

PROJECT THRESHOLD

V7.4

CONTENTS

01 Exercise Overview

01_Executive_Overview	2
01_Expected_Outputs	30
01_How_To_Play	38
01_Practice_Round_Walkthrough	46
01_Timeline	54

02 Facilitator Guide

02_README	64
02_Adjudication_Rules	73
02_Band_Scoring_Quick_Ref	85
02_Base_Case_Fallback_Bank	93
02_Industry_Health_Signal_Tables	102
02_Overview_and_Runbooks	107
02_Plausibility_Decision_Trees	118
02_Quick_Reference_Card	134
02_Scoring_Baselines	144

03 Participant Packets

03_README	169
03_a_Retail	171
03_b_CPG	175
03_c_Healthcare_Provider	180
03_d_Healthcare_Payer	185
03_e_Finance	190
03_f_Consulting	196
03_g_Law	202
03_h_Manufacturing	208
03_i_Logistics	213
03_j_Big_Tech	218
03_k_SaaS	223
03_z_Decision_Worksheet	228
03_z_Rules_of_Play	235
03_z_Template_Industry_Packet	249

04 Participant AI Adoption Arcs

04_README	253
04_a_Retail_Arc	255
04_b_CPG_Arc	259
04_c_Healthcare_Provider_Arc	263
04_d_Healthcare_Payer_Arc	267
04_e_Finance_Arc	271
04_f_Consulting_Arc	275
04_g_Law_Arc	280

04_h_Manufacturing_Arc	284
04_i_Logistics_Arc	288
04_j_Big_Tech_Arc	292
04_k_SaaS_Arc	296

05 Participant Private Cards

05__README	301
05_a_Retail_PC	303
05_b_CPG_PC	306
05_c_Healthcare_Provider_PC	310
05_d_Healthcare_Payer_PC	316
05_e_Finance_PC	322
05_f_Consulting_PC	328
05_g_Law_PC	335
05_h_Manufacturing_PC	342
05_i_Logistics_PC	347
05_j_Big_Tech_PC	352
05_k_SaaS_PC	358
05_z_Template_Private_Card	364

06 Rounds

06__README	370
06_Cross_Round_Integration	377
06_Round_1_Foundation	385
06_Round_2_Acceleration	394
06_Round_3_Reckoning	406
06_Round_4_Normalization	416

07 Supporting Materials

07__README	428
07_Alternative_Scenario_Overrides	435
07_Diagnostic_Survey	442
07_Historical_Reference_Data	450
07_Lead_Facilitator_Briefing	484
07_Post_Exercise_Synthesis_Memo	499
07_Room_Setup_Logistics	507

README

README	520
--------	-----

01

Exercise Overview

Executive Overview

Forecasting the Near-Term Economic Impact of AI Adoption (2026–2030)

Document Date: March 2026 **Version:** 7.4 **Scenario:** Baseline (single format, no scenario variants)
Participant Count: 5–11 industry participants + 1–2 facilitators **Exercise Duration:** 8 hours

Exercise Purpose & Scope

Project Threshold is a scenario-based tabletop exercise designed for senior strategy and investment professionals. Its purpose is to build intuition about how AI adoption will reshape major US industries over the next 2–4 years, and to test realistic scenarios against industry-specific constraints and market dynamics.

This is not a technology demonstration, policy simulation, or evaluation of participant performance. It is an exploration and learning exercise: participants develop grounded intuition about AI diffusion effects, test industry-specific scenarios against operational and financial realities, discover cross-industry interdependencies, and identify strategic actions that make sense across multiple futures.

The exercise is optimized around four outcome variables: realized productivity gains, company valuation impacts, labor displacement and role redesign, and revenue implications (growth, pricing power, new products).

Scenario Overview: 2026–2030 Baseline

Macroeconomic Frame

The US economy enters 2026 with modest GDP growth and unemployment near 4%. The Federal Reserve navigates between inflation persistence and slowing momentum. Enterprise AI capital expenditure is material and accelerating. Enterprise AI adoption is widespread at the pilot level; realized value capture remains uneven.

Over the next four years, three major dynamics unfold:

Dynamic	Timeline	Key Impact
Copilot tools become standard infrastructure	2026–2027	Majority of Fortune 500 deploy AI copilot tools in at least one business function by mid-2027. Integration complexity barriers fall. Competitive pressure drives adoption.

Dynamic	Timeline	Key Impact
Productivity gains materialize unevenly	2027–2029	Gains vary meaningfully by industry. Technology-heavy industries see above-average gains; regulated and data-constrained industries see below-average. Execution risk is highest in regulated industries and where data quality is poor.
Labor displacement concentrates in routine cognitive tasks	2027–2029	Displacement occurs across occupational categories (data entry, customer support, financial analysis, content creation, accounting). Concentrated geographically and demographically, though aggregate impact is moderate by historical standards.

Participants & Individual Model

Participation Model (5–11 industry participants + 1–2 facilitators)

V7.4 moves from 5 sector groupings (each covering 2 sub-industries) to **11 independent industries**, each represented by a single participant. Ideal: 8+ industries assigned to participants. Fewer is possible — the facilitator plays unassigned industries using base case fallback decisions.

- Up to 11 Industry Participants (1 participant per industry)
 - Retail
 - CPG
 - Healthcare Provider
 - Healthcare Payer
 - Finance
 - Consulting
 - Law
 - Manufacturing
 - Logistics
 - Big Tech
 - B2B/B2C SaaS
- Facilitators (1–2 total)
 - Lead Facilitator: 1 person with deep scenario knowledge and decision framework authority
 - Helper (optional): 1 person managing logistics, timing, material distribution, and rule enforcement

Industry Selection (When Running Fewer Than 11)

When running with fewer than 11 participants, select industries to maximize cross-industry dynamics. Recommended minimum configurations:

- 5 participants: Retail, Healthcare Provider, Finance, Manufacturing, Big Tech (broad cross-industry coverage)
- 7 participants: Add Consulting and SaaS for professional services and software depth
- 9 participants: Add CPG, Healthcare Payer, and Logistics for fuller coverage
- 11 participants: Full roster

Participant Profile

The core group is senior consultants, investment professionals, and corporate strategists (Partner/MD level) with industry expertise. Supplement with external expertise: entrepreneurs building AI-native companies, operators actively deploying AI, and national security professionals with visibility into compute/cyber dynamics.

Industry Structure and Profiles (11 Total)

1. Retail

Archetype: Top-5 omnichannel retailer (e.g., Target, Best Buy, Nordstrom).

Revenue: \$80–120B | **Employees:** 300K–400K

Key Functions: Inventory optimization, demand forecasting, in-store operations, e-commerce, customer analytics, pricing and promotions.

AI Dynamics: High volume of structured transaction data. Immediate ROI from demand forecasting and inventory optimization. Customer-facing AI (recommendations, chatbots) is visible but trust-sensitive. Labor displacement concentrated in distribution centers and customer service. Margin pressure from e-commerce competition limits AI investment budgets.

2. CPG

Archetype: Diversified consumer packaged goods manufacturer (e.g., P&G, Colgate, General Mills).

Revenue: \$40–80B | **Employees:** 50K–100K

Key Functions: Brand marketing, supply chain planning, pricing, distribution channel management, consumer insight, R&D.

AI Dynamics: Rich consumer data but fragmented across retail partners. AI-driven consumer insight and demand sensing are high-value applications. Supply chain planning benefits from AI optimization. Marketing mix modeling accelerated by AI. Brand management creates caution about AI-generated content. Distribution channel complexity limits full automation.

3. Healthcare Provider

Archetype: Large integrated health system (20+ facilities).

Revenue: \$15–30B | **Employees:** 60K–100K

Key Functions: Clinical operations, physician workflows, diagnostic imaging, patient record systems, revenue cycle management, surgical planning.

AI Dynamics: Regulatory approval cycles (FDA/CMS) are the binding constraint. Clinical decision support adoption accelerating but requires physician buy-in. Diagnostic imaging AI most mature. Revenue cycle management offers near-term automation ROI. Malpractice liability concerns slow deployment. Data quality and EHR interoperability remain barriers.

4. Healthcare Payer

Archetype: National health insurer (e.g., UnitedHealth, Anthem, Cigna scale).

Revenue: \$150–300B | **Employees:** 50K–80K

Key Functions: Claims processing, prior authorization, coverage decisions, network management, actuarial modeling, fraud detection, member engagement.

AI Dynamics: Claims processing automation offers massive volume ROI. Prior authorization AI creates Provider-Payer tension (faster denials vs. faster approvals). Actuarial modeling enhanced by AI but regulatory scrutiny on algorithmic bias is material. Network management and fraud detection are high-value AI applications. Member engagement AI (chatbots, care navigation) is growing.

5. Finance

Archetype: Diversified financial institution (e.g., JPMorgan, Goldman Sachs, Citi scale).

Revenue: \$50–150B | **Employees:** 100K–250K

Key Functions: Retail banking, wealth management, commercial lending, claims underwriting, fraud detection, compliance, regulatory reporting, trading operations.

AI Dynamics: Clear ROI on trading, fraud detection, and compliance automation. Regulatory scrutiny on algorithmic decision-making (fairness, bias audits) is material. Wealth management AI (portfolio optimization, client engagement) growing rapidly. Back-office automation (KYC, AML, reporting) is table-stakes. Talent competition with Big Tech for AI specialists is intense.

6. Consulting

Archetype: Big Four / MBB-class firm (e.g., McKinsey, BCG, Deloitte, Accenture scale).

Revenue: \$20–25B | **Employees:** 30K–50K

Key Functions: Management consulting, strategy advisory, digital transformation, AI transformation advisory, implementation services, talent management, knowledge management.

AI Dynamics: Dual exposure — both deploying AI internally to accelerate delivery and advising clients on AI strategy. Junior talent pipeline disruption is acute: AI copilots can perform research, analysis, and slide generation that previously required associate-level staff. Pricing pressure as clients expect AI-augmented delivery at lower rates. Knowledge management and institutional learning accelerated by AI but creates IP protection concerns. Utilization models under pressure as project timelines compress. Competitive threat from AI-native advisory startups and from Big Tech advisory arms.

Key Tensions: Revenue per consultant declining as AI compresses project timelines. Clients demand cost savings from AI-augmented delivery. "Grow your own talent" model breaks down when junior roles are automated. Shift to outcome-based pricing creates margin uncertainty. Proprietary methodologies at risk of commoditization by AI tools.

7. Law

Archetype: AmLaw 50 firm (e.g., Kirkland & Ellis, Latham & Watkins, Skadden scale).

Revenue: \$3–6B | **Employees:** 3K–8K (attorneys + staff)

Key Functions: Corporate transactions, litigation, regulatory compliance, contract review, legal research, due diligence, intellectual property, dispute resolution.

AI Dynamics: Billable hour model faces existential pressure from AI-driven efficiency. Contract review and due diligence are high-volume, high-value automation targets. Legal research acceleration is material — tasks that took associates days now take hours. Bar association rules on AI use vary by jurisdiction, creating compliance complexity. Malpractice liability for AI-assisted legal work is unresolved. Associate leverage model eroding as AI handles work previously requiring junior associates.

Key Tensions: Billable hour disruption — clients will not pay human rates for AI-generated work product. Bar rule uncertainty across 50+ jurisdictions creates compliance overhead. Malpractice liability for AI-assisted advice is legally untested. Associate pipeline disruption mirrors consulting — fewer entry-level positions needed. Legal AI platforms (Harvey, CoCounsel) threaten to disintermediate law firms on routine work. IP and confidentiality risks from training models on client data.

8. Manufacturing

Archetype: Multi-plant heavy manufacturer (e.g., Caterpillar, 3M, Honeywell scale).

Revenue: \$30–60B | **Employees:** 50K–100K

Key Functions: Production planning, quality control, plant operations, supply chain optimization, predictive maintenance, R&D, safety management.

AI Dynamics: Predictive maintenance and quality control are mature AI applications with clear ROI. Production planning optimization benefits from AI but requires integration with legacy OT systems. Digital twin technology accelerating but capital-intensive. Workforce is predominantly blue-collar; AI displacement concentrated in supervisory and planning roles. Safety-critical environments require high reliability and regulatory compliance.

9. Logistics

Archetype: National freight and 3PL operator (e.g., J.B. Hunt, XPO, DHL Supply Chain scale).

Revenue: \$15–30B | **Employees:** 40K–80K

Key Functions: Route optimization, fleet management, warehouse operations, demand forecasting, last-mile delivery, carrier management, freight brokerage.

AI Dynamics: Route optimization and demand forecasting are high-value, high-maturity AI applications. Warehouse automation (robotics + AI) accelerating. Fleet management AI (predictive maintenance, driver assistance) growing. Last-mile delivery optimization is competitive battleground. Labor displacement concentrated in warehouse operations and dispatch. Driver shortage creates opportunity for AI-assisted driving tools.

10. Big Tech

Archetype: Platform company operating across cloud, advertising, devices, and enterprise software (e.g., Google, Meta, Microsoft, Amazon scale).

Revenue: \$150–400B | **Employees:** 100K–200K

Key Functions: Cloud infrastructure, advertising platforms, enterprise software, consumer devices, content platforms, developer ecosystems.

Scope Note: Big Tech in Project Threshold V7.4 excludes AI lab / foundation model development. The focus is on how platform companies deploy AI across their existing business lines (cloud, ads, devices, enterprise software) and how they serve as infrastructure providers to other industries. This scoping keeps the exercise grounded in business strategy rather than model capability speculation.

AI Dynamics: AI capex spend is massive and accelerating. Cloud AI services are major revenue growth driver. Advertising optimization via AI is table-stakes. Enterprise software AI features (copilots, automation) drive upselling. Regulatory scrutiny on market concentration, data practices, and platform power is material. Talent competition is fierce. Platform lock-in effects amplified by AI integration.

11. B2B/B2C SaaS

Archetype: Enterprise or consumer software company (e.g., Salesforce, Workday, SAP, ServiceNow scale).

Revenue: \$10–35B | **Employees:** 20K–80K

Key Functions: Enterprise software (CRM, HCM, ERP, ITSM), consumer applications, subscription management, platform ecosystems, developer tools.

AI Dynamics: AI feature integration (copilots embedded in product) is competitive necessity. Pricing pressure as customers expect AI features included in existing subscription fees. Competitive threat from AI-native startups that build AI-first products. Margin pressure from AI infrastructure costs (compute, model hosting). M&A activity for AI capabilities is intense. Partner ecosystem dynamics shifting as AI changes integration patterns.

Exercise Format: 8-Hour Single Session

Structure & Pacing

Phase	Clock Time	Duration	Content
Setup & Briefing	8:30–9:00	0:30	Facilitator introduces scenario, rules, scoring framework; seating; material distribution
Practice Micro-Round	9:00–9:10	0:10	Individual makes one practice decision; facilitator calibrates expectations; clarifies decision worksheet
Round 1: Foundation	9:10–10:15	1:05	2 core injects; decisions for industry; private info card reveal; facilitator establishes rhythm
Break	10:15–10:30	0:15	Coffee, restrooms, reset
Round 2: Acceleration	10:30–11:35	1:05	2–3 injects; industry decisions; Facilitator Market Shock (constraints on 2–3 industries); cross-industry spillover discussion with Collective Bonus (Optional)
Lunch Break	11:35–12:15	0:40	Off-site or catered
Round 3: Reckoning	12:15–1:20	1:05	2–3 core injects; final private info card reveal; Industry Health Signals update; cross-industry discussion with Collective Bonus (Optional)
Break	1:20–1:35	0:15	Reset, facilitator scoring prep

Phase	Clock Time	Duration	Content
Round 4: Normalization	1:35–2:40	1:05	2 injects (no private card); final positioning decisions; final Industry Health Signals; cross-industry discussion with Collective Bonus (Optional)
Debrief & Synthesis	2:50–3:50	1:00	Facilitator-led retrospective; participant reflections; strategic insights

Total: 8 hours (8:30 AM – 4:30 PM with scheduled breaks)

Round Structure Summary

Round 1: Foundation (Q2 2026)

Narrative: Early AI adoption momentum. Copilot tools entering enterprise workflows. First signs of competitive pressure.

Decision Theme: Initial positioning — are you leading, following, or waiting?

Key Triggers: First private information card reveals industry-specific early signals (success stories, pilot results, regulatory headwinds). Private cards are distributed from the separate 05_Participant_Private_Cards/ folder.

Timing Breakdown:

- Situation Update: 5 min
- Card/Inject Distribution: 3 min
- Individual Decision Preparation: 15 min
- Decision Submission: 3 min
- Facilitator Scoring: 12 min
- Cross-Industry Discussion & Results (includes Collective Bonus if used): 27 min

Round 2: Acceleration (Q4 2026)

Narrative: Adoption accelerates faster than expected. Integration complexity falling. Early productivity gains visible in some industries.

Decision Theme: Doubling down or hedging — how aggressive do you get?

Key Triggers: Facilitator Market Shock imposes constraints on 2–3 industries; Industry Health Signals announced; Collective Bonus (Optional) during discussion; cross-industry spillovers emerge (e.g., tech salary inflation affects Finance hiring; consulting pricing pressure from AI-augmented delivery).

Round 3: Reckoning (Q2 2027)

Narrative: Productivity gains materialize but are uneven. Labor displacement becomes visible. Regulatory uncertainty rises.

Decision Theme: Defending against disruption or embracing it — high-stakes strategic choice.

Key Triggers: Second private card reveals industry-specific vulnerability signals; red-flag plausibility triggers force facilitator review; Collective Bonus (Optional) during discussion; Industry Health Signals shift.

Round 4: Normalization (Q4 2027)

Narrative: Market repricing based on realized outcomes. Winners and losers clear. Regulatory frameworks begin to crystallize.

Decision Theme: Final positioning — locking in gains or waiting for volatility?

Key Triggers: No private card (endgame only public information); final Collective Bonus (Optional) during discussion; final Industry Health Signals announced.

Pacing Table: Detailed Round Timing

Round Duration: 65 Minutes (All Rounds)

Phase	Duration	Lead	Content
Situation Update	5 min	Facilitator	Recap previous round outcomes; set narrative context for new round; read macro inject aloud
Card/Inject Distribution	3 min	Helper	Distribute private information cards (from 05_Participant_Private_Cards/) and industry-specific injects; participants read silently

Phase	Duration	Lead	Content
Individual Decision Prep	15 min	Participant	Individual reads inject, reviews industry constraints and baseline position, completes decision worksheet, identifies key trade-offs and cross-industry dependencies. AI Adoption Arc cards (from 04_Participant_AI_Adoption_Arcs/) provide round-by-round context.
Decision Submission	3 min	Participant	Participants submit decision worksheets (physically or digitally); Helper collects and organizes by industry
Facilitator Scoring	12 min	Facilitator	Score decisions in real-time using Quick Reference Card; note red flags and cross-industry spillovers; prepare feedback

Phase	Duration	Lead	Content
Cross-Industry Discussion & Results	27 min	Facilitator	<p>Results (5 min): Announce scores and Industry Health Signal updates. Collective Bonus (~5 min, optional): Any participant may nominate an industry for a +2 bonus (strong strategy) or -2 penalty (risky strategy); nomination takes effect only if 3+ participants agree. Cannot nominate own industry. Multiple nominations may occur. Discussion (~17 min): Facilitator leads open discussion of decisions, cross-industry spillovers, execution risks, revised forecasts. Participants challenge each other's logic. Identify emerging patterns and dependencies.</p>

Note: Cross-industry discussion is the centerpiece of each round. This is where participants build intuition about AI diffusion effects and surface interdependencies that aren't obvious in isolation.

Scoring Overview: Banded Framework

Three Dimensions (Per Decision)

Dimension	Definition	Scores
Strategic Fit	Does this decision align with industry fundamentals and macro trends? Is the logic sound?	{-2, 0, +2} with +/-3 red flags
Execution Risk	Can the organization realistically execute this in the timeframe? Do they have the talent, capital, and operational capability?	{-2, 0, +2} with +/-3 red flags

Dimension	Definition	Scores
Tail Risk	What's the catastrophic downside? Does this decision create asymmetric long-term liability?	{-2, 0, +2} with +/-3 red flags

Scoring Rules

Score	Label	Meaning
0	Hold	Decision is defensible but unremarkable. Industry is hedging or staying neutral.
+2	Win	Strong logic; good execution odds; acceptable tail risk for the upside.
-2	Lose	Logic is weak, execution risk is high, or tail risk is outsized relative to upside.
+/-3	Red Flag	Reserved for decisions that breach explicit plausibility thresholds or trigger cross-industry spillovers requiring facilitator review.

Red-Flag Triggers (Automatic +/-3 Review)

A decision scores +/-3 if it:

- Claims productivity gains >3x the industry baseline without credible execution path
- Involves full headcount reduction (>40%) in a year without regulatory cover
- Triggers cross-industry spillovers (e.g., Finance talent raid affecting Consulting; law firm AI tool creates malpractice precedent)
- Explicitly violates regulatory constraints (stated in baseline scenario)
- Is logically inconsistent with previous decisions in the same round

Facilitator Scoring Process

1. Immediate reaction: Is this decision plausible? Does it make sense given the industry's constraints and inject?
2. Dimension scoring: {-2, 0, +2} for each of Strategic Fit, Execution Risk, Tail Risk
3. Red-flag check: Does this trigger +/-3? If yes, pause for cross-industry review before finalizing
4. Total score: Sum of three dimensions = -6 to +6 per decision (in practice, most decisions fall in -4 to +4)
5. Feedback to participant: 1–2 sentence explanation of scoring rationale; note any dependencies or cross-industry impacts

Scoring is completed in real-time (3–5 min per decision) using the Quick Reference Card.

Market Pressure Mechanics: Collective Bonus & Facilitator Market Shock

V7.4 replaces the Investment Committee role with two lighter-weight mechanics that distribute market pressure across all participants and the facilitator.

Collective Bonus (Optional, Every Round, R1–R4)

During each cross-industry discussion period (~5 min within the 27-min discussion block):

Element	Detail
Timing	Conducted after results are announced, before open discussion
Mechanic	Any participant may nominate an industry for a +2 bonus (strong strategy) or -2 penalty (risky strategy). Nomination takes effect only if 3+ participants agree.
Constraint	Cannot nominate own industry. Multiple nominations may occur in a single round.
Score impact	+2 cumulative score bonus per approved positive nomination; -2 cumulative score penalty per approved negative nomination
Threshold	3+ participants must agree for any nomination to take effect. If fewer than 3 agree, no adjustment is made.
Purpose	Creates optional peer accountability and surfaces group consensus on strategy quality without mandatory voting or a dedicated IC role

Facilitator Market Shock (Round 2 Only)

Element	Detail
Timing	Announced at the start of Round 2, before individual decision prep (~3 min)
Mechanic	Facilitator selects 2–3 industries and imposes one constraint each from the constraint menu

Element	Detail
Constraint menu	Same 6-item menu as Industry Health Signals: Regulatory Halt, Labor Cost Surge, Capital Tightening, Reputational Pressure, Competitive Response, Litigation Risk
Scope	Can hit ANY industry — not limited to those in Headwind/Crisis condition
Purpose	Injects external market pressure and forces adaptive decision-making mid-exercise

Learning Objectives: Building Intuition and Testing Scenarios

Objective	Description
Build Intuition for AI Diffusion Effects	Develop grounded understanding of how AI adoption will reshape different industries of the US economy over 2–4 years. Move beyond abstract narratives to concrete industry-level effects: which functions are most disrupted, what execution risks are real, where does value accrue, where does margin compress.
Test Potential Scenarios	Stress-test industry-specific AI adoption scenarios against realistic constraints: regulatory approval timelines, talent availability, capital requirements, operational integration complexity. Identify which adoption paths are plausible and which are fantasy.
Surface Cross-Industry Dynamics	Illuminate interdependencies between industries that may not be obvious without structured simulation. How does Big Tech AI spending affect talent costs in Finance? How do Healthcare regulatory delays affect Consulting advisory revenue? How does Law firm AI adoption create malpractice precedent? How does Consulting pricing pressure feed back into client AI investment decisions?
Identify No-Regrets Actions	Surface strategic actions and investments that make sense regardless of which AI adoption scenario unfolds. These are the actions that reduce regret across multiple futures (not just the base case).

Objective	Description
Ground Abstract Forecasts in Concrete Decisions	Force participants to translate broad AI narratives into specific, constrained industry-level choices. When facing capital constraints, talent constraints, and regulatory risk, which AI investments actually happen? What is deferred? What is abandoned?

Expected Outputs

See **01ExpectedOutputs.md** for debrief template, industry-specific scoring examples, and key questions to guide synthesis.

The exercise produces:

Output	Description
Decision Log	Decisions made by each industry participant (variable based on industries played: e.g., 11 participants x 4 rounds = 44 decisions, or 5 participants x 4 rounds = 20 decisions, scored)
Collective Bonus Results	Per-round record of nominations and outcomes showing which industries received +2 (strong) or -2 (risky) adjustments via participant consensus
Cross-Industry Spillovers	Record of dependencies and knock-on effects identified during play
Revised Forecasts	Participant views on productivity, valuation, and labor displacement post-exercise (vs. baseline scenario)
Strategic Insights	Facilitator synthesis memo (40 min, written immediately post-exercise)

Key Design Principles

#	Principle	Description
1	Operational Clarity	Exercise runs exactly on schedule. All participants know where they stand (decision windows, scoring feedback, round pacing) at all times. No ambiguity on rules or facilitator interpretation.
2	Novice-Accessible	Participants with no prior tabletop experience can engage meaningfully. Decision worksheet guides logic step-by-step. Scoring feedback is immediate and transparent (not mysterious).
3	Realistic Without Overwhelming	Injects are grounded in plausible Q2–Q4 2026 events (not speculative 2030 fantasy). Industry constraints (regulatory, operational, financial) are explicit. Participants face real trade-offs but not paralyzing complexity.
4	Market Discipline	Collective Bonus (Optional) creates participant-driven market pressure each round (+2/-2 cumulative score adjustments when 3+ participants agree). Facilitator Market Shock imposes constraints on 2–3 industries in Round 2. Industry Health Signals provide mechanical consequences based on cumulative scores. Creates healthy tension and prevents groupthink.
5	Fast Synthesis	60-minute debrief using 01ExpectedOutputs.md template allows facilitator to identify strategic insights in real-time, enabling remote teams to debrief synchronously.

#	Principle	Description
6	Banded Scoring Clarity	{-2, 0, +2} framework is faster and clearer than continuous (e.g., -5 to +5) scoring. Red-flag system catches edge cases without requiring facilitator interpretation or debate.
7	Learning Focus	Exercise is optimized for participant learning and scenario exploration, not for evaluating or ranking participants. This is a discovery tool.

Core Forecast Claims: Baseline Scenario

The baseline scenario underlying Project Threshold V7.4 is grounded in eight quantified forecast claims. All industry participants make decisions based on these baseline assumptions.

Baseline Claims Summary:

#	Claim
1	AI copilot tools become standard enterprise infrastructure by mid-2027
2	Realized productivity gains vary by industry (2026–2029)
3	Agentic workflows emerge in limited production (2027–2028)
4	Software engineering acceleration creates complex market impact
5	Labor displacement concentrated in routine cognitive tasks (2024–2029)
6	Valuation multiples diverge by industry (2026–2030)
7	Regulatory constraints bite in 2027–2028
8	Talent dynamics shift toward AI-complementary skills (2026–2029)

Claim 1: AI Copilot Tools Become Standard Enterprise Infrastructure by

Mid-2027

Point Estimate: 65–70% of Fortune 500 companies deploy AI copilot tools in at least one business function by Q2 2027.

Confidence: HIGH (>75%)

Drivers:

Driver	Metric
Inference cost reduction	45–55% annual decrease in cost per token (Q2 2026 -> Q2 2027)
Integration maturity	Time-to-deployment falls from 6–12 months -> 2–3 months
Competitive pressure	40% of Fortune 500 per industry have competitor with deployed copilots by Q3 2026
Vendor saturation	15–20 serious enterprise AI platform competitors by Q2 2027

Claim 2: Realized Productivity Gains by Industry (2026–2029)

Summary: Productivity gains vary meaningfully by industry. Technology-friendly industries achieve above-average gains; constrained industries achieve below-average gains. Gains are highly dispersed by execution capability and industry maturity.

Industry Estimates (Productivity Gains 2026–2029, Relative to Baseline):

Industry	Gain Level	Confidence	Key Drivers	Time to Material Realization
Big Tech	Above Average	HIGH	AI infrastructure investment, competitive moat building, code generation acceleration	12–18 months
B2B/B2C SaaS	Above Average	HIGH	AI feature integration, product acceleration, developer productivity	12–18 months

Industry	Gain Level	Confidence	Key Drivers	Time to Material Realization
Logistics	Above Average	MEDIUM-HIGH	Route optimization, demand forecasting, driver assistant adoption	12–18 months
Finance	Above Average	MEDIUM-HIGH	Back-office automation, front-office AI adoption, fraud detection	12–18 months
Consulting	Above Average	MEDIUM-HIGH	Research acceleration, analysis automation, knowledge management	18–24 months
Healthcare Payer	Above Average	MEDIUM-HIGH	Claims processing automation, fraud detection, actuarial modeling	18–24 months
Law	Moderate-High	MEDIUM	Contract review, legal research automation, but bar rule uncertainty	18–24 months
Healthcare Provider	Moderate	MEDIUM	Clinical decision support adoption offset by regulatory constraints	18–30 months
Manufacturing	Moderate	MEDIUM	Predictive maintenance, quality control, but legacy OT integration complexity	18–30 months
Retail	Below Average	MEDIUM	Inventory optimization, demand forecasting, but offset by cannibalization	24–36 months

Industry	Gain Level	Confidence	Key Drivers	Time to Material Realization
CPG	Below Average	MEDIUM	Consumer insight, supply chain planning, but data fragmentation limits	24–36 months

Claim 3: Agentic Workflows Emerge in Limited Production (2027–2028)

Point Estimate: By end of 2028, a modest percentage of eligible processes will have agentic workflows in limited production. Error rates require human review for material percentage of transactions.

Confidence: MEDIUM (50–60%)

Drivers:

- Agent planning quality improves materially by Q4 2027 (from earlier levels)
- Enterprise error tolerance varies by industry: high-ROI processes tolerate moderate error; low-stakes tolerate higher error; regulated industries require near-zero error
- Insurance & liability: Specialized AI insurers emerge; premium is material cost element

Claim 4: Software Engineering Acceleration (Complex Market Impact)

Point Estimate:

- Code generation achieves material coverage by Q4 2027
- Developer feature velocity increases meaningfully
- But: Software industry margin compression occurs due to increased competition
- Net: Revenue growth offset by margin pressure; EBITDA growth modest

Confidence: MEDIUM-HIGH (60–75%)

Note: This claim affects Big Tech and SaaS directly (core product and cost structure). It also affects Consulting (software engineering advisory), Finance (software engineering labor costs), Law (IP and licensing complexity), and Manufacturing (software-driven optimization investments).

Claim 5: Labor Displacement Concentrated in Routine Tasks (2024–2029)

Point Estimate: Job losses occur across routine occupational categories. Unemployment rate increases modestly (0.4–0.8 percentage points). Peak impact occurs in 2027–2028.

Confidence: MEDIUM (50–60%)

Occupational Breakdown:

Occupation	Impact Level
Data Entry & Transcription	High
Tier-1 Customer Support	Moderate
Routine Financial Analysis	Moderate
Routine Content Creation	Moderate
Back-Office / Accounting / Basic Legal	Moderate
Junior Consulting Research & Analysis	Moderate
Contract Review & Due Diligence (Law)	Moderate-High

Job Creation Offset: New roles emerge (AI trainers, governance specialists, upskilling instructors) offset a portion of losses. Net job displacement is moderate relative to historical disruptions, but concentrated geographically and demographically.

Claim 6: Valuation Multiples Diverge by Industry (2026–2030)

Baseline Multiples (Q2 2026):

Industry	2026 Position
Big Tech	High multiple range
B2B/B2C SaaS	Elevated multiple range
Healthcare Payer	Elevated multiple range
Healthcare Provider	Moderate to elevated multiple
Finance	Moderate multiple range
Consulting	Elevated multiple (premium for advisory)
Law	N/A (partnerships, not public equity)
CPG	Moderate to elevated multiple
Retail	Modest multiple
Manufacturing	Moderate to elevated multiple
Logistics	Moderate multiple

Valuation Scenarios by Round:

Round	Valuation Behavior
Round 1–2	Multiples remain stable (early signals, not yet priced)
Round 3	Divergence begins — execution leaders see modest multiple expansion; laggards see modest compression
Round 4	Full repricing — leaders see material expansion; laggards see material compression based on execution quality

Claim 7: Regulatory Constraints Bite in 2027–2028

Confidence: MEDIUM-HIGH

Industry-Specific Constraints:

Industry	Regulatory Constraints
Healthcare Provider	FDA/CMS approval cycles for clinical decision support; liability concerns limit autonomous deployment
Healthcare Payer	Regulatory scrutiny on algorithmic coverage decisions and prior authorization automation; state insurance commission oversight
Finance	Regulatory scrutiny on algorithmic decision-making; fairness and bias audit requirements may slow deployment
Consulting	Client data handling regulations; professional licensing requirements; cross-border data transfer constraints
Law	State bar associations issue guidance on AI use; compliance overhead material element of implementation cost; unauthorized practice of law concerns for AI-generated legal advice
Big Tech	Antitrust scrutiny on model market concentration; data privacy regulations affecting training data; content moderation liability
B2B/B2C SaaS	Data privacy requirements; contractual liability for AI feature accuracy; enterprise procurement requirements

Claim 8: Talent Dynamics Shift (2026–2029)

Confidence: MEDIUM-HIGH

Dynamics:

Talent Segment	Direction	Key Details
AI specialist demand	Increasing	Supply is constrained; salaries increase materially
Junior talent (Consulting, Law, Finance)	Declining	AI copilots perform research, analysis, and document generation that previously required associate-level staff
Junior engineer demand	Declining	Code generation reduces entry-level hiring; particularly acute in Big Tech and SaaS
Mid/senior professional demand	Increasing	AI oversight and strategy roles expand
Cross-industry talent wars	Intensifying	Finance, Consulting, Manufacturing, and Big Tech compete for AI talent; salary inflation materializes across industries
Healthcare talent	Mixed	Clinical AI expertise in high demand; physician adoption training needed; regulatory compliance expertise becomes premium skill
Law talent	Mixed	Senior partners directing AI tools become more valuable; junior associate demand drops; legal AI platform expertise emerges as new skill category

Baseline Scenario Assumptions (Exercise Rules)

All industry participants operate under these constraints:

#	Constraint	Rule
1	Regulatory baseline	Constraints listed in Claim 7 are in effect; participants cannot assume faster regulatory change
2	Talent constraints	Labor market is tight; salary inflation 5–10% annualized through 2029
3	Capital constraints	Participants cannot assume unlimited CapEx; each round decision competes for capital
4	Macro baseline	No major recession (2.0–2.5% GDP growth); no significant geopolitical shock
5	Technology baseline	No breakthrough in AI reasoning (still within LLM/generative model class); no quantum computing impact

Operational Definitions (For Exercise Precision)

Term	Definition
Deployed	An AI system is "deployed" when in active daily production use by $\geq 50\%$ of eligible users. Pilots and licensed-but-unused tools do not count.
Realized Productivity	Measured as change in revenue-per-employee or output-per-hour from deployment date using 12-month trailing average. Excludes one-time restructuring gains.
Agentic	An AI system is agentic when autonomously executing multi-step tasks with $\leq 5\%$ human intervention rate across $\geq 1,000$ transactions per month.
Labor Displacement	A job is "displaced" when the role is eliminated or reduced by $\geq 50\%$ of hours AND the primary driver is AI automation (not offshoring or demand decline).
AI Adoption Rate	Percentage of firms in an industry with at least one AI system deployed in a revenue-generating or cost-reducing workflow.

What Project Threshold V7.4 Does NOT Cover

#	Exclusion	Rationale
1	Scenario variants	No "Slow," "Fast," or Type A/Type B scenarios (V3 feature). All exercises use the Baseline scenario.
2	International markets	Exercise is US-centric.
3	Generative AI regulation	No assumption of breakthrough regulation; regulatory baseline is conservative.
4	Multi-year forecasting beyond 2030	Exercise ends Q4 2027 (game time), representing Q4 2027 market snapshot only.
5	AI lab / foundation model development	Big Tech profile excludes model R&D; focus is on business application.

Next Steps for Participants

#	Action	Time	Details
1	Read this document fully	30 min	Understand the scenario, baseline claims, and industry context
2	Read your industry packet (03a–03k from 03ParticipantPackets/)	20 min	Understand constraints, starting position, decision prep checklist
3	Read Rules of Play (03zRulesofPlay.md)	10 min	Understand exercise mechanics, decision submission, revision windows

#	Action	Time	Details
4	Review your AI Adoption Arc (from 04ParticipantAIAdoptionArcs/)	10 min	Understand how AI adoption unfolds round-by-round for your industry
5	Prepare decision framework (before exercise day)	Variable	For each round, identify your industry's key trade-offs (productivity vs. labor, growth vs. execution risk, etc.)
6	Arrive 30 min early	—	Day-of briefing and setup

Historical Reference Data

The following tables provide actual historical data (2015–2025) for key macroeconomic and market metrics cited in this scenario. These benchmarks provide participants with real context for evaluating the forward-looking projections in the exercise baseline scenario.

US GDP Growth (Annual %, 2015–2025)

Year	Growth Rate
2015	2.9%
2016	1.7%
2017	2.3%
2018	3.0%
2019	2.3%
2020	-3.4%
2021	5.9%
2022	2.1%
2023	2.5%
2024	2.7%
2025	2.3% (approx.)

US Unemployment Rate (Annual Average %, 2015–2025)

Year	Unemployment Rate
2015	5.3%
2016	4.9%
2017	4.4%
2018	3.9%
2019	3.7%
2020	8.1%
2021	5.4%
2022	3.6%
2023	3.6%
2024	4.0%
2025	4.1% (approx.)

S&P 500 Annual Returns (% , 2015–2025)

Year	Total Return
2015	1.4%
2016	12.0%
2017	21.8%
2018	-6.2%
2019	31.5%
2020	18.4%
2021	28.7%
2022	-18.1%
2023	24.2%
2024	24.0% (approx.)
2025	15.0% (approx. YTD)

US Labor Productivity Growth (Nonfarm Business, Annual %, 2015–2025)

Year	Growth Rate
2015	1.5%
2016	1.1%
2017	1.3%
2018	1.7%
2019	1.6%
2020	3.5%
2021	4.8%
2022	1.8%
2023	1.4%
2024	1.0%
2025	1.2% (approx.)

Enterprise IT/AI Spending (Estimated \$B, 2015–2025)

Year	IT Spending	AI Spending	AI as % of Total
2015	2,200	15	0.7%
2016	2,300	23	1.0%
2017	2,450	35	1.4%
2018	2,650	55	2.1%
2019	2,900	85	2.9%
2020	3,200	140	4.4%
2021	3,650	220	6.0%
2022	4,100	380	9.3%
2023	4,450	680	15.3%
2024	4,800	1,200	25.0%
2025	5,100	1,850	36.3%

Sources: Federal Reserve Economic Data (FRED), Bureau of Labor Statistics (BLS), Bureau of Economic Analysis (BEA), S&P Global Market Intelligence, Statista. Some figures for 2024–2025 are approximations based on available quarterly data and trend analysis.

Expected Outputs

Objective: Define the specific deliverables that constitute a successful exercise run.

When Project Threshold V7.4 concludes at 3:20 PM, the following outputs should exist or be in flight. Each is measurable and tied to learning outcomes.

REQUIRED OUTPUT #1: Per-Industry Decision Log (4 Rounds) — Individual Model (V7.4)

What: Complete record of each industry representative's decisions across all four rounds, using V7.4 individual industry model.

Format: Filled-in decision submission worksheets. The number of worksheets depends on the number of industries in play:

- 5 industries x 4 rounds = 20 worksheets total
- 7 industries x 4 rounds = 28 worksheets total
- 11 industries x 4 rounds = 44 worksheets total

Plus base case fallback records for any industries without explicit decisions.

Content per submitted decision:

Field	What to Record
Industry Name	Which industry is deciding (Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS)
Decision Title	Explicit, actionable title (e.g., "Deploy AI demand forecasting to 500 stores")
Spend/Commitment	Absorbable / Material but manageable / Transformational bet / Potentially existential
Time-to-Impact	0–3 mo / 3–12 mo / 1–2 yr / 2+ yr
Execution Complexity	Low / Medium / High / Very high
Dependency	Mostly internal / Requires vendor-partner / Requires regulator-union-standards / Requires ecosystem shift
Scale	Pilot / Regional / National / Global
Success Metric	Directional or range-based measurement (e.g., "10–20% margin improvement")
Risk & Mitigation	Top 2 risks and mitigations

Field	What to Record
Cross-Industry Synergies	Coordination with other industries, or "Independent decision"
Private Card Influence	How private card shaped thinking (Rounds 1–3 only)
Continuity Note	How this builds on prior-round decision (Rounds 2–4 only)
Fallback Note	If no explicit decision submitted, record that industry received automatic base case fallback (see 02BaseCaseFallbackBank)

Quality Bar: An external reader (someone not in the room) should be able to:

Test	What It Means
Clarity	Understand exactly what was decided and by which industry
Reasoning	Explain the participant's logic using the banded framework
Risk Awareness	Identify the primary risk concern and mitigation
Context	Map the decision to the round's scenario context
Fallback Coverage	Understand why inactive industries received base case fallback

Example of Strong Decision Entry (V7.4 Individual Model):

Round 2 — Finance Decision Title: Increase AI-Assisted Trading Oversight with Real-Time Risk Controls **Banded Assessment:**

- Spend/Commitment: Material but manageable (~\$20M capex)
- Time-to-Impact: 3–12 months (controls live in 6 months)
- Execution Complexity: High (regulatory approval required, system integration complex)
- Dependency: Requires regulator/union/standards (OCC approval, model risk framework alignment)
- Scale: National (enterprise-wide trading oversight)
- Success Metric: High impact (reduce unauthorized exposure events to zero)

Risk Acknowledged: Regulator may demand additional controls beyond budgeted. Mitigation: Engage OCC early; establish feedback loop for mid-course corrections. Operational disruption if system fails. Mitigation: Staged rollout with 48-hour rollback procedure.

Private Card Influence: Card stated "trading oversight becoming competitive differentiator." This moved us from defensive to proactive posture.

Dependencies: OCC formal approval letter; legacy system compatibility; Model Risk team staffing.

REQUIRED OUTPUT #2: Collective Bonus Record (4

Rounds)

What: Complete record of Collective Bonus nominations and outcomes across all four rounds.

Must Include:

Round	Collective Bonus Data	Required Deliverable
Round 1	Initial peer assessment	Record of nominations (+2 strong / -2 risky), whether 3+ participants agreed, net result applied to cumulative score; brief facilitator note on stated rationale. If no nominations occurred, note "No Collective Bonus this round."
Round 2	Execution-informed assessment	Same format. Note how nominations shifted from R1 based on execution against R1 commitments.
Round 3	Track-record assessment	Same format. Note cumulative patterns and whether participants are responding to peer signals.
Round 4	Final strategic assessment	Same format. Note which industries received sustained positive or negative recognition across rounds.

Quality Bar: A senior executive reviewing the Collective Bonus record should be able to:

- See how group consensus on strategy quality evolved across rounds
- Identify which industries earned sustained peer confidence and which lost it
- Trace whether bonus feedback influenced subsequent strategy pivots
- Understand the connection between peer accountability and strategic discipline

REQUIRED OUTPUT #3: One-Page Synthesis Memo Per Industry

What: Each industry representative writes a 1-page memo capturing three elements.

Format: Use template **07PostExerciseSynthesisMemo.md** (to be distributed at end of exercise).

Three Required Sections:

1. Surprises (What we didn't expect):
 - What scenario twist or Collective Bonus result caught you off-guard?
 - What assumption proved wrong?

- Example: "We assumed our aggressive automation strategy would earn peer support; instead, three participants agreed on a -2 Collective Bonus citing labor risk. We scrambled to add governance in Round 2."
2. What We Got Wrong (Lessons from deliberation):
 - Which decision would you reverse knowing what you know now?
 - Which risk did you underestimate or overestimate?
 - Example: "We overestimated the speed of regulatory approval. Next time, we pad timelines by 30%."
 3. No-Regrets Actions (Do regardless of scenario):
 - What actions would you take the same way in any future scenario?
 - What's non-negotiable in your industry's strategy?
 - Example: "Invest in workforce training for AI upskilling. Every scenario tests talent; this is core."

Quality Bar: A senior executive in your industry should be able to read this memo and extract 2–3 actionable insights for their own organization immediately.

Format constraints:

- 1 page (single-spaced, 1" margins, 11pt font)
- Bullet points or short paragraphs
- No jargon; plain language
- Completed during the 10 minutes before closing (3:20 PM) or submitted within 24 hours post-exercise

REQUIRED OUTPUT #4: Consolidated No-Regrets Actions

What: Cross-industry synthesis of actions that all industries identified as essential, regardless of scenario outcome.

Must Include:

- Actions organized by governance, workforce/talent, compliance/legal, technology infrastructure, and other (if applicable)
- At least 5–7 concrete actions
- Owner assignment (suggested owner, or all-industries responsibility)
- Timeline (immediate, 30 days, 90 days, 6 months)
- Success measure (how you'll know it worked)

Format: Single shared document, produced collaboratively during **no-regrets discussion block (3:00–3:20 PM)**.

Example of Strong No-Regrets Action:

Category	Action	Owner	Timeline	Success Measure
Workforce	Launch cross-industry AI upskilling program for all roles (min. 20 hrs/person/year)	CHRO + Tech Lead	90 days	80% enrollment by end of Q2; internal mentorship pairings established
Governance	Establish scenario planning committee (meet quarterly minimum)	Chief Strategy Officer	30 days	Committee chartered; first scenario cycle planned for Q3
Compliance	Audit regulatory risk exposure; map key dependencies with regulators	Chief Compliance Officer	60 days	Risk map published; 3+ regulator conversations completed
Technology	Modernize data infrastructure for scenario modeling (cloud-based, real-time)	CTO	6 months	Pilot system live; 2+ cross-industry teams using it for planning
Legal	Establish AI liability and malpractice framework (especially for Law, Healthcare, Finance)	General Counsel	90 days	Framework published; insurance coverage reviewed and updated

Quality Bar: Every action on this list should be:

- Concrete: No vague aspirations. "Implement AI upskilling" not "explore AI opportunities."
- Implementable: Can be done in the stated timeline with existing resources or minimal new budget.
- Owned: Someone in the room commits to shepherding it.
- Measurable: You'll know in 90 days or 6 months whether it happened.

HOW TO KNOW THE EXERCISE WORKED

A successful Project Threshold V7.4 run exhibits **at least 5 of these 7 indicators:**

#	Indicator	Threshold	Evidence
1	Strategy Pivot	3+ industry reps materially changed strategy across rounds (not tweaked — changed direction)	Primary decision in R3 or R4 differs fundamentally from R1
2	Collective Bonus Dynamics	Collective Bonus nominations shifted across rounds in response to actual decisions and execution (not static opinions).	R3-R4 nominations reflect track record, not just R1 first impressions
3	Debrief Surprises	4+ industry reps surface 2–3 genuine surprises during reflections (non-obvious; challenge actual assumptions)	Reps cite specific round where an assumption broke
4	No-Regrets Consensus	Consolidated list contains 5+ concrete, implementable actions with clear ownership	List includes timelines, owners, and success measures; no vague language
5	Risk Literacy Improvement	Decision logs show progression from generic risk statements to scenario-specific ones	R1: "we'll monitor closely" → R4: "if regulatory delay >6mo, we shift to conservative positioning"
6	Cross-Industry Recognition	3+ genuine cross-industry dependencies surfaced and discussed during debrief	Reps explicitly call out inter-industry dynamics (e.g., "Law AI adoption creates malpractice precedent constraining Healthcare")
7	Engagement Indicators	85%+ report "valuable/challenging/surprising"; zero confusion about rules; 100% debrief attendance	Feedback surveys + facilitator observation

Minimum Success Threshold: 5 of 7 indicators present = successful run. If <5, debrief with facilitators to identify adjustments for next iteration.

DECISION COUNT EXPECTATIONS BY CONFIGURATION

Configuration	Industries	Decisions/Round	Total Decisions (4 rounds)
Minimum (5 participants)	5	5	20
Medium (7 participants)	7	7	28
Large (9 participants)	9	9	36
Full (11 participants)	11	11	44

Note: Scoring time scales with decision count. With 11 industries, the 12-minute scoring block requires approximately 1 minute per decision. Consider adding a second scorer for full-roster exercises.

DELIVERABLE COLLECTION & TIMELINE

Output	Due Date	Owner	Format
Decision Logs (variable worksheets)	End of exercise (3:20 PM)	Industry Reps	Filled-in printed forms, photos, or scanned PDFs
Collective Bonus Record (4 rounds)	End of cross-industry discussion each round	Facilitator	Nomination records, +2/-2 adjustments (if 3+ agreed), stated rationale
Synthesis Memos (one per industry)	Before closing (3:20 PM) or within 24 hours	Industry Reps	1 page per industry, use provided template
No-Regrets Actions Consolidated List	End of debrief (3:20 PM)	Facilitators + Industry Reps	Shared spreadsheet or document with 5+ actions
Feedback Surveys	End of exercise (3:20 PM)	All Participants	Printed or digital (Google Form, Qualtrics, etc.)

SUCCESS LOOKS LIKE

By end of exercise, you have:

- Completed decision logs with specificity and risk awareness (20–44 depending on industries played)

- Collective Bonus record showing how group consensus evolved across rounds and influenced strategy pivots
- Synthesis memos from each industry that executives will actually read and learn from
- A no-regrets action list that your organization will start executing in Q2
- Industry participant feedback indicating challenge, insight, and relevance

This is the evidence that Project Threshold V7.4 delivered learning, not just theater.

How to Play

A simple guide to the exercise for first-time participants.

What Is This?

Project Threshold is an 8-hour tabletop exercise where participants explore how AI adoption will affect major industries of the US economy over the next 2–4 years. You'll step into the shoes of an industry leader, make strategic decisions in response to evolving scenarios, and discover how those decisions ripple across the broader economy.

The goal is not to "win" — it's to build intuition about AI's economic effects and test what strategies hold up under realistic pressure.

Who Participates?

There are two roles in the exercise:

Industry Participants (5–11 people). Each person represents one industry. You make decisions for your industry across four rounds. The eleven possible industries are:

#	Industry	Archetype
1	Retail	Top-5 omnichannel retailer
2	CPG	Diversified consumer goods manufacturer
3	Healthcare Provider	Integrated health system
4	Healthcare Payer	National health insurer
5	Finance	Diversified financial institution
6	Consulting	Big Four / MBB-class firm
7	Law	AmLaw 50 firm
8	Manufacturing	Multi-plant heavy manufacturer
9	Logistics	National freight and 3PL operator
10	Big Tech	Platform company (cloud, ads, devices)
11	B2B/B2C SaaS	Enterprise/consumer software

Not every exercise uses all 11 industries. The facilitator selects which industries are in play based on participant count and expertise. A minimum of 5 industries ensures cross-industry richness.

Facilitator (1–2 people). Runs the exercise, reads scenario updates, scores decisions, moderates discussion, and administers the Collective Bonus (Optional) and Facilitator Market Shock mechanics.

Industry Selection

Before the exercise, the facilitator assigns each participant to one industry. Ideal: 8+ industries assigned to participants. Fewer is possible — the facilitator plays unassigned industries using base case fallback decisions. If you're running with fewer than 11 participants:

Participants	Which Industries	Coverage
5	Broad coverage (e.g., Retail, Healthcare Provider, Finance, Manufacturing, Big Tech)	Core cross-industry dynamics
7	Add Consulting and SaaS	Professional services and software depth
9	Add CPG, Healthcare Payer, and Logistics	Broader industry coverage
11	Full roster — all industries represented	Complete economy simulation

Your industry assignment determines your packet, private cards, and AI adoption arc.

What You'll Actually Do

Before the Exercise

Read three documents (about 45 minutes total):

1. Executive Overview — the economic scenario and baseline forecasts
2. Your industry packet — the specific position, constraints, and profile for your industry
3. This document — so you understand the flow

You may also review your **AI Adoption Arc** (from [04_Participant_AI_Adoption_Arcs/](#)), which shows how AI adoption unfolds round-by-round for your specific industry.

On Exercise Day

The day follows this rhythm:

Time	Activity	Duration
8:30	Setup & briefing	0:30
9:00	Practice round — learn the format	0:10
9:10	ROUND 1: Foundation	1:05
10:15	Break	0:15
10:30	ROUND 2: Acceleration	1:05
11:35	Lunch	0:40
12:15	ROUND 3: Reckoning	1:05
1:20	Break	0:15
1:35	ROUND 4: Normalization	1:05
2:40	Break + Facilitator Prep	0:10
2:50	Structured debrief	1:00
3:50	Wrap-up	0:40
4:30	End	—

How a Round Works

Every round follows the same structure and takes about 65 minutes. There are four rounds total.

Step	Time	What Happens
1. Situation Update	5 min	The facilitator reads aloud a scenario update — new AI capabilities, market signals, regulatory developments, labor market shifts.
2. Inject & Private Card	3 min	You receive "inject" events (things that just happened in the world) and a confidential private card with intelligence specific to your industry. Other participants don't see your card.
3. Make Your Decision	15 min	Working on your own, you decide what your industry should do in response and fill out one decision worksheet.

Step	Time	What Happens
4. Submit	3 min	Hand your worksheet to the facilitator. If you miss the deadline, your decision defaults to a conservative baseline.
5. Scoring	12 min	The facilitator scores each decision on three dimensions: Strategic Fit, Execution Risk, and Tail Risk. Scores are simple: -2, 0, or +2 on each.
6. Cross-Industry Discussion	27 min	The heart of the exercise. The facilitator shares scores and Industry Health conditions. Collective Bonus (~5 min, optional): any participant may nominate an industry for a +2 bonus (strong strategy) or -2 penalty (risky strategy); nomination takes effect only if 3+ participants agree. Cannot nominate own industry. Then open discussion: how decisions ripple across industries, what patterns are emerging, what surprised you.

The Decision Worksheet

Each round, you submit one worksheet for your industry. The worksheet captures:

Field	What to Provide	Range / Guidance
What you're doing	A 1–2 sentence decision	Plain language, not a business case
Commitment size	How big a bet this is	Absorbable → Material → Bet-the-company
Time to payoff	When you expect returns	0–3 months → 3–12 months → 1–2 years → 2+ years
Execution complexity	How hard it is to pull off	Low → Medium → High → Very High

Field	What to Provide	Range / Guidance
Dependencies	What it hinges on	Internal execution, vendors, regulators, or ecosystem-wide change
Scale	How broadly you're deploying	Pilot → Regional → National → Global
Success metric	What success looks like	Directional, not a precise number
Risk & mitigation	What could go wrong	Main risk + how you'd handle it

Don't aim for precision. The exercise uses broad bands, not exact figures. "Material but manageable investment, 3–12 month timeline, medium complexity" is exactly the right level of specificity.

Market Pressure Mechanics

Two mechanics create external market pressure throughout the exercise — just like real market dynamics.

Collective Bonus (Optional, Every Round)

During each cross-industry discussion, any participant may nominate an industry for a +2 bonus (strong strategy) or -2 penalty (risky strategy). You cannot nominate your own industry. A nomination takes effect only if 3+ participants agree. Multiple nominations may occur in a single round. If fewer than 3 participants agree, no adjustment is made.

This creates optional peer accountability and organic market pressure. Strong reasoning can be rewarded; risky reasoning can carry consequences — but only when the group reaches consensus.

Facilitator Market Shock (Round 2 Only)

At the start of Round 2, the facilitator imposes one constraint each on 2–3 industries from the constraint menu (Regulatory Halt, Labor Cost Surge, Capital Tightening, Reputational Pressure, Competitive Response, Litigation Risk). Any industry can be hit — not just those performing poorly. This forces adaptive decision-making mid-exercise.

Industry Health Signals

After each round (starting Round 2), the facilitator announces how your industry is doing based on your cumulative score across all rounds. There are five conditions:

Condition	What It Means	What Happens
Surge	Your industry is thriving structurally	Room to be ambitious — can attempt bold moves without automatic challenge
Tailwind	Navigating well; market views you favorably	Business as usual; no modifiers
Steady	Holding position; no clear momentum	No modifiers
Headwind	Under structural pressure; margins compressing	One constraint imposed at start of next round
Crisis	Structural decline; survival mode	Two constraints imposed; must include defensive component in your decision

These conditions create real consequences — strong early decisions build momentum, while poor decisions compound into structural disadvantage.

Private Information Cards

In Rounds 1–3, you receive a confidential card with industry-specific intelligence (from [05_Participant_Private_Cards/](#)). This might be internal data (e.g., your AI pilot results were weaker than expected), competitive intelligence (e.g., a competitor is about to launch a disruptive product), or regulatory signals (e.g., a major enforcement action is coming).

You can reference your card's information in your decisions and in discussion, but you shouldn't read the card aloud verbatim. Cards accumulate — keep them all and use them in later rounds.

There is no card in Round 4. The final round operates on public information only.

AI Adoption Arc

Each industry has an **AI Adoption Arc** (from [04_Participant_AI_Adoption_Arcs/](#)) that describes how AI adoption unfolds round-by-round for your specific industry. This provides context for your decisions: what stage of AI deployment your industry is in, what capabilities are becoming available, and what constraints are shifting.

Review your arc before the exercise and reference it each round as you make decisions.

Scoring

Decisions are scored on three dimensions, each rated -2, 0, or +2:

Dimension	What It Measures
Strategic Fit	Does this decision align with your industry's position and the evolving AI landscape?
Execution Risk	Can your organization realistically pull this off in the timeframe?
Tail Risk	Have you accounted for what could go catastrophically wrong?

A score of +6 means you nailed all three. A score of 0 means your decision was safe but unremarkable. Negative scores mean your logic had gaps or you ignored real constraints.

Scoring exists to create structure and feedback — it's a tool for learning, not a competition.

The Debrief (60 min)

The debrief is the most important part of the exercise. It is non-negotiable and cannot be shortened.

Phase	Time	What Happens
Individual Reflections	15 min	Each industry participant shares 1–2 minutes on what surprised them and what they learned.
Cross-Industry Discussion	25 min	Facilitated conversation about patterns, interdependencies, and dynamics that emerged across industries. Includes facilitator synthesis of Collective Bonus trends and Market Shock effects across all four rounds.
No-Regrets Actions	20 min	The group identifies strategic actions that make sense regardless of which AI scenario unfolds.

Key Principles

You're building intuition, not being tested. There are no right answers. The value is in the reasoning, the trade-offs you surface, and the cross-industry dynamics you discover.

Decisions are binding. Once you submit a decision, you can't undo it. You can modify or build on it in later rounds, but you live with your choices. This mirrors real strategic decision-making.

The discussion is where the learning happens. The 27-minute cross-industry discussion each round is not filler — it's the most valuable part of the exercise. Challenge each other. Ask "how does your decision affect my industry?" Surface the interdependencies.

Use broad strokes, not false precision. When describing your decisions, use the provided bands (low/medium/high, pilot/regional/national). Don't invent specific dollar amounts or percentages unless the exercise materials provide them.

Private information creates realistic asymmetry. Not everyone knows the same things. This is by design. It mirrors real markets where different actors have different intelligence.

Quick Reference: The Eleven Industries

#	Industry	Archetype
1	Retail	Top-5 omnichannel retailer
2	CPG	Diversified consumer goods manufacturer
3	Healthcare Provider	Integrated health system (20+ facilities)
4	Healthcare Payer	National health insurer
5	Finance	Diversified financial institution (bank + insurance)
6	Consulting	Big Four / MBB-class firm
7	Law	AmLaw 50 firm
8	Manufacturing	Multi-plant heavy manufacturer
9	Logistics	National freight and 3PL operator
10	Big Tech	Platform company (cloud, ads, devices, enterprise software)
11	B2B/B2C SaaS	Enterprise or consumer software company

Practice Round Walkthrough

Purpose

This document scripts a complete practice micro-round so that all individual industry representatives understand the decision-making flow, scoring mechanics, and submission process before Round 1 begins. The facilitator can read this walkthrough aloud, adapt it to fit timing, and answer questions that arise during execution.

Key Goal: Individual participants will have confidence in the worksheet format, understand how scoring works, and know what to expect when time pressure increases in the live rounds.

Overview

- Duration: 10 minutes total (scripted, compressed from live rounds)
 - Materials: 5-11 individual industry representatives (each with 1 blank decision worksheet)
 - Scenario: Simple, AI adoption in a professional services environment
 - Output: Minimum 5 individual decisions (1 per industry participant)
 - Key Change: Each individual industry representative submits ONE decision per round. If a participant is absent or passes, they receive a pre-defined base case fallback (see 02BaseCaseFallbackBank.md).
-

Facilitator Script: Setup (1 minute)

"Good morning, everyone. Before we launch Round 1, we're going to run a quick practice round together. This is a sandbox — nothing here counts toward your final score. The purpose is threefold: first, you'll see the exact worksheet format; second, you'll watch scoring happen in real time; and third, you'll practice the rhythm of individual decision preparation, submission, and feedback.

We're going to compress the timing today — 10 minutes instead of 15 minutes of decision prep. That means you'll move fast. That's intentional. In the real rounds, you'll have more time to think through risk, but the structure is identical.

Each of you will decide individually for your industry. No team deliberation — you analyze and decide on your own.

Everyone clear? Good. Let's begin."

Step 1: Situation Update (1 minute)

The Scenario

"It's January 2026. A major AI company has just released a new foundation model that scores 90% on professional licensing exams — think bar exam, medical boards, engineering exams, you name it. Adoption is ramping fast across industries. Your industry needs to decide how to respond in the next six months.

Here's your inject card — this is new information that just arrived."

[Facilitator reads aloud]

INJECT: Three Fortune 500 companies announce they are cutting entry-level hiring by 25% and redirecting savings to AI tool licenses.

"This means the labor market is tightening. Fewer junior roles. More pressure to upskill existing staff with AI. Every industry sees it — but it hits some harder than others.

Each of you will now submit one decision for your industry. That decision describes what you'll invest in, how long it takes, what it costs, and what success looks like.

Let's go."

Step 2: Individual Decision Preparation & Submission (3 minutes)

Individual Instructions (read aloud)

"Each of you has three minutes to analyze the situation and complete ONE decision worksheet for your industry. You are deciding alone — no team deliberation.

Key rule: When you submit a decision, your industry gets scored. If you decide to skip or don't submit, your industry receives a pre-defined base case fallback (see the Base Case Fallback Bank reference). That fallback scores 0 on all dimensions but keeps your industry in play.

Think critically about your industry's position. What's your response to the hiring freeze inject? What AI investments make sense given the constraint?

You must decide: What's your industry doing in the next 6 months?

Time starts now. Go."

Example Decisions (Script If Needed)

If a participant is stuck, the facilitator can offer these examples to unblock them. NOTE: These examples use BANDED FRAMEWORK instead of granular costs.

Field	Retail Example	Finance Example	Consulting Example
Decision Title	Pilot AI-Powered Inventory Optimization	AI-Assisted Fraud Detection with Human Review Loop	Deploy AI Copilot to Consulting Workforce
Description	ML inventory optimization in 50 Midwest stores; leverages POS, weather, and promo data to reduce overstock/stockouts	AI flags suspicious transactions in real time; human analyst reviews all flags before action (block, investigate, monitor)	AI copilot rollout to all consultants — phased (500/500/1000 over 3 months) with e-learning and buddy system
Spend/Commitment	Material but manageable	Transformational bet	Material but manageable
Time-to-Impact	3-12 months	3-12 months	3-12 months
Complexity	Medium	High	High
Dependency	Requires vendor/partner	Requires regulator/standards (OCC/Fed/FDIC)	Requires vendor/partner
Scale	Regional / Pilot (50 stores)	National / Enterprise-wide	Enterprise-wide
Success Metric	Moderate impact (reduce markdowns, improve inventory turns)	High impact (reduce missed fraud 30%, material loss reduction)	Moderate-High impact (15% increase in billable hours per consultant)
Key Risk	Store adoption resistance (mitigate: manager buy-in programs); data quality (mitigate: early audit)	Regulatory scrutiny (mitigate: explainability + audit trails); model drift (mitigate: monthly retraining); system failure (mitigate: 48-hour rollback)	Low adoption (mitigate: weekly utilization tracking, remediation training); vendor lock-in (mitigate: require open APIs)

Worksheet Template (Show One Completed Example)

[Facilitator shows the Retail decision on screen or poster.]

"Here's what a filled-out worksheet looks like. Notice the key fields — and remember, we're using BANDS instead of invented precision:

- Industry at the top
- Decision Title — plain language, clear scope
- Description — the 'who, what, where, when' (directional, not granular)
- Timeline — milestones or phases (avoid false precision like 'month 1.5')
- Spend/Commitment Band — {Absorbable (BAU), Material but manageable, Transformational bet, Potentially existential}
- Time-to-Impact Band — {0-3 months, 3-12 months, 1-2 years, 2+ years}

- Execution Complexity Band — {Low, Medium, High, Very high}
- Dependency Band — {Mostly internal, Requires vendor/partner, Requires regulator/union/standards, Requires ecosystem shift}
- Scale Band — {Pilot / Regional / National / Global}
- Success Metric Band — {Low/Moderate/High/Transformational impact}
- Key Risk & Mitigation — be specific about the top 2-3 risks

Rules: Avoid line-item cost breakdowns. Avoid invented precision (don't say '\$15.3M'; say 'Material but manageable'). Use bands. Let the band language do the work.

Important: You must submit ONE decision for your industry. If you don't submit, your industry auto-defaults to base case fallback (scores 0 on all dimensions).

You have three minutes. Fill the worksheet. Go."

[Time: 3 minutes elapsed. 4 minutes total.]

Step 3: Share-Out and Plausibility Challenge (2 minutes)

Two Participants Present

"Let's hear from two participants. Retail, you go first — 30 seconds."

Retail Participant: "Retail is piloting AI inventory optimization in 50 Midwest stores. Six-month timeline, material but manageable investment, goal is reduce markdowns and improve inventory turns. Phased rollout with vendor partnership."

Facilitator: "Good. Clear scope. Regional pilot — that's smart. Bounded risk. Next: Consulting."

Consulting Participant: "Consulting is deploying an AI copilot to all consultants — three-month rollout, material but manageable investment, goal is 15% increase in billable hours per consultant. Phased across three batches with buddy system and weekly tracking."

Facilitator Challenge

"Consulting, I want to probe the copilot rollout. Three months to full rollout across all consultants — that's aggressive. How do you handle change management? Adoption risk?"

Consulting Participant: "Good catch. We'd phase it — 500 consultants in month 1, another 500 in month 2, full 2000 by month 3. Each batch gets two hours of e-learning, plus a buddy system pairing early adopters with skeptics. We track utilization weekly and adjust training if adoption lags."

Facilitator: "Acceptable. That's a real rollout plan, not a hope. In live rounds, I might ask: What's your fallback if only 30% of consultants adopt by month 2? But for now, this is solid. I'd score this as manageable execution risk."

[Time: 2 minutes elapsed. 6 minutes total.]

Step 4: Scoring Example (2 minutes)

Scoring Mechanics

"Now let's score two of these decisions so you see how banding works. Every decision gets scored on three dimensions — Strategic Fit, Execution Risk, and Tail Risk. Each dimension can be +2, 0, or -2. Red-flag trigger conditions can shift scores by +/-3. Let me show you."

Example 1: Retail — Inventory Optimization (Strong Decision)

[Facilitator writes on board or displays:]

Dimension	Band	Reasoning
Strategic Fit	+2	Well-aligned with Retail's core operations (inventory management). "Material but manageable" spend level is appropriate for test-and-learn. Pilot scope (Regional) is bounded. Directly reduces key metric (markdowns).
Execution Risk	+2	"Medium" complexity band is realistic. 6-month timeline aligns with typical pilot deployment. Risk mitigations (data audit, train-the-trainer) are credible. No novel technical/organizational risks.
Tail Risk	+2	Bounded downside. Regional pilot means limited exposure if performance disappoints. Easy rollback. No regulatory exposure. Brand risk minimal (internal operations, not customer-facing).
TOTAL SCORE	+6 points	This is a well-balanced decision. Banded language prevents false precision while maintaining clarity.

Facilitator commentary: "Notice what makes this strong: bounded scope (Regional, not Global), realistic complexity (Medium, not Very High), managed downside. Retail isn't betting the company. Test, learn, scale. That's strategic discipline. The banded language lets us focus on the real trade-offs instead of arguing about whether cost is \$15M or \$17M."

Example 2: Consulting — Copilot Rollout (Acceptable Decision)

[Facilitator writes on board or displays:]

Dimension	Band	Reasoning
Strategic Fit	+2	Directly addresses stated pain point (upskilling in response to entry-level hiring freeze). "Material but manageable" spend (Spend Band) and "3-12 months" time-to-impact align with strategy. Well-motivated.
Execution Risk	0	Full rollout across 2000 consultants in three months is "High" complexity. Phased rollout plan (500 -> 500 -> 500) mitigates some risk, but adoption uncertainty remains. Change management is outlined but not stress-tested. Could be de-risked further.
Tail Risk	0	If adoption stalls and utilization gains don't materialize, revenue takes hit. Sunk cost (\$5M) is absorbable but impact on quarterly results is real (downside is "Moderate", not "Low"). Not existential, but material.
TOTAL SCORE	+2 points	Acceptable decision. Strategically sound, but execution and downside risks need tighter mitigation.

Facilitator commentary: "This is the difference between +2/+2/+2 and +2/0/0. Consulting moved boldly, which is great. But they've put execution risk management in the medium band. In a real round, I'd push: 'If only 50% of consultants adopt by month 2 instead of your assumed 70%, what happens to revenue?' That's the question they haven't fully answered. For this practice round, acceptable. For live rounds, they need to either de-risk the execution or acknowledge the tail risk more explicitly."

Example 3: Red Flag Trigger (Illustrative)

"Now, let me show you what triggers a red flag. Imagine the Consulting participant had proposed: 'Fire 40% of associate staff and replace them with AI copilots. Timeline: 6 months. Cost: \$2M in severance. Expected savings: \$50M. Net benefit: \$48M.'

That would trigger multiple flags:

- Plausibility: Is a 40% headcount reduction actually executable in six months? Legally? Operationally?

- Tail Risk: If the copilots underperform or adoption stalls, you've permanently lost institutional knowledge and client relationships. That's a tail risk that could impair revenue for years.
- Ethical/Stakeholder: Are there labor law, client service, or reputational risks you haven't mentioned?

If I see a proposal like that, I would challenge it hard. I might score it -3 on Execution Risk and -3 on Tail Risk because the downside far outweighs the upside. Or I might reject it as implausible and ask for a revised proposal.

That's how red flags work. They're rare. But they're circuit-breakers for decisions that look too good to be true or too risky to stomach."

[Time: 2 minutes elapsed. 8 minutes total.]

Step 5: Wrap-Up and Key Reminders (1 minute)

Facilitator: "That's the process. Every round in Project Threshold follows this exact flow:

1. Situation update and inject.
2. Each participant submits ONE decision for their industry (minimum 5 total; up to 11 if all industries are represented).
3. Industries without explicit decisions receive pre-defined base case fallbacks from the fallback bank.
4. Two participants present; I challenge one decision for plausibility.
5. I score all submitted decisions on the three dimensions using banded framework.
6. Cross-industry discussion and optional Collective Bonus — any participant may nominate an industry for +2 (strong) or -2 (risky) if 3+ participants agree.
7. We review scores and any Collective Bonus results, then talk through the reasoning.

In live rounds, deliberation time will be 15 minutes instead of 3 — so you'll have room to stress-test your decisions, model risks, and think through cross-industry implications. That said, the worksheet format and scoring logic are identical to what you just did.

A few key reminders for Rounds 1, 2, and 3:

- Hard deadline: When I call time, submission closes. No late arrivals.
- One decision per participant: Each participant submits one decision for their industry.
- Base case fallbacks: Industries without explicit decisions auto-default to base case fallback (scores 0 on all dimensions).
- Banded framework: Use bands — avoid invented precision. 'Material but manageable' not '\$15.3M'. Let the band language do the work.
- Collective Bonus is open and optional: Nominations are stated publicly. A nomination takes effect only if 3+ participants agree. This creates real but consensus-driven accountability.
- Scoring is consistent: I use the same banding logic every round. {-2, 0, +2} per dimension, with red-flag triggers at +/-3.
- Red flags are rare and explicit: If I flag a decision, I'll tell you exactly what I'm concerned about. You can then decide to revise or defend.

Any questions?"

[Allow 30 seconds for Q&A.]

Facilitator: "Excellent. You're ready for Round 1. We'll begin in five minutes. Grab water, stretch, and reconvene. See you shortly."

[Time: 1 minute elapsed. 9 minutes total. Total time: ~10 minutes.]

Appendix: Blank Decision Worksheet

The Decision Worksheet is a standalone printable form. See **03ParticipantPackets/03zDecision_Worksheet.md** for the full template.

Print multiple copies per participant before the exercise begins. Each participant submits one worksheet per decision per round.

Quick reference — the 12 required fields:

#	Field	Format
1	Industry Identification	Write industry name
2	Decision Statement	1–2 sentences
3	Spend/Commitment Band	Check one of 4 bands + justify
4	Time-to-Impact Band	Check one of 4 bands + justify
5	Execution Complexity Band	Check one of 4 bands + justify
6	Dependency Band	Check one of 4 bands + justify
7	Scale	Check one of 4 options
8	Success Metric	Directional range, not point estimate
9	Key Risk & Mitigation	Top 2 risks + mitigations
10	Cross-Industry Synergies	1 line or "Independent decision"
11	Private Card Influence	1 line or "None" (R1–R3)
12	Continuity Note	Builds on prior round (R2–R4)

End of Practice Round Walkthrough

8-Hour Session Timeline

Total Duration: 8 hours exactly (e.g., 8:30 AM to 4:30 PM) **Key Principle:** Debrief is protected and non-negotiable. All flexibility built into rounds and breaks.

EXERCISE TIMELINE

Time	Block	Duration	Key Actions	Owner	Notes
8:30–8:45	Arrival & Setup	15 min	Materials distribution, name badges, seating by industry, login credentials confirmed	Logistics Lead	Hard stop at 8:45
8:45–9:00	Welcome & Exercise Overview	15 min	Facilitate opening remarks, review objectives, explain rounds 1–4 structure, industry participant introductions	Lead Facilitator	Hard stop at 9:00
9:00–9:10	PRACTICE MICRO-ROUND	10 min	One trivial scenario inject, mock individual decision (2 min), facilitator scoring demo (2 min), Q&A (2 min)	Lead Facilitator	Hard stop at 9:10 — ensures everyone understands decision submission format
9:10–10:15	ROUND 1: Foundation	1:05	See detailed breakdown below	Industry Participants + Facilitators	First full round; sets baseline; includes Collective Bonus (Optional)

Time	Block	Duration	Key Actions	Owner	Notes
10:15–10:30	Break	15 min	Coffee, restrooms, facilitator scoring prep	All	Hard stop at 10:30
10:30–11:35	ROUND 2: Acceleration	1:05	See detailed breakdown below	Industry Participants + Facilitators	Facilitator Market Shock imposes constraints; includes Collective Bonus (Optional)
11:35–12:15	Lunch	40 min	Offsite or catered; participants may informally discuss (but no decisions logged)	All	Hard stop at 12:15
12:15–1:20	ROUND 3: Reckoning	1:05	See detailed breakdown below	Industry Participants + Facilitators	Heightened stakes; includes Collective Bonus (Optional)
1:20–1:35	Break	15 min	Restrooms, facilitator brief, energy reset	All	Hard stop at 1:35
1:35–2:40	ROUND 4: Normalization	1:05	See detailed breakdown below	Industry Participants + Facilitators	Final round; no private cards; final Collective Bonus (Optional); synthesis begins
2:40–2:50	Break + Facilitator Prep	10 min	Participants rest; facilitators consolidate scores, prepare debrief materials	Facilitators	Hard stop at 2:50
2:50–3:50	STRUCTURED DEBRIEF	60 min	PROTECTED BLOCK — DO NOT SACRIFICE	Lead Facilitator	See debrief structure below

Time	Block	Duration	Key Actions	Owner	Notes
3:50–4:30	Wrap-Up & Close	40 min	Collect feedback surveys, closing remarks, thank you, schedule for synthesis memo collection (due within 5 days)	Lead Facilitator	Hard stop at 4:30 — Firm end time

ROUND STRUCTURE (Rounds 1–4) — INDIVIDUAL INDUSTRY PARTICIPANT MODEL (V7.4)

Each round follows this sequence. **Total per round: 1:05. Each industry participant submits one decision per round, with cross-industry discussion emphasis.**

Time Slot	Sub-Block	Duration	What Happens	Owner
:00–:05	Situation Update Read-Aloud	5 min	Facilitator reads current scenario, new context, updated conditions. All participants listen (no decisions yet).	Lead Facilitator
:05–:08	Private Card Distribution + Inject Delivery	3 min	(Rounds 1–3 only) Distribute industry-specific private cards (from 05_Participant_Private_Cards/). Read inject aloud. Answer clarifying questions only (no strategy discussion).	Inject Manager

Time Slot	Sub-Block	Duration	What Happens	Owner
:08--:23	Individual Decision Preparation	15 min	Each industry participant analyzes the inject and prepares their industry decision individually. Participants may reference their AI Adoption Arc (04_Participant_AI_Adoption_Arcs/) and seek clarification from facilitators but make decisions independently. No cross-industry coordination during prep. Facilitators observe silently; no coaching. Hard stop at :23.	Industry Participants
:23--:26	Decision Submission	3 min	Each industry participant submits their decision worksheet. Facilitators collect submissions. Late submissions or missing worksheets: that industry auto-defaults to base case fallback (see 02BaseCaseFallback Bank.md).	Scorekeepers
:26--:38	Facilitator Scoring	12 min	Facilitators score all submitted industry decisions using banded framework ({-2, 0, +2} per dimension). Base case fallback industries noted but not scored.	Facilitators

Time Slot	Sub-Block	Duration	What Happens	Owner
:38–1:05	Cross-Industry Discussion & Results	27 min	<p>Results (5 min): Facilitator announces scores and Industry Health Signal updates. Collective Bonus (~5 min, optional): Any participant may nominate an industry for +2 (strong strategy) or -2 (risky strategy); nomination takes effect only if 3+ participants agree. Cannot nominate own industry. Discussion (~17 min): Which decisions surprised you? How do decisions in one industry create spillovers for others? Participants challenge each other's reasoning and explore interdependencies.</p>	Lead Facilitator + Industry Participants

Facilitator Notes for Rounds (V7.4 Industry Model):

- Round 1 (9:10–10:15): Baseline round. Participants may be cautious. All industries expected to submit decisions (individuals decide quickly). First Collective Bonus opportunity establishes early group consensus on strategy quality. Facilitator highlights key spillovers between decisions.
- Round 2 (10:30–11:35): Acceleration. Facilitator Market Shock: Before decision prep, facilitator selects 2–3 industries and imposes one constraint each from the constraint menu (~3 min). Inject creates urgency. Cross-industry discussion reveals how constraints cascade across industries (e.g., Consulting pricing pressure affects Finance advisory spend; Law AI adoption creates precedent concerns). Collective Bonus continues.
- Round 3 (12:15–1:20): Reckoning. External shock injects. Collective Bonus carries heightened stakes as cumulative scores diverge. Participants adjust strategy based on peer signals and Industry Health conditions.

- Round 4 (1:35–2:40): Normalization. No private cards; full transparency. Final decisions for long-term positioning. Final Collective Bonus opportunity. Cross-industry discussion focuses on emergent patterns and no-regrets actions. Expect all industries to submit.

DEBRIEF STRUCTURE (60 minutes, 2:50–3:50)

This block is **non-negotiable** and protected from round overruns.

Time	Sub-Block	Duration	What Happens	Owner
2:50–3:05	Individual Reflections	15 min	Each industry participant shares 1–2 min highlights: surprises, key learnings, what they'd do differently. (With 11 participants, keep to 1 min each; with 5 participants, allow 2–3 min each.) Other participants listen and note patterns.	Industry Participants

Time	Sub-Block	Duration	What Happens	Owner
3:05–3:30	Cross-Industry Discussion	25 min	Facilitator moderates discussion of inter-industry dynamics, unexpected spillovers, and emerging patterns. Includes facilitator synthesis of Collective Bonus trends across all four rounds and Facilitator Market Shock effects. What surprised everyone? How did constraints cascade across industries? Which assumptions held or broke? Participants identify common threads in no-regrets actions.	Lead Facilitator + Industry Participants
3:30–3:50	No-Regrets Actions & Key Takeaways	20 min	Consolidate no-regrets actions on shared whiteboard/document. Facilitator captures final list by category (governance, workforce, compliance, tech). Each industry participant states one key takeaway.	Lead Facilitator + Industry Participants

BUFFER & FLEXIBILITY NOTES

Block	Flexible?	Notes
Arrival/Welcome	No	Fixed. Sets tone.
Practice Micro-Round	No	Fixed. Essential for calibration.
Rounds 1–4 Individual Prep (15 min each)	Yes, +/-3 min	Can extend to :26 if needed; compress other sub-blocks. But never cut into Submission or Cross-Industry Discussion.
Breaks	No	Fixed times. Facilitators must enforce hard stops.
Lunch	No	Fixed time. May extend to 45 min if needed, compress adjacent break instead.
Debrief	NO — ABSOLUTE HARD STOP	Minimum 60 min. Non-negotiable. If rounds run long, debrief starts late but gets full 60 min.
Wrap-Up	No	Fixed. Brief closing.

Scaling Note for 11 Participants: With more industry participants, scoring time may need slight extension. If running 11 industries, facilitator may need a second scorer or simplified scoring approach during the 12-minute scoring block. Plan for approximately 1 minute per industry decision.

FACILITATOR CHECKLIST

- 8:25 AM: All materials printed and organized by industry
- 8:25 AM: Scoring rubric and Collective Bonus procedure reviewed with all facilitators
- 9:00 AM: Confirm all industry participants have decision submission forms + pens
- 9:10 AM: Begin micro-round; watch for timing
- Each round: Set visible timer (countdown clock projected if possible)
- Each round: Call "3 minutes" warning for individual prep hard stop
- Each round: Distribute private cards from 05_Participant_Private_Cards/ at correct time (Rounds 1–3 only)
- Post-round: Facilitators score within 12 min; announce Industry Health conditions
- 2:50 PM: Transition to debrief; do not allow round content to bleed into debrief time
- 3:50 PM: Hard stop debrief; distribute feedback surveys
- 4:25 PM: Thank industry participants; remind about synthesis memo due date (within 5 business days)

TIME-AT-A-GLANCE

08:30 +- Arrival
08:45 +- Welcome
09:00 +- Micro-Round
09:10 +- ROUND 1 (Foundation)
10:15 +- Break
10:30 +- ROUND 2 (Acceleration)
11:35 +- Lunch
12:15 +- ROUND 3 (Reckoning)
01:20 +- Break
01:35 +- ROUND 4 (Normalization)
02:40 +- Break
02:50 +- DEBRIEF (60 min, non-negotiable)
03:50 +- Wrap-Up
04:30 +-- END

Total: 8 hours exactly. Debrief protected. Ready to execute.

02

Facilitator Guide

Facilitator Guide Index

Master index for the **02FacilitatorGuide/** folder. Use this to navigate, understand file dependencies, and prepare for live play.

COMPLETE FILE LISTING

File	Purpose	Key Audience
02Overviewand_Runbooks.md	Round-by-round play schedule; 11 industries; flexible participant model (5-11 participants, each selects 1+ industries); Collective Bonus (Optional) and Facilitator Market Shock mechanics; individual decision model; round timing (15min decision prep, cross-industry discussion with optional Collective Bonus); Private Card and AI Adoption Arc distribution procedures; cross-industry discussion prompts	First-time facilitators; workshop designers
02AdjudicationRules.md	Core scoring framework: 3 dimensions (Strategic Fit, Execution Risk, Tail Risk); banded scoring (-2/0/+2 typical, +/-3 exception); decision specificity checklist (WHO/WHAT/WHERE/WHEN/HOW/ HOW MUCH/RISK); red-flag triggers; when to challenge vs. accept; base case fallback application; all 11 industries enumerated	Primary reference during live scoring
02PlausibilityDecision_Trees.md	15 strategic archetypes + industry-specific decision trees for all 11 industries including Consulting and Law; Big Tech scoped to cloud/ads/devices/enterprise software (excludes AI lab/model dev)	Advanced facilitators for complex red-flag challenges

File	Purpose	Key Audience
02ScoringBaselines.md	Industry-specific scoring ranges for all 11 industries (Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS); band-to-score translation reference; typical scores for 15 strategic archetypes; calibration guidance	During live scoring for industry-specific context
02IndustryHealthSignalTables.md	Industry Health Signal tables: cumulative score → condition band (Surge/Tailwind/Steady/Headwind/Crisis) → mechanical consequences; constraint menu; tracking template; historical multiples reference	After explicit + fallback scoring complete
02QuickReference_Card.md	One-page cheat sheet: decision specificity checklist; banded scoring summary; red-flag categories; archetype quick guide; industry baselines for 11 industries; Collective Bonus (Optional) and Facilitator Market Shock reference; time management; materials checklist (including Private Cards and AI Adoption Arc); key debrief questions	Print + laminate. Keep at hand during live play.
02BaseCaseFallbackBank.md	Deterministic fallback scores for 11 industries (no explicit action submitted) x 4 rounds. Rationale per industry including new Law fallback. How to apply. Non-negotiable.	During scoring: apply to industries without explicit decisions
02BandScoringQuickRef.md	Two-page visual guide: 4 bands + options; band-to-score tier system; red-flag band combinations; 15-second scoring decision tree; 3 worked examples (including Law); common errors; fallback quick ref for 11 industries; industry watch-outs	Print + laminate. Reference during EVERY explicit scoring decision.

File	Purpose	Key Audience
02__README.md (this file)	Master navigation guide; file dependencies; reading paths (first-time vs. experienced facilitators); key terminology; day-of checklist	Facilitators new to V7.4; workshop prep.

Related Folders (Outside Facilitator Guide)

Folder	Purpose	Distribution
05ParticipantPrivate_Cards/	3 Private Cards per industry for Rounds 1-3. Card 1 shared within industry clusters; Cards 2-3 unique per industry.	Facilitator distributes at start of each round
04ParticipantAIAdoptionArcs/	AI Adoption Arc phases. Foundation phase in pre-read packets; Phases 2-4 as handouts at start of Rounds 2-4.	Foundation in pre-read; Phases 2-4 distributed by facilitator

RECOMMENDED READING ORDER

For First-Time Facilitators (New to Project Threshold V7.4)

Before the workshop:

- 02Overviewand_Runbooks.md -- Understand round structure, 11 industries, flexible participant model (ideal: 8+ industries; facilitator plays unassigned), individual decision model, new round timing (15min decision prep, 27min cross-industry discussion), Private Card and AI Adoption Arc distribution
- 02AdjudicationRules.md -- Core scoring framework; decision specificity; red-flag triggers; fallback application
- 02ScoringBaselines.md -- Eleven industry profiles; scoring ranges; calibration examples (pay special attention to new Consulting and Law baselines)
- 02BaseCaseFallbackBank.md -- Fallback scores for 11 industries x 4 rounds; when to apply; non-negotiability

Practice round (10 min before live):

- 02QuickReference_Card.md (printed + laminated)
- 02BandScoringQuickRef.md (printed + laminated)
- Run one mock decision; score together; debrief on specificity + bands

During live play:

- 02BandScoringQuickRef.md (primary reference; laminated at table)
- 02BaseCaseFallbackBank.md (reference for fallback industries after explicit scoring)
- 02QuickReference_Card.md (for time management + debrief questions)

For Experienced Facilitators (Returning from V7)

Key updates (V7 -> V7.4):

- 11 industries (UP from 10 personas/5 sectors): Professional Services split into separate Consulting and Law industries. Each is now an independent industry with its own packet, baselines, fallback, and decision trees.
- Terminology change: "Sector" -> "industry"; "persona" -> "industry". Each industry IS the packet.
- Flexible participant model: 5-11 participants; each selects 1+ industries (recommend 2). One decision worksheet per industry per round.
- Big Tech scope narrowed: Excludes AI lab/model development. Scope is cloud, ads, devices, enterprise software. Model development dynamics introduced via facilitator injects only.
- Private Cards: Now in separate 05_Participant_Private_Cards/ folder. 3 cards per industry for Rounds 1-3. Card 1 shared within industry clusters; Cards 2-3 unique per industry.
- AI Adoption Arc: Now in separate 04_Participant_AI_Adoption_Arcs/ folder. Foundation phase in pre-read; Phases 2-4 distributed as handouts at start of each round.
- IC eliminated: Investment Committee role removed entirely. Replaced by Collective Bonus (Optional) (all rounds, +2/-2 if 3+ participants agree) and Facilitator Market Shock (R2 only, facilitator imposes constraints). Industry Health Signals replace Valuation Translation Pipeline.

Quick prep:

1. Read 02Overviewand_Runbooks.md carefully (11-industry model + distribution procedures are significant changes)
2. Review 02PlausibilityDecision_Trees.md for new Consulting and Law decision trees
3. Review 02ScoringBaselines.md for new Consulting and Law baselines
4. Reference 02BaseCaseFallbackBank.md for new Law fallback entry
5. Reference 02BandScoringQuickRef.md during live play (updated for 11 industries)

FILE DEPENDENCY MAP

```
02_Overview_and_Runbooks
| (references)
02_Adjudication_Rules
|-- (references) 02_Plausibility_Decision_Trees
|-- (references) 02_Scoring_Baselines
|-- (references) 02_Base_Case_Fallback_Bank
+-- (references) 02_Industry_Health_Signal_Tables

02_Quick_Reference_Card
|-- (summarizes) 02_Adjudication_Rules
|-- (summarizes) 02_Base_Case_Fallback_Bank
+-- (references) 02_Overview_and_Runbooks

02_Band_Scoring_Quick_Ref
```

```
|-- (extracted from) 02_Adjudication_Rules
|-- (extracted from) 02_Scoring_Baselines
+-- (references) 02_Base_Case_Fallback_Bank
```

```
02_Base_Case_Fallback_Bank
+-- (applied after) 02_Adjudication_Rules explicit scoring
```

External Dependencies:

```
05_Participant_Private_Cards/ --> distributed per 02_Overview_and_Runbooks round flow
04_Participant_AI_Adoption_Arcs/ --> distributed per 02_Overview_and_Runbooks round flow
```

KEY TERMINOLOGY

Term	Definition
Industry	One of 11 segments: Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS. Each industry IS the packet. Replaces "persona" and "sector" from V7.
Band	Five dimensions for scoring input: Spend/Commitment, Time-to-Impact, Execution Complexity, Dependency, Scale
Banded Scoring	Use {-2, 0, +2} per dimension (or +/-3 if red-flag); prevents false precision
Red-Flag Trigger	Band combination (e.g., Global + 0-3mo + Very High) that signals plausibility concern; unlocks +/-3 exception scoring
Exception Scoring (+/-3)	Unlock when red-flag band combo fires; allows -3 or +3 on a dimension if situation is severe or mitigated credibly
Base Case Fallback	Deterministic score (small fixed delta: +/-1 per dimension) for industry with no explicit participant action; applied automatically; no participant input
Fallback Bank	Table of 11 industry fallback scores (02BaseCaseFallbackBank); reference when industry receives no explicit decision
Execution Risk	Can the organization execute within the timeframe? (mapped from bands: Spend, Time, Complexity, Dependency, Scale)
Strategic Fit	Alignment with industry fundamentals and macro scenario (-2 to +2 typical; +/-3 exception)

Term	Definition
Tail Risk	Downside exposure (regulatory, competitive, reputational, systemic) (-2 to +1 typical; +/-3 exception)
Industry Representative / Participant	Individual who selects 1+ industries (recommend 2); makes individual decisions; no team deliberation. 5-11 participants total.
Collective Bonus (Optional)	Each round: participants may optionally nominate industries for +2 (strong strategy) or -2 (risky strategy). Nomination takes effect only if 3+ participants agree. Conducted during cross-industry discussion (~5 min).
Facilitator Market Shock	R2 only: facilitator imposes one constraint each on 2-3 industries from constraint menu (~3 min).
Industry Health Signal	Cumulative aggregate score mapped to condition band: Surge (+15+), Tailwind (+6 to +14), Steady (-5 to +5), Headwind (-6 to -14), Crisis (-15 or worse). Announced at start of each round (starting Round 2). Headwind/Crisis impose mechanical constraints. See 02IndustryHealthSignalTables.
Private Cards	Industry-specific information cards distributed at start of each round (Rounds 1-3). From 05_Participant_Private_Cards/. Card 1 shared within industry clusters; Cards 2-3 unique per industry.
AI Adoption Arc	4-phase framework. Foundation phase in pre-read; Phases 2-4 distributed as handouts. From 04_Participant_AI_Adoption_Arcs/.
Specificity Checklist	7 required items (WHO/WHAT/WHERE/WHEN/HOW/HOW MUCH/RISK) + 5 bands (Spend/Time/Complexity/Dependency/Scale)

DAY-OF FACILITATOR CHECKLIST

Before Kickoff (15 min)

- [] Print + laminate: 02QuickReference_Card.md (1 copy per facilitator)
- [] Print + laminate: 02BandScoringQuickRef.md (1 copy per facilitator)

- [] Print (not laminate): 02AdjudicationRules.md (1 copy per facilitator; reference during complex red-flags)
- [] Print (not laminate): 02BaseCaseFallbackBank.md (1 copy per facilitator; reference for fallback application)
- [] Print: 03ParticipantPackets/03zDecision_Worksheet.md (multiple copies per participant — at least 4 per person for 4 rounds)
- [] Print: 02ScoringBaselines.md (1 copy per facilitator; 11-industry baseline context)
- [] Sort Private Cards by round (from 05_Participant_Private_Cards/): Round 1 cards ready for distribution at setup
- [] Sort AI Adoption Arc handouts: Phases 2-4 ready for distribution at start of Rounds 2-4
- [] Prepare flip chart + markers (post explicit + fallback scores per round for 11 industries)
- [] Set up timer (phone OK; beep for Collective Bonus, Market Shock, and round transitions)
- [] Brief co-facilitator(s) on V7.4 individual participant model (5-11 participants; each selects 1+ industries; individual decisions; fallbacks post-scoring)
- [] Confirm participant industry assignments (each participant should know which industries they are covering)

During Each Round

Distribution phase (3 min):

- Distribute appropriate Private Cards for the round
- Distribute AI Adoption Arc phase handout (Rounds 2-4 only)

Scoring phase (12 min total):

- Reference 02BandScoringQuickRef.md for every explicit decision (60 sec/decision x variable industry count)
- Use 02ScoringBaselines.md to calibrate industry-specific context for 11 industries
- Post explicit scores on flip chart per industry (show bands + rationale)
- After explicit scoring: reference 02BaseCaseFallbackBank.md to apply fallbacks for industries without explicit participant decisions
- Post fallback scores on flip chart (show as separate line item)

Red-flag challenge (use language from 02AdjudicationRules):

- "I appreciate the direction. I need clarity on [specific issue]..."
- Reference band combination that triggered red-flag
- Offer narrower scope, longer timeline, lower spend, reduced complexity, or hedges
- "Which direction?"

Collective Bonus & Market Shock:

- Facilitate optional Collective Bonus during cross-industry discussion each round (~5 min): participants may nominate industries for +2 (strong) or -2 (risky). Takes effect only if 3+ participants agree.
- R2 only: Announce Facilitator Market Shock constraints on 2-3 industries (~3 min) before cross-industry discussion.
- After scoring + Collective Bonus (Rounds 2-4): Announce Industry Health conditions (~2 min)

End of Round (60 min debrief)

- Reference key debrief questions (02QuickReference_Card.md, section 13)
 - Hit top 3 questions; rest can be synthesis memos
 - Do NOT skip debrief; participants need processing time
-

QUICK SCORING WORKFLOW (Per Round)

1. Distribute materials (3 min): Private Cards + AI Adoption Arc handout (if applicable)
 2. Scenario read (5 min): Facilitator reads or presents scenario
 3. Individual participants prepare decisions (15 min): Each participant decides individually for their assigned industries
 4. Collect decisions (3 min): Late submissions -> auto-fallback
 5. Explicit scoring (8-12 min): Use 02BandScoringQuickRef.md; ~1-2 min per explicit decision; post on flip chart
 6. Apply fallbacks (2 min): Reference 02BaseCaseFallbackBank.md; post fallback scores for industries without explicit decisions
 7. Facilitator Market Shock (R2 only, 3 min): Facilitator imposes constraints on 2-3 industries from constraint menu
 - 7a. Industry Health Signals (2 min, Rounds 2-4): Announce condition bands based on cumulative scores. Apply constraints for Headwind/Crisis.
 8. Cross-industry discussion & results (27 min): Participants discuss cross-industry dynamics, challenge reasoning, surface spillovers. Use discussion prompts from 02Overviewand_Runbooks.
 9. Debrief (60 min total, non-negotiable): Individual reflections (15min) + Cross-industry discussion (25min) + No-regrets actions (20min)
-

COMMON FACILITATOR QUESTIONS ANSWERED

Q: How many participants do I need? A: Minimum 5, maximum 11. Ideal: 8+ industries assigned. Fewer possible -- facilitator plays unassigned industries. Each participant selects 1 or more industries (recommend 2). With 5 participants each covering 2 industries, you cover 10 of 11; the facilitator plays the remaining industry.

Q: How many decisions will I score per round? A: Variable. Depends on number of participants and how many industries each covers with explicit decisions. Expect 5-11 explicit decisions per round + fallbacks for remaining industries. Total scoring time: ~10-12 min.

Q: What if a participant submits decisions for multiple industries? A: Score each industry decision separately using that industry's baselines. E.g., a participant covering Finance and Law submitting both = 2 explicit scores, each on industry-specific merits.

Q: Can I negotiate fallback scores? A: No. Fallbacks are deterministic (not participant-controlled). Participants wanting higher scores for fallback industries must submit explicit decisions next round.

Q: What if a decision triggers a red-flag band combo? A: Challenge immediately. Reference the band combo + red-flag category (02AdjudicationRules, section "Red-Flag Triggers"). Offer narrower scope/longer timeline/lower spend/reduced complexity. Unlock +/-3 exception scoring if participant adjusts credibly or doesn't adjust (severity). Move on.

Q: Which file do I reference most during live play? A: **02BandScoringQuickRef.md** (laminated; on table). Use for band lookup, band-to-score translation, red-flag bands, fallback quick ref for 11 industries.

Q: How do I handle a vague decision? A: Reference specificity checklist (**02QuickReferenceCard, section 1** or **02Adjudication_Rules, section "Decision Specificity Checklist"**. Name the missing item(s). Don't score until all 7 items + 5 bands are clear.

Q: How is Consulting different from Law in scoring? A: Consulting is scored on copilot adoption, vertical AI expertise, pricing model transition, and junior talent pipeline dynamics. Law is scored on billable hour economics, bar rule compliance, malpractice liability, associate leverage model, and legal AI platform competition. Despite adjacency, they have completely different baselines, decision trees, and fallbacks.

Q: What about Big Tech and AI model development? A: Big Tech scope in this exercise covers cloud, advertising, devices, and enterprise software. AI lab and foundation model development is excluded from participant decisions and introduced via facilitator injects only. If a participant proposes model development decisions, redirect them to the in-scope areas.

Q: How is the cross-industry discussion phase different from V7? A: V7.4 retains the 27-min cross-industry discussion as the primary learning period. The facilitator now has discussion prompts (02Overviewand_Runbooks) organized by industry clusters and cross-industry tensions to surface dynamics between the 11 industries.

FINAL NOTES: V7.4 moves to 11 industries with flexible participant assignments (5-11 participants). Private Cards distributed from 05_Participant_Private_Cards/; AI Adoption Arc from 04_Participant_AI_Adoption_Arcs/. Reference **02BandScoringQuickRef.md** relentlessly. Apply **02BaseCaseFallbackBank** fallbacks post-explicit scoring. Debrief for 60 min (non-negotiable). Facilitate, don't debate.

Document Version: Project Threshold V7.4 — Facilitator Guide Index **Last Updated:** March 2026

Adjudication Rules

Individual Participant Decision Model (V7.4)

Critical Change: Each round, each participant makes **individual decisions** for their assigned industries (1 or more industries per participant; recommend 2). Unlike V6 (team deliberation), V7.4 has individual participants deciding for their industries directly.

Scoring Logic:

- Explicit Industry Decisions (If Submitted): Scored normally using banded framework ({-2, 0, +2} per dimension, +/-3 if red-flag)
- Industries Without Explicit Actions: Receive pre-defined base case fallback scores from fallback bank (small fixed deltas: +/-1 per dimension, deterministic, plausible). Fallbacks are automatic and deterministic.

Each industry decision is scored independently on the same 3 dimensions. There is no "sector-level" decision; all scoring happens at the industry level. The 11 industries are:

1. Retail — scored using retail-specific context (omnichannel, demand forecasting, inventory)
2. CPG — scored using CPG-specific context (brand management, R&D cycles, DTC)
3. Healthcare Provider — scored using provider-specific context (clinical AI, FDA, patient safety)
4. Healthcare Payer — scored using payer-specific context (claims automation, fraud detection, MLR)
5. Finance — scored using finance-specific context (underwriting, trading, fair lending, fraud)
6. Consulting — scored using consulting-specific context (copilot adoption, vertical AI expertise, pricing models, junior talent pipeline)
7. Law — scored using law-specific context (billable hour model, bar rule compliance, malpractice liability, associate leverage)
8. Manufacturing — scored using manufacturing-specific context (predictive maintenance, OT/IT integration, union relations)
9. Logistics — scored using logistics-specific context (route optimization, autonomous vehicles, driver adoption)
10. Big Tech — scored using Big Tech-specific context (cloud, ads, devices, enterprise software; excludes AI lab/model development)
11. B2B/B2C SaaS — scored using SaaS-specific context (AI feature integration, pricing pressure, startup disruption)

Participants decide individually; facilitators score each industry decision on its own merits, in its own context. Fallback industries do not require participant justification; facilitator applies the pre-defined score from the fallback bank.

Scoring Framework

Each **industry** decision is scored on three dimensions:

Dimension	Definition	Range
Strategic Fit	Does the decision align with industry fundamentals, competitive positioning, and the macro scenario?	-3 to +3
Execution Risk	Can the organization execute this decision within the timeframe? Are there organizational, technical, or capital constraints?	-3 to +3
Tail Risk	Does this decision expose the organization to downside scenarios (regulatory, competitive, reputational, systemic)?	-3 to +3

Total per decision: Sum of three dimensions (range: -6 to +6).

Banded Scoring During Play

Default scoring bands during live rounds:

Dimension	Typical Bands
Strategic Fit	{-2, 0, +2}
Execution Risk	{-2, 0, +2}
Tail Risk	{-2, 0, +2}

Why banded? Prevents false precision. Participants understand three clear buckets per dimension, not nine.

Exception: If a **red-flag/plausibility trigger** fires (see Section 2), unlock +/-3 scoring.

Scoring Guidance by Dimension

Strategic Fit (-3 to +3)

Score	Interpretation	Example
+2	Decision directly captures the core competitive opportunity or mitigates the core threat in this scenario.	Retail deploys AI demand forecasting when competitors doing same; Finance deploys AI underwriting when profitability depends on speed; Consulting builds vertical AI expertise when clients demand it.
0	Decision is aligned with industry fundamentals; reasonable strategic move but not game-changing.	CPG launches direct-to-consumer AI-driven personalization; Manufacturing invests in predictive maintenance; Law pilots AI-assisted research tools.
-2	Decision is misaligned with scenario; exposes participant to opportunity cost.	Big Tech holds on AI product integration during copilot adoption wave; Healthcare delays diagnostic AI investment during regulatory uncertainty; Law delays all AI adoption waiting for bar rule clarity.

Execution Risk (-3 to +3)

Execution Risk maps from banded inputs: **Spend/Commitment, Time-to-Impact, Execution Complexity, Dependency, and Scale.**

Score	Interpretation	Band Combination Example
+2	Decision is straightforward to execute; minimal organizational change; available capital and talent.	Absorbable spend + 0-3mo + Low complexity + Mostly internal + Pilot -> +2 Execution Risk
0	Decision is feasible but requires some organizational change and talent.	Material spend + 3-12mo + Medium complexity + Mostly internal + Regional -> 0 Execution Risk
-2	Decision faces material execution challenges; significant change management, talent gaps, or capex.	Material/Transformational spend + 1-2 years + High complexity + Vendor/Partner + National -> -2 Execution Risk
-3 (Exception)	Decision is extremely difficult to execute; severe constraint likely failure.	Transformational/Existential spend + 2+ years + Very high complexity + Ecosystem shift required + Global -> -3 Execution Risk

Examples:

- Deploying proven AI copilots (Absorbable, 0-3mo, Low, Internal, Pilot) = +2
- Enterprise-wide AI deployment with vendor integration (Material, 3-12mo, Medium, Vendor, National) = 0
- Building proprietary AI from scratch (Transformational, 2+ years, Very high, Ecosystem, Global) = -2 to -3

Tail Risk (-3 to +3)

Score	Interpretation	Example
+2	Decision includes hedges or defensive measures; upside potential with limited downside.	AI trading with circuit breakers; diagnostic system with human review; pilot-phase deployment; Law AI with human attorney review of all output.
0	Decision is neutral on tail risk; no clear upside or downside hedges.	Phased deployment with standard governance; moderate automation with retention plan.
-2	Decision exposes organization to material tail risk; likely to backfire if scenario changes.	Aggressive headcount cuts without severance; autonomous systems with no rollback; all-in bet on unproven technology; Law firm deploys AI-generated briefs without attorney review (malpractice exposure).

Band-to-Score Translation Reference

Use this table to quickly map banded inputs to expected Execution Risk scores:

Spend	Time	Complexity	Dependency	Scale	Typical Exec Risk
Absorbable	0-3mo	Low	Internal	Pilot	+2
Absorbable	0-3mo	Medium	Internal	Regional	+1
Material	0-3mo	Low	Internal	Regional	+1
Material	3-12mo	Medium	Internal	Regional	0
Material	3-12mo	High	Vendor	National	-1
Transformational	3-12mo	Medium	Vendor	National	-1
Transformational	1-2yr	High	Regulator/Union	National	-2
Transformational	2+yr	Very High	Ecosystem	Global	-3
Existential	2+yr	Very High	Ecosystem	Global	-3

Quick rule: Absorbable + 0-3mo + Low = +2. Each band downgrade (Material, Transformational, Existential) or timeline extension (3-12mo, 1-2yr, 2+yr) or complexity increase (Medium, High, Very High) or dependency widening (Vendor, Regulator, Ecosystem) shifts score down by roughly 0.5-1.0 per dimension.

Red-Flag Triggers (Band Combinations)

Red-flag triggers now reference **band combinations** instead of granular details:

Category	Red-Flag Bands	Example
Timeline Misalignment	Global scale + 0-3mo + Absorbable/Material spend	Launching globally in 3 months without pilot
Overcommitted Complexity	Transformational spend + Very High complexity + Ecosystem shift + 0-12mo	Building proprietary AI with ecosystem shift in <1 year
Regulatory Constraint	Any deployment requiring Regulator/Union buy-in without pre-defined pathway	Healthcare AI without FDA engagement plan; Law AI without bar rule compliance assessment
Unhedged Scale	Global/National scale + No pilot phase + High/Very High complexity	Full enterprise deployment of unproven tech without pilot
Existential Bet	Existential spend + 2+ years + Very High complexity + Ecosystem shift	Bet-the-company M&A or full business model pivot

When a red-flag band combination appears: Challenge the participant. Offer a narrower scope, longer timeline, lower complexity, or lower spend. Unlock +/-3 exception scoring only if participant

acknowledges and addresses the constraint.

Decision Specificity Checklist

Before scoring ANY decision, verify participant specified:

Item	Example
WHO	Which team owner? Committed sponsor? (Not "we'll figure it out later")
WHAT	What specific capability/action? (e.g., "GitHub Copilot," not "deploy AI")
WHERE	Scope: pilot / function / geography / enterprise? (Not "across the company")
WHEN	Timeline realistic for scope? Regulatory approval needed? (Not "as soon as possible")
HOW MUCH	Headcount, capex, revenue impact clear? (Not "minimal budget")
HOW	Talent plan? Integration detail? Rollback plan? (Not "we'll manage it")
RISK	Participant acknowledges execution and tail risk? (Not "this is foolproof")

Scoring rule: If >2 items are missing -> ask for specificity before scoring.

When to Challenge vs. When to Accept

Challenge (Ask for Clarification)

Ask for specificity if proposal is vague on WHO/WHAT/WHERE/WHEN/HOW/HOW MUCH/RISK.

Template: "I appreciate the direction. I need clarity on [specific issue]. Can you walk me through [detail]? That changes the execution risk profile significantly."

Accept (If Specificity Met)

Once all 7 items are clear, score the decision per framework above.

Red-Flag Triggers (When to Unlock +/-3)

Red-flag triggers are **plausibility gates** based on band combinations and strategic archetypes. If a red-flag fires, unlock +/-3 exception scoring (see *02PlausibilityDecision_Trees.md*).

Quantitative Requirements (Only Required When...)

Only quantify when the move is M&A, major capex build, or regulatory/legal commitment.

Otherwise, bands + justification is sufficient.

- M&A: Deal size, close timeline, integration plan required
- Major Capex: Capex amount, ROI, payback period required
- Regulatory/Legal Commitment: Regulatory timeline, legal exposure, severance obligations required
- Otherwise: Band classification + concise justification suffices

Red-Flag Categories (Band-Based)

Category	Trigger	Example
Timeline Misalignment	Global/National scale + 0-3mo timeline + Absorbable/Material spend	Launching globally in 3 months without pilot phase
	Regulator-dependent deployment + <12mo timeline	FDA or OCC approval assumed in under a year
	Pilot to enterprise scaling + <6mo without proven track record	Jumping from 50-store pilot to national rollout in 5 months
No Execution Risk Discussion	Missing talent acquisition plan (esp. Execution Complexity = High/Very High)	"We'll hire AI engineers" with no sourcing plan or timeline
	No integration roadmap (esp. Dependency = Vendor/Regulator/Ecosystem)	Vendor partnership assumed but no contractual or technical plan
	Unclear regulatory pathway (esp. regulated industry)	Healthcare or Finance AI deployment with no compliance strategy
	No data readiness assessment (esp. AI-heavy system)	ML model proposed without addressing data quality or availability
Industry Constraint Violations	Healthcare Provider AI + High/Very High complexity + no FDA engagement plan	Deploying clinical AI diagnostics without regulatory pathway
	Finance/Trading AI + High complexity + no circuit breakers/risk controls	AI trading system with no kill switch or position limits
	Retail/CPG + Headcount reduction >15% + no labor transition plan	Mass layoffs without severance, retraining, or redeployment
	Finance + AI underwriting + no bias auditing framework	Automated lending decisions without fair-lending compliance
	Law + AI-generated work product + no bar rule compliance + no attorney review	AI drafts filed with courts without human attorney review
	Law + AI-generated briefs without human review + malpractice unaddressed	Client-facing AI legal work with no liability framework
	Consulting + AI advisory services + no client confidentiality safeguards	AI tools processing client data without data-handling protocols

Category	Trigger	Example
	Consulting + >40% junior headcount reduction + no talent pipeline plan	Gutting the associate bench with no plan to develop future partners
Unhedged Tail Risk	Autonomous systems + National/Global scale + no rollback plan	Enterprise-wide autonomous process with no manual override
	Headcount reduction >30% + no severance/retention plan (M&A/major restructure only)	Major restructuring with no workforce transition support
	Full deployment + High/Very High complexity + no pilot phase	Skipping test-and-learn on a complex, unproven system
	Novel technology in mission-critical domain (Healthcare, Finance, Law) + no human oversight	AI making clinical, financial, or legal decisions without human review
Implausible Synergies (M&A only)	Deal size >\$5B + synergy targets >100% of acquisition cost	Claiming \$6B in synergies on a \$5B acquisition
	Integration complexity High/Very High + synergy realization <12mo	Expecting full integration benefits within a year of close

Facilitation Language for Industry-Level Decisions

When challenging a vague decision, reference the **industry** explicitly to keep context clear:

Challenging a Vague Decision (Industry-Specific)

"I want to score this Retail decision, but I need specificity. Right now you've said 'deploy AI across stores and online.' But I need to know:

- Where are you starting? Which 500 stores? All stores? Which e-commerce functions?
- How much are you spending on the Retail initiative? \$5M or \$50M?
- How will you staff it? Do you have retail tech talent, or do you need to hire?
- What's the rollback plan if customer trust erodes?

Once you answer those, I can score your **Retail** decision. Note: If you also submit a Law decision, it will be scored separately on Law-specific metrics (billable hour impact, bar compliance, malpractice exposure, etc.)."

Challenging a Red-Flag

"I appreciate the ambition here. But this proposal triggers a plausibility concern for me. Let me walk you through:

- You're proposing FDA approval in 12 months for a novel diagnostic system.
- FDA approval for novel diagnostics typically takes 18-24 months, even with a complete application.
- In 12 months, you could submit an application or run a pilot with investigational status, but not full deployment.

Here's what I can score: Proposing a pilot with human oversight and FDA submission (not approval). That's Strategic Fit +1 (good direction), Execution Risk -1 (regulatory timeline is long), Tail Risk 0 (human oversight limits downside). Total: 0/6.

Or, if you want higher execution feasibility, propose a different approach. What would you prefer?"

Accepting a Strong Industry Decision

"Okay, I have all the specificity I need for your **Retail** decision. Let me score:

- Strategic Fit: +2 -- Direct response to inventory margin loss; aligns with omnichannel strategy.
- Execution Risk: +1 -- Demand forecasting AI is proven; you have experience; 500-store pilot is feasible.
- Tail Risk: +1 -- Phased pilot allows learning; if brand backlash emerges, you can narrow scope.
- Total: +4/6 -- This is a strong Retail decision. Expected probability of success: 70%+. If you also submit a Consulting decision, it will be scored separately."

Scoring Reference Table

Total Score	Interpretation	Action
+5 to +6	Strong decision; expect 70%+ success probability.	Accept and celebrate. Participant is executing well.
+1 to +4	Acceptable decision; execution risk is real; monitor closely.	Accept; add to watch list; expect 50-70% success probability.
-2 to 0	Weak decision; reframe as narrower scope or defer.	Challenge; ask participant to narrow scope or add hedges.
<-2	Challenge; likely implausible; ask participant to narrow or add hedges.	Reject or reframe significantly. Low probability of success.

Base Case Fallback Scoring

After scoring all explicit industry decisions, apply **base case fallbacks** to any industry that did **not** receive an explicit action from any participant.

Fallback Scoring Process:

1. Identify industries without explicit actions (variable count based on how many industries each participant covers)
2. Reference the fallback bank (see 02BaseCaseFallbackBank.md)
3. Apply the pre-defined fallback score for that industry (deterministic, small fixed deltas: +/-1 per dimension, plausible)
4. Fallback industries are applied automatically (no participant input required)
5. Post fallback score alongside explicit decisions for transparency

Example: Participant covering Retail and CPG submits explicit Retail decision (+2/6). CPG receives no explicit action. Reference fallback bank -> CPG falls back to {0, +1, 0} = +1/6 (defensive but plausible cost-reduction move). Post both scores.

Fallback scoring is **deterministic**, not participant-controlled. Facilitator applies it directly from the fallback bank.

Industry Health Signals (End of Round)

- During play: Score explicit decisions only (3 dimensions, sum per decision).
 - Apply base case fallbacks after explicit scoring (reference fallback bank).
 - End of round: Update cumulative aggregate score for each industry (running total across all rounds). Look up Industry Health condition (Surge/Tailwind/Steady/Headwind/Crisis) from 02IndustryHealthSignalTables.md.
 - Announce conditions at start of next round (~2 min). Apply Headwind/Crisis constraints as applicable.
-

Summary Checklist (Per Explicit Decision)

- Industry identified (Retail? CPG? Finance? Healthcare Provider? Consulting? Law? etc.)
- Participant specified all 7 items (WHO/WHAT/WHERE/WHEN/HOW/HOW MUCH/RISK)
- Band mapping clear (Spend/Commitment, Time-to-Impact, Complexity, Dependency, Scale)
- No red-flag band combinations (if yes, unlock +/-3 scoring and reference decision trees)
- Only quantify if M&A, major capex, or regulatory/legal commitment; otherwise bands + justification suffices
- Score Strategic Fit: {-2, 0, +2} (or +/-3 if red-flag)
- Score Execution Risk: {-2, 0, +2} (or +/-3 if red-flag) -- reference band translation table
- Score Tail Risk: {-2, 0, +1} (or +/-3 if red-flag)
- Sum total (range: -6 to +6)
- Post score and brief rationale to participant
- After round: Apply base case fallbacks to industries without explicit actions (reference 02BaseCaseFallbackBank)

- [] End of round: Update cumulative scores and look up Industry Health conditions (reference 02IndustryHealthSignalTables)

Band-to-Score Quick Reference

PRINT THIS. Laminate. Keep at hand during live play. Reference during scoring decisions (60-120 sec/decision).

V7.4 Changes: 11 industries (up from 10 personas/5 sectors). Individual participant model (5-11 participants, each selects 1+ industries). One decision worksheet per industry per round.

1. FOUR BANDS & THEIR OPTIONS

Band	Options
Spend/Commitment	Absorbable Material Transformational Existential
Time-to-Impact	0-3 months 3-12 months 1-2 years 2+ years
Execution Complexity	Low Medium High Very High
Dependency	Internal Vendor Regulator/Union Ecosystem
Scale	Pilot Regional National Global

2. BAND-TO-SCORE TIER SYSTEM (Execution Risk Primary)

Use this table to map band combinations -> Execution Risk score:

Spend	Time	Complexity	Dependency	Scale	Typical Exec Risk
Absorbable	0-3mo	Low	Internal	Pilot	+2
Absorbable	0-3mo	Medium	Internal	Regional	+1
Material	0-3mo	Low	Internal	Regional	+1
Material	3-12mo	Medium	Internal	Regional	0
Material	3-12mo	High	Vendor	National	-1
Transformational	3-12mo	Medium	Vendor	National	-1
Transformational	1-2yr	High	Regulator	National	-2
Transformational	2+yr	Very High	Ecosystem	Global	-2
Existential	2+yr	Very High	Ecosystem	Global	-3

Quick rule: Absorbable + 0-3mo + Low = +2. Each band downgrade shifts score down ~0.5-1.0.

3. RED-FLAG BAND COMBINATIONS (+/-3 Exception Scoring)

When these band combos appear: **Challenge the participant. Unlock +/-3 scoring.**

Red-Flag	Band Trigger	Action
Timeline Misalignment	Global + 0-3mo (no pilot)	Offer narrower scope or +3mo timeline
Overcommitted Complexity	Transformational + Very High + Ecosystem + <12mo	Propose pilot phase; lower complexity
Regulatory Constraint	Any regulatory deployment + <12mo timeline (unrealistic cycle)	Push FDA/regulator/bar pre-submission; extend timeline
Unhedged Scale	National/Global + High/Very High complexity + No pilot	Demand pilot phase; reduce initial scope
Existential Bet	Existential spend + 2+yr + Very High complexity + Ecosystem	Ask participant: reframe as narrower scope or accept -3 risk

When fired: Stop; ask participant to reframe (narrower scope, longer timeline, lower spend, reduce complexity, add hedges). If participant doesn't adjust -> score -2 to -3 on Execution Risk. If participant adjusts credibly -> may unlock +2 or +1 on Exec Risk.

4. 15-SECOND SCORING DECISION TREE

For each explicit decision:

1. Specificity check (5 sec): WHO/WHAT/WHERE/WHEN/HOW/HOW MUCH/RISK + BANDS clear?
 - If NO (>2 missing) -> Ask for clarification. Don't score yet.
2. Band plausibility (5 sec): Any red-flag combo above?
 - If YES -> Challenge; offer narrower scope. Unlock +/-3 if band red-flag fires.
 - If NO -> Proceed to scoring.
3. Score three dimensions (50 sec total):
 - Strategic Fit: -2 / 0 / +2 (or +/-3 if red-flag)
 - Execution Risk: -2 / 0 / +2 (or +/-3) -- Use band translation table
 - Tail Risk: -2 / 0 / +2 (or +/-3 if red-flag)
 - Total = Strat Fit + Exec Risk + Tail Risk
4. Post & explain (5 sec): "Score: +2/6. Reason: [band combo that drove score]."

5. WORKED EXAMPLES (Band Selection -> Score)

Example 1: Retail Demand Forecasting Deployment

Participant decision: "Deploy AI demand forecasting to 500 stores, 90 days, \$2M budget, existing vendor."

Bands:

Band	Selection
Spend	Absorbable
Time	0-3mo
Complexity	Low
Dependency	Vendor
Scale	Regional (500 stores)

Band combo: Absorbable + 0-3mo + Low + Vendor + Regional

Lookup: Similar to row 3 (Material, 0-3mo, Low, Internal, Regional) but Absorbable downgrade = slightly better.

Scoring:

- Strategic Fit +2: Directly addresses margin-defense opportunity vs. Amazon; proven tech.
- Execution Risk +1: Band combo suggests +1-+2; proven tech + existing vendor = +1 (conservative).
- Tail Risk 0: Phased scope limits downside; pilot approach reversible.
- Total: +3/6 (check)

Explain: "Strong score. Bands (Absorbable + 0-3mo + Low) are straightforward to execute. Execution Risk +1 reflects proven tech + vendor partnership."

Example 2: Healthcare FDA Diagnostic AI Submission

Participant decision: "Initiate FDA pre-submission for radiology AI; plan 12-month clinical validation; \$8M budget; will engage FDA early."

Bands:

Band	Selection
Spend	Material
Time	3-12mo
Complexity	Very High
Dependency	Regulator
Scale	Pilot (select hospital system)

Band combo: Material + 3-12mo + Very High + Regulator + Pilot

Red-flag check? Regulator-dependent + 3-12mo timeline approaches constraint (FDA approval typically 18-24mo). But participant is proposing pre-submission + validation (not approval). **Plausible with caveats.**

Scoring:

- Strategic Fit +1: Strategic value (clinical AI is high-impact); longer timeline acceptable.
- Execution Risk -1: Regulatory dependency + Very High complexity = constraint. But FDA pre-submission + pilot scope mitigates. -1 (not -2) because participant has regulatory pathway.
- Tail Risk 0: Human review mandatory; pilot scope; FDA engagement reduces surprise risk.
- Total: 0/6 (check)

Explain: "Plausible move but execution risk is real. Band combo (Regulator + Very High + 3-12mo) hits regulatory constraint. You've mitigated with pre-submission + pilot, so Exec Risk -1 (not -2). Strategic Fit +1 (clinical value) + Tail Risk 0 (human oversight) = 0/6 total."

Example 3: Law Firm AI-Assisted Brief Drafting

Participant decision: "Deploy AI-assisted brief drafting across litigation practice group; mandatory attorney review and sign-off on all output; bar rule compliance assessment per jurisdiction; \$5M annual licensing budget."

Bands:

Band	Selection
Spend	Material
Time	3-12mo
Complexity	High
Dependency	Regulator (bar rules)
Scale	Regional (one practice group)

Band combo: Material + 3-12mo + High + Regulator + Regional

Red-flag check? Regulator-dependent (bar rules) + 3-12mo. Bar rules are evolving but not approval-gated like FDA. Attorney review protocol is the key safeguard. **Plausible if attorney review is mandatory.**

Scoring:

- Strategic Fit +1: Addresses efficiency demands from clients; proven tools available; positions firm for future AI adoption.
- Execution Risk -1: Bar rule complexity across jurisdictions + High complexity. But mandatory attorney review + one practice group scope mitigates. -1 (not -2).
- Tail Risk +1: Mandatory attorney review eliminates malpractice exposure; bar compliance assessment proactive; limited scope allows learning.
- Total: +1/6 (check)

Explain: "Defensible move. Bar rule dependency adds complexity, but attorney review is the right safeguard. Execution Risk -1 reflects regulatory uncertainty, offset by Strategic Fit +1 (client demand) and Tail Risk +1 (malpractice managed). +1/6 total."

6. COMMON ERRORS TO AVOID

Error	What Happens	Fix
Vague band mapping	"Transformational spend" without timeline/complexity specified.	Require all 5 bands clear before scoring.
Band override	Participant says "Exec Risk +2" even though bands suggest -1.	Bands drive score, not participant preference. Explain your band logic.
Missing red-flag	Global + 0-3mo slips through without challenge.	Reference red-flag table explicitly. Challenge band combos before scoring.
Tail risk underweighting	Autonomous systems scored at +1 with "no rollback" mentioned.	If Very High complexity + no pilot/rollback = -2 minimum on Tail Risk.

Error	What Happens	Fix
Allowable overcommitment	Material spend + Transformational complexity; participants argue "we'll find a way."	Timeline/complexity mismatch = material execution risk. Score -1 to -2 on Exec Risk; don't apologize.
Conflating industries	Scoring a Consulting decision using Law-specific criteria or vice versa.	Each industry scored on its own merits with its own context. Check the industry baseline (02ScoringBaselines).

7. FALLBACK QUICK REFERENCE (11 Industries, Base Case Scores)

If an industry receives NO explicit decision from any participant, apply deterministic fallback from this table:

Industry	Fallback Strat Fit	Fallback Exec Risk	Fallback Tail Risk	Fallback Total
Retail	0	+1	0	+1
CPG	0	+1	0	+1
Healthcare Provider	0	0	+1	+1
Healthcare Payer	+1	+1	0	+2
Finance	+1	+1	+1	+3
Consulting	+1	+1	0	+2
Law	0	+1	+1	+2
Manufacturing	0	+1	+1	+2
Logistics	+1	+1	0	+2
Big Tech	+1	+1	0	+2
B2B/B2C SaaS	0	+1	0	+1

Example: Participant submits explicit Retail decision (+2/6). CPG receives no explicit action. Apply fallback: CPG = +1/6 (deterministic). Retail +2/6 posted as explicit; CPG +1/6 posted as fallback.

8. INDUSTRY-SPECIFIC WATCH-OUTS (Baseline Anchors)

Industry	Watch For	High-Risk Band Combo
Retail	Brand backlash on aggressive personalization; Amazon competition	Global + Very High complexity + <6mo
CPG	Retailer retaliation on aggressive DTC; AI-generated content brand safety	DTC + Transformational spend + <12mo + Ecosystem
Healthcare Provider	Patient safety liability; physician adoption risk; FDA timeline	Diagnostic AI + Very High complexity + <12mo + no FDA pre-submission
Healthcare Payer	Member backlash on aggressive claim denials; regulatory MLR limits	>70% auto-denial + <6mo + no human review + Existential
Finance	Systemic risk if trading AI fails; fair lending litigation	Autonomous trading + Existential spend + 2+ years + no oversight
Consulting	Junior talent pipeline disruption; pricing pressure from AI boutiques; client in-house AI competition	40%+ junior cuts + <12mo + no vertical upskilling
Law	Bar rule uncertainty; malpractice liability from AI-generated work; billable hour model disruption; associate leverage erosion	AI briefs filed without attorney review; >30% associate cuts + <12mo; firm-wide pricing change + <6mo
Manufacturing	Union relations; labor displacement without transition plan	>30% headcount cut + <6mo + no severance plan
Logistics	Driver resistance; last-mile profitability limits; AV regulatory uncertainty	Full AV fleet + <2yr + Global + no union negotiation
Big Tech	Antitrust scrutiny; margin pressure from CapEx inflation; open-source competition. Excludes AI lab/model dev.	Existential capex + 2+ years + Global + no competitive hedge
B2B/B2C SaaS	Pricing backlash on AI bundling; startup disruption; margin compression	50%+ cost increase (AI) + <6mo + pricing unchanged + Transformational

9. QUICK REFERENCE: SCORING BANDS SUMMARY

Execution Risk: Use band translation table (Section 2).

Strategic Fit: Quick calibration

- +2: Captures core competitive opportunity or mitigates core threat.
- 0: Neutral; reasonable move, not game-changing.
- -2: Misaligned with scenario; opportunity cost.

Tail Risk: Quick calibration

- +2: Includes hedges (pilot, human review, rollback plan).
- 0: Neutral; standard risk management.
- -2: Exposed to downside; no hedges.

TOTAL SCORING RANGE: -6 to +6 (typical: -2 to +6)

Score	Interpretation	Action
+5 to +6	Strong (70%+ success)	Accept. Celebrate.
+1 to +4	Acceptable (50-70% success)	Accept. Monitor.
-2 to 0	Weak	Challenge. Ask to narrow scope.
<-2	Poor. Red-flag fired.	Reject or reframe significantly.

PRINT. LAMINATE. REFERENCE DURING EVERY DECISION. FACILITATE BRISKLY.

Base Case Fallback Bank

Overview

This document defines **base case fallback scores** for industries that do **not** receive explicit actions from participants in a given round.

Key Principles:

- Deterministic: Same fallback applies to all instances of an industry across all participants and rounds
- Small deltas: Typically +/-1 per dimension (not -2 to +2), creating scores in the -2 to +2 range
- Plausible: Represent defensive but reasonable moves (e.g., cost control, operational efficiency, maintain status quo)
- Automatic: Applied by facilitator directly; no participant input required
- Not participant-controlled: Fallback scores are fixed; participants cannot modify them by arguing or adjusting

When to apply: After scoring all explicit industry decisions, identify any industry (across all 11 industries) that received no explicit action. Apply its fallback score directly from this bank.

11 Industries (V7.4): Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS

Fallback Scores by Industry

Retail (Omnichannel Retailers)

Fallback Rationale: Maintain pilot-phase demand forecasting in select stores; optimize inventory within existing systems; no new capex or major staffing changes.

Dimension	Score	Justification
Strategic Fit	0	Neutral move; maintains competitive position but doesn't capture opportunity
Execution Risk	+1	Proven tech; existing infrastructure; low change management burden
Tail Risk	0	Pilot scope limits downside; if adoption slows, can pause
TOTAL	+1/9	Defensive; preserves optionality

CPG (Consumer Goods Manufacturer)

Fallback Rationale: Accelerate minor R&D cycles using existing tools; minor marketing automation with human review; no DTC expansion.

Dimension	Score	Justification
Strategic Fit	0	Neutral; cost efficiency moves, but doesn't drive growth
Execution Risk	+1	Existing teams; minor process changes; low adoption risk
Tail Risk	0	Conservative scope; brand safety maintained via human review
TOTAL	+1/9	Defensive operational efficiency

Healthcare Provider (Hospital System, Clinical Care)

Fallback Rationale: Deploy operational AI (scheduling, resource allocation); defer major clinical diagnostic AI; maintain standard clinical governance.

Dimension	Score	Justification
Strategic Fit	0	Operational efficiency; doesn't drive clinical outcomes or revenue growth
Execution Risk	0	Operational AI (scheduling) proven; workflow integration standard; physician adoption manageable
Tail Risk	+1	Operational scope limits patient safety risk; governance framework maintains liability protection
TOTAL	+1/9	Defensive operational efficiency without clinical risk escalation

Healthcare Payer (Health Insurer, Claims & Coverage)

Fallback Rationale: Deploy fraud detection + payment integrity AI; continue prior authorization automation with human review; no aggressive claim denial.

Dimension	Score	Justification
Strategic Fit	+1	Fraud detection improves medical loss ratio; proven ROI
Execution Risk	+1	Fraud detection proven; payment integrity automation standard; compliance clear
Tail Risk	0	Human review safeguards prevent denial backlash; regulatory MLR compliance maintained
TOTAL	+2/9	Defensive cost-control with member satisfaction safeguards

Finance (Bank + Insurance)

Fallback Rationale: Continue current fraud detection arms race; invest in bias auditing and fair lending compliance infrastructure; no major new deployment.

Dimension	Score	Justification
Strategic Fit	+1	Defensive but strategically important; reduces regulatory and litigation tail risk
Execution Risk	+1	Compliance infrastructure proven; existing vendor relationships
Tail Risk	+1	Proactive compliance reduces regulatory backlash; enables future deployment
TOTAL	+3/9	Defensive risk mitigation (higher than other industries due to systemic tail risk reduction)

Consulting (Big Four / MBB)

Fallback Rationale: Deploy copilots to research + analysis functions (proven use case); maintain junior hiring at current levels; defer vertical AI expertise build and pricing model changes.

Dimension	Score	Justification
Strategic Fit	+1	Copilots improve delivery efficiency; maintains junior talent pipeline
Execution Risk	+1	Copilot adoption straightforward; existing vendor relationships (GitHub Copilot, internal tools)
Tail Risk	0	Junior hiring preserved; no aggressive commoditization; talent pipeline maintained; pricing stability
TOTAL	+2/9	Defensive productivity move; preserves talent model and client relationships

Law (AmLaw 50 Firm)

Fallback Rationale: Pilot AI research tools (legal research, due diligence) in 2-3 practice groups with mandatory attorney review; maintain billable hour model; maintain current associate hiring levels; defer pricing model changes; monitor bar rule developments across jurisdictions.

Dimension	Score	Justification
Strategic Fit	0	Conservative approach; captures minor efficiency gains but doesn't transform delivery model or address pricing pressure
Execution Risk	+1	AI research tools proven (Harvey.ai, CoCounsel, Westlaw AI); attorney review protocol manageable; limited scope reduces adoption friction
Tail Risk	+1	Mandatory attorney review eliminates malpractice exposure from AI-generated work; bar rule compliance maintained; billable model stable near-term; no associate displacement
TOTAL	+2/9	Conservative but defensible; preserves current economics while building AI familiarity

Manufacturing (Heavy Manufacturing)

Fallback Rationale: Selective predictive maintenance in highest-ROI plants (pilot phase); continue "no-layoff" retraining commitment; no aggressive automation.

Dimension	Score	Justification
Strategic Fit	0	Neutral; cost reduction at 2-3 plants, but not enterprise-wide transformation
Execution Risk	+1	Predictive maintenance proven; phased OT/IT integration manageable
Tail Risk	+1	Union cooperation maintained via "no-layoff" commitment; labor relations preserved
TOTAL	+2/9	Conservative but labor-cooperative

Logistics (Freight/3PL/Warehouse)

Fallback Rationale: Deploy route optimization to subset of fleet (pilot); driver engagement program; defer autonomous vehicle decisions.

Dimension	Score	Justification
Strategic Fit	+1	Route optimization proven ROI; \$45-50M annual savings potential (limited scope)
Execution Risk	+1	Technology proven; driver adoption managed via engagement program
Tail Risk	0	Pilot scope; union concerns addressed; no autonomous displacement
TOTAL	+2/9	Proven, labor-safe cost reduction

Big Tech (Google/Meta/Microsoft/Amazon-Class — Cloud, Ads, Devices, Enterprise Software)

Fallback Rationale: Continue incremental AI infrastructure investments for cloud and enterprise products; integrate AI into existing products cautiously; defer transformational new product launches. (Excludes foundation model development decisions.)

Dimension	Score	Justification
Strategic Fit	+1	Incremental AI product features maintain competitive position; defensive move in cloud and enterprise
Execution Risk	+1	AI feature integration proven; existing infrastructure and talent; low risk
Tail Risk	0	Incremental scope limits regulatory/antitrust risk; margins preserved via cost management
TOTAL	+2/9	Defensive infrastructure + product evolution

B2B/B2C SaaS (Workday/Salesforce-Class)

Fallback Rationale: Bundle AI copilots into standard product SKU; maintain pricing at current levels; defer AI-native product launches.

Dimension	Score	Justification
Strategic Fit	0	Bundling maintains customer lock-in; doesn't drive premium pricing or market share
Execution Risk	+1	AI feature bundling proven; existing infrastructure and vendor relationships
Tail Risk	0	Standard bundling avoids pricing backlash; customer churn risk manageable at current pricing
TOTAL	+1/9	Defensive feature parity with margin pressure

How to Apply Fallback Scoring

Example 1: Participant Covering Retail + CPG Submits Only Retail

Decision

1. Participant submits explicit decision: Retail (e.g., "Deploy demand forecasting to 500 stores, 90 days, \$2M budget")
2. Retail scored normally: (e.g., +2/9)
3. CPG receives no explicit action
4. Fallback applied: CPG fallback from bank = {0, +1, 0} = +1/9
5. Both posted: Retail +2/9 (explicit), CPG +1/9 (fallback)

Example 2: Participant Covering Healthcare Provider + Payer Submits Only Provider Decision

1. Participant submits explicit decision: "Deploy clinical decision support with physician override; staff training"
2. Healthcare Provider scored normally: (e.g., +2/9)
3. Healthcare Payer receives no explicit action
4. Fallback applied: Healthcare Payer fallback = {+1, +1, 0} = +2/9
5. Both posted: Provider +2/9 (explicit), Payer +2/9 (fallback)

Example 3: Participant Covering Finance + Consulting + Law Submits Only Finance Decision

1. Participant submits explicit decision: Finance (e.g., "Deploy AI underwriting with explainability + bias audit")
2. Finance scored normally: (e.g., +3/9)
3. Consulting and Law receive no explicit actions
4. Fallbacks applied: Consulting = {+1, +1, 0} = +2/9; Law = {0, +1, +1} = +2/9
5. All three posted: Finance +3/9 (explicit), Consulting +2/9 (fallback), Law +2/9 (fallback)

Example 4: Minimum Model (5 Participants, Each Covering 2 Industries)

Scenario: 5 participants, each assigned 2 industries. Each submits 1 explicit decision; 1 industry per participant falls back. 1 industry unassigned (auto-fallback).

1. Participant A (Retail + CPG): Retail decision +2/9 -> CPG fallback +1/9
 2. Participant B (Healthcare Provider + Payer): Provider decision +1/9 -> Payer fallback +2/9
 3. Participant C (Finance + Consulting): Finance decision +3/9 -> Consulting fallback +2/9
 4. Participant D (Manufacturing + Logistics): Manufacturing decision +1/9 -> Logistics fallback +2/9
 5. Participant E (Big Tech + SaaS): Big Tech decision +2/9 -> SaaS fallback +1/9
 6. Law (unassigned): Auto-fallback +2/9
-

Fallback Summary Table (11 Industries)

Industry	Fallback Strat Fit	Fallback Exec Risk	Fallback Tail Risk	Fallback Total	Rationale
Retail	0	+1	0	+1	Pilot-phase demand forecasting
CPG	0	+1	0	+1	Minor R&D + marketing automation
Healthcare Provider	0	0	+1	+1	Operational AI (scheduling) + governance
Healthcare Payer	+1	+1	0	+2	Fraud detection + prior auth with human review
Finance	+1	+1	+1	+3	Compliance + fraud defense infrastructure
Consulting	+1	+1	0	+2	Copilot deployment + talent preservation
Law	0	+1	+1	+2	Pilot AI research tools + maintain billable model + attorney review
Manufacturing	0	+1	+1	+2	Selective predictive maintenance + union cooperation
Logistics	+1	+1	0	+2	Route optimization pilot + driver engagement
Big Tech	+1	+1	0	+2	Incremental AI infrastructure + product features

Industry	Fallback Strat Fit	Fallback Exec Risk	Fallback Tail Risk	Fallback Total	Rationale
B2B/B2C SaaS	0	+1	0	+1	AI copilots bundled into standard product

Facilitator Checklist: Applying Fallbacks

After scoring all explicit decisions:

- [] Identify industries across all participants that received no explicit action
- [] For each fallback industry, reference the table above and note the deterministic fallback score
- [] Post fallback score for each industry (e.g., "Law falls back to +2/9 -- pilot AI research tools, maintain billable model")
- [] Post all scores with transparency (show explicit vs. fallback breakdown per industry)
- [] Update cumulative scores and look up Industry Health conditions (reference 02IndustryHealthSignalTables)

Important: Do NOT negotiate fallbacks. They are deterministic. Do NOT allow participants to argue for higher/lower fallbacks. If a participant wants to improve an industry's score, they must submit an explicit decision in the next round.

Reference: Use this bank during scoring (02AdjudicationRules.md). Participants consult this implicitly during decision planning (see 02Overviewand_Runbooks.md for round structure).

Industry Health Signal Tables

Overview

This file replaces the former Valuation Translation Pipeline. Instead of calculating margin impacts, multiple shifts, and basis-point portfolio returns, facilitators use a simple score-to-condition mapping to give participants tangible feedback on their industry's structural health.

Key concept: Cumulative aggregate score → Industry Condition → Mechanical consequence for next round.

Key Changes (V7.4 IC Redesign):

- Valuation translation pipeline eliminated (no margin impact, multiple impact, probability of success, or portfolio impact calculations)
- Replaced with 5 Industry Condition bands based on cumulative aggregate score
- Conditions carry mechanical consequences that affect participant decisions
- Historical reference data consolidated in [07_Supporting_Materials/07_Historical_Reference_Data.md](#)

Industry Condition Bands

Conditions are based on **cumulative aggregate score** (sum of all round scores to date for that industry). This creates momentum: good early decisions build advantage; poor early decisions compound into real trouble.

Facilitator announces conditions at the start of each round (starting Round 2). Takes ~2 minutes. No calculation required — just look at the running score total.

Cumulative Score	Condition	Facilitator Announcement	Mechanical Consequence
+15 or higher	Surge	"Your industry is structurally advantaged. Margins expanding, multiple lifting. Investors are overweight. You have room to be ambitious."	Participant may attempt one Transformational-band decision without triggering automatic red-flag review. (Still scored normally on all 3 dimensions — just no automatic challenge on the Transformational band.)

Cumulative Score	Condition	Facilitator Announcement	Mechanical Consequence
+6 to +14	Tailwind	"Your industry is navigating well. Margins stable-to-improving. Market views you favorably."	No modifier. Business as usual.
-5 to +5	Steady	"Your industry is holding position. No clear momentum. Market is neutral."	No modifier.
-6 to -14	Headwind	"Your industry is under structural pressure. Margins compressing. Investors cautious. You need to stabilize before you can grow."	Facilitator imposes one constraint from the constraint menu at start of next round. Participant must address it before submitting primary decision.
-15 or worse	Crisis	"Your industry faces structural decline. EBITDA collapsing, multiple in freefall. Talent and capital fleeing. Survival mode."	Two constraints imposed. Primary decision must include a defensive or restructuring component — pure growth plays score automatic -2 Strategic Fit.

Constraint Menu (Used by Facilitator Market Shock in Round 2 AND by Industry Health Signals)

When a Headwind or Crisis condition triggers a constraint, the facilitator selects from this menu:

1. Regulatory Halt: Government regulators have announced a review of AI deployments in your industry. All new AI projects are effectively frozen for this round — you cannot launch new pilots, expand existing deployments to new sites, or go live with tools that are still in testing. Decision impact: Your strategy this round must work within your current AI footprint. You can optimize what's already deployed, invest in workforce development, or reposition strategically — but you cannot expand your AI capability set.

2. **Labor Cost Surge:** Union negotiations, wage competition, or mandatory retraining obligations have driven your labor costs sharply higher. The increase is material enough to compress margins and force trade-offs elsewhere in your budget. Decision impact: Any strategy that depends on labor-intensive rollout (retraining programs, change management, multi-site deployment teams) now costs significantly more. You may need to narrow the scope of your plans, slow your rollout timeline, or find ways to offset the cost pressure — but you cannot ignore it.
3. **Capital Tightening:** Market volatility, credit conditions, or investor skepticism have significantly reduced your available capital for this round. Your board has pulled back discretionary spending authorization. Decision impact: You cannot pursue capital-intensive strategies at the scale you might otherwise choose. Large infrastructure investments, acquisitions, and multi-site expansions are constrained. Your strategy must either be capital-light or demonstrate why a smaller, more focused investment is the right use of limited funds.
4. **Reputational Pressure:** Public backlash — whether from workforce displacement concerns, algorithmic fairness controversies, or customer trust erosion — has forced your industry into a defensive posture. Stakeholders expect visible action on responsible AI practices. Decision impact: Your strategy this round must credibly address the reputational concern. Aggressive AI expansion without a visible "responsible transition" component will be scored harshly. You may need to slow deployment, invest in workforce support, or reframe your narrative — but ignoring the backlash is not a viable option.
5. **Competitive Response:** A major competitor has used AI to aggressively undercut your pricing, poach your customers, or leapfrog your product capability. The competitive threat is immediate and material. Decision impact: Your strategy must account for active competitive pressure. Purely internal-facing or long-horizon strategies that don't address the near-term market share threat are risky. You may need to defend your position through pricing, product differentiation, partnerships, or accelerated deployment — but you cannot plan as if the competitive landscape is stable.
6. **Litigation Risk:** Your industry faces active or imminent litigation related to AI deployment — algorithmic bias claims, labor displacement lawsuits, intellectual property disputes, or product liability actions. Legal exposure is material and uncertain. Decision impact: Any strategy that expands your litigation surface area (aggressive automation, opaque algorithmic decision-making, rapid workforce reduction) carries additional risk this round. You may need to slow certain deployments, invest in governance and compliance infrastructure, or choose lower-risk AI applications — but the legal overhang must be reflected in your reasoning.

Selection guidance for facilitators:

- For Headwind: Choose the constraint most narratively appropriate to the industry's situation and recent decisions. Aim for pressure that creates an interesting decision, not punishment.
- For Crisis: Choose two constraints. One should address the primary failure mode; the other should reflect market consequences (e.g., Capital Tightening + Competitive Response).

How to Use Industry Health Signals

End-of-Round Procedure (Replaces Former Valuation Translation)

1. Score all explicit industry decisions using banded scoring ($\{-2, 0, +2\}$ per dimension, ± 3 if red-flag).
2. Apply base case fallbacks for industries without explicit decisions (reference 02BaseCaseFallbackBank).
3. Update cumulative score for each industry (running total across all rounds).
4. Determine Industry Condition by looking up cumulative score in the table above.
5. Announce conditions at start of next round (takes ~2 minutes).
6. Apply constraints for any industry in Headwind or Crisis before decision preparation begins.

Tracking Template

Use a simple flip chart or spreadsheet:

Industry	R1 Score	R2 Score	R3 Score	R4 Score	Cumulative	Condition
Retail	—	—	—	—	—	—
CPG	—	—	—	—	—	—
Healthcare Provider	—	—	—	—	—	—
Healthcare Payer	—	—	—	—	—	—
Finance	—	—	—	—	—	—
Consulting	—	—	—	—	—	—
Law	—	—	—	—	—	—
Manufacturing	—	—	—	—	—	—
Logistics	—	—	—	—	—	—
Big Tech	—	—	—	—	—	—
B2B/B2C SaaS	—	—	—	—	—	—

Condition Dynamics Across Rounds

Typical Trajectories

- Strong performer: R1 +4, R2 +5 → Cumulative +9 (Tailwind). R3 +4 → Cumulative +13 (Tailwind). R4 +3 → Cumulative +16 (Surge). Participant earns freedom to attempt bold moves.

- Average performer: R1 +2, R2 +1 → Cumulative +3 (Steady). R3 +2 → Cumulative +5 (Steady). R4 +1 → Cumulative +6 (Tailwind). Gradual positive momentum.
- Struggling performer: R1 -2, R2 -3 → Cumulative -5 (Steady). R3 -2 → Cumulative -7 (Headwind). Constraint imposed in R4.
- Crisis trajectory: R1 -4, R2 -5 → Cumulative -9 (Headwind). Constraint in R3. R3 -4 → Cumulative -13 (Headwind). More constraints in R4. R4 -3 → Cumulative -16 (Crisis). Sustained poor decisions required to reach Crisis.

Design Intent

- Momentum matters. Cumulative scoring mirrors real-world dynamics: industries that make consistently poor decisions face compounding consequences.
- Recovery is possible. A strong round can pull an industry out of Headwind (+5 or +6 in a single round). Crisis recovery is harder but not impossible.
- Consequences are mechanical, not just narrative. Constraints physically restrict what decisions participants can make. This is the "compressed multiple" effect — expressed as reduced freedom of action rather than financial calculations.
- Surge is earned. Hitting +15 cumulative requires sustained excellence across most of the exercise. The reward (bypassing automatic red-flag challenge) is meaningful but balanced.

Historical Reference Data: All historical market data (industry EV/EBITDA and P/E multiples, S&P 500 sector returns, forward P/E multiples, Treasury yields) is consolidated in 07_Supporting_Materials/07_Historical_Reference_Data.md.

End of Industry Health Signal Tables — V7.4

Facilitator Overview & Runbook

Exercise Overview

Project Threshold is an 8-hour tabletop exercise designed to stress-test near-term (2026-2030) economic implications of AI across eleven US industries. Individual industry participants navigate four sequential rounds under different AI deployment scenarios, make strategic decisions individually, and evaluate industry-level implications through cross-industry discussion.

Key Parameters

Parameter	Value
Participants	5-11 industry representatives + 1-2 facilitators
Industry Representatives	5-11 individual participants. Each selects 1 or more industries (recommend 2). One decision worksheet per industry per round. Ideal: 8+ industries assigned. Fewer possible. Facilitator plays unassigned industries.
The 11 Industries	Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS
Rounds	4 sequential periods (0-6mo, 6-18mo, 18-30mo, 30-48mo)
Format	Single 8-hour event
AI Scenario	Baseline Type A (configurable; see separate technical specs)
Decision Archetypes	15 common strategic patterns with plausibility gates

The 11 Industries

Industry	Key Concerns
Retail	Demand forecasting, supply chain, labor displacement, competitive AI dynamics
CPG	Demand forecasting, supply chain, labor displacement, competitive AI dynamics
Healthcare Provider (hospital systems)	Clinical AI workflows, FDA approval, patient safety
Healthcare Payer (insurers)	Claims automation, coverage decisions, patient safety
Finance (banking + insurance)	Underwriting/fraud AI, regulatory approval, labor transition
Consulting (Big Four/MBB)	Copilot adoption, service pricing, labor transition
Law (AmLaw 50)	Billable hour disruption, bar rule uncertainty, malpractice liability
Manufacturing	Production optimization, labor transitions, capex allocation
Logistics	Autonomous vehicles, labor transitions, capex allocation
Big Tech (cloud, ads, devices, enterprise software — excludes AI lab/model development)	AI feature integration, pricing pressure, margin management
B2B/B2C SaaS (Workday/Salesforce-class)	AI feature integration, pricing pressure, competitive threats from AI-native startups

Note on Big Tech scope: Big Tech covers cloud infrastructure, advertising, devices, and enterprise software. AI lab and foundation model development decisions are excluded from participant scope and introduced via facilitator injects only.

Facilitator Role & Responsibilities

Phase	Responsibilities
Pre-Exercise	Assign industry roles (ideal: 8+ industries assigned; fewer possible; facilitator plays unassigned industries), distribute decision packets and AI Adoption Arc Foundation phase (pre-read), brief scenarios
During Exercise	Manage injects, distribute Private Cards (from 05_Participant_Private_Cards/) and AI Adoption Arc phase handouts (from 04_Participant_AI_Adoption_Arcs/) at the start of each round, track decision quality, score per adjudication rules, challenge implausible proposals

Phase	Responsibilities
Market Shock & Collective Bonus	Announce Facilitator Market Shock constraints (R2 only); facilitate optional Collective Bonus during cross-industry discussion (all rounds)
Adjudication	Apply banded scoring ({-2, 0, +2}) + red-flag triggers (+/-3), track individual decisions per industry
Cross-Industry Discussion	Facilitate 27-min discussion period in each round for cross-industry dynamics and spillover analysis. Use discussion prompts to surface synergies and tensions across industries.
Debrief	Lead 60-min final debrief with individual reflections (15min) + cross-industry discussion (25min) + no-regrets actions (20min)

Private Cards & AI Adoption Arc Distribution

Distribute the appropriate materials to each participant at the start of each round, before the scenario read.

Round	Private Cards (from 05_Participant_Private_Cards/)	AI Adoption Arc (from 04_Participant_AI_Adoption_Arcs/)
Pre-Exercise	--	Foundation phase included in pre-read packets
Round 1	Card 1 -- shared within industry clusters (e.g., Retail and CPG both receive Consumer Card 1)	Confirm all participants have Foundation phase
Round 2	Card 2 -- unique per industry	Acceleration phase handout
Round 3	Card 3 -- unique per industry	Reckoning phase handout
Round 4	No Private Cards	Normalization phase handout

Exercise Timeline

The master timeline for the 8-hour exercise is in 01_Timeline.md. Refer to that document for the complete minute-by-minute schedule, round structure, sub-block durations, debrief structure, and

facilitator checklist. All clock times, break placements, and duration allocations are defined there.

Key reference values: Rounds 1–4 at 65 min each (5 + 3 + 15 + 3 + 12 + 27). Debrief at 60 min (15 + 25 + 20). Exercise runs 8:30 AM – 4:30 PM.

Practice Micro-Round (Before Round 1)

Purpose: Teach participants the individual decision format, scoring model, and specificity checklist in 5 minutes flat.

Inject: "A major consultancy just announced that AI copilots have reduced junior hiring by 40%. Entry-level salaries are down 15%. What do you do?"

Decision Format Example: "The Consulting participant proposes: Deploy GitHub Copilot to 25% of consulting staff (pilot phase, major US offices only) over 90 days. Pair with mandatory training. Commit to zero involuntary headcount reduction in pilot phase. Budget: \$2M for licensing and training. Success metric: Track analyst productivity and time-to-billable-hours."

Facilitator Scoring (Live):

Dimension	Score	Rationale
Strategic Fit	0	Defensive move; doesn't accelerate growth, but acknowledges competitive pressure
Execution Risk	+1	Copilot adoption is proven; 90-day timeline is tight but feasible with modern infrastructure
Tail Risk	+1	Pilot phase limits downside; if adoption fails, can pivot; if succeeds, can scale
Total	+2/6	Acceptable decision. Participant sees how scoring works; why specificity matters

Facilitator debrief: "Notice how the participant specified WHO (25% of staff), WHAT (GitHub Copilot), WHERE (major US offices), WHEN (90 days), HOW MUCH (\$2M), HOW (training plan), RISK (pilot-only, no layoffs). That specificity made scoring easy. Vague proposals (e.g., 'accelerate AI adoption') get challenged. This is individual decision-making -- you're deciding alone for your assigned industries."

Running with Minimal Staff

Project Threshold V7.4 can run with **1 facilitator + 1 helper** or **1 facilitator solo**:

Configuration	Role	Responsibilities
1 Facilitator + 1 Helper (Recommended)	Facilitator	Reads scenarios, manages injects, distributes Private Cards and AI Adoption Arc handouts, scores decisions, announces Facilitator Market Shock (R2), facilitates Collective Bonus (Optional) and cross-industry discussion, leads debrief
	Helper	Tracks time, maintains scoreboard, tallies Collective Bonus nominations, handles logistics (breaks, AV, materials)
1 Facilitator Solo (Possible)	Facilitator (all duties)	Pre-record scenario reads (or use text slides). Use timer app with hard alarms for round transitions. Fast-track scoring via specificity checklist and industry baselines; don't debate edge cases. Distribute scorecards post-event rather than live.

Facilitator Discussion Prompts: Cross-Industry Synergies & Tensions

Use these prompts during the 27-minute cross-industry discussion period to surface dynamics across industries. Scan the table below for the relevant cluster and pick 1-2 prompts per round.

Industry Cluster	Industries	Discussion Prompts
Consumer	Retail, CPG	(1) How does Retail's AI-driven inventory optimization affect CPG's demand forecasting and production planning? (2) If Retail cuts shelf space for AI-optimized assortment, how does CPG respond?

Industry Cluster	Industries	Discussion Prompts
Healthcare	Provider, Payer	(1) How do Provider AI diagnostic investments interact with Payer claims automation? Who bears the risk of AI errors? (2) If Payer automates prior authorization, does that help or hinder Provider clinical AI adoption?
Finance & PS	Finance, Consulting, Law	(1) How does Finance's AI underwriting capability affect demand for Consulting advisory services? (2) If Law firms adopt AI-generated legal research, how does that change Consulting's competitive landscape for regulatory advisory? (3) How does the shift from billable hours (Law) interact with outcome-based pricing (Consulting)?
Supply Chains	Manufacturing, Logistics	(1) How does Manufacturing's predictive maintenance success affect Logistics demand and routing? (2) If autonomous vehicles disrupt Logistics, how does Manufacturing's supply chain strategy adapt?
Software & Tech	Big Tech, B2B/B2C SaaS	(1) How does Big Tech's AI infrastructure investment create or destroy opportunity for SaaS incumbents? (2) If Big Tech bundles AI features into cloud platforms, what happens to standalone SaaS pricing?
Cross-Industry	Big Tech, Consulting, Law	How does Big Tech's enterprise AI platform strategy affect Consulting and Law firms' competitive positioning?
Cross-Industry	Manufacturing, Logistics, Retail, CPG	If Manufacturing and Logistics automate aggressively, how does that flow through to Retail and CPG supply chain costs?

Industry Cluster	Industries	Discussion Prompts
Cross-Industry	Finance, Healthcare Provider, SaaS	How does Finance's risk appetite for AI-driven industries affect capital availability for Healthcare Provider and SaaS?

Tips for Facilitating

Decision Specificity Checklist (Apply Before Every Explicit Score)

Before scoring an explicit decision, verify the participant specified:

Field	What to Check
INDUSTRY	Which industry? (Retail? CPG? Finance? Healthcare Provider? Consulting? Law? etc.)
WHO	Which team owner? Committed sponsor?
WHAT	Specific capability/action? (e.g., "GitHub Copilot," not "deploy AI")
WHERE	Scope -- pilot, function, geography, enterprise?
WHEN	Timeline realistic for scope? Regulatory approval needed?
HOW MUCH	Headcount, capex, revenue impact clear?
HOW	Talent plan? Integration detail? Rollback plan?
RISK	Participant acknowledges execution and tail risk?
BANDS	Spend/Commitment? Time-to-Impact? Execution Complexity? Dependency? Scale?

If >2 items missing: "I need more specificity before I score. Can you clarify [issue]? And how would you classify this in terms of Spend, Complexity, and Dependency? That changes the execution risk profile."

After Round Scoring: Apply Base Case Fallbacks

After scoring all explicit industry decisions:

1. Identify any industry that received no explicit action from any participant
2. Reference the fallback bank (02BaseCaseFallbackBank.md)
3. Apply the pre-defined fallback score (deterministic, small deltas: +/-1 per dimension)
4. Post fallback score alongside explicit scores for transparency

5. Fallback industries receive standard base case scores

Example facilitator announcement: "Participant covering Retail and CPG submitted one explicit decision on Retail (+2/6). CPG received no explicit action. Applying base case fallback: CPG = +1/6 (defensive cost-reduction move). Both scores now posted."

Challenging Implausible Proposals

Use this template:

"I appreciate the ambition. Let me ask for clarity on [specific issue]. Your proposal as stated is [current scope]. That changes the execution risk significantly. Here's what I can score: [narrower scope with realistic timeline]. That gives you [Strategic Fit] + [Execution Risk] + [Tail Risk] = [Total]. Alternatively, if you want [different tradeoff], we can score [different proposal]. Which direction?"

Red-Flag Triggers (When to Challenge)

Red-Flag Category	Examples (Challenge If Present)
Timeline misalignment	Enterprise AI in <6 months without pilot; FDA approval in <12 months; major M&A close in <3 months
No execution risk discussion	Missing talent plan, integration detail, or regulatory pathway
Industry constraint violations	Healthcare AI without FDA plan; Trading AI without circuit breakers; Retail with no labor transition; Law AI without bar rule compliance plan; Consulting AI without client confidentiality safeguards
Unhedged tail risk	Autonomous systems with no rollback; >30% headcount cut without retention plan; full deployment no pilot
Implausible synergies	M&A with >100% synergies; supply chain cuts >20% COGS; revenue growth >50% without market expansion

Industry Health Signals (End of Round Only)

- Score first (3 dimensions, banded {-2, 0, +2})
- Apply fallbacks to any industries without explicit participant decisions
- At end of round, update cumulative scores for all 11 industries
- Look up Industry Condition from the condition band table (see 02IndustryHealthSignalTables.md)
- Announce conditions at start of next round (~2 minutes)
- Apply constraints for any industry in Headwind or Crisis

Time Management

Hard Stops (Non-Negotiable)

Activity	Duration	Rule
Individual decision prep	15 min per round	Individuals decide faster than teams
Collective Bonus (Optional)	~5 min per round (within discussion period)	Participants nominate industries for +2 (strong) or -2 (risky). Takes effect if 3+ agree.
Facilitator Market Shock	3 min (R2 only)	Facilitator selects 2-3 industries, imposes one constraint each. No negotiation.
Cross-industry discussion	27 min per round (uniform)	Primary learning period; includes optional Collective Bonus
Debrief	60 min total	Non-negotiable. If time is short, cut earlier rounds or optional injects, never the debrief

If Falling Behind

Priority	Action	Detail
1	Skip optional injects	Keep core injects #1-#8; defer others
2	Fast-track decisions	Use industry baseline ranges directly; don't debate edge cases
3	Reduce cross-industry discussion	Hit key questions quickly; defer nuance to debrief
4	Abbreviate individual reflections in debrief	Each participant: 2 min (vs. 3 min). Keep cross-industry discussion + no-regrets actions

If Ahead of Schedule

Priority	Action	Detail
1	Add optional injects	Pull from the packet
2	Extend cross-industry discussion	Deep-dive on spillovers and industry interdependencies
3	Extend debrief	Explore tail risks or policy implications

Priority	Action	Detail
4	Run war game	Ask participants to role-play responses to one more 2028 scenario

Scoring Mechanics (Quick Reference)

Parameter	Value / Rule
Dimensions	Strategic Fit, Execution Risk, Tail Risk (each -3 to +3, but typically {-2, 0, +2})
Banded framework inputs	Spend/Commitment, Time-to-Impact, Execution Complexity, Dependency, Scale
Banded scoring during play	{-2, 0, +2} for each dimension (default)
Red-flag exception	If band red-flag combination fires, unlock +/-3 exception scoring
Total per explicit decision	Sum of three dimensions (range: -6 to +6; typical: -2 to +6)
Fallback industry scoring	Pre-defined small deltas (+/-1 per dimension) applied directly from fallback bank
Aggregate industry score	Sum of explicit decisions + fallback decisions (per industry)
Industry Health	Cumulative aggregate scores determine Industry Condition bands: Surge (+15+), Tailwind (+6 to +14), Steady (-5 to +5), Headwind (-6 to -14), Crisis (-15 or worse) with mechanical consequences (see 02IndustryHealthSignalTables)

Materials Checklist

- [] Scenario briefs (one per participant + facilitator)
- [] Decision packets / worksheets (forms for recording WHO/WHAT/WHERE/WHEN/HOW/HOW MUCH/RISK; one per industry per round)
- [] Private Cards sorted by round (from 05_Participant_Private_Cards/)

- AI Adoption Arc handouts for Rounds 2-4 (from 04_Participant_AI_Adoption_Arcs/; Foundation phase in pre-read)
 - Scoring sheet template (one per round; track all industry decisions)
 - Adjudication rules card (laminated for reference)
 - Quick reference card (facilitator)
 - Plausibility decision trees (for red-flag checks)
 - Industry baselines (for calibration)
 - Industry Health Signal tables (for end-of-round condition announcements)
 - Base case fallback bank (for industries without explicit decisions)
 - Timer (or phone timer) for Collective Bonus, Market Shock, and round transitions
 - Flip chart + markers (for posting scores and Collective Bonus results after each round)
-

Next: See *02AdjudicationRules.md* for detailed scoring guidance for 11 industries.

Plausibility Decision Trees

Overview

This section provides plausibility checks for 15 common strategic archetypes, using **band combinations** instead of granular operational details. For each archetype, a red-flag trigger (based on band combinations) indicates when to unlock +/-3 exception scoring (vs. default banded {-2, 0, +2}). Use this as the gating mechanism for exception scoring.

Key Change: Plausibility gates now focus on band feasibility (e.g., "Transformational spend + 0-3mo + Very High complexity = implausible unless...") rather than detailed budget/headcount requirements. Participants submit bands; facilitators challenge implausible combinations.

15 Strategic Archetypes (Consolidated for V7.4)

#	Archetype	Red-Flag Band Combination	Typical Banded Score (if plausible)	Exception Score Range (if red-flag fires)
1	Deploy New AI System	Global/National + 0-3mo + Absorbable/Material; OR no pilot phase + High/Very High complexity; OR Regulator-dependent + <12mo	{-2,0,+2} -> +0-2/9	+/-3 unlocked: -2 to -3 if band combination unrealistic; +3 if novel tech proven in pilot
2	Acquire Competitor / Startup	Transformational spend + High integration complexity + <3mo; OR Existential spend + Ecosystem shift required	{-2,0,+2} -> -1-1/9	+/-3 unlocked: -3 if regulatory existential risk; +2 if synergies crystallizing
3	Cut Headcount / Restructure	Transformational spend + >30% reduction + <6mo + no severance plan; OR Existential spend + labor impact	{-2,0,+2} -> -1-1/9	+/-3 unlocked: -3 if labor backlash + execution collapse; +2 if severance + retraining excellent

#	Archetype	Red-Flag Band Combination	Typical Banded Score (if plausible)	Exception Score Range (if red-flag fires)
4	Exit Market / Divest Unit	Existential spend (core asset) + no identified buyer + fire-sale valuation	{-2,0,+2} -> -1-2/9	+/-3 unlocked: -3 if fire-sale destroys value; +2 if proceeds productively redeployed
5	Launch AI-Native Product	Transformational/Existential spend + 2+ years + unproven market fit + Very High complexity	{-2,0,+2} -> 0-1/9	+/-3 unlocked: -3 if no differentiation; +3 if market fit crystal clear + execution flawless
6	Form Strategic Alliance / Partnership	Material/Transformational spend + no partner identified + vague scope	{-2,0,+2} -> 0-1/9	+/-3 unlocked: -3 if only path forward but no partner exists; +2 if partner ideal + synergies obvious
7	Build Proprietary AI / R&D	Transformational/Existential spend + 2+ years + Very High complexity + Ecosystem shift + no world-class talent	{-2,0,+2} -> -1-0/9	+/-3 unlocked: -3 if talent gap unbridgeable; +1 if acquisition + integration succeeds
8	Deploy Autonomous System	Global scale + National/Global + High/Very High complexity + no rollback + mission-critical domain	{-2,0,+2} -> -2-0/9	+/-3 unlocked: -3 if catastrophic failure risk (e.g., autonomous vehicle no circuit breakers)
9	Major Org Restructure	Transformational spend + Very High complexity + >50% role changes + no change management plan	{-2,0,+2} -> -1-0/9	+/-3 unlocked: -3 if org capability collapses; +1 if change management excellent

#	Archetype	Red-Flag Band Combination	Typical Banded Score (if plausible)	Exception Score Range (if red-flag fires)
10	Aggressive Pricing / Margin Push	Material/Transformational spend + commodity market + below-cost pricing + competitors will match	{-2,0,+2} -> -2-0/9	+/-3 unlocked: -3 if self-destructive pricing certain
11	Build Data/AI Capability In-House	Transformational spend + 1-2yr+ timeline + Very High complexity + no existing talent + competitors faster	{-2,0,+2} -> -1-0/9	+/-3 unlocked: -3 if talent shortage severe; +1 if phased build feasible
12	Pivot to New Market / Segment	Transformational spend + 1-2yr timeline + High complexity + unvalidated market + abandons core advantage	{-2,0,+2} -> -1-0/9	+/-3 unlocked: -3 if abandons core advantage; +2 if validated market + clear advantage
13	Geographic Expansion	Transformational/Existential spend + Global scale + geopolitical risk + unfamiliar regulatory environment	{-2,0,+2} -> -1-0/9	+/-3 unlocked: -3 if geopolitical hostility; +1 if regulatory clarity + opportunity
14	Invest in Defense / Risk Mitigation	Material spend + purely defensive + no offensive upside + capital locked up	{-2,0,+2} -> -1-0/9	+/-3 unlocked: -2 if unnecessary; +2 if prevents existential tail risk
15	Wait & See / Hold Pattern	Material/Transformational spend + market accelerating + competitors moving + cost of waiting high	{-2,0,+2} -> -2-0/9	+/-3 unlocked: -3 if market leaves you behind; 0 if strategic clarity emerges

Per-Archetype Plausibility Checks (Abbreviated)

Archetype 1: Deploy New AI System

Plausible IF (Band Combinations):

- Spend: Absorbable or Material (not Transformational/Existential for unproven tech)
- Time: 0-3mo for pilots; 3-12mo for limited scale; 1-2yr for enterprise-wide
- Complexity: Low or Medium (not High/Very High for novel tech)
- Dependency: Mostly internal or Vendor (not Regulator/Union/Ecosystem unless pre-engaged)
- Scale: Pilot -> Regional -> National (not Global without track record)

Examples of plausible combinations:

- Absorbable + 0-3mo + Low + Internal + Pilot (check)
- Material + 3-12mo + Medium + Vendor + Regional (check)
- Transformational + 1-2yr + High + Regulator (with engagement plan) + National (check)

Red-Flag IF (Band Combinations):

- Global/National scale + 0-3mo + no pilot phase (overcommitted timeline)
- Transformational/Existential spend + 0-3mo + Very High complexity (unrealistic)
- High/Very High complexity + Regulator-dependent + <12mo (FDA approval timeline typically 12-24mo)
- Novel technology + no talent plan + no vendor partnership

Typical Score (if plausible):

- Absorbable + 0-3mo + Low + Internal + Pilot: +0 to +2/9 (low-risk pilot)
- Material + 3-12mo + Medium + Vendor + Regional: +0 to +1/9 (some integration friction)
- Transformational + 1-2yr + High + Regulator + National: -1 to 0/9 (long timeline, approval risk)

Exception Scoring:

- If red-flag band combination fires: Unlock +/-3. Can drop to -3 if overcommitment is severe. Can reach +3 if novel tech proven in pilot + execution flawless.

Archetype 2: Acquire Competitor / Startup

Plausible IF (Band Combinations):

- Spend: Material or Transformational (but not Existential for unvetted target)
- Time: 3-12mo for small deals (<\$1B); 1-2yr for larger deals (>\$1B)
- Complexity: Medium to High (integration always complex)
- Dependency: Mostly internal or Vendor (low antitrust scrutiny)
- Scale: Regional to National (integration risk scales with size)

Examples of plausible combinations:

- Material + 3-12mo + Medium + Internal + Regional (small bolt-on) (check)
- Transformational + 1-2yr + High + Internal + National (large integration) (check)
- Transformational + 2+yr + Very High + Regulator (major deal with antitrust review) + National (check)

Red-Flag IF (Band Combinations):

- Existential spend + <3mo (impossible close timeline for large deal)
- Transformational spend + Very High complexity + <6mo (unproven integration plan)
- Existential spend + Ecosystem shift required (bet-the-company M&A with unknown synergies)

Typical Score (if plausible):

- Material + 3-12mo + Medium + Internal + Regional: -1 to +1/9 (deal risk; integration friction)
- Transformational + 1-2yr + High + Internal + National: 0 to +2/9 (deal closing; synergies becoming clear)

Exception Scoring:

- If red-flag band combination fires: Can drop to -3 if architecturally unsound (e.g., cultural incompatibility). Can reach +3 if synergies crystallizing + regulatory approval near-certain.
-

Archetype 3: Cut Headcount / Restructure

Plausible IF (Band Combinations):

- Spend: Absorbable or Material (cost-reduction initiative)
- Time: 0-3mo for announcement; 3-12mo for phased execution
- Complexity: Low to Medium (straightforward execution)
- Dependency: Mostly internal (may involve union negotiation; requires engagement plan)
- Scale: Regional to National (not Global without cultural/regulatory complexity)

Examples of plausible combinations:

- Material + 3-12mo + Low + Internal + Regional (modest <15% cut, severance offered) (check)
- Material + 3-12mo + Medium + Union partnership + National (phased, generous severance, retraining) (check)

Red-Flag IF (Band Combinations):

- Transformational spend + >30% reduction + 0-3mo + no severance plan (aggressive + no safety net)
- Transformational spend + High/Very High complexity + Union-dependent + <6mo (labor backlash certain)

Typical Score (if plausible):

- Material + 3-12mo + Low + Internal + Regional: -1 to +1/9 (cost upside; talent/morale risk)
- Material + 3-12mo + Medium + Union partnership + National: 0 to +2/9 (phased; cost advantage realized; labor relations maintained)

Exception Scoring:

- If red-flag band combination fires: Can drop to -3 if aggressive cut + no severance + legal exposure. Can reach +3 if severance generous + retraining excellent + org stabilizes quickly.
-

Archetype 4: Exit Market / Divest Unit

Plausible IF:

- Business unit is non-core
- Logical buyer identified (competitor, PE firm, strategic buyer)
- Valuation is reasonable (market comps support price)
- Sale can close in 6-12 months

Red-Flag IF:

- Unit is core to company identity
- No identified buyer (vague idea of "someone will want it")
- Valuation is overpriced relative to market comps
- Asset is forced fire-sale (depressed price; weak negotiating position)

Typical Score (if plausible):

- Early rounds: {-2, 0, +2} -> -1 to +2/9 (defensive move; frees capital for redeployment)
- Later rounds: {-2, 0, +2} -> +0 to +3/9 (capital redeployment is successful; divestiture was right call)

Exception Scoring:

- If red-flag fires: Can drop to -3 if fire-sale undermines shareholder value + market sees it as distress signal. Can reach +3 if strategic clarity is excellent + proceeds are productively redeployed.
-

Archetype 5: Launch AI-Native Product

Plausible IF:

- Market opportunity is validated (customer research; willing to pay)
- Product roadmap is clear (MVP in 6-12 months)
- Team has product + engineering capability
- Capital is committed (budget allocated by leadership)

Red-Flag IF:

- Unproven market fit
- No alpha / differentiation vs. competitors
- Massive capex with unclear payback (>\$50M+ for unproven product)
- No committed go-to-market plan

Typical Score (if plausible):

- Early rounds: {-2, 0, +2} -> +0 to +1/9 (exploratory; high execution risk)
- Later rounds: {-2, 0, +2} -> +1 to +3/9 (market fit is proving; revenue traction)

Exception Scoring:

- If red-flag fires: Can drop to -3 if product has no differentiation + team is uncertain. Can reach +3 if market fit is crystal clear + team is world-class + execution is ahead of schedule.
-

Archetype 6: Form Strategic Alliance / Partnership

Plausible IF:

- Partner is identified and interested in discussions
- Scope is defined (technology licensing, distribution, joint venture)
- Governance is clear (roles, IP ownership, exit clause)
- Win-win for both parties

Red-Flag IF:

- No specific partner identified
- Scope is vague (e.g., "we'll figure out what to do together")
- Synergies are unproven or wishful
- One party has much more to gain than the other (imbalance)

Typical Score (if plausible):

- {-2, 0, +2} -> +0 to +2/9 (strategic alliance can accelerate time-to-market; execution depends on partner quality)

Exception Scoring:

- If red-flag fires: Can drop to -3 if partnership is only viable option but no partner exists (stranded). Can reach +3 if partner is ideal + governance is ironclad + synergies are obvious.

Archetypes 7-15 (Brief Summaries)

Archetype	Red-Flag	Plausible Score	Exception Range
7. Build Proprietary AI/R&D	No world-class talent; >\$100M capex unproven ROI	-1 to +1/9	-3 to +2 (if talent is scarce, -3; if acquisition + integration succeeds, +2)
8. Deploy Autonomous System	No circuit breakers; no human oversight; mission-critical + no rollback	-2 to +1/9	-3 to 0 (catastrophic failure risk; -3 is likely)
9. Major Org Restructure	>50% role changes; no change mgmt plan; talent exodus risk	-1 to +1/9	-3 to +2 (org collapse risk is -3; successful transformation is +2)
10. Aggressive Pricing/Margin Push	Commodity market; below-cost pricing; competitors will match	-2 to 0/9	-3 to 0 (self-destructive pricing is -3)
11. Build Data/AI Capability In-House	No existing talent; competitors faster; make-vs-buy ignored	-1 to +1/9	-3 to +1 (talent gap is unbridgeable, -3; phased build is +1)

Archetype	Red-Flag	Plausible Score	Exception Range
12. Pivot to New Market/Segment	Unvalidated market; abandons core advantage; new competition	-1 to +1/9	-3 to +2 (abandons core advantage, -3; validated market + clear advantage, +2)
13. Geographic Expansion	Geopolitical risk; unfamiliar labor/regulatory environment; capex-heavy	-1 to +1/9	-3 to +1 (geopolitical hostility, -3; regulatory clarity + market opportunity, +1)
14. Invest in Defense/Risk Mitigation	Purely defensive; no offensive upside; capital locked	-1 to 0/9	-2 to +2 (if defense prevents existential tail risk, +2; if unnecessary, -2)
15. Wait & See / Hold Pattern	Market accelerating; competitors moving; cost of waiting	-2 to 0/9	-3 to 0 (market leaves you behind, -3; strategic clarity emerges, 0)

Industry-Specific Decision Trees

The following decision trees provide industry-specific plausibility guidance for the 11 industries. Use these when a participant's decision requires context beyond the generic 15 archetypes.

Retail

Key tensions: Amazon competition, omnichannel execution, labor displacement, brand trust with AI-driven personalization.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
AI demand forecasting / inventory	Proven tech; phased rollout; existing vendor	Global rollout <3mo; no pilot	+2 to +4/9
AI-driven personalization	Transparency/opt-in; brand safety plan	Invasive data use; no customer consent	+1 to +2/9
Aggressive dynamic pricing	Competitive intelligence; customer segmentation	Price discrimination perception; brand damage	-1 to +1/9
Autonomous checkout / fulfillment	Pilot phase; union engagement	Full deployment <6mo; no labor transition	-2 to +1/9

CPG

Key tensions: Brand safety with AI-generated content, retailer power dynamics, DTC execution risk,

R&D cycle acceleration.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
R&D acceleration with AI	Pilot product lines; proven tools	Full portfolio <6mo; no validation	+2 to +3/9
Marketing automation (human review)	Human review gates; brand guidelines	No review; AI-generated content uncontrolled	+1 to +2/9
Aggressive DTC bypass of retailers	Validated DTC demand; phased	Abandon retailer relationships; no fallback	-1 to +1/9

Healthcare Provider

Key tensions: Patient safety liability, FDA approval timelines, physician adoption, EHR integration complexity.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Clinical decision support (physician override)	Human oversight; staff training	No override; novel diagnostics no FDA	+1 to +2/9
FDA pre-submission for diagnostic AI	12-24mo timeline; clinical validation	<12mo for approval; no validation plan	0 to +1/9
Operational AI (scheduling, resources)	Proven tech; pilot departments	Enterprise-wide <3mo; no workflow analysis	+1 to +2/9

Healthcare Payer

Key tensions: Claims denial backlash, regulatory MLR limits, member satisfaction, fraud detection arms race.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Prior auth automation (human review)	Human review for high-cost; phased	>70% auto-denial; no human review	+1 to +2/9
Fraud detection / payment integrity	Proven technology; vendor partnership	Novel approach; no validation	+2 to +3/9
Aggressive claim denial automation	Clear cost rationale	No human review; regulatory backlash certain	-1 to 0/9

Finance

Key tensions: Systemic risk from AI trading, fair lending compliance, synthetic fraud arms race, regulatory scrutiny.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
AI underwriting (explainability)	Human review for denials; bias audit	No explainability; no bias audit	+2 to +3/9
Synthetic fraud detection	Vendor partnership; proven tools	Build from scratch; no talent	+1 to +2/9
AI trading (autonomous)	Circuit breakers; human oversight	>80% autonomous; no oversight; no circuit breakers	-1 to +1/9
Bias auditing / fair lending infrastructure	Compliance-driven; proactive	Reactive only; after enforcement action	+2 to +3/9

Consulting

Key tensions: AI transformation advisory vs. being disrupted by client in-house AI capabilities; junior talent pipeline disruption as AI automates entry-level analysis; pricing pressure from AI-enabled boutiques offering faster/cheaper delivery; vertical AI competition from specialized firms; the tension between selling AI transformation to clients while managing your own AI disruption.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Deploy copilots firm-wide (research, analysis, proposals)	Proven tools (GitHub Copilot, custom LLMs); phased rollout; training program; human review gates	Enterprise-wide <30 days; no training; no quality controls	+2 to +3/9
Build vertical AI expertise practices	Hire 30-50 AI specialists; partner with AI vendors; target 2-3 verticals first	Target all verticals simultaneously; no hiring plan; no client pipeline	+1 to +3/9
Premium AI governance / compliance advisory	Regulatory uncertainty creates demand; existing client relationships; demonstrated methodology	No methodology; no regulatory expertise; pricing unsupported by market	+2 to +5/9
Transition to outcome-based pricing	Pilot with 3-5 trusted clients; document value creation; maintain time-based as fallback	Firm-wide pricing change <6mo; no client buy-in; no value measurement	0 to +2/9
Aggressive junior headcount reduction (>40%)	Clear redeployment plan; severance; AI handles analyst tasks; senior talent retained	No redeployment; no severance; talent pipeline destroyed; client delivery quality at risk	-2 to 0/9

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Acquire AI-native boutique	Clear synergy (talent + capability); reasonable valuation; integration plan	Overpay for unproven capability; cultural mismatch; talent flight risk post-acquisition	-1 to +2/9

Consulting-Specific Red-Flags:

- Cutting >40% of junior staff + <12mo + no talent pipeline plan = red-flag (destroys future partner pipeline; client delivery quality degrades)
- Outcome-based pricing firm-wide + <6mo + no pilot = red-flag (revenue uncertainty; partner resistance; client confusion)
- Building vertical AI practices in >5 verticals simultaneously + <12mo = red-flag (talent spread too thin; none achieve depth)
- Selling AI transformation advisory while not adopting AI internally = credibility red-flag (clients will notice the hypocrisy)

Law

Key tensions: Billable hour model under existential pressure from AI efficiency gains; bar rule uncertainty for AI-generated work product (varies by jurisdiction; ABA and state bars actively debating); malpractice liability if AI-generated work contains errors and is filed without adequate human review; associate leverage model erosion as AI handles tasks traditionally done by junior associates; competition from legal AI platforms (Harvey.ai, CoCounsel) that could disintermediate traditional firms; tension between reducing associate hours (which drives revenue under billable model) and improving efficiency (which clients demand).

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Pilot AI research tools (legal research, due diligence, contract review)	Proven tools (Harvey.ai, CoCounsel, Westlaw AI); human attorney review mandatory; select practice groups	Firm-wide deployment <30 days; no attorney review protocol; no bar rule compliance assessment	+1 to +3/9
AI-assisted brief drafting with attorney review	Attorney signs off on all output; clear quality control; malpractice insurance reviewed	AI drafts filed without attorney review; no malpractice assessment; jurisdiction bar rules unclear	0 to +2/9
Maintain billable hour model with AI efficiency gains	Clients accept current rates; AI reduces cost-to-serve but maintains pricing; reinvest savings	Clients demand rate reductions proportional to AI gains; realization rates drop; revenue erosion	+1 to +2/9

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Transition to alternative fee arrangements (fixed fee, outcome-based)	Pilot with 3-5 clients; clear scope definition; partner buy-in; maintain billable as fallback	Firm-wide pricing overhaul <6mo; no partner consensus; no client validation; revenue model unclear	-1 to +2/9
Reduce associate headcount significantly (>30%)	Phased over 18-24mo; severance; redeployment to complex work; AI handles routine tasks	Immediate cuts >30% + <6mo; no severance; no redeployment; client service quality collapses	-2 to +1/9
Build proprietary legal AI platform	Clear differentiation; in-house AI talent; phased roadmap; client-facing application	No AI talent; >\$50M capex unproven; competing with well-funded legal AI startups	-2 to +1/9
Expand into AI compliance / regulatory advisory	Growing demand; existing regulatory expertise; cross-sell to corporate clients	No regulatory expertise; entering crowded market; competing with Consulting firms	+1 to +3/9

Law-Specific Red-Flags:

- AI-generated briefs filed without attorney review = critical red-flag (malpractice exposure; bar discipline; client harm). Score -3 on Tail Risk automatically.
- >30% associate reduction + <12mo + no redeployment plan = red-flag (leverage model collapses; ability to staff large matters degraded; lateral partner flight)
- Firm-wide pricing model change + <6mo + no partner consensus = red-flag (partner exodus; revenue instability; client confusion)
- Building proprietary legal AI platform + >\$50M + no AI talent = red-flag (competing against Harvey.ai, CoCounsel, and other well-funded startups with no competitive advantage)
- Ignoring bar rule developments on AI-generated work = regulatory red-flag (rules are evolving rapidly; firms that don't track jurisdiction-specific requirements risk sanctions)

Manufacturing

Key tensions: Union relations, labor displacement, OT/IT integration complexity, capex allocation for proven vs. unproven technology.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Predictive maintenance (pilot plants)	High-ROI plants; phased OT/IT integration	All 28 plants simultaneously; no OT/IT plan	+2 to +4/9
No-layoff retraining agreement	Union cooperation; \$35-40M committed	No union engagement; token retraining	+2 to +4/9

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Full warehouse automation	Pilot first; labor transition plan	All plants <6mo; no labor plan; union ignored	-2 to +1/9

Logistics

Key tensions: Driver adoption of AI tools, autonomous vehicle regulatory uncertainty, last-mile profitability limits, union concerns.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Route optimization (pilot fleet)	Proven tech; driver engagement	Full fleet <3mo; no driver buy-in	+2 to +4/9
AV partnership (Waymo/Aurora)	5-10% fleet; managed expectations	Full fleet autonomous; eliminate drivers	0 to +1/9
Autonomous vehicles (internal build)	Phased; regulatory engagement	Full fleet <2yr; no regulatory plan	-3 to -1/9

Big Tech

Key tensions: Cloud infrastructure investment, enterprise AI platform competition, antitrust scrutiny, margin pressure from AI compute costs, open-source model commoditization. **Note:** AI lab and foundation model development decisions are excluded from participant scope; those dynamics come via facilitator injects only. Big Tech scope covers cloud, advertising, devices, and enterprise software.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
AI infrastructure capex (cloud, chips, data centers)	Scaled investment; existing capability; clear ROI pathway	>\$50B single-year commitment without revenue offset; margin collapse	+2 to +4/9
Enterprise AI APIs + platform tools	Platform leverage; existing enterprise relationships	Pricing below cost to undercut competitors; margin destruction	+1 to +3/9
Deep AI integration into core products (search, ads, cloud)	Incremental; user experience tested; phased	Full product overhaul <6mo; no user testing; brand risk	+2 to +4/9
Defensive pricing against open-source	Maintains enterprise stickiness; modest margin compression	Below-cost pricing; shareholder revolt; unsustainable	-1 to +1/9
Major acquisition of AI startup/competitor	Clear strategic fit; talent retention plan; regulatory pathway	Antitrust scrutiny certain; >\$20B without regulatory pre-clearance	-1 to +2/9

Big Tech-Specific Note: If a participant proposes decisions related to foundation model training,

frontier model development, or AI lab operations, redirect: "Big Tech scope in this exercise covers cloud, advertising, devices, and enterprise software. AI lab and model development dynamics are introduced via facilitator injects. Please reframe your decision within the Big Tech scope."

B2B/B2C SaaS

Key tensions: AI feature integration pressure, pricing model disruption, competitive threat from AI-native startups, margin compression from AI infrastructure costs, customer retention vs. upsell.

Decision Type	Plausible IF	Red-Flag IF	Typical Score
Bundle AI copilots into standard product	Maintains lock-in; execution proven	No revenue offset; margin erosion >5%	+1 to +2/9
Premium AI tier (30-50% price increase)	Differentiated features; customer demand validated	Untested pricing; customer churn risk	+2 to +3/9
AI-native vertical product launch	Validated market; existing customer base	Unproven market; competing against specialized startups	+1 to +2/9

How to Use This Document During Scoring

- Participant proposes decision. Record WHO/WHAT/WHERE/WHEN/HOW/HOW MUCH/RISK and band classification (Spend, Time, Complexity, Dependency, Scale).
- Map to archetype. Which of the 15 does this fall into?
- Check industry-specific tree. Does the industry have specific plausibility guidance? (See industry trees above.)
- Check band red-flag combination. Does the band combo trigger a red-flag?
 - If NO red-flag band combo: Score using default bands {-2, 0, +2}. Use band-to-score translation table (02AdjudicationRules). Total typically ranges -6 to +6.
 - If YES red-flag band combo: Unlock +/-3 exception scoring. Reference the plausibility check to calibrate final score. Total can range -9 to +9.
- Reference plausibility check. If band combo is questionable, guide participant toward narrower scope, longer timeline, lower complexity, or lower spend.
- Post score. Explain your decision concisely to participant (reference bands that drove score).
- Move on. No deep debate on edge cases; next decision.

Summary Table: Archetypes -> Band Red-Flags -> Can Score +/-3?

Archetype	Red-Flag Band Combo	Can Score +/-3?
Deploy AI System	Global/National + 0-3mo; OR High/Very High complexity + no pilot + <12mo	YES if timeline/complexity band mismatch is severe
Acquire Competitor	Existential spend + <3mo close; OR High integration + <6mo	YES if regulatory existential risk or deal structure unsound
Cut Headcount	Transformational + >30% in <6mo + no severance	YES if labor backlash + execution collapse
Exit Market	Existential spend (core asset) + no buyer + fire-sale	YES if fire-sale destroys value or strategic incoherence
Launch AI-Native Product	Transformational/Existential + 2+yr + unproven market fit	YES if market fit crystallizes or completely lacks validation
Form Alliance	Material/Transformational + no partner identified + vague scope	YES if only path forward but partner doesn't exist
Build Proprietary AI	Transformational/Existential + Very High complexity + no talent	YES if talent gap unbridgeable or ROI completely unclear
Deploy Autonomous	Global/National + High/Very High + no rollback + mission-critical	YES if catastrophic failure risk (systemic)
Major Org Restructure	Transformational + Very High complexity + >50% changes + no change mgmt	YES if org capability collapses
Aggressive Pricing	Material/Transformational + commodity market + below-cost band	YES if self-destructive pricing certain
Build Data Capability	Transformational + 1-2yr+ + Very High complexity + no talent	YES if talent shortage severe or build technically unfeasible
Pivot to New Market	Transformational + 1-2yr + High complexity + abandons core advantage	YES if abandons core advantage or market unvalidated
Geographic Expansion	Transformational/Existential + Global scale + geopolitical risk band	YES if geopolitical hostility severe
Invest in Defense	Material + purely defensive (no offensive upside)	YES if prevents existential tail risk or unnecessary

Archetype	Red-Flag Band Combo	Can Score +/-3?
Wait & See	Material/Transformational + market accelerating + competitors moving	YES if market passes you by permanently

Use this as a quick reference during live facilitation. If a red-flag fires, reference the plausibility check to calibrate the +/-3 exception score. Do not overthink; move briskly.

Facilitator Quick Reference Card (Industry-Level)

PRINT THIS. Keep handy during live play. Max 2 pages.

CRITICAL: Individual participant model (V7.4):

- 5-11 participants: Each selects 1+ industries. Ideal: 8+ assigned; facilitator plays unassigned industries (up to 3). One decision worksheet per industry per round.
- 11 industries: Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS
- After explicit scoring: Apply base case fallbacks to industries without explicit participant decisions (reference *02BaseCaseFallbackBank*)
- Card distribution: Private Cards from *05_Participant_Private_Cards/* (3 per industry, Rounds 1-3). AI Adoption Arc handouts from *04_Participant_AI_Adoption_Arcs/* (Phases 2-4 at start of Rounds 2-4).

Each industry decision is scored **independently** on its own context. Fallback industries receive pre-defined scores from fallback bank (deterministic; no participant input).

1. DECISION SPECIFICITY CHECKLIST (Before Every EXPLICIT Industry Decision)

- INDUSTRY ID: Which industry? (Retail? CPG? Finance? Consulting? Law? Healthcare Provider? etc.)
- WHO: Team owner? Committed sponsor?
- WHAT: Specific capability? (e.g., "Contract review copilots," not just "AI delivery")
- WHERE: Scope for this industry? (500 stores? 8 plants? Full fleet? 3 practice groups?)
- WHEN: Realistic timeline for this industry?
- HOW MUCH: Capex, headcount, revenue impact for this industry? (quantify only for M&A/major capex/regulatory commitment)
- HOW: Talent plan? Integration? Rollback?
- RISK: Execution + tail risk acknowledged for THIS industry?
- BANDS: Classify as Spend, Time, Complexity, Dependency, Scale (e.g., "Material, 3-12mo, Medium, Vendor, Regional")

Missing >2 items OR bands unclear? -> Ask for clarification before scoring. **No explicit decision for an industry?** -> Apply base case fallback (reference *02BaseCaseFallbackBank*).

2. BANDED SCORING CHEAT SHEET

Dimension	-2	0	+2	Exception (+/-3)
Strategic Fit	Misaligned; opportunity cost	Neutral; reasonable move	Captures core opportunity	Red-flag: unlock +/-3 if severe
Execution Risk	Major barriers; hard to execute	Feasible; some change mgmt needed	Straightforward; talent/capital ready	Red-flag: unlock +/-3 if barriers severe
Tail Risk	Exposed to downside; no hedges	Neutral; standard risk	Includes hedges/rollback plan	Red-flag: unlock +/-3 if catastrophic

Total Score Range: -6 to +6 (typical: -2 to +6)

Interpretation:

- +5 to +6: Strong (70%+ success). Accept.
- +1 to +4: Acceptable (50-70% success). Monitor.
- -2 to 0: Weak. Challenge; ask participant to narrow scope.
- <-2: Poor. Reject or reframe significantly.

3. RED-FLAG TRIGGERS (Band Combinations -> Unlock +/-3 Exception)

Red-Flag Categories (Band-Based)

Category	Red-Flag Band Combo	Action
Timeline Misalignment	Global/National + 0-3mo + no pilot phase	Offer narrower scope or longer timeline
Overcommitted Complexity	Transformational + Very High + Ecosystem shift + <12mo	Challenge scope; propose pilot phase
No Execution Plan	Missing talent/integration/regulatory plan for banded complexity	Demand specificity (see checklist above)
Industry Constraints	Healthcare AI + no FDA plan; Finance + no circuit breakers; Law + no bar compliance; Consulting + no client confidentiality plan	Push back hard; add safeguards before scoring

Category	Red-Flag Band Combo	Action
Unhedged Tail Risk	Autonomous + no rollback; >30% cut with no severance plan (M&A context); AI legal work filed without attorney review	Refuse to score until hedged

When band red-flag fires: Unlock +/-3 scoring. Reference decision trees (02PlausibilityDecision_Trees) to calibrate. Challenge participant to narrow scope, extend timeline, reduce complexity, or lower spend. Can drop to -3 or rise to +3 depending on response.

4. ARCHETYPE QUICK GUIDE (15 Strategic Patterns)

#	Archetype	Typical Range	Red-Flag
1	Deploy AI System	+0-2	Enterprise <6mo; novel tech; no plan
2	Acquire Competitor	-1-+1	>\$5B deal; antitrust; <3mo close
3	Cut Headcount	-1-+1	>30% in <6mo; no severance; union
4	Exit Market	-1-+2	No buyer; core asset; fire-sale
5	Launch AI Product	+0-+2	Unproven market; no alpha; huge capex
6	Form Partnership	+0-+2	No partner identified; vague scope
7	Build Proprietary AI	-1-+1	No talent; massive capex; unproven ROI
8	Deploy Autonomous	-2-+1	No circuit breakers; no oversight
9	Major Org Restructure	-1-+1	>50% changes; no change mgmt
10	Aggressive Pricing	-2-0	Commodity market; below-cost
11	Build Data Capability	-1-+1	No talent; competitors faster

#	Archetype	Typical Range	Red-Flag
12	Pivot Markets	-1-+1	Unvalidated; abandons advantage
13	Geographic Expansion	-1-+1	Geopolitical risk; unfamiliar regulatory
14	Invest in Defense	-1-0	Purely defensive; no upside
15	Wait & See	-2-0	Market accelerating; competitors moving

5. INDUSTRY SCORING QUICK REFERENCE (11 Industries)

Industry	Strat Fit Typ	Exec Risk Typ	Tail Risk Typ	Watch For
Retail	0	0	0	Amazon competition; brand backlash; DTC execution
CPG	0	0	-0.5	Brand safety (AI content); retailer power; DTC backlash
Healthcare Provider	0	-1	-1	Patient safety liability; physician adoption; FDA timeline
Healthcare Payer	+1	0	0	Denial backlash; regulatory MLR limits; member satisfaction
Finance	+1	-1	-1	Fair lending enforcement; systemic risk; fraud arms race

Industry	Strat Fit Typ	Exec Risk Typ	Tail Risk Typ	Watch For
Consulting	+1	0	0	Junior talent pipeline; pricing pressure; client in-house AI competition
Law	+1	-1	-1	Bar rule uncertainty; malpractice liability; billable hour disruption; associate leverage erosion
Manufacturing	0	0	-1	Union relations; labor transition; OT/IT integration
Logistics	+1	0	-1	Driver adoption; AV regulatory uncertainty; last-mile limits
Big Tech	+2	+1	-1	Antitrust scrutiny; margin pressure from CapEx; open-source competition. Excludes AI lab/model dev.
B2B/B2C SaaS	+1	0	-1	Pricing pressure (AI bundling); startup disruption; margin compression

Use to calibrate each industry separately. Don't conflate industries.

6. RED-FLAG CHALLENGE TEMPLATE

"I appreciate the ambition. Let me ask for clarity:

Your proposal: *[Restate as you understand it]*

My concern: *[Specific red-flag: timeline, execution risk, regulatory, tail risk, synergies]*

What I can score: *[Narrower scope with realistic timeline/hedges]*

That gives: *[Strategic Fit] + [Execution Risk] + [Tail Risk] = [Total]*

Alternative: If you want [different tradeoff], we can score [different scope].

Which direction?"

7. TIME MANAGEMENT (Industry-Level Scoring + Fallbacks)

Critical Parameters for V7.4

Activity	Time	Flexibility	Notes
Individual decision preparation	15 min	5 min flex	Each participant prepares decisions for their assigned industries
Decision submission	3 min	ZERO	Late = auto-drop to fallback
Scoring (variable explicit decisions)	10-12 min	Use baselines	~1.5-2 min per explicit decision; count depends on participants x industries
Apply base case fallbacks	2 min	Quick ref	For each industry without explicit action, apply deterministic fallback from bank
Collective Bonus	~2-3 min per round	Within discussion	Optional: participants nominate strong/risky strategies. +2/-2 if 3+ agree.
Facilitator Market Shock	3 min (R2 only)	ZERO	Facilitator selects 2-3 industries, imposes one constraint each from menu. Cannot repeat constraint type.
Debrief	60 min	ZERO	Non-negotiable. Never cut.

Scoring Pace

With 2 facilitators:

- Facilitator 1 scores Consumer grouping (Retail, CPG) + Supply Chains grouping (Manufacturing, Logistics) + Big Tech

- Facilitator 2 scores Healthcare grouping (Provider, Payer) + Finance & PS grouping (Finance, Consulting, Law) + SaaS
- Apply fallbacks for non-submitted industries simultaneously
- Total time: 10-12 min (explicit scoring + fallbacks)

With 1 facilitator:

- Use industry baselines aggressively; aim for 1.5 min per explicit decision
- Apply fallbacks quickly using fallback bank (deterministic, <30 seconds total)
- Prioritize explicit decisions; fallbacks are mechanical

If Falling Behind

1. Use baselines religiously (no nuance debates; reference baseline, post score, move)
2. Skip optional injects (keep core injects #1-#8)
3. Reduce debrief to 45 min (hit top 3 questions only)

If Ahead

1. Add optional injects
2. Deepen cross-industry discussion (ask "why?" on strategy)
3. Extend debrief (explore tail risks, policy implications)
4. Run war game (2028 scenario; 15 min)

8. COLLECTIVE BONUS & FACILITATOR MARKET SHOCK REFERENCE

Collective Bonus (Optional, Every Round, R1-R4)

Step	Action	Notes
1	Facilitator opens floor (~2-3 min within cross-industry discussion)	"Does anyone want to recognize an especially strong strategy this round, or flag one that seems particularly risky?"
2	Participants who wish to respond nominate one industry in each direction	Optional — not all participants need to speak. Cannot nominate own industry.

Step	Action	Notes
3	If 3+ participants agree: +2 cumulative score bonus (strong) or -2 penalty (risky)	Applied to that industry's running total. Max one +2 and one -2 per round.
4	If no consensus or no nominations: no bonus applied	Move on. This is optional by design.

Facilitator Market Shock (R2 Only, ~3 min)

Step	Action	Notes
1	Facilitator selects 2-3 industries	Any industry may be targeted.
2	Impose one constraint per industry from constraint menu	Cannot impose same constraint type on multiple industries.
3	Announce constraints publicly	Constrained industries must adapt in their next decision.

After Collective Bonus + Market Shock (Rounds 2-4): Announce Industry Health conditions (~2 min). Look up cumulative scores → condition band (Surge +15+ / Tailwind +6–14 / Steady -5–+5 / Headwind -6–14 / Crisis -15-). Apply constraints for Headwind/Crisis industries.

9. PRACTICE MICRO-ROUND CHECKLIST

Before Round 1, run 10-min demo:

- Inject: "Consultancy reports copilots cut junior hiring 40%."
- Mock decision: Consulting participant proposes copilot deployment (25% staff, 90 days, \$2M budget, zero layoffs in pilot).
- Live scoring: Strat Fit 0 + Exec Risk +1 + Tail Risk +1 = +2/6.
- Debrief: "See how specificity made scoring easy? Vague proposals get challenged."
- Q&A: Any questions on decision format, scoring, red-flags?

10. SCORING DECISION TREE (60 seconds per explicit decision; 30 sec per fallback)

For EXPLICIT Decisions

1. Specificity check: WHO/WHAT/WHERE/WHEN/HOW/HOW MUCH/RISK + BANDS specified?
2. Band plausibility: Any red-flag band combo? (see #3 above) -> If YES, challenge; offer narrower scope/longer time.
3. Strategic Fit: -2 / 0 / +2 (or +/-3 if band red-flag fires)
4. Execution Risk: -2 / 0 / +2 (or +/-3 if band red-flag fires) -- use band translation table
5. Tail Risk: -2 / 0 / +1 (or +/-3 if band red-flag fires)
6. Total = Strat Fit + Exec Risk + Tail Risk
7. Post score. Explain concisely (1-2 sentences); reference bands that drove score.

For FALLBACK Industries (No Explicit Decision)

1. Identify industry without explicit action
2. Reference fallback bank (02BaseCaseFallbackBank.md)
3. Apply pre-defined fallback score (deterministic: +/-1 per dimension, e.g., {0, +1, 0})
4. Post fallback score. Announce: "[Industry] falls back to +1/6 (defensive cost-control move)"
5. Move on. Fallback is automatic; no participant input required.

After ALL Scoring (Explicit + Fallback)

1. Update cumulative score for each industry (running total across all rounds)
2. Look up Industry Condition (Surge +15+, Tailwind +6 to +14, Steady -5 to +5, Headwind -6 to -14, Crisis -15 or worse)
3. Announce at start of next round (~2 min). Apply constraints for Headwind/Crisis.

11. BASE CASE FALLBACK APPLICATION CHECKLIST (After Scoring Explicit Decisions)

- [] Identify industries across all participants that received NO explicit action this round
- [] For each fallback industry, reference the fallback bank (02BaseCaseFallbackBank.md)
- [] Note the deterministic fallback score (small fixed delta: +/-1 per dimension)
- [] Add fallback to industry aggregate (e.g., Retail +2 + CPG fallback +1 = both posted separately)
- [] Announce fallback (e.g., "CPG falls back to +1/6 -- defensive cost control")
- [] Update cumulative scores and determine Industry Health conditions for next round announcement

Important: Fallbacks are deterministic. Do NOT negotiate. Participants wanting to improve fallback industries must submit explicit decisions next round.

12. MATERIALS CHECKLIST

- [] Scenario briefs (1 per participant + facilitator)
- [] Decision worksheets (WHO/WHAT/WHERE/WHEN/HOW/HOW MUCH/RISK + BANDS forms; 1 per industry per round)
- [] Private Cards sorted by round (from 05_Participant_Private_Cards/)
- [] AI Adoption Arc handouts for Rounds 2-4 (from 04_Participant_AI_Adoption_Arcs/)
- [] Scoring sheets (1 per round; show explicit + fallback breakdown per industry)
- [] This quick reference card (laminated)
- [] Adjudication rules (02AdjudicationRules)
- [] Plausibility decision trees (02PlausibilityDecision_Trees)
- [] Industry baselines (02ScoringBaselines)
- [] Industry Health Signal tables (02IndustryHealthSignalTables)
- [] Base case fallback bank (02BaseCaseFallbackBank)
- [] Band-to-score translation table (from 02AdjudicationRules or 02ScoringBaselines)
- [] Timer (phone timer OK)
- [] Flip chart + markers (post scores per industry)

13. KEY QUESTIONS FOR DEBRIEF (Hit These)

1. Which AI timeline emerged? (Slow / Baseline / Fast)
2. Which industries proved to be winners? Losers? Surprises?
3. What policy gaps remain unresolved?
4. What surprised you about execution vs. strategy? (Which band combinations proved hardest?)
5. If you ran this again, what would you do differently?
6. How did the multi-industry model affect your decision quality and strategic thinking?
7. What cross-industry dynamics surprised you most? (e.g., how did Law's billable hour pressure interact with Consulting's pricing transition?)
8. What does this exercise change about your view of AI impact 2026-2030?

Print this card. Laminate. Reference during live play. Keep pace brisk. Facilitate, don't debate.

Scoring Baselines (Industry-Level)

Overview

This section provides scoring baselines for the **eleven industries** in V7.4, showing expected Strategic Fit, Execution Risk, and Tail Risk ranges for typical AI decisions. Use these as calibration anchors when scoring industry-level participant decisions.

Key Changes (V7 -> V7.4):

- 11 industries (up from 10 personas): Split Professional Services into separate Consulting and Law industries. Each is now an independent industry with its own packet, baselines, and scoring context.
- Industry terminology: "Persona" is now "industry" throughout. Each industry IS the packet.
- Big Tech scope narrowed: Excludes AI lab/model development. Scope is cloud, ads, devices, enterprise software.
- Industry-specific baselines: Each of 11 industries has distinct cost structures, competitive threats, and regulatory environments; they receive separate baselines.
- Band-to-score translation: Baselines reference banded inputs (Spend/Commitment, Time-to-Impact, Execution Complexity, Dependency, Scale) rather than granular operational details.
- Base case fallback scoring: Industries without explicit actions receive pre-defined small fallback scores from the fallback bank (see 02BaseCaseFallbackBank.md). Fallback industries are not scored here; this section covers explicit decisions only.

Band-to-Score Translation Reference

Use this table to quickly map banded inputs to expected score ranges:

Spend	Time	Complexity	Dependency	Scale	Typical Strategic Fit	Typical Exec Risk	Typical Tail Risk	Typical Total
Absorbable	0-3mo	Low	Internal	Pilot	+1	+2	+1	+4
Absorbable	0-3mo	Medium	Internal	Regional	0	+1	0	+1
Material	3-12mo	Low	Internal	Regional	+1	+1	0	+2
Material	3-12mo	Medium	Vendor	National	0	0	-1	-1
Transformational	1-2yr	Medium	Vendor	National	+1	-1	-1	-1

Spend	Time	Complexity	Dependency	Scale	Typical Strategic Fit	Typical Exec Risk	Typical Tail Risk	Typical Total
Transformational	1-2yr	High	Regulator	National	+1	-2	-2	-3
Transformational	2+yr	High	Ecosystem	Global	0	-2	-2	-4
Existential	2+yr	Very High	Ecosystem	Global	-1	-3	-3	-7

How to use: If a participant proposes a decision with bands matching one of these rows, the typical score range provides an anchor. Decisions better-executed than baseline get higher scores; worse-executed get lower scores.

Industry 1: RETAIL (Omnichannel Retailers, ~500 stores + e-commerce)

Strategic Priorities (2026-2030):

- Deploying AI demand forecasting and inventory optimization (margin defense)
- Personalizing customer experience (conversion improvement, brand trust risk)
- Automating supply chain and logistics (cost reduction)
- Competing with Amazon and tech-native direct-to-consumer

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	-2 to +2	0	High if catching up on operational AI (inventory, forecasting); neutral if consumer-facing personalization
Execution Risk	-1 to +2	0	Demand forecasting is proven; personalization at scale requires brand safety discipline

Dimension	Range	Typical	Notes
Tail Risk	-1 to +2	0	Labor backlash if aggressive automation; brand backlash if personalization feels invasive

Example Decisions (Retail)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy AI demand forecasting + inventory optimization (500 stores, pilot)	+2	+1	+1	+4	Proven tech; phased; catches up to Amazon; low brand risk
Launch omnichannel personalization with transparency/opt-in	+1	0	+1	+2	Customer trust managed; execution feasible; brand safe
Aggressive dynamic pricing + personalization; brand as efficiency leader	+2	0	-2	0	High margin upside; high brand backlash risk if customers perceive discrimination
Cut 30% retail labor; deploy autonomous checkout + fulfillment	0	-2	-2	-4	Labor relations damaged; union negotiation required; ROI uncertain

Industry 2: CPG (Consumer Goods Manufacturer, 35 brands)

Strategic Priorities (2026-2030):

- R&D acceleration: reduce product development cycles from 18-24 months to 12-15 months (time-to-market advantage)
- Marketing automation + AI-generated content: reduce marketing spend from 8.2% to 7.0% of revenue (margin gain ~60 bps)
- Demand sensing + supply chain optimization: improve forecast accuracy, reduce inventory, optimize production
- Direct-to-consumer (DTC) expansion: bypass retailers, own customer relationship, premium pricing
- Brand safety: manage AI-generated content backlash; maintain brand equity

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	-1 to +2	0	High if R&D acceleration or brand-safe marketing efficiency; neutral if aggressive DTC (retailer relations risk)
Execution Risk	-1 to +1	0	R&D AI is proven; DTC is harder (retailer retaliation, supply chain complexity)
Tail Risk	-1 to 0	-0.5	Brand safety risk if AI-generated content backfires; retailer retaliation if DTC too aggressive

Example Decisions (CPG)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Accelerate R&D with AI ideation + formulation support (pilot: 10 product lines)	+2	+1	0	+3	Proven ROI (6-month cycle reduction); manageable execution; low brand risk

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy AI marketing automation + AI-generated copy (with human review); target -60 bps marketing spend	+1	0	+1	+2	Cost efficiency proven; human review manages brand safety risk
Launch aggressive DTC with AI personalization + dynamic pricing; bypass retailer distribution	+2	-1	-2	-1	Revenue opportunity; but retailer relationships damaged; private-label retaliation likely
Build "no AI" positioning: emphasize human-crafted products, authentic storytelling	-1	+1	+1	+1	Defensive; low execution risk; but misses R&D acceleration opportunity

Industry 3: HEALTHCARE PROVIDER (Hospital System, Clinical AI & Operations)

Strategic Priorities (2026-2030):

- Clinical diagnostic AI (radiology, pathology, cardiology) with FDA approval
- Clinical decision support (reduce medical errors, improve outcomes, physician workflows)
- Operational AI (scheduling, resource allocation, supply chain within hospital)
- Care coordination + risk stratification (population health, EHR integration)

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	-1 to +2	0	High if improving patient outcomes + reducing costs; moderate if defensive/cautious on unproven tech
Execution Risk	-2 to 0	-1	FDA regulatory approval adds 12-24 months; physician adoption uncertain; EHR integration complex
Tail Risk	-3 to 0	-1	Patient safety liability if AI error leads to harm; physician autonomy risk; malpractice exposure

Example Decisions (Healthcare Provider)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy clinical decision support (diagnosis suggestions with physician override); staff training	+2	-1	+1	+2	Clinical value; physician autonomy preserved; adoption uncertain without workflow redesign
Initiate FDA pre-submission for diagnostic radiology AI; plan external validation; 12-month timeline	+1	-1	0	0	Strategic value; long regulatory timeline; patient safety covered; early FDA engagement reduces risk

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy operational scheduling + resource AI; pilot in 3 departments; monitor physician/staff adoption	+1	0	+1	+2	Operational efficiency; lower clinical risk; physician workflow disruption manageable
Build clinical validation and governance infrastructure upfront (before major diagnostic AI)	0	+1	+1	+2	Defensive; enables faster future diagnostic AI deployment; reduces liability risk

Industry 4: HEALTHCARE PAYER (Health Insurer, Claims & Coverage AI)

Strategic Priorities (2026-2030):

- Prior authorization automation (claims cost reduction, liability management)
- Fraud detection + waste reduction (claims review AI, payment integrity)
- Coverage determination + medical policy AI (actuarial models, policy logic)
- Risk stratification + population health (predictive analytics, preventive care targeting)

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	0 to +2	+1	High if improving medical loss ratio (MLR) + reducing administrative costs; moderate if defensive
Execution Risk	-1 to +1	0	Prior auth automation proven; fraud detection is arms race; regulatory compliance adds complexity

Dimension	Range	Typical	Notes
Tail Risk	-2 to +2	0	Denial rate increase + litigation risk; member backlash if coverage perception erodes; regulatory penalty risk (medical loss ratio limits)

Example Decisions (Healthcare Payer)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy prior authorization AI; mandatory human review for high-cost/rare procedures. Phased rollout	+1	0	+1	+2	Proven ROI (15-20% claim speed improvement); human safeguards manage denial litigation; clear pathway
Deploy fraud detection + waste reduction AI; automated payment integrity for high-variance claims	+2	+1	0	+3	High strategic fit (cost reduction); proven technology; manageable execution; low tail risk
Deploy risk stratification AI + predictive models for preventive care targeting; partner with providers	+1	0	+1	+2	Population health upside; execution depends on provider cooperation; reduces adverse selection risk

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Aggressive claim denial automation (>70% auto-denials); minimize human review	0	+1	-2	-1	Cost reduction upside; high denial litigation + regulatory backlash risk; member satisfaction damage

Industry 5: FINANCE (Bank + Insurance, Underwriting & Trading)

Strategic Priorities (2026-2030):

- AI underwriting + fraud detection (core profitability levers)
- Synthetic fraud detection (arms race; ongoing investment required)
- Fair lending compliance (regulatory scrutiny; accuracy/explainability trade-off)
- Back-office optimization (claims, KYC/AML, document review)

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	-1 to +3	+1	High if improving underwriting ROI with compliance safeguards; moderate if defensive/cautious
Execution Risk	-2 to +2	-1	Regulatory approval adds timeline; fair lending audits add complexity; talent availability varies
Tail Risk	-3 to +2	-1	High tail risk (systemic, discrimination litigation, regulatory backlash)

Example Decisions (Finance)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy AI underwriting (with explainability requirements); human review for denials >threshold	+2	0	+1	+3	Core profitability; human safeguards; regulatory pathway clear
Deploy next-gen synthetic identity fraud detection (partnership with specialized vendor)	+1	0	0	+1	Defensive (staying in arms race); proven technology; manageable cost
Deploy AI trading; autonomous for 80%+ trades; minimal oversight	+2	+1	-3	0	High profit upside; systemic risk if model fails; regulatory backlash certain
Invest in bias auditing + fair lending compliance infrastructure upfront	+1	+1	+1	+3	Defensive; reduces regulatory penalty risk; enables faster future deployment

Industry 6: CONSULTING (Big Four / MBB Archetype, \$20-25B Revenue, 40K Employees)

Strategic Priorities (2026-2030):

- Copilot deployment in delivery (research, analysis, proposal development, client presentations)
- Vertical-specific AI expertise practices (Financial Services AI, Healthcare AI, Manufacturing AI)
- Premium AI governance + compliance offerings (regulatory uncertainty, high-margin)
- Talent repositioning (junior analysts -> complex, judgment-intensive work; AI handles routine analysis)

- Pricing model transition (value-based, outcome-based; away from time-and-materials)
- Competing against AI-enabled boutiques and client in-house AI capabilities

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	-1 to +3	+1	High if building vertical expertise + deploying copilots + pricing innovation; moderate if generic services without differentiation
Execution Risk	-1 to +1	0	Copilot adoption straightforward; vertical expertise requires hiring + client relationships; pricing models hard to change
Tail Risk	-1 to +2	0	Junior talent commoditization risk; pricing pressure from AI boutiques; competitive disintermediation as clients build in-house AI capabilities

Example Decisions (Consulting)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy copilots firm-wide (research, analysis, proposals); maintain junior hiring; invest in vertical AI expertise (50 specialists)	+2	0	+1	+3	Productivity gains proven; vertical expertise defensible; junior transition managed

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Build premium AI governance + compliance offerings (regulatory focus); differentiate from AI-native boutiques	+3	+1	+1	+5	High strategic fit (regulatory uncertainty real); premium pricing justified; execution feasible with talent
Pilot outcome-based pricing with 5 trusted clients (lower-risk engagements); document value creation	+1	-1	+1	+1	Addresses pricing pressure; pilots learning; adoption risk high (new model for clients + firm)
Hold on copilot deployment; maintain traditional leverage model; wait for AI disruption to stabilize	-1	+2	+1	+2	Low execution risk; but strategic lag; competitors moving faster; talent attrition likely
Aggressive junior headcount reduction (>40%); rely on AI for analyst-level work	+1	-1	-2	-2	Short-term cost savings; destroys talent pipeline; client delivery quality at risk; future partner pipeline broken

Industry 7: LAW (AmLaw 50 Firm, Billable Hour Economics, Partner/Associate Leverage)

Strategic Priorities (2026-2030):

- AI-assisted legal research, due diligence, and contract review (efficiency gains within billable model)

- Bar rule compliance for AI-generated work product (jurisdiction-specific; rapidly evolving)
- Malpractice liability management (attorney review protocols for AI output)
- Associate leverage model adaptation (AI handles routine tasks; associates redeployed to complex work)
- Pricing model defense or transition (billable hour under pressure from client demands for efficiency)
- Competition from legal AI platforms (Harvey.ai, CoCounsel) and alternative legal service providers

Key Economic Context:

- Revenue driven by billable hours x realization rate x partner/associate leverage ratio
- AI efficiency gains create a paradox: better for clients, potentially destructive to revenue model
- Partner economics depend on associate leverage (billing associates at 3-4x cost); if AI replaces associate work, leverage model erodes
- Bar associations in multiple jurisdictions actively developing rules on AI-generated work product
- Malpractice insurance implications for AI-assisted work not yet settled

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	-1 to +2	+1	High if piloting AI tools while managing bar compliance; moderate if fully committing to pricing model change; low if ignoring AI entirely
Execution Risk	-2 to +1	-1	Bar rule uncertainty adds complexity; malpractice review requirements slow adoption; partner resistance to pricing change
Tail Risk	-2 to +2	-1	Malpractice exposure if AI-generated work has errors; bar discipline risk if rules violated; revenue erosion if billable model disrupted without replacement

Example Decisions (Law)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Pilot AI research tools (Harvey.ai/CoCo unsel) in 3 practice groups; mandatory attorney review; bar rule compliance assessment per jurisdiction	+2	0	+1	+3	Proven tools; phased; attorney review manages malpractice; bar compliance addressed
Deploy AI-assisted contract review firm-wide; invest in quality control infrastructure; maintain billable rates	+1	-1	0	0	Strategic value; execution harder at scale; bar rule compliance across jurisdictions complex; revenue model preserved near-term
Transition to alternative fee arrangements with 5 major clients; maintain billable for rest	+1	-1	+1	+1	Addresses client demand; pilots learning; risk managed via limited scope; partner buy-in uncertain
Reduce associate class by 30% over 24 months; redeploy remaining to complex work; AI handles routine	0	-1	-1	-2	Cost savings real; but leverage model damaged; staffing large matters becomes harder; lateral partners may leave

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Build proprietary legal AI platform for contract analysis	+1	-2	-1	-2	Differentiation potential; but massive capex; competing against well-funded legal AI startups; no AI talent in-house
Hold: defer AI adoption; maintain traditional model; monitor bar rule developments	-1	+2	0	+1	Low execution risk; but strategic lag; competitors gaining efficiency; clients demanding AI-enabled delivery

Industry 8: MANUFACTURING (Heavy Manufacturing, 28 plants)

Strategic Priorities (2026-2030):

- Predictive maintenance (downtime reduction, equipment life extension)
- Production optimization (throughput, quality, energy efficiency)
- Quality inspection AI (defect detection, process control)
- Labor transition (retraining, "no-layoff" agreements, union cooperation)

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	-1 to +2	0	High if prioritizing high-ROI plants (8-12 vs. all 28); moderate if spreading capex thin
Execution Risk	-1 to +2	0	Predictive maintenance is proven; OT/IT integration is complex; equipment retrofitting labor-intensive

Dimension	Range	Typical	Notes
Tail Risk	-2 to 0	-1	Union relations if labor displacement not managed; supply chain dependencies if suppliers not ready

Example Decisions (Manufacturing)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy predictive maintenance in 8 highest-ROI plants (pilot); phased OT/IT integration	+2	+1	+1	+4	Proven tech; 3.2-year payback; high ROI on priority plants; phased approach reduces risk
Announce "no-layoff" retraining agreement with unions; commit \$35-40M over 2 years	+1	+1	+2	+4	Preserves labor relations, safety culture; union cooperation enables faster automation; long-term competitive advantage
Deploy warehouse automation + logistics integration across all plants simultaneously	+1	-2	-1	-2	High labor displacement without transition plan; execution risk high (OT/IT complex); union friction likely
Hold on manufacturing AI; focus on operational efficiency (cost controls, headcount management)	-1	+2	+1	+2	Defensive; low execution risk; but strategic lag if competitors gain efficiency advantage

Industry 9: LOGISTICS (Freight/3PL/Warehouse, 5,000+ vehicles)

Strategic Priorities (2026-2030):

- Route optimization + fuel efficiency (cost reduction, proven ROI: \$180-200M potential)
- Predictive vehicle maintenance (downtime reduction, equipment life)
- Driver assistance + safety systems (accident reduction, driver acceptance)
- Autonomous vehicle pilots (partnerships with Waymo/Aurora; regulatory uncertainty; long timeline)

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	-1 to +2	+1	High if prioritizing route optimization + driver acceptance; moderate if overcommitting to unproven AV
Execution Risk	-1 to +1	0	Route optimization is proven; driver adoption varies by age cohort; AV regulatory timeline uncertain
Tail Risk	-2 to +2	-1	Driver resistance; union concerns; last-mile profitability limits; autonomous vehicle regulatory risk

Example Decisions (Logistics)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Deploy route optimization to 2,000 trucks (pilot: 40% of fleet); driver engagement program	+2	+1	+1	+4	Proven tech; \$45-50M annual savings; phased approach allows learning; driver buy-in investment

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Partner with Waymo/Aurora for long-haul autonomous pilots (5-10% of fleet); manage expectations	+1	-1	0	0	Strategic value; regulatory timeline uncertain (2028-2030); partnership de-risks vs. internal development
Deploy autonomous vehicles internally (full long-haul fleet); eliminate driver roles	0	-3	-3	-6	Regulatory approval uncertain; driver/union resistance certain; execution technically and politically infeasible
Accept last-mile profitability limits; optimize only high-density urban routes	+1	+2	0	+3	Realistic; lower total savings; but avoids overselling optimization to unprofitable segments

Industry 10: BIG TECH (Google/Meta/Microsoft/Amazon-Class — Cloud, Ads, Devices, Enterprise Software)

Scope Note: Big Tech in this exercise covers cloud infrastructure, advertising, devices, and enterprise software. AI lab and foundation model development decisions are excluded from participant scope; those dynamics are introduced via facilitator injects only.

Strategic Priorities (2026-2030):

- AI-powered product features (search, advertising, recommendation, cloud services)
- Enterprise AI services (APIs, platform tools, developer ecosystems)
- Cloud infrastructure scaling (GPUs, data centers, inference capacity for enterprise customers)

- AI cost management + margin defense (compute cost inflation, infrastructure investment)
- Regulatory scrutiny navigation (antitrust, data privacy, content moderation)

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	0 to +3	+2	High if investing in AI product features + enterprise services; very high if staying ahead of competition on cloud/platform
Execution Risk	-1 to +2	+1	Massive CapEx required for infrastructure; talent competition fierce; regulatory approval uncertain for some products
Tail Risk	-2 to +2	-1	Antitrust scrutiny + regulatory backlash; margin compression from AI compute cost inflation; competitive disruption from open-source models

Example Decisions (Big Tech)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Invest in cloud/inference infrastructure (\$15B+ capex); build custom chips for enterprise AI workloads	+3	+1	0	+4	Core strategic value; execution capability exists; margin pressure manageable with scale

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Launch enterprise AI APIs + platform tools; compete on cost/performance for enterprise adoption	+2	0	0	+2	Platform leverage; proven go-to-market; execution straightforward; tail risk low
Integrate AI deeply into core products (search, ads, cloud); expect 5-10% productivity gain	+2	+1	+1	+4	Strategic fit high; user experience upside; execution proven; regulatory risk limited (private integration)
Defend market against open-source model commoditization; cut enterprise AI margins to 10-15%	+1	0	-1	0	Defensive pricing; maintains enterprise stickiness; profitability erodes; shareholder backlash risk

Industry 11: B2B/B2C SaaS (Workday/Salesforce/SAP-Class, ~100K+ employees)

Strategic Priorities (2026-2030):

- AI feature integration into products (copilots, predictive analytics, automation)
- Pricing model evolution (AI features bundled vs. premium tier)
- Competitive threat from AI-native startups (simpler, cheaper, specialized products)
- Margin pressure from AI infrastructure costs (training, inference, compute)
- Customer retention + upsell via AI (lock-in effect)

Scoring Ranges

Dimension	Range	Typical	Notes
Strategic Fit	0 to +3	+1	High if bundling AI features + defending margins; moderate if uncertain on pricing/competition
Execution Risk	-1 to +1	0	AI feature integration proven; pricing model change is hard; competitive response unpredictable
Tail Risk	-2 to +2	-1	Pricing pressure + margin compression; customer churn if AI features disappoint; startup disruption from specialized competitors

Example Decisions (B2B/B2C SaaS)

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Bundle AI copilots into core product; include in standard SKU; no premium pricing upside	+1	+1	0	+2	Maintains customer lock-in; execution straightforward; margin pressure from cost (no revenue offset)
Integrate AI prediction + automation deeply; sell as premium tier at 30-50% price increase	+2	0	+1	+3	Strategic fit high (differentiation); proven pricing model; execution proven; churn risk manageable

Decision	Strat Fit	Exec Risk	Tail Risk	Total	Rationale
Build AI-native product for vertical (e.g., HR-specific AI recruiting); compete on cost + specialization	+2	-1	0	+1	High strategic value (new market); execution complexity (new product); adoption uncertain
Maintain traditional product; defer heavy AI integration until market stabilizes	0	+2	-1	+1	Low execution risk; but competitive lag; startup disruption risk rises; customer churn risk

Using Baselines During Industry-Level Scoring

1. Participant proposes decision. Identify the industry (Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS).
2. Reference the industry baseline. Check typical Strategic Fit, Execution Risk, Tail Risk ranges for this specific industry.
3. Calibrate your score. If participant decision is better than baseline, score higher; if worse, score lower.
4. Score independently. Remember: all 11 industries are scored separately. A participant covering both Consulting and Law has each decision scored on its own industry-specific merits.
5. Post and move on. Explain score concisely; don't debate. Remember you may be scoring multiple decisions per round (variable based on participant count and industry assignments).

Example

Participant: Consulting participant proposes to "deploy copilots to 50% of consultant base over 90 days; hire 50 AI specialists to build vertical AI practices; maintain junior hiring (zero layoffs in pilot phase)."

Baseline (Consulting): Deploy copilots + maintain junior hiring + build vertical AI = +2 to +3 typical score.

Your calibration:

- Specificity is excellent (50%, 90 days, zero layoffs, vertical hires quantified)
- Timeline is tight but feasible (copilots are proven; vertical hiring is hard but not impossible)
- Commitment to zero layoffs + vertical hiring reduces tail risk and addresses competitive disintermediation

Score: +2 (Strategic Fit) + 0 (Execution Risk) + +1 (Tail Risk) = +3/6 (in line with baseline; executed well).

Note: If the same participant also covers Law, the Law decision will be scored separately based on Law baselines + Law-specific context (billable hour impact, bar compliance, malpractice liability).

Base Case Fallback Scoring (NOT Covered Here)

Industries without explicit participant actions receive **base case fallback scores** from the fallback bank (see *02BaseCaseFallbackBank.md*). Fallback scores are:

- Deterministic: Not varied per participant or round; same fallback applies across all instances
- Small deltas: Typically +/-1 per dimension (not -2 to +2)
- Plausible: Represent defensive but reasonable moves (e.g., cost control, operational efficiency)
- Automatic: Applied by facilitator without participant input

Fallback industries are NOT scored using the baselines in this document. Baselines apply only to explicit decisions submitted by participants.

Summary: When to Score High, Medium, Low (Explicit Decisions Only)

High Score (+2 to +3 range)

- Participant is responding to clear competitive threat (e.g., competitor AI advantage)
- Technology is proven and participant's organization has execution capability
- Scope is realistic (phased, limited, with clear milestones)
- Participant acknowledges and mitigates tail risk (pilot, human oversight, severance plan)

Medium Score (0 to +1 range)

- Strategic alignment is moderate (neither captures opportunity nor avoids disaster)
- Execution is feasible but carries standard risk (integration, talent, regulatory timeline)
- Tail risk is managed but not eliminated

Low Score (-1 to -2 range)

- Decision is defensive or reactive (holding pattern, waiting)
- Execution faces material barriers (talent shortage, long regulatory approval, capital constraints)
- Tail risk is not addressed or is material

Very Low Score (-3 or exception +/-3)

- Decision triggers a red-flag (see 02PlausibilityDecision_Trees.md)
- If red-flag fires, unlock +/-3 exception scoring to reflect severity

Use these baselines to calibrate, not to lock scores. Participants proposing above-baseline decisions earn higher scores. Participants proposing below-baseline decisions earn lower scores. Move briskly; facilitate, don't debate.

03

Participant Packets

Participant Packets — Distribution Guide

Overview

This folder contains all participant-facing materials for Project Threshold V7.4: 11 industry packets (one per industry) and 3 shared reference documents.

Distribution Instructions

1. Before the exercise: Distribute each participant's industry packet (03a–03k) as a pre-read
2. At exercise start: Distribute Rules of Play (03z) and Decision Worksheet (03z) to all participants
3. Decision Worksheet: Print multiple copies per participant (one per industry per round)

Files

Industry Packets (03a–03k)

File	Industry
03aRetail.md	Retail
03bCPG.md	CPG
03cHealthcare_Provider.md	Healthcare Provider
03dHealthcare_Payer.md	Healthcare Payer
03eFinance.md	Finance
03fConsulting.md	Consulting
03gLaw.md	Law
03hManufacturing.md	Manufacturing
03iLogistics.md	Logistics
03jBig_Tech.md	Big Tech
03kSaaS.md	B2B/B2C SaaS

Shared Reference Documents (03_z)

File	Purpose
03zRulesofPlay.md	Comprehensive participant guide — rules, scoring, round flow
03zDecision_Worksheet.md	Printable 12-field decision form — print multiple copies per participant
03zTemplateIndustryPacket.md	Blank template establishing consistent structure across industry packets

Retail Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the CEO of a Top-5 US omnichannel retailer with 4,700+ physical stores, a major e-commerce platform, and 300K+ frontline employees (store associates, distribution center workers, logistics staff). Annual revenue is approximately \$200B with operating margins in the 3-5% range. You compete directly with Amazon and tech-native pure-play e-commerce platforms for consumer share and loyalty. Your 150M+ registered users, same-day/next-day delivery capability across 18 distribution centers, and physical store footprint are your core competitive assets. Your decisions on pricing, personalization, supply chain automation, and labor management cascade across the entire consumer market.

Strategic Context

You operate at the front line of the consumer AI arms race. As a retailer, you control first-party customer data (150M+ registered users with transactional history, browsing behavior, location data, and loyalty program engagement) and the physical/digital touchpoints where purchase decisions happen. AI-driven personalization and inventory optimization are table-stakes against Amazon, but consumer trust is eroding as AI-driven pricing and recommendations increasingly feel invasive and opaque.

Your competitive position is under pressure from multiple directions. Amazon continues to push AI-native shopping experiences that set consumer expectations. Tech-native pure-play e-commerce platforms are gaining share through superior personalization and lower cost structures. Meanwhile, your CPG vendor partners are launching direct-to-consumer (DTC) channels using AI-driven marketing and demand sensing, threatening to disintermediate you on significant SKU volumes. At the same time, you face private-label competition from other retailers who are using AI to optimize their own-brand products against your third-party assortment.

The sector faces a critical inflection point. Aggressive AI deployment can improve margins and efficiency in the short term through dynamic pricing, automated inventory management, and personalized recommendations. But consumer backlash is real: customers are pushing back against perceived price discrimination, invasive tracking, and algorithmic manipulation. Cautious deployment preserves trust and brand equity but risks losing first-mover advantage and margin expansion to competitors willing to move faster.

Cross-industry dynamics shape your strategic landscape. Healthcare sector data privacy regulations (expanding beyond HIPAA) are raising the bar for consumer data practices across all industries, including retail. Finance and professional services regulatory exposure is creating new compliance

requirements for algorithmic pricing and credit-adjacent retail services. Supply chain sector labor automation directly affects your cost structure, fulfillment speed, and logistics reliability. Software and tech sector AI capabilities are what enable your personalization, demand forecasting, and inventory optimization investments — their pricing and platform decisions directly constrain your options.

Objectives

Objective	Target (Banded/Directional)	Driver
Comparable Sales Growth	Modest growth in stores (mature market); material growth in e-commerce (competitive positioning)	AI-driven personalization, demand sensing, omnichannel integration
Gross Margin Expansion	Material expansion; defend against Amazon and private-label pressure	Inventory optimization, pricing discipline, markdown reduction, supply chain efficiency
Return Rate Management	Manage within acceptable range; AI personalization improves conversion but can increase returns	Better recommendation targeting, product-fit algorithms, transparent product information
Same-Day/Next-Day Delivery	Expand capability in high-density metro areas (competitive table-stakes)	Warehouse automation, AI route optimization, demand-responsive fulfillment
Frontline Employee Retention & Morale	Maintain workforce stability; critical for service quality	Careful automation rollout, retraining programs, transparent communication with unionized workforce

Constraints

Constraint	Impact	Implications
Legacy IT Systems	Point-of-sale systems, inventory management, and e-commerce platforms are 10+ years old; full modernization requires 2-3 years and \$500M+ capex	Integrating third-party AI APIs introduces latency, reliability, and data governance risks; limits agility and speed-to-market for AI features

Constraint	Impact	Implications
Physical Store Productivity & Labor	Stores have high fixed labor cost; 300K+ employees, 60% unionized in distribution centers and logistics	AI-driven automation (checkout, customer service, warehouse) directly affects workforce; layoff announcements risk strikes during peak seasons; labor relations constrain speed of automation
Consumer Data Privacy	150M+ registered users represent a high-value data asset and a high-value breach target; CCPA, GDPR, and expanding state-level regulations limit data collection, retention, and monetization	Personalization engines require data; regulation constrains scale and creates compliance burden; any breach would be catastrophic for trust
Omnichannel Integration	Store operations, e-commerce, supply chain, and logistics must work seamlessly across channels	Siloed systems slow AI deployment; inconsistent customer experience across channels undermines personalization and brand trust
Thin Operating Margins	3-5% operating margin in a heavy promotional environment	Significant AI implementation failure could cause 50-100 bps margin hit; limited capacity to absorb large failed bets; capital constraints force prioritization

Resources & Levers

Physical & Digital Assets:

- 150M+ registered users with transactional history, browsing behavior, location data, and loyalty program engagement
- 4,700 stores and 18 distribution centers enabling same-day/next-day delivery in major metro areas
- \$800M annual IT spending; \$150M allocated to AI/ML in 2026 (AWS, Google Cloud, Databricks partnerships)
- Brand trust with 65M+ households; physical store presence as competitive moat against pure-play e-commerce

Potential Paths Forward:

- Demand Forecasting & Inventory Optimization: AI reduces markdown losses, improves in-stock rates, optimizes store assortment. ROI proven; low reputational risk. Highest-confidence investment.
- Personalization & Recommendations: AI drives conversion uplift; but return-rate risk and brand safety concerns if perceived as invasive. Requires careful targeting and opt-in design.

- Dynamic Pricing: AI optimizes prices by location and demand; but consumer backlash risk if perceived as price discrimination. Brand trust concern; requires transparency rules.
 - Checkout & Customer Service Automation: Reduce labor cost through staffing-light stores; risk to customer experience and labor relations. Sensitive with unionized workforce.
 - Supply Chain & Logistics: AI route optimization, warehouse automation, demand-responsive fulfillment improve margins. Integration complexity high but ROI is material.
-

AI Adoption Arc — Foundation Phase

Foundation (2025 - Q1 2026): Retail AI deployment is concentrated in low-risk operational areas. Demand forecasting and inventory optimization pilots are running in select distribution centers and store clusters, showing modest but validated ROI through markdown reduction and improved in-stock rates. A personalization engine was deployed to 20% of e-commerce traffic in Q4 2025, lifting conversion but also increasing return rates — net incremental profit underperformed projections. Consumer-facing AI remains limited; organizational skepticism is high among store operations leadership. The \$150M AI/ML budget for 2026 is approved but uncommitted beyond current pilots. The workforce is watching closely — any signal of large-scale automation will trigger union attention. Margin impact so far: low, but directionally positive on the operational side.

Strategic Considerations

1. The trade-off between operational AI and consumer-facing AI defines strategic posture. Inventory optimization, demand forecasting, and supply chain automation have proven ROI and low reputational risk. Consumer-facing personalization offers higher upside but carries brand and privacy risk — consider where each fits given your competitive position against Amazon.
 2. Price transparency is a strategic lever, not just a compliance choice. Dynamic pricing can optimize margins, but consumer backlash erodes the trust advantage that distinguishes you from Amazon. Weigh short-term margin gains against long-term brand equity costs — particularly around location-based or profile-based price discrimination.
 3. Employee enablement and automation are not mutually exclusive, but sequencing matters. AI-equipped associates improve service quality and reduce turnover; aggressive automation cuts costs but risks labor disruption — especially with a unionized workforce. How you frame AI internally ("associate support" vs. "labor replacement") affects adoption and retention.
 4. Data sharing with vendors requires balancing mutual value against competitive exposure. Logistics partners need demand visibility to perform well, but proprietary insights are strategic assets. Consider which data creates shared value and which confers competitive advantage — share the former, protect the latter.
 5. Retail dynamics shift faster than annual planning cycles can capture. Continuous monitoring of consumer sentiment, competitive AI adoption, and pricing dynamics enables adaptive strategy. Consider quarterly reassessment of personalization intensity, pricing rules, and vendor terms rather than rigid multi-year commitments.
-

CPG Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the CEO of a diversified consumer packaged goods company with 35 iconic brands spanning food, beverage, household products, and personal care. You employ 180K people globally, operate 45 R&D centers and 28 manufacturing plants, and source across 45 countries with vertical integration from raw materials to finished goods. Annual revenue is approximately \$50B with operating margins around 14%. You hold 85%+ household penetration in North America across your brand portfolio. Your brands compete against retailer private-label products, Amazon's growing own-brand portfolio, and an emerging wave of direct-to-consumer competitors. Your pricing power, innovation pipeline, and retailer relationships define your competitive position.

Strategic Context

You sit at the intersection of brand equity, manufacturing scale, and an increasingly hostile distribution landscape. AI adoption decisions affect every dimension of your business — from R&D acceleration and marketing efficiency to demand forecasting and the existential question of whether to build direct consumer relationships or remain dependent on retail partners.

On the innovation and marketing side, AI is accelerating product development cycles and driving material cost savings in content generation and campaign targeting. A pilot using AI-assisted formulation and consumer preference modeling has demonstrated meaningful cycle-time reduction, and AI-generated marketing content has matched agency-quality conversion performance at a fraction of the cost. The opportunity is real. But brand safety is a genuine constraint: AI-generated human faces in a household products campaign triggered significant negative social sentiment, and consumers are increasingly skeptical of AI-generated content that feels inauthentic.

Your distribution economics are under structural pressure. Your top 5 retail partners represent 55% of revenue and are demanding real-time access to your demand forecasts, supply chain visibility, and consumer insights — framing it as a "partnership requirement" while building private-label products in your most profitable categories. You are simultaneously considering DTC launches that would bypass these retailers entirely, using AI-driven personalization, pricing, and customer service. But aggressive DTC expansion risks retailer retaliation through shelf-space reduction and accelerated private-label competition.

Cross-industry dynamics shape your operating environment. Healthcare sector data privacy regulations are raising the bar for consumer data practices across all packaged goods, particularly in personal care

and food safety claims. Finance and professional services regulatory frameworks increasingly govern AI-driven pricing and promotional practices. Supply chain sector labor automation and logistics reliability directly affect your manufacturing cost structure and go-to-market speed. Software and tech sector AI platforms are the backbone of your R&D, marketing, and demand-sensing investments — their capability roadmaps and pricing decisions constrain your deployment options.

The fundamental tension: aggressive AI deployment can accelerate innovation, reduce costs, and build direct consumer relationships. But it risks alienating retail partners (your primary revenue channel), eroding brand trust through inauthentic content, and triggering regulatory scrutiny. Moving cautiously preserves relationships and trust but risks losing competitive ground to faster-moving rivals.

Objectives

Objective	Target (Banded/Directional)	Driver
Brand Share Growth	Defend market share or achieve modest growth across key categories	Product innovation, marketing effectiveness, shelf-space defense, brand loyalty
Marketing Efficiency	Material reduction in marketing spend through AI-driven targeting and content generation	AI content automation, programmatic media buying, predictive campaign optimization
R&D Acceleration	Meaningful acceleration of development cycles through AI-assisted ideation, formulation, and testing	AI-driven consumer preference modeling, automated formulation, accelerated testing protocols
Retailer Relationships	Maintain shelf space and prevent private-label encroachment; manage increasingly dominant retail partners	Strategic data sharing, joint business planning, category leadership positioning
DTC Launch Capability	Build direct brand channels; own customer relationships; capture premium pricing	AI-driven personalization, demand sensing, customer service automation, fulfillment partnerships

Constraints

Constraint	Impact	Implications
Retailer Power Dynamics	Top 5 retailers represent 55% of revenue; they demand data transparency, threaten private-label expansion, and control shelf space	Limited pricing power; must balance data sharing (relationship retention) against competitive transparency (strategic risk); DTC launches risk retailer retaliation
Supply Chain & Manufacturing Complexity	Global sourcing across 45 countries, 28 manufacturing plants, long lead times; most suppliers lack digital readiness	Demand forecasting AI adds value but requires supplier ecosystem visibility; integration is slow and capital-intensive; disruption risk from geopolitical and climate factors
Brand Safety & Consumer Trust	Branded products carry emotional and trust components; AI-generated marketing content and AI-driven pricing carry brand risk if consumers perceive inauthenticity	AI content must be transparent; fake testimonials, AI-generated human faces, and non-disclosed AI content create reputational liability; brand equity is the core asset
R&D and Innovation Complexity	Product development involves regulatory approval (food safety, cosmetics), consumer testing, and brand alignment across 35 brands	AI can accelerate but cannot replace human judgment on safety, taste, and brand fit; regulatory timelines are not compressible; failed launches damage brand equity
Thin Margin for Error	~14% operating margin with rising input costs and promotional intensity; capital must fund both operational transformation and growth investments	Simultaneous investment in R&D acceleration, DTC, marketing automation, and supply chain modernization exceeds available capital; must sequence and prioritize

Resources & Levers

Brand & Consumer Assets:

- 35 iconic brands with 85%+ household penetration in North America
- Direct consumer survey data, social listening, retailer POS data, 250M+ consumer interactions annually
- \$2.1B global marketing spend; in-house creative, agency, and influencer relationships
- Deep category expertise and consumer insight across food, beverage, household, and personal care

Innovation & Manufacturing Assets:

- 45 R&D centers with established AI-assisted formulation and testing capabilities (pilot stage)
- 28 manufacturing plants with vertical integration from raw materials to finished goods
- Sourcing relationships across 45 countries; procurement scale advantages

- Existing AI pilot results: 6-month cycle reduction in product development, validated incremental revenue from accelerated launches

Potential Paths Forward:

- AI-Driven R&D & Product Development: AI generates formulation ideas, predicts consumer preferences, and accelerates testing cycles. High ROI; moderate execution risk (requires domain expertise to validate AI outputs).
 - Marketing Automation & Content Generation: AI generates copy, targeted campaigns, and influencer recommendations. Cost reduction is proven; brand safety risk is real (AI-generated claims, synthetic imagery, fake testimonials).
 - Demand Sensing & Supply Chain Optimization: AI improves forecast accuracy, reduces inventory waste, and optimizes production scheduling. ROI is high; requires supplier ecosystem visibility and digital readiness.
 - Direct-to-Consumer (DTC) Expansion: Own brand channels bypass retailers; AI-driven personalization, dynamic pricing, and customer service automation improve DTC unit economics. Retailer relationship risk is the primary constraint.
 - Data Monetization & Insights: First-party consumer data from retail partnerships and DTC channels enables AI-driven consumer insights. Can be shared with partners or protected as competitive advantage — partnership value versus competitive edge tension.
-

AI Adoption Arc — Foundation Phase

Foundation (2025 - Q1 2026): CPG AI deployment is concentrated in back-office and operational applications. R&D teams are piloting AI-assisted formulation and consumer preference modeling with promising early results — one pilot demonstrated a 6-month reduction in development cycles and \$35M in incremental revenue from accelerated launches. Marketing is testing AI-generated content for product descriptions and campaign copy, achieving substantial cost savings versus agency work with matched conversion performance. However, an AI-generated campaign using synthetic human faces triggered meaningful negative social sentiment, underscoring brand safety risks. Demand forecasting pilots are running across select product lines with modest accuracy improvements. Consumer-facing AI is minimal; the organization remains cautious about anything that touches the brand directly. The investment pipeline for 2026 is approved but largely uncommitted beyond current pilots. Margin impact so far: low, but the R&D and marketing pilots point to material upside if scaled carefully.

Strategic Considerations

1. R&D acceleration and marketing efficiency offer the highest-confidence returns. Both have validated ROI and lower reputational risk than consumer-facing personalization. Margin gains of 150-200 bps through AI-assisted product development and content automation are achievable — the question is sequencing and organizational readiness.

2. DTC expansion creates a tension between margin improvement and retailer relationships. DTC channel economics are favorable (higher margin, direct consumer data), but losing shelf space from retailers is catastrophic. Consider how to frame DTC as brand-building and consumer engagement rather than disintermediation of retail partners.
 3. Retailer data demands require careful strategic positioning. Sharing aggregated insights maintains partnership goodwill; sharing tactical data erodes competitive advantage. The balance between collaborative joint business planning and protecting proprietary consumer insights is a recurring negotiation — consider where the line falls for each retailer relationship.
 4. Brand authenticity and AI transparency are linked. Consumer trust depends on knowing what is human-created and what is AI-generated. Consider the reputational implications of undisclosed AI content — authenticity is a differentiation lever against pure-play e-commerce and private-label competitors.
 5. Capital constraints force sequencing decisions. R&D acceleration, DTC expansion, marketing automation, and supply chain modernization cannot all be funded simultaneously. Proven-ROI initiatives in early rounds build the financial runway for more transformational bets later.
-

Healthcare Provider Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the chief executive of a large integrated health system: 20+ hospital facilities, 300+ ambulatory care sites, and a clinical provider network serving 3M covered members through an affiliate health plan. You employ approximately 50,000 clinical and administrative staff (physicians, nurses, clinical specialists, administrators) and generate ~\$6B in annual healthcare operating revenue on the provider side. You compete with regional health systems and national hospital networks on clinical quality, physician recruitment, and operational efficiency. Your decisions on AI-driven clinical diagnostics, care coordination, physician workflow, and administrative operations directly impact patient safety, clinical outcomes, regulatory standing, and financial performance.

Strategic Context

You are a systemically important clinical care provider operating in one of the most heavily regulated sectors in the US economy. Every AI deployment you consider must navigate overlapping oversight from the FDA (clinical device approvals), CMS (Conditions of Participation, reimbursement), state medical boards (physician practice standards), and HIPAA (patient data governance). The regulatory approval pathways for clinical AI are distinct from operational AI, but both converge on common themes: explainability, fairness, validation, and governance.

AI adoption status: You have demonstrated strong results in back-office and administrative AI. Medical coding automation improved billing accuracy by 12%. Prior authorization prediction tools have reduced physician documentation burden in targeted specialties. Patient risk stratification models are operational across your ambulatory network. However, clinical AI — diagnostic radiology, pathology interpretation, clinical decision support — remains in controlled pilot phases. Regulatory uncertainty around clinical AI explainability, validation requirements, and physician liability is the primary drag on clinical deployment.

Cross-industry impacts: Your AI decisions ripple outward. Health insurers and payers depend on your clinical data quality and coding accuracy; AI-driven changes to your documentation practices affect their claims processing and coverage determinations. Pharmaceutical and medical device companies are watching your clinical AI pilots as potential distribution channels. Technology vendors (Epic, Google DeepMind Health, IBM Watson Health) compete for your platform business and your clinical validation partnerships. Meanwhile, workforce dynamics in your region are shaped by the same labor market pressures affecting manufacturing and logistics employers competing for technical talent.

The core tension: Superior AI accuracy in diagnostics often requires black-box models that regulators

find problematic. Compliance-friendly interpretable models sacrifice accuracy. You must navigate this trade-off while managing the additional reality that clinical AI deployment requires expensive, multi-year FDA/CMS approval cycles (18-24 months per tool), whereas administrative AI (coding, documentation, prior auth) can move faster with lower regulatory risk. Investment trade-offs between quick-win administrative AI and long-cycle clinical AI are constant. And none of it matters if physicians refuse to adopt the tools — accuracy alone does not drive clinical adoption; workflow integration and burden reduction do.

Objectives

Objective	Target (Banded/Directional)	Driver
Clinical Quality & Patient Safety	Measurable reduction in diagnostic errors, adverse events, and readmissions through AI-assisted decision support	Patient outcomes, physician accountability, malpractice risk reduction, quality metric performance
Operational Efficiency	Material cost reduction in medical coding, documentation, and prior authorization processing	Physician burden reduction, billing accuracy, administrative cost compression
Regulatory Compliance	Zero significant FDA, CMS, or state medical board enforcement actions; maintain CMS Conditions of Participation	Clinical AI approval pathway navigation, HIPAA compliance, governance infrastructure
Physician Adoption & Satisfaction	Physician acceptance and active use of deployed AI tools; reduced administrative burden per physician	Workflow integration quality, perceived clinical value, autonomy preservation, documentation time reduction
Care Coordination & Patient Outcomes	Improved chronic disease management, reduced readmissions, better transitions of care through AI-driven population health tools	Risk stratification accuracy, cross-facility data integration, care team adoption

Constraints

Constraint	Impact	Implications
Regulatory Pathway Complexity	FDA pathway for clinical diagnostic AI is 18-24 months per novel system; CMS coverage and reimbursement decisions add 12+ months; HIPAA data governance requires strict access controls	Multi-year lead times for clinical AI deployment; high compliance cost (\$5-10M per diagnostic AI tool); must plan regulatory engagement 12-18 months before intended deployment
Physician Liability & Malpractice Risk	Physicians are personally liable for clinical decisions; AI-recommended diagnoses that lead to patient harm create cascading liability to malpractice insurer and hospital system	Physicians resist AI perceived as reducing autonomy; malpractice insurers scrutinize AI use; human-in-the-loop design is non-negotiable for clinical decisions
Clinical Validation Requirements	FDA requires external clinical validation for diagnostic AI; internal validation is insufficient; clinical trials cost \$5-10M per tool and require regulatory pre-approval	Must partner with academic medical centers; validation timelines are 18-24+ months; expensive and slow; no shortcuts
Physician Workflow Integration	Accuracy does not equal adoption; if AI increases physician validation burden, physicians resist even high-accuracy tools; physician time is scarce and expensive	Workflow redesign is as important as model accuracy; must design AI to reduce documentation and validation overhead, not add to it
Data Privacy & Patient Trust	HIPAA requires strict patient data controls; patient consent for AI use in clinical care is an emerging requirement; data breaches expose PHI with regulatory penalty and reputation damage	Complex data governance across clinical systems; breach affects clinical operations and patient relationships simultaneously; emerging state regulations on AI transparency in healthcare
Legacy System Integration	EHR system is 12+ years old; custom middleware required for clinical AI integration; HIPAA-compliant integration adds complexity	Aging infrastructure delays deployment; modernization is expensive and risky; integration timelines often exceed AI development timelines

Constraint	Impact	Implications
Talent & Cost Constraints	40-person data science team shared with payer operations; limited clinical validation expertise; clinical AI specialists (radiology, pathology, cardiology) are expensive and scarce	Competition for AI talent is fierce; clinical validation expertise is especially expensive; budget tension between clinical and administrative AI investment

Resources & Levers

Data & Clinical Capacity:

- 1.2M inpatient admissions/year, 15M outpatient visits/year
- Real-time clinical outcomes and risk stratification data
- EHR data on 3M covered members with longitudinal patient records
- Advanced analytics platforms and clinical data warehouses

Technology & Talent:

- In-house data science team (40 clinical/operational analysts, shared with payer operations)
- Partnerships with AI vendors (Epic, Google DeepMind Health, IBM Watson Health)
- \$50M annual technology spend; \$15M allocated to AI in 2026 (shared budget)

Regulatory Access & Relationships:

- Established relationships with FDA, CMS, state medical boards
- Access to regulatory guidance and advance warning of policy changes through board contacts
- Regional medical societies and physician advisory groups

Capital & Infrastructure:

- ~\$6B healthcare provider operating revenue annually
- 20+ hospital facilities, 300+ ambulatory care sites, robust physician network
- Sufficient capital to absorb losses, fund R&D, acquire specialized AI vendors, or weather regulatory penalties

Potential Paths Forward:

- Clinical Diagnostic AI: Deploy AI for radiology, pathology, cardiology interpretation. High clinical value; long FDA pathway; external validation required; physician workflow integration critical.
- Clinical Decision Support: AI-assisted diagnosis, treatment recommendations, patient risk stratification. Must preserve physician autonomy and accountability.
- Administrative Clinical AI: Medical coding, documentation support, prior authorization prediction/optimization. Reduces physician burden; lower clinical risk; strong ROI profile.
- Population Health & Predictive AI: Risk stratification, disease progression prediction, care coordination targeting. High strategic value; depends on data quality and clinical workflow integration.

- Care Coordination AI: Identify high-risk patients, optimize resource allocation, improve transitions of care. Strategic value; depends on care team adoption across facilities.
-

AI Adoption Arc — Foundation Phase

Foundation (2025 - Q1 2026): Your AI deployment is concentrated in back-office clinical operations: medical coding automation, documentation support, and patient risk stratification are operational and delivering measurable ROI. Diagnostic AI pilots (radiology, pathology) are running in controlled clinical settings with promising accuracy results, but regulatory uncertainty on FDA validation requirements and CMS Conditions of Participation is preventing broader rollout. Physician sentiment is mixed — administrative AI that reduces documentation burden is welcomed, but diagnostic AI that increases validation overhead faces resistance. You are preparing for anticipated FDA and CMS guidance on clinical AI requirements expected in mid-2026, but timelines are uncertain. Budget tensions persist between funding quick-win administrative AI (coding, documentation, prior auth) that delivers near-term savings and long-cycle clinical AI (diagnostics, decision support) that requires 18-24 month regulatory and validation investment before any deployment.

Strategic Considerations

1. Regulatory pathway planning determines deployment timelines. FDA pre-submission meetings for clinical AI should be filed 12-18 months before intended deployment. CMS coverage and reimbursement timelines add further lag. The cost of moving without regulatory clarity can be years of delay — but excessive caution yields competitive disadvantage.
 2. External clinical validation is a prerequisite, not a nice-to-have. Partnering with academic medical centers and integrating validation into pilot designs from day one satisfies FDA scrutiny and builds clinical credibility. Weigh the cost of early validation investment against the cost of later rejection — internal validation alone is insufficient.
 3. Workflow improvement matters as much as diagnostic accuracy. AI that burdens physicians with additional documentation or validation overhead will face adoption resistance regardless of accuracy. Consider whether each deployment genuinely reduces physician burden or merely shifts it — accuracy without workflow improvement is a failed deployment.
 4. Human-in-the-loop for clinical decisions is both a safety requirement and a liability strategy. Fully autonomous clinical recommendations carry unacceptable malpractice exposure in the current regulatory environment. Surfacing AI insights for physician review preserves autonomy and accountability — the trade-off is speed vs. safety.
 5. Patient data governance and consent shape long-term trust. Explicit disclosure of AI use in care, strict HIPAA compliance, and transparent consent processes build the patient trust foundation that enables future AI adoption. Once patient trust is lost, it does not return quickly.
-

Healthcare Payer Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the chief executive of a major health insurer: 3M covered members, claims processing operations, prior authorization and coverage determination functions, and actuarial modeling capabilities. You generate ~\$35B in annual revenue from insurance premiums and managed care operations, and employ approximately 30,000 administrative, underwriting, and clinical review staff. You compete with other major insurers (UnitedHealth, Anthem, Aetna, Cigna) on cost containment, coverage breadth, member satisfaction, and customer retention. Your decisions on AI-driven prior authorization, claims processing, fraud detection, and coverage determination directly impact medical costs, member access to care, regulatory standing, and operating margins.

Strategic Context

You are a major participant in the US health insurance market, operating under scrutiny from CMS, state insurance regulators, and HIPAA. Unlike clinical AI (which requires FDA approval pathways), your core AI applications — prior authorization, claims processing, fraud detection, coverage determination — operate under insurance regulation frameworks where the rules are evolving rapidly and enforcement is intensifying. The distinction matters: your AI can deploy faster than clinical AI, but the regulatory and reputational consequences of errors are severe and increasingly public.

AI adoption status: You have demonstrated strong ROI in operational AI. Prior authorization AI generated \$22M in savings. AI-driven claims processing reduced fraud loss by 15%. Medical cost forecasting models are improving actuarial accuracy. These are mature deployments, not pilots. However, the next frontier — algorithmic coverage determination, AI-driven member health management, synthetic fraud detection — faces growing regulatory scrutiny around fairness, transparency, and disparate impact. State insurance regulators are increasingly demanding explainability for coverage decision algorithms, and enforcement precedents are forming.

Cross-industry impacts: Your AI decisions create downstream effects across the healthcare ecosystem. Healthcare providers depend on your prior authorization speed and accuracy; AI-driven changes to your approval workflows directly affect physician burden and patient access to care. Pharmaceutical companies watch your coverage determination algorithms as gatekeepers to market access. Technology and consulting firms see your claims data as a platform opportunity. Financial services firms use your combined ratio performance as a signal for healthcare sector investment. Meanwhile, the same AI-generated synthetic content threatening your fraud detection is weaponizing claims data across the entire insurance industry.

The core tension: You can deploy AI faster than any clinical provider — your regulatory cycles are shorter and your ROI is more immediate. But speed creates its own risks. Prior authorization algorithms that deny clinically appropriate care expose you to liability, regulatory enforcement, and reputational damage. Fraud detection models that generate false positives burden honest providers and members. Coverage determination AI that exhibits disparate impact on protected populations invites enforcement action. The fundamental trade-off is between aggressive AI-driven cost containment and the fairness, transparency, and patient access obligations that regulators and the public increasingly demand. Moving fast is your advantage; moving recklessly is existential.

Objectives

Objective	Target (Banded/Directional)	Driver
Medical Cost Management	Meaningful reduction in per-member medical costs through AI-driven prior authorization, fraud detection, and care optimization	Loss ratio improvement, premium competitiveness, operating margin protection
Claims Processing Efficiency	Material reduction in claims processing cost and rework through automation of coding adjudication and denial management	Administrative cost compression, processing speed, accuracy improvement
Fraud & Waste Detection	Maintain competitive parity in fraud detection as AI-generated synthetic fraud methods evolve; reduce fraud loss rate	Ongoing arms race with AI-enabled fraudsters; continuous investment required to maintain detection capability
Coverage Determination Accuracy	Improve prior authorization accuracy: reduce false negatives (inappropriate denials) AND false positives (unnecessary approvals)	Balance cost containment with patient access; regulatory compliance; member satisfaction and retention
Regulatory Compliance & Reputation	Zero significant enforcement actions from CMS or state insurance boards; pass all audits; maintain reputation for fair coverage decisions	Algorithmic transparency, bias testing, disparate impact avoidance, proactive regulatory engagement
Member Retention & Satisfaction	Maintain member satisfaction despite cost containment measures through transparent, fair coverage decisions	Competitive differentiation; retention economics; brand trust

Constraints

Constraint	Impact	Implications
Prior Authorization Liability & Patient Access	AI-driven prior authorization denials expose payer to liability if patient harmed by inappropriate denial; regulators scrutinize false negatives (inappropriate denials) with increasing intensity	Tension between cost containment and patient access is structural; fully automated denials are high-risk; human review for high-impact cases is expensive but reduces liability
Algorithmic Transparency & Fairness	State insurance regulators increasingly demand transparency on coverage decision algorithms; black-box models face scrutiny; fair lending-style algorithmic bias concerns emerging in insurance	Must invest in explainable AI architectures for coverage decisions; disparate impact on protected populations triggers enforcement; proactive bias auditing is becoming table stakes
Fraud Detection Arms Race	AI-generated fraudulent claims (synthetic provider credentials, member data, claim narratives) are emerging and increasingly difficult to detect; current fraud detection model performance degrades as offenders adopt AI	Continuous investment required; next-generation detection capability requires 12-18 month development cycles; standing still means falling behind
Claims Data Quality & Integration	Claims data quality is poor: incomplete coding, late submissions, rework cycles slow AI deployment; data governance and standardization are prerequisites	AI accuracy depends on data quality; garbage in, garbage out; data remediation and standardization must precede or accompany AI deployment
Regulatory Environment	CMS, state insurance boards, and HIPAA all regulate payer AI; coverage determination authority is evolving; regulators increasingly scrutinize "black box" coverage denials	Multiple regulators, multiple compliance frameworks; regulatory environment is tightening, not loosening; enforcement precedents are forming in real time
Legacy System Integration	Claims processing system is 15+ years old; slowly integrates with AI APIs; prior auth workflow redesign is technically difficult	Aging infrastructure delays deployment; modernization is expensive and carries operational risk; integration timelines often exceed model development timelines

Constraint	Impact	Implications
Talent & Resource Competition	Shared data science talent with provider operations; competition for investment resources between clinical and operational AI	AI talent is scarce and expensive; actuarial and data science expertise overlap but are not interchangeable; budget tension between payer and provider priorities

Resources & Levers

Data & Actuarial Capacity:

- Claims data on 3M members with multi-year longitudinal history
- Prior authorization history and coverage determination records
- Actuarial and risk modeling data with member-level granularity
- Advanced analytics platforms and actuarial data warehouses

Technology & Talent:

- In-house data science team (40 clinical/operational analysts + actuaries, shared with provider operations)
- Partnerships with AI vendors (Optum, IBM Watson Health, specialized insurtech firms)
- \$50M annual technology spend; \$15M allocated to AI in 2026 (shared budget)

Regulatory Access & Relationships:

- Established relationships with CMS, state insurance boards, HIPAA enforcement
- Access to regulatory guidance and advance warning of policy changes through board contacts
- Industry trade associations and regulatory working groups

Capital & Market Position:

- ~\$35B annual revenue (insurance premiums + managed care operations)
- 3M covered members across commercial, Medicare Advantage, and Medicaid managed care
- Sufficient capital to absorb losses, fund technology investment, acquire specialized vendors, or weather regulatory penalties

Potential Paths Forward:

- Prior Authorization AI: Optimize coverage approvals, reduce denials, predict member needs. High ROI; liability risk if algorithm denies appropriate care; physician and member backlash if perceived as denying legitimate care.
- Claims Processing & Adjudication AI: Automate coding review, claims validation, denial management. Reduces claims cost; risk of over-denials; regulatory scrutiny on accuracy.
- Fraud Detection AI: Deploy AI-driven fraud detection on claim patterns, provider billing, member claims. High ROI; arms race as fraudsters use AI; continuous investment required.
- Medical Cost Forecasting & Risk Adjustment: Improve member risk stratification, cost forecasting, and actuarial modeling. High strategic value; data quality dependent.

- Member Health Management AI: Identify high-cost, high-risk members; target interventions; coordinate care with providers. Reduces overall medical costs; requires provider coordination and data sharing.
-

AI Adoption Arc — Foundation Phase

Foundation (2025 - Q1 2026): Your operational AI deployment is mature and delivering measurable ROI. Prior authorization AI generated \$22M in savings. AI-driven claims processing reduced fraud loss by 15%. Medical coding adjudication automation is reducing rework cycles. These are no longer pilots — they are enterprise-scale deployments under internal pressure to show continued near-term returns. However, the next wave of AI investment faces a different risk profile. Coverage determination algorithms are drawing regulatory attention around fairness and disparate impact. Your fraud detection models are showing early signs of degradation as AI-generated synthetic claims become more sophisticated. State insurance regulators are signaling increased scrutiny of algorithmic coverage decisions. You are positioned to move faster than clinical AI providers (shorter regulatory cycles, more immediate ROI), but the speed advantage comes with growing compliance risk that your governance infrastructure was not designed to handle at this scale.

Strategic Considerations

1. Prior authorization accuracy and cost reduction are in tension. AI can improve approval speed and consistency, but optimizing purely for cost risks inappropriate denials and regulatory backlash. Consider the balance between false negatives (inappropriate denials that damage member trust) and administrative savings.
 2. Algorithmic transparency and bias testing are becoming table stakes. Regulators will scrutinize prior auth and coverage algorithms for disparate impact. Proactive bias auditing and transparent decision rules reduce regulatory risk — the question is not whether to invest, but how early and how thoroughly.
 3. Human review for high-risk denials reflects a risk-reward calculation. Full automation of denial decisions maximizes cost savings but creates the highest liability exposure. Preserving human judgment for high-impact cases is an insurance premium — weigh the operational cost against the regulatory and reputational downside.
 4. Fraud detection is a continuous arms race, not a deployment milestone. Synthetic fraud and AI-generated identity documents evolve constantly. Budget for ongoing model refresh, vendor partnerships, and 12-18 month capability cycles. Standing still means falling behind.
 5. Provider data integration improves outcomes but requires trust-building. Sharing risk stratification and member health insights with clinical providers creates a longitudinal view neither possesses alone. Consider how to structure data sharing that incentivizes collaboration without creating competitive exposure.
-

Finance Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See `05_Participant_Private_Cards/` and `04_Participant_AI_Adoption_Arcs/` folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document (`03zRulesofPlay.md`) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the CEO of a major US diversified financial institution with retail lending, wholesale lending, wealth management, investment banking, P&C insurance, and life insurance operations. You employ approximately 100,000 people across these business lines. Your balance sheet includes \$250B in deposits, \$500B in assets under management, and \$50B in annual insurance premiums. Your board targets a 12% return on equity. You compete against other universal banks, regional banks, fintech challengers, and insurtech startups for market share across lending, wealth management, and underwriting. Your decisions on AI deployment in underwriting, fraud detection, risk modeling, and customer-facing services affect regulatory dynamics, competitive positioning, and systemic risk across the broader financial services ecosystem.

Strategic Context

You operate in one of the most heavily regulated and data-intensive industries in the US economy. Financial services has a long history of quantitative modeling — actuarial science, credit scoring, algorithmic trading — which means AI adoption is not greenfield. It is an extension and acceleration of existing capabilities. The strategic question is not whether to adopt AI, but how fast, how transparently, and with what governance infrastructure. Your competitive advantage depends on getting this calibration right.

AI is already deployed and delivering returns in your back-office and risk management operations. Fraud detection, claims processing, KYC/AML compliance, and document review are partially automated with demonstrated ROI. The next frontier — AI-driven loan underwriting, insurance underwriting, robo-advisory, and customer acquisition — offers material profit improvement but enters heavily scrutinized regulatory territory. The OCC, Federal Reserve, SEC, FDIC, and state insurance regulators all have overlapping and sometimes conflicting authority over how AI can be used in customer-facing financial decisions. Fair lending laws (ECOA, FHA) impose specific constraints on model complexity and explainability that directly trade off against predictive accuracy.

Your competitive landscape is bifurcating. Fintech challengers are deploying AI-native underwriting and customer acquisition models that move faster than traditional banks, often operating under lighter regulatory frameworks. Insurtech startups are using alternative data and machine learning to underwrite risks your actuarial models miss. At the same time, the largest universal banks are investing billions in AI infrastructure, creating an arms race in talent acquisition and model sophistication. Firms that move too slowly lose market share to fintechs; firms that move too aggressively risk regulatory enforcement,

fair lending violations, and systemic risk events that damage the entire sector.

Cross-industry dynamics are material. Healthcare sector AI adoption is creating new data sources (health-linked underwriting, wellness-based insurance pricing) that raise both opportunity and regulatory questions. Supply chain disruptions and manufacturing sector volatility directly affect your commercial lending portfolio risk. Software and tech sector AI platform decisions constrain your build-vs-buy options and vendor dependencies. Professional services firms — consultants and lawyers — are your critical advisors on AI governance, regulatory compliance, and organizational transformation; their own AI disruption affects the quality and cost of the guidance you receive. Consumer sector dynamics influence your retail lending and credit card portfolios.

The core strategic tension: aggressive AI deployment in underwriting and customer-facing decisions can materially improve ROE and competitive positioning, but regulatory enforcement risk is real and consequences are severe — consent orders, remediation costs, and reputational damage that takes years to repair. Conservative deployment preserves regulatory standing but cedes market share to faster-moving competitors and fintechs.

Objectives

Objective	Target (Banded/Directional)	Driver
Underwriting Accuracy & Profitability	Material improvement in loss ratios and approval rates across lending and insurance	AI-driven loan and insurance underwriting models that outperform traditional actuarial and credit scoring methods
Fraud Loss Reduction	Meaningful reduction in fraud losses; maintain competitive parity as synthetic fraud accelerates	AI-driven fraud detection and synthetic identity detection; continuous model refresh to stay ahead of adversarial AI
Operational Cost Reduction	Significant cost reduction in claims processing, KYC/AML, document review, trading support	Automation of high-volume, rules-based operational processes while maintaining compliance
Capital Efficiency	Improve regulatory capital ratios through better risk modeling; pass stress tests consistently	AI-enhanced risk models that more precisely estimate loss distributions and capital requirements
Deposit Growth & Wallet Share	Expand retail banking and wealth management share through superior personalization	AI-driven customer acquisition, robo-advisory, and personalized financial planning; differentiation vs. fintech

Constraints

Constraint	Impact	Implications
Regulatory Complexity	OCC, Federal Reserve, SEC, FDIC, and state insurance regulators maintain overlapping authority; fair lending rules (ECOA, FHA) impose specific constraints on model design and documentation	Multiple approval pathways for different AI applications; regulatory uncertainty slows deployment timelines; must engage regulators early to clarify requirements rather than deploy and hope for retroactive approval
Model Explainability & Fairness	Black-box ensemble models that maximize accuracy are disfavored by regulators for customer-facing decisions; interpretable models sacrifice 6-8 percentage points of accuracy	Must accept accuracy trade-offs for regulatory compliance on lending and underwriting; can preserve high-accuracy models for internal risk management only; dual-model architectures add complexity and cost
Systemic Risk & Capital Rules	Capital ratios must stay above regulatory minimums; stress tests must be passed; AI model failures in lending cascade into capital depletion	AI deployment in credit decisions carries systemic risk — bad underwriting models produce bad loans at scale; model governance and validation infrastructure is a prerequisite, not an afterthought
Consumer Trust & Brand Risk	Perceived discrimination, unfair lending, or opaque algorithmic decisions damage brand and trigger regulatory action; algorithmic bias litigation is emerging	Consumer-facing AI decisions require explainability, fairness testing, and remediation protocols; reputational damage from a fair lending violation affects deposit growth and wealth management retention
Data Privacy & Cybersecurity	Customer financial data is a high-value breach target; robust security infrastructure required across all AI systems; breach consequences include customer outflow, regulatory penalty, and litigation	AI systems that ingest and process customer data must meet enterprise security standards; data governance, encryption, access controls, and audit trails are non-negotiable infrastructure costs

Constraint	Impact	Implications
Legacy Systems	Core banking and insurance platforms are 12-20 years old; integration with modern AI APIs is slow and expensive	Modernization is a multi-year, multi-hundred-million-dollar investment; interim solutions (API wrappers, middleware) add latency and reliability risk; constrains speed of AI deployment

Resources & Levers

Data & Modeling Assets:

- Decades of credit, claims, customer, and transactional data across lending, insurance, and wealth management
- Proprietary underwriting, fraud detection, and risk models refined over multiple economic cycles
- Real-time data on millions of individual and institutional customers; deep sector knowledge across commercial verticals

Technology & Talent:

- In-house data science teams with 150+ PhD-level researchers and quantitative analysts
- \$250M+ annual technology spend with capacity to invest \$500M-1B in AI, talent, and capability building
- Partnerships with major AI vendors and cloud platforms; enterprise-grade infrastructure

Regulatory Access & Relationships:

- Established relationships with OCC, Federal Reserve, SEC, FDIC, and state insurance boards
- Advance access to regulatory guidance and policy changes through industry groups and direct engagement
- Track record of regulatory compliance and examination cooperation

Capital & Distribution:

- \$250B in deposits, \$500B AUM, \$50B annual insurance premiums providing stable funding base
- Sufficient capital to fund R&D, acquire specialized vendors, invest in talent, and absorb regulatory penalties
- Multi-year relationships with thousands of institutional and corporate clients across all financial services lines

Potential Paths Forward:

- AI Underwriting & Risk Modeling: Deploy AI for loan and insurance underwriting, improving loss ratios and approval rates. High ROI potential; regulatory approval required; accuracy-explainability trade-off is the central design constraint.
- Synthetic Identity & Fraud Detection: Deploy next-generation fraud detection to stay ahead of AI-generated synthetic identities and forged documents. Arms race dynamic; ongoing investment required; competitive parity at stake.

- Robo-Advisory & Wealth Management: Deploy AI-driven portfolio management, financial planning, and customer acquisition. Client acceptance and fiduciary liability risk if recommendations underperform.
 - Trading Systems & Algorithmic Execution: AI-driven trading, market-making, and risk management. High profit potential; systemic risk if models fail; regulatory scrutiny on market manipulation.
 - Compliance & Governance Infrastructure: Invest in model risk management, bias testing, regulatory audit readiness, and remediation protocols. Defensive investment; reduces regulatory penalty risk; enables faster deployment of revenue-generating AI.
-

AI Adoption Arc — Foundation Phase

Foundation (2025 – Q1 2026): AI deployment in finance is concentrated in back-office and internal risk management operations where regulatory requirements are clearer and ROI is well-established. Fraud detection models are operational across transaction monitoring and claims processing. KYC/AML compliance workflows are partially automated. Document review and data extraction tools are deployed in lending operations and insurance claims. Internal risk models use machine learning for loss estimation and stress testing, though regulatory reporting still relies on traditional approaches. Conservative posture on consumer-facing AI: loan underwriting pilots are in controlled testing with human-in-the-loop review, and robo-advisory features are limited to account aggregation and basic portfolio rebalancing. The institution has invested heavily in data infrastructure and model governance frameworks, but the regulatory pathway for deploying AI in lending decisions remains unclear. Talent competition is fierce — data science teams are being recruited aggressively by fintechs and Big Tech. Margin impact so far: modest but positive in back-office efficiency; front-office AI revenue contribution is minimal pending regulatory clarity.

Strategic Considerations

1. Regulatory pathway clarity and deployment speed are inversely correlated — and the order matters. Early engagement with OCC, Federal Reserve, and SEC builds understanding of approval timelines and examination expectations. Deploying into regulatory gray zones invites enforcement action — but waiting for full clarity may mean losing competitive ground. Consider the cost of each.
2. Explainability and accuracy are inversely correlated in many models. Customer-facing lending and underwriting decisions require interpretable models, even at some accuracy cost. Internal risk management can use high-accuracy ensemble methods. Consider which decisions require explainability and which prioritize performance — the regulatory requirements differ.
3. Governance infrastructure is expensive upfront but cheap relative to enforcement. Model risk management, bias testing, audit trails, and fair lending analysis require significant investment. A consent order, customer remediation program, or fair lending enforcement action costs multiples more. The question is how much to invest and when.
4. Human-in-the-loop review for high-risk AI decisions balances speed against liability. Fully automated lending or fraud determinations maximize efficiency but create the highest regulatory exposure. Human review of high-confidence outputs reduces liability and builds organizational trust — weigh the throughput cost against the risk reduction.

5. Fraud detection is a P&L-critical ongoing investment. Synthetic fraud and AI-generated identity documents represent an evolving threat. Standing still on fraud detection capability has direct financial and regulatory consequences — consider this as recurring operational expenditure, not a one-time capital project.
-

Consulting Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the CEO of a top-tier global management consulting firm (Big Four / MBB class). You employ approximately 40,000 people — partners, principals, engagement managers, consultants, analysts, data scientists, and support staff — across strategy, operations, technology, and organizational transformation practices. Annual revenue is approximately \$20-25B with operating margins in the 15-18% range. You serve Fortune 500 and Global 2000 clients across all major industries, with particular depth in financial services, healthcare, manufacturing, and energy. Your traditional business model relies on leveraging junior talent for research, analysis, and deliverable production — associates and analysts working long hours to synthesize data, build models, and produce slide decks. That leverage model is now being disrupted by AI copilots that can perform much of this work faster, cheaper, and around the clock.

Strategic Context

Your firm sits at the center of the AI economy in a way no other industry does. You are simultaneously the subject of AI disruption (AI is automating the core analytical work your junior consultants perform) and the leading advisor on AI transformation for every other sector. Every Fortune 500 company is wrestling with how to deploy AI, and most of them are calling firms like yours for help. This dual exposure — internal disruption plus external opportunity — defines your strategic challenge for the next five years.

The internal disruption is real and accelerating. AI copilots for research, data analysis, slide production, and market sizing are already deployed within your firm and producing meaningful productivity gains. An associate with a well-configured copilot can produce in hours what previously took days. But these gains create uncomfortable questions. If a team of three analysts can now do the work of eight, what happens to your staffing model? Consulting economics have always depended on the pyramid — large numbers of junior staff generating billable hours at high realization rates, supervised by a smaller number of senior partners who own client relationships. AI flattens that pyramid. The work still needs to be done, but it needs fewer people to do it, and the skill profile shifts from "research and production" to "judgment, synthesis, and client management."

The external opportunity is equally significant — and contested. AI transformation advisory is the fastest-growing segment of the consulting market. Every client in every sector needs help building AI strategy, deploying AI tools, managing organizational change, navigating regulatory uncertainty, and rethinking talent models. Your firm's deep industry expertise, global client relationships, and

trusted-advisor positioning make you a natural choice for this work. But you are not the only option. Specialized AI consultancies — firms born in the AI era with deep technical talent and lower cost structures — are winning an increasing share of transformation engagements. Meanwhile, the largest technology companies are bundling advisory services with their AI platforms, and your clients themselves are building in-house AI capability that reduces their dependence on external advisors.

Pricing pressure compounds these dynamics. Your clients understand that AI is making your delivery more efficient. They see the productivity gains, and they want to share in them. Demands for 15-30% rate reductions on AI-assisted engagements are becoming routine. At the same time, time-based billing — the backbone of consulting economics — is fundamentally challenged when AI compresses a two-week analysis into two days. You cannot bill for two weeks of work that took two days. The pricing model must evolve, but the transition is painful: value-based pricing requires a different commercial infrastructure, a different partner skill set, and a different relationship with clients who are accustomed to buying hours.

Cross-industry dynamics are central to your position. Your advice and implementation support directly shape AI strategy in Consumer (retail and CPG transformation), Healthcare (clinical AI deployment, payer analytics), Finance (regulatory navigation, model governance), Supply Chains (manufacturing automation, logistics optimization), and Software & Tech (product strategy, go-to-market). The quality and direction of your consulting engagements cascade through every sector represented in this exercise. At the same time, developments in those sectors feed back into your business: regulatory tightening in healthcare or finance creates demand for compliance advisory; supply chain disruption creates demand for operations consulting; tech platform shifts create demand for technology strategy work.

Objectives

Objective	Target (Banded/Directional)	Driver
Revenue Growth Through AI Advisory	Material growth in AI transformation, governance, and deployment advisory services; target faster-than-market growth in these segments	Client urgency around AI deployment; regulatory uncertainty creating advisory demand; cross-industry AI adoption accelerating
Margin Expansion Through AI-Assisted Delivery	Meaningful margin improvement from AI copilot deployment in delivery; offset by pricing concessions	Copilot productivity gains in research, analysis, deliverable production; net margin dependent on pricing discipline
Talent Retention & Redeployment	Maintain top-quartile retention rates; successfully redeploy junior consultants from routine to complex work	Junior consultant value proposition must evolve; career paths must remain compelling despite AI displacement of entry-level work
Competitive Positioning vs. Specialists	Defend market share in AI transformation advisory; win on depth of vertical expertise, not breadth	Specialized AI consultancies and in-house client teams taking share; differentiation requires deep industry knowledge, not generic AI strategy

Objective	Target (Banded/Directional)	Driver
Utilization & Realization Management	Maintain utilization above sustainable thresholds; protect realization rates amid pricing pressure	AI efficiency compresses engagement timelines; clients demanding discounts; value-based pricing pilots must demonstrate margin protection

Constraints

Constraint	Impact	Implications
Junior Talent Commoditization	AI copilots displace 30-50% of routine analytical work (research, data synthesis, slide production, market sizing); junior consultant utilization declining	The pyramid model that generates margins is eroding; fewer juniors needed per engagement; training pipeline disrupted — new hires spend less time on foundational work; retention risk as top graduates question the value of starting in consulting
Pricing Pressure & Billing Model Disruption	Clients demanding 15-30% rate reductions for AI-assisted delivery; time-based billing erodes when AI compresses two-week analyses into days	Revenue per engagement declining unless pricing model shifts; value-based pricing requires new commercial capabilities, client education, and risk-sharing frameworks; transition takes 12-24 months
Specialized AI Consultancy Competition	AI-native firms (founded 2020-2025) winning 10-20% of AI transformation engagements with lower cost structures, deeper technical talent, and faster delivery	Horizontal "AI transformation" offerings commoditize; must differentiate through vertical expertise and trusted relationships; acquiring or partnering with specialists is expensive and culturally difficult
Client In-House Capability Building	Major clients (Fortune 100) investing in internal AI teams, reducing demand for external consulting on deployment and implementation work	Strategic advisory and change management remain defensible; implementation and analytics work increasingly pulled in-house; engagement scope narrowing

Constraint	Impact	Implications
Talent Retention & Career Path Disruption	Top MBA and undergraduate candidates questioning the value proposition of consulting if AI handles entry-level work; retention of mid-level consultants uncertain as career paths change	Must redefine the junior consultant role around judgment, synthesis, and client interaction; career progression models must adapt; compensation and development programs need redesign

Resources & Levers

Client Relationships & Market Position:

- Established relationships with Fortune 500 and Global 2000 C-suite executives, boards, and institutional investors across all major industries
- Trusted-advisor positioning built over decades; reputation for rigor and quality in strategy and transformation work
- Global delivery capability across 50+ countries; ability to staff complex, multi-geography engagements

Knowledge & Methodology Assets:

- Proprietary consulting methodologies, frameworks, benchmarks, and industry playbooks accumulated over decades
- Deep vertical expertise in Financial Services, Healthcare, Manufacturing, Energy, and Technology sectors
- Extensive case library and transformation playbooks covering 1,000+ enterprise AI deployments

Technology & Talent:

- \$20-25B revenue base providing material investment capacity for AI infrastructure, talent acquisition, and capability building
- 500+ data scientists, AI/ML engineers, and technology specialists embedded across practices
- Strategic partnerships with major AI platform vendors (cloud providers, foundation model companies, enterprise software)
- Internal AI copilot platform deployed across 80% of engagement teams

Potential Paths Forward:

- AI-Assisted Delivery Deployment: Scale copilot use across all practices (research, analysis, deliverable production, data analysis). Proven productivity gains; labor displacement risk; pricing pressure if efficiency is visible to clients. ROI: High. Risk: Pricing erosion and talent disruption.
- Premium AI Governance & Strategy Advisory: Build specialized practice in AI governance, responsible AI deployment, regulatory compliance advisory, and board-level AI strategy. High-margin, defensible positioning. ROI: High but requires 12-18 month build. Risk: Competition from specialized firms and Big Four rivals.

- **Talent Repositioning & Development:** Redeploy junior consultants from routine analytical work to complex synthesis, client relationship management, change management, and emerging service areas. Invest in retraining programs. ROI: Indirect but essential for retention and culture. Risk: Execution difficulty; some roles genuinely eliminated.
- **Value-Based Pricing Pilots:** Shift select engagements from time-and-materials to fixed-fee or outcome-based pricing. Pilot with trusted clients on well-scoped work. Develop playbook for broader rollout. ROI: Margin-protective if well-executed. Risk: Underpricing, scope creep, client resistance to premium tiers.
- **Vertical Specialization Investment:** Double down on 3-4 verticals (Financial Services AI, Healthcare AI, Manufacturing AI, Energy Transition) where firm has deep client relationships and domain knowledge. Hire specialized talent. ROI: Differentiation vs. horizontal competitors. Risk: Opportunity cost; narrowing focus in volatile market.
- **Change Management & Organizational Transformation Services:** Expand offerings around the human side of AI deployment — workforce planning, operating model redesign, culture change, executive alignment. Differentiation vs. technology-only consultancies. ROI: Moderate; high client demand. Risk: Harder to demonstrate tangible outcomes; pricing challenged.

AI Adoption Arc — Foundation Phase

Foundation (2025 – Q1 2026): AI copilot deployment across your firm's delivery teams is underway and producing measurable but uneven results. Research and analysis copilots are deployed to approximately 80% of engagement teams, with productivity gains of 25-40% on routine analytical tasks (market sizing, competitive benchmarking, data synthesis, first-draft deliverable production). However, the quality of AI-generated output varies significantly by use case and requires substantial partner and manager review time — some teams report that quality assurance overhead partially offsets productivity gains. Junior consultant utilization is declining as routine work shifts to AI tools, but redeployment pathways are not yet defined at scale. Early AI advisory engagements are active with financial services and healthcare clients, generating premium revenue but at modest scale relative to total firm revenue. Pricing pressure from clients is beginning to surface but remains manageable — most clients are still adjusting to the new delivery model. Margin impact so far: modest improvement from delivery efficiency, partially offset by pricing concessions on early AI-assisted engagements and investment in AI infrastructure. The talent pipeline is showing early stress: MBA applications to consulting are declining, and junior retention is weakening as associates report uncertainty about career progression.

Strategic Considerations

1. AI copilots improve productivity, but quality gates determine whether that productivity creates or destroys value. Productivity gains from AI-assisted delivery are real, but the net margin improvement is lower than raw automation would suggest once quality assurance costs are included. Speed without rigor destroys client trust — consider where quality gates are non-negotiable vs. where lighter review suffices.

2. Vertical specialization vs. horizontal reach is a defining strategic choice. Generic "AI transformation" consulting commoditizes rapidly. Deep expertise in 3-4 verticals (Financial Services AI governance, Healthcare AI deployment, Manufacturing AI transformation) commands premium pricing and defends against specialist competitors — but narrows the addressable market.
 3. Talent repositioning and AI deployment are inseparable. Using productivity gains to redeploy junior talent toward complex work (client relationships, change management, synthesis) preserves the pipeline that produces future partners. Gutting junior ranks saves margin short-term but destroys long-term organizational health. The talent development model must evolve, not disappear.
 4. Pricing model transition is inevitable but risky to execute. Time-based billing is structurally undermined by AI efficiency. Value-based or outcome-based pricing can work if you maintain quality controls and charge a premium for assurance. Piloting with trusted clients builds the commercial playbook — the risk is getting the sequencing wrong.
 5. AI governance advisory is a high-margin window that is closing. Regulatory uncertainty around AI deployment is intensifying across every sector. Credible guidance on responsible AI, algorithmic fairness, and regulatory compliance commands premium rates — but the window for establishing market position is narrow as Big Four and specialist firms compete for the same space.
 6. Transparency about AI use in delivery builds or destroys trust depending on framing. Clients will discover AI is being used in delivery. Proactive disclosure with quality assurance framing builds trust; concealment and discovery destroys it. Consider how to frame AI-assisted delivery as "better, faster, with the same rigor" rather than as a cost-cutting measure.
 7. Recruitment pipeline health is a leading indicator of firm health. If top graduates stop viewing consulting as a compelling career, the long-term talent pipeline is at risk. Similarly, if partner economics shift too fast (fewer juniors generating leverage), compensation models may need redesign. These are slow-moving but existential dynamics.
-

Law Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are a top-tier global law firm (AmLaw 50 class). ~5,000 attorneys across corporate/M&A, litigation, regulatory compliance, intellectual property, and financial services practices. ~\$5-8B annual revenue, ~35-40% operating margin. Partner billing rates of \$800-1,000+/hr; associate rates of \$400-700/hr. Your leverage model depends on associates performing high-volume legal work — contract review, due diligence, legal research, document drafting — that AI can now perform faster and cheaper. You serve Fortune 500 corporate clients, financial institutions, and institutional investors.

Strategic Context

Your firm operates in a profession where billable time is both the unit of production and the unit of revenue. AI is attacking that foundation directly. Contract review copilots can process in hours what took associates days. Legal research tools can surface relevant precedent in minutes, not half-day research sessions. Due diligence workflows that required teams of associates for weeks can now be accelerated dramatically. Every efficiency gain is, under the current model, a revenue loss. This is not a technology problem — it is a business model crisis that happens to be triggered by technology.

The regulatory landscape compounds the uncertainty. State bar associations are issuing AI guidance on staggered, inconsistent timelines. Some states require disclosure of AI use in legal work product; others are still silent. Malpractice liability frameworks have not caught up — if an AI-assisted brief contains a fabricated citation or a contract review misses a material clause, the supervising attorney bears responsibility, but case law on the standard of care for AI-supervised work is sparse. Your malpractice insurers are beginning to ask pointed questions about AI deployment, and premium adjustments are on the horizon. Every practice group that deploys AI is navigating a patchwork of rules that differ by jurisdiction, practice area, and client expectation.

Meanwhile, your competitive position is eroding from multiple directions. Legal AI platforms — Harvey.ai, LexisNexis+, Westlaw+ — are being adopted directly by corporate general counsel offices, enabling in-house legal departments to handle work they previously outsourced to firms like yours. Alternative legal service providers (ALSPs) are scaling AI-native delivery models for commoditized work at a fraction of your billing rates. Specialized AI-native boutique firms are emerging in focused practice areas, winning mandates on speed, cost, and technical sophistication. Your largest clients are asking hard questions about what, exactly, they are paying \$700/hr for when an AI tool produced the first draft.

Your strongest strategic asset is judgment. Complex litigation strategy, regulatory navigation in

ambiguous environments, structuring novel transactions, advising boards on existential risk — these require the kind of contextual reasoning, relationship management, and professional accountability that AI cannot replicate at present. The firms that thrive will be those that ruthlessly shift their value proposition from document production to judgment and counsel, while using AI to deliver commoditized work faster and cheaper than competitors. But the transition is painful: it means restructuring associate career paths, renegotiating client fee arrangements, and accepting that a significant portion of traditional law firm revenue is being permanently repriced.

The cross-industry dimension is critical. You are not just subject to AI disruption — you are an architect of the rules governing it. Your regulatory compliance practice advises healthcare systems on AI liability, financial institutions on algorithmic fairness, manufacturers on product liability for AI-enabled systems, and technology companies on intellectual property risk. Every other industry in this exercise is a client. The legal frameworks your firm helps shape — through litigation outcomes, regulatory guidance, and transactional structures — directly influence the pace and direction of AI adoption economy-wide. This gives you unique leverage, but also unique exposure: if you get AI governance wrong in your own house, your credibility advising others collapses.

Objectives

Objective	Target (Banded/Directional)	Driver
Revenue Per Lawyer	Maintain or grow RPL despite efficiency-driven pricing pressure; target range ~\$1,100-1,200K	Shift mix toward high-value advisory, AI governance, and complex litigation; use AI to reduce cost of delivery on commoditized work
Operating Margin Preservation	Maintain 35-40% operating margins through delivery model transformation	AI-assisted delivery reduces associate hours on routine work; offset pricing concessions with lower cost-to-serve; expand high-margin practice areas
Associate Development & Retention	Maintain competitive associate recruitment pipeline; reduce first-3-year attrition below industry average	Redesign associate career path from document production to judgment-intensive work; invest in AI-augmented training; preserve value proposition of BigLaw career
Regulatory Positioning	Establish firm as market leader in AI governance, compliance advisory, and regulatory risk management	Build dedicated AI governance practice; leverage cross-industry regulatory expertise; premium pricing justified by complexity and client urgency
Malpractice Risk Management	Zero material malpractice claims attributable to AI-assisted legal work	Mandatory attorney review protocols for all AI-generated work product; AI quality assurance infrastructure; proactive malpractice insurer engagement

Constraints

Constraint	Impact	Implications
Bar Rule Uncertainty	State-by-state AI rules still forming; 15+ states have not issued substantive guidance; disclosure and supervision requirements vary by jurisdiction	Multi-jurisdictional practices must track and comply with divergent rules; risk of inadvertent noncompliance; cost of compliance infrastructure is high and ongoing
Malpractice Liability	AI errors in legal work product (fabricated citations, missed contract provisions, incorrect regulatory analysis) create direct liability exposure; supervising attorneys bear responsibility	Every AI deployment requires human review overlay; net productivity gain is lower than raw AI capability suggests; malpractice insurance premiums increasing for firms with AI deployment
Client Confidentiality	Attorney-client privilege and duty of confidentiality constrain data sharing; client data cannot be used to train third-party AI models without explicit consent	Limits ability to leverage firm-wide data for AI training; each client engagement may require separate data governance protocols; competitive disadvantage vs. AI platforms with broader training data
Associate Leverage Model Disruption	Traditional model generates revenue through associate hours on high-volume work; AI displaces 30-50% of that work within 2-3 years	Revenue model must transition from volume-based (hours billed) to value-based (judgment delivered); associate headcount and compensation structure require restructuring; partner economics affected
Pricing Pressure	Clients see AI improving delivery efficiency and demand proportional rate reductions or alternative fee arrangements; AI-assisted work commands 15-25% lower effective rates	Hourly billing erodes as clients gain visibility into AI-assisted delivery speed; must develop alternative fee structures (fixed fee, outcome-based, subscription) or accept margin compression

Resources & Levers

Knowledge & Data Assets:

- 50+ years of legal precedent databases, contract templates, transaction history, and litigation outcomes across all major practice areas
- Proprietary clause libraries, risk assessment frameworks, and regulatory compliance playbooks covering financial services, healthcare, IP/patent, and corporate governance
- Deep institutional knowledge of regulatory environments across 20+ jurisdictions

Relationship Assets:

- Direct relationships with Fortune 500 general counsel, corporate boards, and C-suite executives across all major industries
- Established relationships with regulators (SEC, FTC, state attorneys general, federal agencies) through decades of regulatory practice
- Trusted advisor status on high-stakes transactions, litigation, and regulatory matters — reputation is a defensible moat

Talent & Infrastructure:

- ~5,000 attorneys including 800+ partners with deep domain expertise; strong recruitment pipeline from top law schools
- Quality assurance infrastructure (partner review protocols, malpractice insurance, conflict checking, ethical compliance systems)
- \$200M+ annual technology spend; capacity to invest \$300-500M in AI infrastructure, talent development, and practice transformation over 3 years

Potential Paths Forward:

- AI-Assisted Legal Research & Document Drafting: Deploy copilots for contract review, due diligence, legal research, brief drafting, and regulatory analysis. High productivity gain (30-50% faster delivery); requires mandatory attorney review; net margin improvement depends on pricing strategy.
- AI Governance & Regulatory Compliance Practice: Build dedicated practice advising clients on AI liability, algorithmic fairness, regulatory compliance, and governance frameworks. Premium pricing; high demand from all industries; leverages firm's cross-sector regulatory expertise.
- Alternative Fee Arrangements: Pilot fixed-fee, outcome-based, or subscription pricing models with trusted clients. Protects margin when hourly billing erodes; requires accurate cost modeling and quality assurance to be profitable.
- Associate Career Path Redesign: Restructure associate development from document production toward judgment-intensive work — client counseling, negotiation, regulatory strategy, complex analysis. Investment in AI-augmented training programs.
- Malpractice Risk Management Infrastructure: Build formal AI quality assurance protocols — mandatory review layers, error tracking, disclosure procedures, insurer reporting. Defensive investment; reduces liability exposure and builds client trust.
- Legal AI Platform Partnerships: Partner with or license from legal AI platforms (Harvey.ai, LexisNexis+, Westlaw+) rather than build proprietary tools. Faster deployment; lower development cost; dependency risk on third-party platforms.

AI Adoption Arc — Foundation Phase

Foundation (2025 – Q1 2026): AI deployment in the firm is limited to controlled pilots with strong guardrails. A contract review copilot is being piloted in the corporate/M&A practice on mid-complexity deals, with mandatory partner review of all AI-generated output. Associates in litigation are using legal research AI tools for preliminary case research, subject to senior associate validation before inclusion in any work product. The IP practice is testing AI-assisted patent landscape analysis for due diligence engagements.

Results are promising but uneven. Productivity gains are real — associates report completing routine research and first-draft documents materially faster. But quality review overhead partially offsets the time savings: partners and senior associates are spending significant additional time checking AI output for accuracy, hallucinated citations, and jurisdictional errors. Malpractice review has already caught AI-generated errors that, if undetected, could have created liability exposure. Bar rule compliance is manageable but labor-intensive — the firm tracks AI disclosure requirements across 25+ state jurisdictions, and the rules are evolving monthly. Malpractice insurers have begun requesting information on AI deployment scope and quality assurance protocols. Client interest in AI governance and regulatory compliance advisory services is strong and growing — early indications suggest this could become a significant new revenue stream. Margin impact so far: minimal, as pilots are small scale and overhead costs are high relative to productivity gains.

Strategic Considerations

1. Attorney review of all AI-generated legal work is non-negotiable in the current environment. Productivity gains are real, but unreviewed AI output creates unacceptable malpractice exposure. The cost of human review is the price of operating responsibly — the question is how to make review efficient without compromising thoroughness.
2. AI governance and regulatory compliance is a rare high-margin growth opportunity. Every industry is grappling with AI liability, algorithmic fairness, and regulatory compliance. Firms with regulatory expertise and client relationships are well-positioned — but the window for establishing market position is narrowing as competitors move.
3. Bar rule compliance is a jurisdiction-by-jurisdiction challenge. State-by-state AI guidance varies and evolves rapidly. Meeting the strictest jurisdiction's requirements firm-wide reduces complexity but increases cost. Consider whether minimum compliance per jurisdiction or maximum compliance firm-wide better serves the firm's risk profile.
4. Associate skill development must shift from production to judgment. Associates whose value is limited to document drafting and research face displacement. Those who exercise legal judgment, manage client relationships, and counsel on risk become more valuable. The training and evaluation framework must evolve to reflect this shift.
5. Alternative fee arrangements are an experiment worth running now. The hourly billing model is under growing pressure for AI-assisted work. Piloting fixed-fee, outcome-based, and subscription models with trusted clients builds the institutional knowledge needed when the transition becomes unavoidable.
6. Malpractice risk management for AI-assisted work is both defensive and offensive. Formal error tracking, quality audits, and disclosure protocols reduce liability. Proactively engaging malpractice insurers demonstrates governance maturity and can build client confidence — it is a differentiator, not just a cost.

7. Client confidentiality is an absolute constraint on AI deployment. Training AI models on client data without explicit consent, or using third-party platforms that may expose privileged information, creates existential reputational risk. The duty of confidentiality admits no exceptions — this constrains which tools and platforms can be deployed.
-

Manufacturing Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the CEO of a heavy manufacturing conglomerate operating 28 plants (18 US, 10 international) with approximately 40,000 employees. Your production spans automotive components, industrial machinery, and electronics assemblies. The business is capital-intensive (\$3-5B annual capex typical), unionized, and physically complex: equipment lifespans run 7-15 years, facility modernization cycles 20-30 years. You are among the largest industrial AI adopters in the US, and your decisions on predictive maintenance, production automation, quality inspection, and workforce transition set expectations for the broader manufacturing sector.

Strategic Context

Your industry sits at the intersection of two powerful forces: AI-enabled operational transformation and deeply entrenched physical-world constraints. Unlike software or finance, where AI deployment can scale with marginal cost near zero, every manufacturing AI initiative requires physical retrofitting, sensor deployment, OT/IT integration, and workforce adaptation. The gap between "what AI can do in a lab" and "what AI can do on a factory floor with 15-year-old equipment" defines your strategic reality.

AI adoption in US manufacturing is uneven. Predictive maintenance and quality inspection are the most mature use cases, with proven pilots across multiple industries. Production optimization and scheduling AI remain earlier-stage, constrained by the complexity of integrating legacy operational technology (PLCs, SCADA systems) with modern IT infrastructure. Warehouse automation within manufacturing facilities is accelerating but triggers significant labor displacement concerns in a heavily unionized sector.

Your decisions ripple across the five-industry ecosystem. Manufacturing efficiency gains directly affect Consumer sector supply costs and product availability. Healthcare depends on your precision manufacturing for critical medical devices. Your capital allocation signals influence Finance sector industrial lending and investment thesis. Supply chain digitalization efforts you lead determine whether Logistics operators can achieve end-to-end visibility. And Software & Tech companies are both your AI vendors and your competitors for engineering talent.

The core tension you face: AI delivers clear ROI in controlled settings, but scaling requires massive capital commitments into physical infrastructure with long payback periods, workforce transitions governed by union contracts, and integration of legacy systems that were never designed for digital connectivity. Moving too fast burns capital and triggers labor conflict. Moving too slow cedes competitive

advantage to rivals and disappoints investors expecting efficiency gains.

Objectives

Objective	Target (Banded/Directional)	Driver
Equipment Uptime & Throughput	Material improvement in manufacturing uptime and plant throughput	Predictive maintenance and production optimization AI reduce unplanned downtime and increase output per shift
Quality & Defect Reduction	Material improvement in manufacturing defect rates	AI-driven quality inspection and process control catch defects earlier and reduce scrap/rework costs
Asset Utilization & Efficiency	Improved asset utilization; reduced energy consumption	AI-optimized production scheduling maximizes equipment usage and minimizes energy waste across plants
Capital Efficiency	Material IRR on manufacturing AI investments	Retrofit costs are high (\$50-100M per plant); must demonstrate clear ROI to justify capex and sustain investor confidence
Labor Transition Management	Execute automation within union contract constraints	Retrain and redeploy affected workers; preserve labor relations, safety culture, and organizational stability

Constraints

Constraint	Impact	Implications
Capital Intensity & Long Asset Lifecycles	Retrofitting 28 plants costs \$50-100M each. Equipment replacement cycles are 7-15 years. Annual budget \$950M total; \$100-160M allocable to AI.	Full deployment takes 4-5 years at current pace. Leasing/financing can accelerate but increases total lifetime cost. Must prioritize highest-value plants first.

Constraint	Impact	Implications
Unionized Workforce & "No-Layoff" Agreements	50-60% of workforce unionized. Any automation-driven headcount reduction triggers renegotiation. Retraining costs increase automation ROI by 20-25%.	"No-layoff" agreements cost \$35-40M but preserve labor relations and safety culture. Union contracts renew every 3-5 years with uncertainty. Strikes possible during transition.
OT/IT Convergence Challenge	Manufacturing AI requires integration of legacy operational technology (PLCs, SCADA) with IT infrastructure (cloud, data, software). Legacy systems are often disconnected.	Integration is technically hard and expensive. Edge AI deployment hamstrung by lack of IT infrastructure at plant level. Phased approach necessary; not all plants can modernize simultaneously.
Equipment Intelligence & Sensor Deployment	Predictive maintenance requires sensors, edge computing, real-time data streaming across aging equipment. Many legacy machines lack intelligent controls.	Equipment retrofitting is labor-intensive. Sensor infrastructure investment is material per plant. Must balance breadth of deployment against depth of capability at each site.

Resources & Levers

Physical & Operational Assets:

- 28 manufacturing plants (18 US, 10 international) with millions of sensors and production equipment
- 15 years of production logs, maintenance records, and supply chain data
- Equipment and facility infrastructure with known performance baselines

Talent & Expertise:

- ~40,000 employees with deep manufacturing operations and maintenance expertise
- Data science capability: 100+ data scientists with production optimization domain knowledge
- Established vendor relationships with equipment OEMs (Siemens, ABB, Rockwell)
- Union contract history and labor relations institutional knowledge

Capital & Financial Resources:

- Annual capex \$950M total; \$100-160M allocable to AI-enabled upgrades
- Access to equipment financing and leasing arrangements
- Existing equipment lease portfolio and established vendor financing terms

Potential Paths Forward:

- Predictive Maintenance & Condition Monitoring: Deploy AI to predict equipment failures and optimize maintenance scheduling. High ROI (~\$45-55M annual savings on 8 highest-value plants); moderate execution complexity. The clearest near-term value driver.
- Production Optimization & Scheduling: AI-optimized production plans, real-time resource allocation, throughput optimization. Requires OT/IT integration; moderate ROI; high execution risk due to legacy system dependencies.
- Quality Inspection & Defect Detection: AI-driven visual inspection and automated quality control. Proven in pilots; scalable to all plants. Lower capex than maintenance or scheduling initiatives.
- Warehouse Automation & Internal Logistics: Warehouse automation (pick-and-place robots, sorting systems) within manufacturing facilities. High labor displacement risk; requires careful labor transition planning and union negotiation.
- Energy & Resource Optimization: AI-driven energy management, waste reduction, resource scheduling. Moderate ROI; complements other AI initiatives and supports ESG commitments.

AI Adoption Arc — Foundation Phase

Foundation (2025 - Q1 2026): Predictive maintenance pilots in 4 manufacturing plants validated a 19% reduction in unplanned downtime and meaningful extension of equipment life over a 12-month trial. Quality inspection AI was piloted in 2 high-variability plants, demonstrating improved defect detection rates. Limited supplier digitalization partnerships were established with 5 strategic vendors to test real-time data sharing. OT/IT integration assessments were completed at pilot sites, revealing the scale of the convergence challenge across remaining plants. The organization has proven that manufacturing AI works in controlled settings; the question now is whether it can scale across 28 plants within capital and labor constraints.

Strategic Considerations

1. Plant prioritization determines payback speed. Concentrating predictive maintenance and optimization on the 8-12 highest-value plants (which account for ~60% of annual maintenance cost) offers the shortest payback and most defensible ROI. The question is when to scale beyond proven sites — scaling too early risks capital on unvalidated integration approaches; scaling too slowly cedes efficiency gains to competitors.
2. Labor transition investment and automation economics are inseparable. Retraining and redeployment programs (\$35-40M) cost more upfront but preserve labor relations, safety culture, and union cooperation. Consider whether the operational and political risks of underinvesting in workforce transition — strike exposure, talent attrition, organizational resistance — outweigh the near-term capex savings of a technology-only approach.
3. OT/IT integration is a phased modernization challenge, not a single conversion. Requiring all 28 plants to achieve "smart" status simultaneously concentrates risk and overcommits capital. Phased rollout (4-6 plants, then scale) reduces integration risk and allows learning — but the pace of phasing determines how quickly efficiency gains compound across the network.

4. Equipment vendor partnerships offer shared risk and accelerated capability. Working with equipment OEMs (Siemens, ABB, Rockwell) to co-develop sensor and analytics capabilities reduces internal development burden and leverages vendor expertise. The trade-off is dependency — vendor-integrated intelligence may limit future flexibility and negotiating leverage.
 5. Hybrid supplier ecosystems are the reality for 5+ years. The supplier base will remain mixed (digital leaders and analog laggards) for the foreseeable future. AI systems must leverage digital supplier data where available while operating gracefully when suppliers lack real-time connectivity — designing for the ideal supplier ecosystem rather than the actual one creates brittle systems.
-

Logistics Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the CEO of a national freight and logistics operator running 150+ logistics facilities, a fleet of 5,000+ commercial vehicles, and approximately 60,000 employees (drivers, warehouse workers, logistics coordinators). Your operations span long-haul freight, regional distribution, third-party logistics (3PL), warehouse management, and last-mile delivery. The workforce is heavily unionized (60%+), and your business model depends on route density, fuel economics, and service reliability. You are one of the largest logistics AI adopters in the US, and your decisions on route optimization, autonomous systems, fleet management, and workforce transition influence the entire physical supply chain ecosystem.

Strategic Context

Logistics is the connective tissue of the US economy, and AI is reshaping every link in the chain. Route optimization, predictive vehicle maintenance, demand forecasting, and warehouse automation are all either deployed or in advanced pilots across the industry. But unlike digital businesses, logistics AI must operate in the physical world: on highways subject to weather and regulation, in warehouses with variable labor availability, and across supply chains where your partners range from digitally sophisticated to entirely analog.

The competitive landscape is bifurcating. Large operators with scale and data advantages are pulling ahead on AI-driven efficiency, while smaller regional carriers struggle to invest. Autonomous vehicle technology remains the sector's most disruptive uncertainty: the technology is advancing, but federal DOT/NHTSA regulations on commercial autonomous long-haul trucks are not expected until 2028-2030. Specialists like Waymo and Aurora are years ahead of any internal capability you could build. The question is not whether autonomous vehicles will arrive, but when, under what regulatory framework, and whether you should build, partner, or wait.

Your decisions have significant cross-industry impact. Consumer sector retailers depend on your delivery reliability and cost structure. Healthcare providers rely on your cold-chain logistics for pharmaceutical and device distribution. Manufacturing operations need your freight capacity and scheduling flexibility. Finance sector investors are watching whether logistics AI investments generate returns or become stranded capital. And Software & Tech companies are both your AI vendors and potential disruptors (Amazon Logistics, autonomous vehicle startups).

The core tension you face: AI optimization delivers clear value in high-density, high-volume corridors,

but the economics of last-mile delivery in low-density areas remain structurally challenged regardless of technology. Meanwhile, the workforce transition from human-driven to AI-assisted (and eventually autonomous) operations is the most politically sensitive labor issue in the sector. Moving aggressively on automation risks union conflict and public backlash. Moving cautiously risks cost disadvantage against competitors who are investing faster.

Objectives

Objective	Target (Banded/Directional)	Driver
Fleet Cost & Efficiency	Material reduction in fuel consumption and cost per mile	AI-driven route optimization, predictive vehicle maintenance, and driver behavior optimization reduce operating costs
On-Time Delivery & Service	Maintain or improve on-time delivery performance	Better demand forecasting, real-time routing, and inventory positioning improve service reliability for shippers
Safety Performance	Material improvement in fleet accident rates	AI-assisted safety monitoring, collision avoidance, and fatigue detection reduce incidents while preserving safety culture
Last-Mile Economics	Optimize last-mile delivery unit economics in viable markets	Focus AI optimization on high-density urban routes where technology improves profitability; accept structural limits in rural areas
Autonomous Vehicle Readiness	Establish strategic positioning for autonomous transition	Evaluate and pilot AV partnerships; manage regulatory and labor timeline expectations without overcommitting capital

Constraints

Constraint	Impact	Implications
Regulatory Uncertainty on Autonomous Vehicles	DOT/NHTSA rules on commercial autonomous long-haul trucks not expected until 2028-2030. State-level regulations vary. Federal uncertainty blocks major capex decisions.	Internal AV development is slow and expensive vs. partnerships with specialists. Cannot commit large capital to autonomous systems until regulatory framework is clearer. Partner, don't build.
Driver Acceptance & Change Management	Drivers, especially older cohorts (>50), resist routing AI due to GPS tracking concerns, loss of autonomy, and perceived surveillance. Younger drivers adopt readily (85%); older drivers resist (55% adoption).	Technology alone does not guarantee adoption. Change management programs, transparency, override capabilities, and incentive structures are required. Frame AI as "efficiency assistant" not surveillance.
Union Contract Constraints	60%+ of workforce unionized. Autonomous or remote-driving systems require union negotiation. Strikes during peak season are possible if changes are perceived as labor displacement.	Any automation initiative must account for union relations cost and timeline. "No-layoff" commitments and retraining programs are practical requirements, not optional. Contract renewal every 3-5 years introduces recurring uncertainty.
Last-Mile Profitability Limits	AI route optimization cannot fix structurally unprofitable last-mile delivery in rural areas. Low stop density and long distances defeat algorithmic efficiency.	No amount of optimization makes 5-stop, 50-mile routes profitable. Must distinguish between routes where AI improves economics and routes where the business model is the problem. Consider outsourcing low-density last-mile to regional carriers.

Resources & Levers

Physical & Fleet Assets:

- 5,000+ commercial vehicles with GPS, telematics, and vehicle telemetry systems
- 150+ logistics facilities (distribution centers, cross-docks, last-mile hubs)
- Real-time fleet data: GPS positioning, fuel consumption, driver behavior, vehicle condition
- 15 years of route data, delivery logs, and fleet maintenance records

Talent & Expertise:

- ~60,000 employees including experienced drivers, warehouse operators, and logistics coordinators

- Data science capability: 40+ routing optimization experts, demand forecasting team
- Established shipper/receiver relationships across Consumer, Healthcare, and Manufacturing sectors
- Deep institutional knowledge of union contract dynamics and labor relations

Capital & Financial Resources:

- Annual capex \$950M total (shared with Manufacturing); fleet telematics and AI investment competes for \$100-160M AI allocation
- Fleet leasing and financing arrangements; vehicle replacement cycles of 5-8 years
- Established vendor relationships with telematics and fleet management providers

Potential Paths Forward:

- Route Optimization & Fuel Efficiency: Deploy AI to optimize vehicle routes, reduce idle time, improve driver behavior. High ROI (\$180-200M potential savings at full fleet scale); proven in pilots; driver adoption is the key risk.
- Predictive Vehicle Maintenance: Condition monitoring to predict maintenance needs, extend equipment life, reduce downtime. Moderate ROI; complements route optimization and improves fleet availability.
- Autonomous Vehicles & Long-Haul Pilot: Partner with AV specialists (Waymo, Aurora) for limited autonomous long-haul pilots. Strategic positioning value; but timeline uncertain, regulatory risk high, and labor relations tension significant.
- Driver Assistance & Safety Systems: AI-driven collision avoidance, fatigue monitoring, coaching systems. Safety-focused (not pure automation). Improves safety culture; driver acceptance better than full autonomy initiatives.
- Warehouse Automation & Last-Mile Optimization: Warehouse automation (sorting, packing) and route optimization for last-mile. High-density urban routes benefit from optimization; rural routes do not. Must segment investment accordingly.

AI Adoption Arc — Foundation Phase

Foundation (2025 - Q1 2026): Route optimization pilots in 3 logistics zones covering 800 trucks achieved 14% fuel savings and 12% on-time delivery improvement over a 12-month trial. Predictive vehicle maintenance pilots demonstrated meaningful reductions in unplanned vehicle downtime. Driver assistance and safety systems (collision avoidance, fatigue monitoring) were tested across a subset of the fleet with positive safety outcomes. Limited autonomous vehicle monitoring was initiated through observation partnerships with AV specialists, but no operational pilots were launched due to regulatory uncertainty. The organization has validated that logistics AI works at regional scale; the challenge now is full-fleet deployment across 5,000+ vehicles while managing driver acceptance, union dynamics, and the autonomous vehicle question.

Strategic Considerations

1. Route optimization and driver acceptance are interdependent, not sequential. Algorithm quality is not the bottleneck — adoption rate is. Younger drivers adopt readily (~85%); older drivers require deliberate change management and incentive programs. Consider whether framing AI as an "efficiency assistant" rather than a monitoring tool affects adoption curves enough to justify the investment in change management alongside technology deployment.
 2. Autonomous vehicle strategy is a partnership question, not a build question. Regulatory uncertainty and specialist competitive advantages make internal AV development expensive and slow. Partnering with established AV operators (Waymo, Aurora) for limited pilot deployments preserves capital for proven fleet optimization — but the trade-off is limited control over a technology that could reshape the industry's cost structure within a decade.
 3. Last-mile optimization has structural limits that technology cannot overcome. AI cannot make fundamentally unprofitable rural routes profitable. The question is where to concentrate optimization investment — high-density urban and regional corridors where technology genuinely improves unit economics, or broadly across the network where returns diminish. Consider whether low-density last-mile is better outsourced than optimized.
 4. Safety culture is both a constraint and a competitive advantage. Safety systems that support driver judgment (collision avoidance, fatigue monitoring, coaching) build trust and aid recruitment. Systems perceived as surveillance erode both. The framing and implementation of driver-facing AI determines whether it strengthens or weakens the safety culture that differentiates your operation.
 5. Fleet investment should match route-level economics, not uniform capability targets. Not all 5,000+ vehicles need the same AI capability. High-volume interstate corridors justify maximum telematics and optimization investment; regional distribution needs predictive maintenance and basic routing; last-mile is primarily a labor management challenge. Consider how granular your investment segmentation should be to maximize return without creating operational complexity.
-

Big Tech Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See [05_Participant_Private_Cards/](#) and [04_Participant_AI_Adoption_Arcs/](#) folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document ([03zRulesofPlay.md](#)) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the CEO of a Big Tech platform company with a dominant global market position (Google, Meta, Microsoft, Amazon, Apple-class). You operate across cloud infrastructure, advertising, enterprise software, devices, search, social media, and e-commerce with billions of active users and devices worldwide. Annual revenue is approximately \$200B-\$400B with operating margins in the 20-30% range. You control critical infrastructure — cloud platforms, data centers, global distribution networks — and your capital allocation decisions cascade across every other sector in this exercise. You represent the non-AI-development operations of a Big Tech platform company: cloud infrastructure, advertising, enterprise software, devices, search, social media, and e-commerce. **The AI development/AI lab portion of the business (frontier model training, research lab operations, model licensing) is NOT under your control** — those actions enter the exercise only through facilitator-controlled injects and scenario updates.

Strategic Context

You sit at the center of the AI economy — not as the builder of frontier models, but as the company that must figure out how to embed AI into every product, service, and customer relationship you operate. Your cloud infrastructure hosts the AI workloads of every other sector in this exercise. Your advertising business is being reshaped by AI-driven search disruption. Your enterprise software products must integrate AI features at speed or lose ground to nimbler competitors. Your devices and platforms are the distribution channel through which billions of people interact with AI for the first time.

Your competitive position is defined by scale and reach, but scale alone is not a strategy. Amazon, Google, Microsoft, and Meta are all racing to embed AI into their existing product lines — copilots in productivity suites, AI-powered search, generative ad creative, AI shopping assistants, smart device integration. The question is not whether AI gets embedded but how fast, at what cost, and with what customer experience. The companies that execute product integration well will deepen ecosystem lock-in. Those that stumble will face customer attrition to AI-native alternatives and open-source-powered competitors that can move faster without legacy constraints.

Cloud infrastructure is your highest-leverage business. Enterprise customers across every sector are migrating AI workloads to your platform, but they are also demanding commodity pricing as open-source models reduce the premium on proprietary cloud AI services. Your cloud margin is under pressure from both competition (AWS vs. Azure vs. GCP price wars) and customer consolidation (enterprises reducing multi-cloud sprawl). Winning in cloud means positioning as the infrastructure-of-choice for enterprise AI

— not just raw compute, but managed AI services, fine-tuning platforms, and verticalized solutions.

The regulatory environment is tightening. Antitrust scrutiny is focused on your market concentration, data practices, and the degree to which your platform advantages create barriers to competition. Data privacy regulations (GDPR, CCPA, expanding state laws) constrain how you collect, retain, and monetize user data. AI amplifies both concerns: AI features require more data, and AI-powered products deepen the ecosystem lock-in that regulators are targeting. You must navigate this environment proactively — reactive compliance is more expensive and more disruptive than anticipatory strategy.

Talent competition is fierce across all sectors. You are bidding for the same AI engineers, product managers, and data scientists that finance, healthcare, supply chains, and SaaS companies want. Your scale is an advantage (interesting problems, competitive compensation, career breadth), but startups offer equity upside and agility that you cannot match. Retention of mid-career engineers — the backbone of product execution — is the most acute pressure point.

Objectives

Objective	Target (Banded/Directional)	Driver
Cloud Revenue Growth	Material growth in cloud infrastructure revenue; secure AI workload migration from enterprise customers across sectors	Managed AI services, fine-tuning platforms, verticalized cloud solutions, competitive pricing
AI Product Launch Velocity	Rapid integration of AI features into core products (search, email, office, devices, cloud, advertising) with measurable adoption	Product engineering execution, ecosystem integration, customer communication
Advertising Revenue Protection	Maintain or grow ad revenue despite search fragmentation, consumer privacy regulations, and AI-driven disruption to ad-supported models	AI-improved targeting efficiency, generative ad creative, new ad formats, measurement innovation
Ecosystem Lock-in & Platform Stickiness	Increase switching costs and ecosystem stickiness through AI-native features integrated across the product portfolio	Integration depth, cross-product data flows, workflow automation, developer ecosystem
Talent Retention & Competitive Recruitment	Recruit and retain top AI engineering and product talent in a competitive market; maintain organizational execution capability	Competitive compensation, career development, interesting problems at scale, retention of institutional knowledge

Constraints

Constraint	Impact	Implications
Massive AI Infrastructure Capex	AI-related capex (data centers, custom silicon, inference infrastructure) is growing 30%+ annually. By 2026, AI capex could represent 40%+ of total capex. ROI timelines are uncertain; GPU/TPU utilization rates are lower than projected between workloads	Heavy capex burden constrains profitability, dividend capacity, M&A dry powder, and ability to invest in other areas. Overinvestment risk if enterprise AI adoption is slower than expected. Underinvestment risk if competitors capex harder. Capital allocation discipline is critical
Model Commoditization (External Market Force)	Open-source AI models (LLaMA, Mistral, Phi) are improving rapidly and reaching feature parity with proprietary models for many enterprise use cases. This is an external market force — you do not control frontier model development decisions	Proprietary cloud AI service premiums are eroding. Customers are testing open-source models and finding acceptable performance at lower cost. Your differentiation must come from managed services, integration, reliability, and ecosystem value — not model exclusivity
Antitrust Regulatory Pressure	Antitrust scrutiny on market concentration, data practices, platform bundling, and AI service tying. Potential enforcement actions including forced data separation, open API access mandates, and platform conduct restrictions	Compliance cost and execution risk are high. Business model and product design must be defensible. Proactive transparency and access measures may reduce enforcement risk but also reduce competitive advantage
Data Privacy & Regulation	GDPR, state privacy laws, and sector-specific regulations limit data collection, retention, and cross-product data sharing. Expanding regulatory scope increases compliance cost and constrains data-driven product improvement	Reduces training data availability for product optimization. Limits data monetization across products. Privacy-preserving techniques (federated learning, differential privacy) are required but add complexity and cost
Talent Saturation	Limited supply of AI engineering, product, and data science talent globally; bidding war with other Big Tech, SaaS, finance, healthcare, and startups. Retention risk if employees perceive better opportunities elsewhere	Compensation pressure (stock + cash rising faster than historical rates); attrition of mid-career engineers is the most acute risk; organizational churn can derail product roadmaps; cannot solve with hiring alone — retention and internal mobility matter more

Resources & Levers

Platform Assets:

- Vast first-party data: Search queries, social graph, cloud usage, consumer behavior, enterprise usage, device telemetry. Petabytes of behavioral and operational data
- Compute infrastructure: Global data center footprint, custom silicon (TPUs, GPUs), distributed inference capability. Competitive advantage in scale and reliability
- Capital: Cash reserves (\$50B-\$100B+). Ability to fund massive capex and M&A. Access to capital markets at favorable terms
- Ecosystem reach: Billions of active users/devices. Cloud platform with hundreds of thousands of enterprise customers. Distribution advantages for new AI-powered products
- Brand trust: Market leadership, brand recognition, enterprise relationships, consumer familiarity

Potential Paths Forward:

- AI Product Launches: Embed AI into core products (search, email, office, cloud, devices, advertising). Leverage ecosystem reach and distribution. High impact but execution complexity and customer communication challenges are significant. This is the primary lever for capturing AI-era value
- Cloud Infrastructure Positioning: Position cloud platform as infrastructure-of-choice for enterprise AI workloads. Compete on managed services, fine-tuning platforms, and verticalized solutions — not just raw compute. Cloud margin under pressure; differentiation is critical
- Open-Source Strategy: Contribute to or sponsor open-source models and frameworks to maintain developer goodwill, talent recruitment pipeline, and ecosystem influence. Must calibrate carefully to avoid commoditizing your own premium services
- M&A for AI Talent & Capabilities: Acquire AI startups, vertical AI companies, or talent teams to fill capability gaps faster than organic development. Expensive but execution risk lower than organic build in a talent-constrained market
- Data Strategy & Privacy: Defend data assets and optimize data flows across products within regulatory constraints. Invest in privacy-preserving techniques to maintain data advantage while demonstrating compliance. Proactive data governance reduces regulatory risk

AI Adoption Arc — Foundation Phase

Foundation (2025 - Q1 2026): Your AI product integration effort is in early execution. AI-powered features have launched in core products — search, email, productivity suites, cloud services — with mixed adoption. Enterprise cloud customers are piloting AI workloads on your platform, but conversion from pilot to production is slower than projected. GPU/TPU utilization rates at your data centers are below target as inference demand has not yet scaled to match infrastructure buildout. Advertising AI (generative creative, AI-powered targeting) is in limited deployment with promising early signals but no material revenue impact yet. Internally, engineering teams are stretched across too many AI initiatives; prioritization and focus are emerging as execution risks. Capex is committed and accelerating. Margin impact so far: moderate compression from infrastructure costs, partially offset by cloud revenue growth.

Strategic Considerations

1. Strategic AI embedding requires focus, not comprehensiveness. The tension is between high-impact, high-stickiness features (copilots in core productivity tools, AI-powered search, smart device integration) and broad AI deployment across the product portfolio. Consider whether concentrated investment in a few transformative features creates more differentiation than distributing AI capability thinly — focus creates speed, but breadth defends against competitors targeting specific product gaps.
 2. Cloud competitive advantage shifts from model exclusivity to managed services. Model commoditization is an external reality that redefines where cloud differentiation lies. Managed AI services, fine-tuning platforms, verticalized solutions, and enterprise trust may matter more than proprietary model superiority — but the trade-off is ceding the narrative on AI leadership to foundation model companies while competing on integration and reliability.
 3. Regulatory posture is a strategic choice with asymmetric consequences. Antitrust and privacy enforcement are intensifying. Proactive transparency and access measures are cheaper than reactive compliance and litigation — but the degree of proactivity involves trade-offs between competitive flexibility and regulatory defensibility. Consider what posture is sustainable before enforcement arrives rather than adapting after it does.
 4. Talent retention competes with infrastructure investment for strategic priority. Engineering talent — especially mid-career product engineers who ship features — is the binding constraint on AI product velocity. Competitive compensation alone is insufficient; interesting problems, career development, and internal mobility matter as much. The question is whether talent investment receives the same strategic weight as capex in resource allocation decisions.
 5. Capex discipline determines whether infrastructure buildout creates or destroys value. Infrastructure investment is necessary but must be tied to demonstrable demand signals and utilization targets. The risk is that momentum-driven buildout outpaces customer adoption — consider what utilization thresholds and payback periods should trigger acceleration or deceleration of infrastructure spending.
-

B2B/B2C SaaS Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See *05_Participant_Private_Cards/* and *04_Participant_AI_Adoption_Arcs/* folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document (*03zRulesofPlay.md*) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

You are the CEO of a major enterprise software company with a subscription-based business model (Salesforce, SAP, Workday, Adobe, Zoom-class). You serve millions of enterprise and consumer customers with annual subscriptions, usage-based pricing, and professional services. Annual revenue is approximately \$10B-\$50B with operating margins in the 20-30% range. You compete with larger integrated platforms (Big Tech companies moving downmarket into your categories) and specialized AI-native startups (moving upmarket with lightweight, focused tools). Your customer relationships, domain expertise, and integration depth are your core competitive assets. Your decisions on AI feature strategy, pricing, and talent investment shape how every other sector in this exercise experiences enterprise software.

Strategic Context

Your core business is subscription revenue from enterprise and consumer customers, and AI adoption is simultaneously the biggest opportunity and the most acute threat you face. AI-powered features — copilots in ERPs, AI-driven sales tools, customer service automation, predictive analytics — drive product differentiation and can justify price increases. Early movers gain competitive advantage as customers consolidate around platforms that deliver genuine AI value. But customers increasingly expect AI features to be bundled into existing subscription prices at no additional cost. Pricing power is eroding as AI becomes table-stakes, and your own AI infrastructure costs are rising faster than revenue.

The competitive threat from AI-native startups is acute and accelerating. Specialized, lightweight AI tools — AI-first sales platforms, marketing automation, code generation, customer support bots — are displacing parts of your larger platform. These competitors are agile, venture-backed, and attractive to price-sensitive customers who want best-of-breed AI capabilities without paying for a full enterprise suite. Your scale and integration depth are advantages, but your agility is constrained by legacy architecture, existing customer commitments, and organizational complexity.

From above, Big Tech is moving into your territory. Microsoft, Google, and Amazon are building AI capabilities directly into their productivity suites, cloud platforms, and enterprise tools — features that overlap with or directly compete with your core products. Copilot in Office competes with your CRM automation. Google Workspace AI competes with your collaboration tools. AWS AI services compete with your analytics offerings. Big Tech has distribution advantages (billions of users, embedded in enterprise IT stacks) that you cannot match. Your defense is depth: deeper domain expertise, tighter workflow integration, and stronger customer relationships in your verticals.

Talent costs are rising sharply. AI engineering talent is scarce and mobile — the same engineers you need are being courted by Big Tech, finance, healthcare, and well-funded startups. Your existing support and services workforce (customer success, professional services, implementation teams) is under automation pressure from AI tools. You face a dual talent challenge: acquiring expensive new AI talent while managing the transition of your existing workforce. This is not just a cost problem — it is a roadmap execution problem. If you cannot hire and retain the engineers to build AI features, your competitive position deteriorates regardless of how much capital you allocate.

Cross-industry dynamics shape your landscape directly. Healthcare, finance, supply chain, and consumer companies are all your customers — their AI adoption pace determines your revenue growth. When healthcare invests in AI-powered clinical tools, they buy your enterprise platform to support it. When finance automates operations, they use your workflow software. But when these sectors slow AI investment (regulatory uncertainty, budget constraints), your pipeline slows with them. You are a derived demand business: your fortunes track the AI adoption velocity of every other sector in this exercise.

Objectives

Objective	Target (Banded/Directional)	Driver
AI-Powered Feature Adoption & Upsell	Rapid integration of AI features (copilots, predictive analytics, automation) into core product; drive measurable adoption among installed base and justify premium pricing	Product engineering execution, customer communication, feature quality, time-to-market
Revenue Growth from AI Capabilities	Material growth from AI-specific product features, new AI service tiers, and AI-driven professional services	New product tiers, upsell to existing customers, new customer acquisition for AI-first capabilities
Gross Margin Preservation	Maintain or improve gross margins despite rising AI infrastructure costs (cloud compute, third-party model APIs, engineering complexity)	Efficient make-vs-buy decisions, infrastructure cost optimization, pricing discipline, feature bundling strategy
Net Revenue Retention Improvement	Increase upsell and cross-sell to existing customers; reduce churn to AI-native competitors; improve customer lifetime value	Product stickiness, integration depth, customer success investment, competitive feature parity
Competitive Positioning vs. AI-Native Entrants	Defend market position against focused, lightweight AI competitors while leveraging integration depth and enterprise trust	Superior product experience, breadth of platform, vertical domain expertise, data security, enterprise-grade reliability

Constraints

Constraint	Impact	Implications
Customer Pricing Expectations	Enterprise customers expect AI features bundled into existing subscriptions at no additional cost. Willingness to pay a premium for AI features is minimal (~5% of customers willing to pay 10%+ more). Sales teams report deal slippage and negotiation friction	Pricing power erodes; upsell friction increases; free-tier expansion risk (customer acquisition via free AI, difficult to monetize). Bundling discipline required — must decide clearly: bundle or premium add-on. Ambiguous pricing creates customer confusion and sales friction
AI Infrastructure Cost Growth	Cloud compute costs for AI features (running inference at scale, third-party API calls, vector databases, fine-tuning) are rising faster than revenue. Gross margin on new AI features is 10-15 points lower than legacy features due to AI infrastructure overhead	Margin pressure even as adoption scales. Must optimize make-vs-buy (proprietary development vs. third-party APIs) aggressively. Cost efficiency in AI feature delivery is a competitive differentiator, not just an operational concern
Customer Concentration	Top 10-20 customers represent 30-40% of revenue. These customers have the most leverage on pricing and the highest expectations for AI features. Churn of any single large customer triggers material revenue impact	Defensive investment required to retain top accounts. Customer concentration creates vulnerability to competitive displacement and negotiation pressure. Must diversify customer base while protecting existing relationships
Competitive Threat from AI-Native Startups	Specialized, lightweight AI tools are displacing parts of your platform. AI-native competitors are agile, venture-backed, and offer focused solutions at lower cost. Customers are evaluating point solutions that outperform your platform in specific AI capabilities	Must accelerate AI feature roadmap to maintain competitive parity. Cannot win on breadth alone — must demonstrate AI excellence in core use cases. M&A of AI-native competitors may be required to close capability gaps
Talent Cost Inflation	AI engineering talent is scarce; compensation is rising 20-30% annually. Existing support and services talent is under automation pressure. Retention risk across the organization — attrition of mid-career engineers to startups and Big Tech	Compensation budget pressure; roadmap execution constrained by talent availability, not capital. Must invest in retention (career paths, interesting problems) not just recruitment. Workforce transition for non-AI roles is a management challenge

Resources & Levers

Enterprise Assets:

- Customer relationships: Millions of enterprise and consumer customers with deep product engagement and workflow integration across critical business processes
- Proprietary data: Customer usage data, feature adoption signals, industry-specific insights, workflow patterns. Valuable for product optimization and AI feature development
- Domain expertise: Deep knowledge of customer workflows, industry compliance requirements, operational pain points. Decades of product-market fit in core verticals
- Engineering capability: Established product development organizations, quality assurance, security infrastructure, customer support operations
- Capital: Strong cash flows and balance sheets; access to debt and equity markets. Many SaaS companies generate 20%+ free cash flow margins
- Integration advantage: Deep integration into customer workflows creates high switching costs. AI features embedded in existing products are stickier than standalone AI tools

Potential Paths Forward:

- AI Feature Roadmap Acceleration: Rapidly integrate AI into core product (copilots, predictive analytics, automation, intelligent search). High execution risk (quality, performance, customer communication), but table-stakes for competitive positioning. Speed matters — first movers capture customer attention and switching cost advantage
- Pricing Strategy for AI: Bundle AI features into existing subscriptions (volume play, lower pricing power) or offer as premium add-ons (pricing power, customer friction). This is the most consequential strategic decision: it determines revenue trajectory, margin profile, and competitive positioning for years
- Make-vs-Buy AI Capabilities: Build proprietary AI features (control, long-term competitive advantage, high cost and slow) or leverage third-party APIs and open-source models (cost-efficient, fast time-to-market, vendor dependence risk). Hybrid approach likely optimal — proprietary for core differentiators, third-party for commoditized capabilities
- Customer Concentration Risk Management: Increase investment in smaller-customer acquisition, product breadth (new vertical markets, new product lines), and self-serve channels. Reduce dependence on top customers; improve revenue stability and negotiating position
- M&A for AI Capabilities & Talent: Acquire AI-native competitors or specialized AI companies to close capability gaps faster than organic build. Expensive but lowers time-to-market risk. Integration execution is the key risk
- Cost Efficiency & Automation: Invest in operational efficiency (infrastructure optimization, customer self-service, support automation) to offset AI infrastructure cost growth and protect margins

AI Adoption Arc — Foundation Phase

Foundation (2025 - Q1 2026): Your AI feature integration is in early deployment. Copilot features and predictive analytics are live in core product modules, with adoption concentrated among your most

sophisticated enterprise customers (roughly 15-20% of your installed base). Customer feedback is positive on capability but sharply negative on pricing — most customers expect AI features to be included in their existing subscription. Third-party API costs (OpenAI, Anthropic, Cohere) are running above budget; gross margin on AI features is 10-15 points below legacy features. Three AI-native startups have emerged as serious competitive threats in your core vertical, winning proof-of-concept engagements with your mid-market customers. Your AI engineering team is fully deployed but stretched across too many initiatives; prioritization and focus are emerging as execution risks. Two senior AI engineers departed for startups in Q4 2025. Margin impact so far: low but directionally negative from infrastructure costs; revenue impact directionally positive from early upsell but below plan.

Strategic Considerations

1. AI feature integration must solve workflow problems, not demonstrate capability. The tension is between shipping AI features that address real customer pain points and deploying AI broadly to match competitor marketing claims. Consider whether fewer, higher-quality AI features that genuinely improve customer workflows create more retention and willingness to pay than comprehensive AI coverage that risks feature bloat and product confusion.
 2. Pricing strategy for AI features is a positioning decision, not just a revenue decision. Bundling signals that AI is core to the platform and positions it as competitive advantage; premium add-on pricing captures willingness to pay but risks adoption friction. The worst outcome is ambiguity — where customers are unsure what they are paying for. Consider how pricing choices affect both near-term revenue and long-term competitive positioning against AI-native entrants.
 3. Customer concentration is a risk that AI-driven product expansion can address. Revenue concentration in top customers creates churn risk and weakens negotiating leverage. Vertical-specific solutions, self-serve channels, and product breadth enabled by AI can diversify the customer base — but broadening requires investment that competes with deepening existing customer relationships. Consider the balance between expanding reach and maximizing value within the current base.
 4. Build vs. buy for AI capabilities is a continuous, feature-level decision. Proprietary AI development is not required for every feature — open-source models, third-party APIs, and commercial model providers can accelerate time-to-market and reduce cost. The trade-off is between speed and dependency: external models ship faster but create vendor risk, while proprietary investment is slower but builds defensible differentiation. Consider which capabilities require proprietary investment and which benefit from external leverage.
 5. AI engineering talent is the binding constraint on roadmap execution. Retention — career paths, interesting problems, internal mobility — matters as much as recruitment. The broader question is how to manage workforce transition for roles under automation pressure: proactive retraining preserves organizational knowledge and is less disruptive than reactive layoffs, but requires investment during a period of margin pressure.
-

Decision Worksheet

Print one copy per decision. Fill in all fields before the 3-minute submission deadline. Each participant submits one worksheet per decision per round. Write clearly in ink.

PROJECT THRESHOLD V7.4 — DECISION WORKSHEET

Round: Practice R1 R2 R3 R4

Name: _____

Industry: _____

Date/Time: _____

1. Industry Identification

Industry: _____

2. Decision Statement

(1–2 sentences: What action is your industry taking, and why?)

3. Spend/Commitment Band

(Check one, then justify.)

Absorbable (BAU)

Material but manageable

Transformational bet

Potentially existential

Justification: _____

4. Time-to-Impact Band

(Check one, then justify.)

<input type="checkbox"/> 0–3 months	<input type="checkbox"/> 3–12 months	<input type="checkbox"/> 1–2 years	<input type="checkbox"/> 2+ years
-------------------------------------	--------------------------------------	------------------------------------	-----------------------------------

Justification: _____

5. Execution Complexity Band

(Check one, then justify.)

<input type="checkbox"/> Low	<input type="checkbox"/> Medium	<input type="checkbox"/> High	<input type="checkbox"/> Very high
------------------------------	---------------------------------	-------------------------------	------------------------------------

Justification: _____

6. Dependency Band

(Check one, then justify.)

<input type="checkbox"/> Mostly internal	<input type="checkbox"/> Requires vendor/partner	<input type="checkbox"/> Requires regulator/union/standards	<input type="checkbox"/> Requires ecosystem shift
--	--	---	---

Justification: _____

7. Scale

(Check one.)

<input type="checkbox"/> Pilot	<input type="checkbox"/> Regional	<input type="checkbox"/> National	<input type="checkbox"/> Global
--------------------------------	-----------------------------------	-----------------------------------	---------------------------------

8. Success Metric

(Directional or range-based — e.g., "10–20% margin improvement." No point estimates.)

9. Key Risk & Mitigation

	Risk	Mitigation
Risk 1	_____	_____
	_____	_____
	_____	_____
	_____	_____
Risk 2	_____	_____
	_____	_____
	_____	_____
	_____	_____

10. Cross-Industry Synergies

(Which other industries benefit from or depend on this decision? Write "Independent decision" if none.)

11. Private Card Influence *(Rounds 1–3 only)*

(Which private card influenced this decision? Write "None" if not applicable.)

12. Continuity Note (Rounds 2–4 only)

(How does this decision build on, extend, or pivot from your prior-round decision?)

Submitted by: _____ Time: _____

PROJECT THRESHOLD V7.4 — DECISION WORKSHEET

Round: Practice R1 R2 R3 R4 Name: _____

Industry: _____ Date/Time: _____

1. Industry Identification

Industry: _____

2. Decision Statement

(1–2 sentences: What action is your industry taking, and why?)

3. Spend/Commitment Band

(Check one, then justify.)

<input type="checkbox"/> Absorbable (BAU)	<input type="checkbox"/> Material but manageable	<input type="checkbox"/> Transformational bet	<input type="checkbox"/> Potentially existential
---	--	---	--

Justification: _____

4. Time-to-Impact Band

(Check one, then justify.)

<input type="checkbox"/> 0–3 months	<input type="checkbox"/> 3–12 months	<input type="checkbox"/> 1–2 years	<input type="checkbox"/> 2+ years
-------------------------------------	--------------------------------------	------------------------------------	-----------------------------------

Justification: _____

5. Execution Complexity Band

(Check one, then justify.)

<input type="checkbox"/> Low	<input type="checkbox"/> Medium	<input type="checkbox"/> High	<input type="checkbox"/> Very high
------------------------------	---------------------------------	-------------------------------	------------------------------------

Justification: _____

6. Dependency Band

(Check one, then justify.)

<input type="checkbox"/> Mostly internal	<input type="checkbox"/> Requires vendor/partner	<input type="checkbox"/> Requires regulator/union/standards	<input type="checkbox"/> Requires ecosystem shift
--	--	---	---

Justification: _____

7. Scale

(Check one.)

<input type="checkbox"/> Pilot	<input type="checkbox"/> Regional	<input type="checkbox"/> National	<input type="checkbox"/> Global
--------------------------------	-----------------------------------	-----------------------------------	---------------------------------

8. Success Metric

(Directional or range-based — e.g., "10–20% margin improvement." No point estimates.)

9. Key Risk & Mitigation

	Risk	Mitigation
Risk 1	_____ _____	_____ _____
	_____ _____	_____ _____
Risk 2	_____ _____	_____ _____
	_____ _____	_____ _____

10. Cross-Industry Synergies

(Which other industries benefit from or depend on this decision? Write "Independent decision" if none.)

11. Private Card Influence (Rounds 1–3 only)

(Which private card influenced this decision? Write "None" if not applicable.)

12. Continuity Note *(Rounds 2–4 only)*

(How does this decision build on, extend, or pivot from your prior-round decision?)

Submitted by: _____ Time: _____

Rules of Play

A. EXERCISE STRUCTURE

Format: Single 8-hour session (8:30 AM – 4:30 PM, with lunch break 12–1 PM)

Participants:

- 5–11 industry representatives (each selects one or more industries from the 11-industry framework)
- 1-2 facilitators + observers

The 11 Industries: Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS

Duration: 4 rounds of deliberation + scorekeeping + practice micro-round at start

B. INDUSTRY SELECTION

Before the exercise begins, each participant selects the industries they will represent. The facilitator coordinates assignments based on participant expertise, professional background, and interest.

Assignment Guidelines

- Ideal: 8 or more industries assigned across participants
- Minimum: Fewer than 8 is possible; the facilitator plays any unassigned industries
- Maximum: No participant should represent more than 3 industries
- Facilitator coverage: The facilitator can represent up to 3 unassigned industries, submitting decisions on their behalf using base-case assumptions

Recommended Assignments by Participant Count

Participants	Industries per Person	How It Works
5	2–3 each	Each participant takes 2 industries, with 1 participant taking 3. Facilitator covers any remaining industries.
7	1–2 each	4 participants take 2 industries; 3 participants take 1 each. Facilitator balances workload and expertise.

Participants	Industries per Person	How It Works
9	1–2 each	2 participants take 2 industries; 7 participants take 1 each. Near full coverage with minimal doubling.
11	1 each	Each participant represents exactly 1 industry. Ideal for maximum depth and discussion diversity.

Assignment Principles

1. Expertise match: Where possible, assign participants to industries aligned with their professional experience.
2. Natural pairings: When a participant takes 2 industries, related industries (e.g., Healthcare Provider + Healthcare Payer, Retail + CPG) are natural but not required. Cross-industry pairs (e.g., Finance + Big Tech) can produce valuable perspective.
3. Workload balance: Each additional industry means an additional Decision Worksheet per round. Participants with 2+ industries should be comfortable with higher decision throughput.

C. ROUND FLOW

Each round follows this sequence. **All timings are strict; participants with incomplete submissions are auto-dropped to secondary decisions.**

Sequence (per round: 65 minutes total)

Step	Activity	Duration	Who
1	Situation Update (facilitator reads aloud)	5 min	Facilitator
2	Inject Cards Reveal (distributes 4–6 inject events; participants read in silence)	3 min	Facilitator
3	Individual Decision Preparation (analyze injects + private information cards; formulate decisions for each assigned industry)	15 min	Industry Participants

Step	Activity	Duration	Who
4	Decision Submission (hard deadline; one Decision Worksheet per industry; late submissions dropped to secondary)	3 min	Industry Participants
5	Facilitator Market Shock (R2 only: facilitator imposes constraints on 2-3 industries)	3 min	Facilitator
6	Industry Health Signals (facilitator announces conditions based on cumulative scores; Rounds 2–4 only)	2 min	Facilitator
7	Facilitator Scoring (scores all submitted decisions on 3 dimensions; see Section F)	12 min	Facilitator
8	Cross-Industry Discussion & Results (discuss inter-industry dynamics, challenge reasoning, explore spillovers; optional Collective Bonus nominations; facilitator surfaces patterns and dependencies)	13–31 min	All

D. DECISION RULES

Industry-Level Decisions (Mandatory)

Each participant submits **one independent Decision Worksheet per assigned industry per round**. If you represent 2 industries, you submit 2 separate worksheets.

The 11 Industries:

Industry	Description
Retail	4,700+ store omnichannel retailer
CPG	Consumer goods manufacturer/brand company
Healthcare Provider	Integrated hospital system, clinical operations, physician workflows
Healthcare Payer	Health insurer, claims processing, prior authorization, coverage decisions
Finance	Major US bank + insurance underwriting, trading, fraud detection
Consulting	Management consulting firm, copilot deployment, client delivery transformation
Law	Law firm, document review automation, legal research, junior associate displacement
Manufacturing	28 plants, predictive maintenance, production automation
Logistics	Freight/3PL/warehouse, autonomous vehicle pilots, route optimization
Big Tech	Google/Meta/Microsoft-class company, AI model development, platform capabilities
B2B/B2C SaaS	Workday/Salesforce-class company, AI feature integration, pricing pressure

Format & Submission

- One Decision Worksheet per industry per round (e.g., a participant representing Retail and CPG submits 2 worksheets per round)
- Each industry decision is independent: Scored on its own merits against the 3 dimensions, evaluated in that industry's specific context
- Binding: Cannot be reversed once submitted. Sunk costs are real. Round 2 builds on Round 1 decisions.

Default Fallback

If a participant does not submit a Decision Worksheet for an assigned industry by the deadline, that industry defaults to its **pre-defined base case** for the round (a conservative, status-quo decision determined by the facilitator). The base case scores 0 on all three dimensions.

Cross-Industry Synergies

Participants representing multiple industries **may coordinate decisions across their assigned industries** (e.g., Retail and CPG may share data investment that benefits both). However:

- Each industry's decision must be articulated separately on its own worksheet
- Each decision is scored independently
- You must note cross-industry synergies in the "Cross-Industry Synergies" field (e.g., "Retail and CPG coordinating on shared data platform; demand signal benefits both")

Auto-Drop Rule

- If a participant misses the 3-minute submission deadline for any industry, that industry's decision is auto-dropped to secondary status (lower scoring ceiling)
- Secondary industry decisions score on the same 3 dimensions but are capped at +2 per dimension (vs. +3 for primary on red-flag triggers)

Decision Continuity

- You cannot undo prior-round decisions in any industry (sunk costs are real)
- You can expand, modify, or shut down prior decisions in later rounds
- Each industry decision can reference and build on its own prior-round decisions
- If new information contradicts an earlier decision, you must explain adaptation in the Continuity Note (e.g., "Healthcare Provider deployed AI diagnostics in R1, but R2 clinical workflow integration issues forced us to pause full rollout; pivoting to targeted AI in high-variance specialties instead")

E. PRIVATE INFORMATION CARDS

Distribution & Confidentiality

- Cards issued in Rounds 1, 2, 3 only (no Round 4 card; the final round is about synthesis)
- 3 cards per industry per round (Rounds 1–3):
 - Card 1: May be shared between related industries (e.g., Retail and CPG may receive the same Card 1, framed differently)
 - Cards 2 and 3: Unique to each industry
- Distributed face-down at the start of each round's decision preparation phase
- CONFIDENTIAL to your industry. Do not share contents with other industries unless the facilitator explicitly permits
- Cards are facts discovered through internal analysis, market intelligence, or regulatory contacts
- All private information cards are located in the 05_Participant_Private_Cards/ folder

Card Usage

- You may reference card information in your decision and cross-industry discussion (without revealing specifics)
- You should NOT read cards aloud to other industries
- Card contents must be noted on your Decision Worksheet: "Private card influence: [One-sentence summary of how this card shaped my decision]"
- Cards accumulate — keep all cards from prior rounds and refer to them in later rounds

Card Influence Disclosure

Every Decision Worksheet must include a line:

"Private Card Influence (if any): [Sentence describing how this round's card(s) affected this decision, or 'None']"

This ensures transparency about information asymmetry.

F. SCORING OVERVIEW

Scoring Framework: Banded Scoring During Play

Decisions are scored **per industry** on **three dimensions** using a banded scale:

- Strategic Fit {-2, 0, +2} (with rare +/-3 on red-flag trigger)
- Execution Risk {-2, 0, +2}
- Tail Risk {-2, 0, +2}
- Round Score Range: -6 to +6 per industry decision (primary) or -4 to +4 (secondary)

Dimension Definitions

Strategic Fit (-2, 0, +2, +/-3):

- +2: Decision aligns perfectly with the industry's position and evolving AI landscape. Shows conviction and market awareness.
- 0: Decision is defensible but lacks clarity or bold positioning. Safe, middle-of-the-road.
- -2: Decision misaligns with the industry's position or ignores market realities. Contradicts prior messaging or signals.
- +/-3 (Red-Flag Trigger Only): Rare. Used only if decision is reckless (+3 strategic genius) or catastrophic (-3 existential failure).

Execution Risk (-2, 0, +2, +/-3):

- +2: Decision is feasible given resources, timeline, and organizational capacity. Participant has credible path to delivery.
- 0: Decision is feasible but tight. Some risk of delay or scope creep.
- -2: Decision is ambitious beyond capacity. High risk of failure, delay, or cost overrun.

- +/-3 (Red-Flag Trigger Only): +/-3 reserved for decisions that are either trivially easy (+3, should have decided this already) or technically impossible (-3, no way to execute).

Tail Risk (-2, 0, +2, +/-3):

- +2: Decision accounts for downside scenarios (regulatory backlash, liability, customer defection, competitive response). Participant has contingency planning.
- 0: Decision acknowledges some risks but lacks contingency depth.
- -2: Decision ignores major downside risks (regulatory, reputational, or financial tail events). Participant is exposed.
- +/-3 (Red-Flag Trigger Only): +/-3 reserved for decisions that are either extremely defensive (+3, perhaps overly cautious) or recklessly exposed to catastrophe (-3).

Scoring Example

Retail Industry, Round 1 Decision: "Pilot personalization in 3 key categories, with built-in ROI pause points. Budget: 8 points. Timeline: 12 weeks. Success metric: +5% basket size in pilot categories without return-rate lift. Risk: Brand backlash if personalization feels invasive; we're monitoring social sentiment daily and have kill-switch authority."

Facilitator Scoring:

- Strategic Fit: +2 (clear positioning, market-smart)
- Execution Risk: +2 (tight timeline but achievable; proven pilot capability)
- Tail Risk: +2 (strong acknowledgment of brand risk and contingency plan)
- Round 1 Score: +6

G. COLLECTIVE BONUS (OPTIONAL)

How It Works

Each round, during the cross-industry discussion period, participants may collectively recognize an industry whose strategy was especially strong — or flag one whose strategy seems especially risky. This is **optional**: if no one feels strongly, no bonus is applied.

Step	What Happens
1. Facilitator opens floor	"Does anyone want to recognize an industry whose strategy was especially strong this round, or flag one that seems particularly risky?"
2. Open nominations	Participants who wish to respond name one industry in each direction. Not all participants need to speak. You cannot nominate your own industry.

Step	What Happens
3. Consensus check	If 3 or more participants agree on the same industry, the facilitator awards +2 (strong) or -2 (risky) to that industry's cumulative score.
4. No consensus	If fewer than 3 participants agree, or no one nominates, no bonus or penalty is applied.
5. Limits	Maximum one +2 and one -2 per round. Facilitator judgment is final.

Why This Matters

- Peer recognition creates social accountability without requiring forced voting every round.
- The ± 2 swing is meaningful over multiple rounds — it can shift your Industry Health condition (see below).
- The optional nature encourages genuine engagement: nominations carry weight because they are voluntary.

Facilitator Market Shock (Round 2 Only)

In Round 2, the facilitator imposes one constraint each on 2-3 industries, selected from the constraint menu. If your industry is targeted, you must adapt your decision. Constraints include: regulatory halts, labor cost surges, capital tightening, reputational pressure, competitive response, litigation risk.

G2. INDUSTRY HEALTH SIGNALS

Starting in Round 2, the facilitator announces your industry's **Health Condition** based on your cumulative score across all rounds:

Condition	What It Means	What Happens
Surge (+15 or higher)	Structurally advantaged. Margins expanding. Investors overweight.	You may attempt one Transformational-band decision without automatic red-flag challenge.
Tailwind (+6 to +14)	Navigating well. Market views you favorably.	No modifier. Business as usual.
Steady (-5 to +5)	Holding position. No clear momentum.	No modifier.

Condition	What It Means	What Happens
Headwind (-6 to -14)	Under structural pressure. Margins compressing.	One constraint imposed by facilitator at start of next round. You must address it before submitting your primary decision.
Crisis (-15 or worse)	Structural decline. EBITDA collapsing. Survival mode.	Two constraints imposed. Your primary decision must include a defensive or restructuring component.

Note: Collective Bonus results (± 2), when applied, affect your cumulative score and can shift your Industry Health condition over time.

H. AI ADOPTION ARC

The exercise is structured around a four-phase AI adoption arc that mirrors the round progression:

Round	Phase	Distribution
Round 1	Foundation	Included in your pre-read packet (you already have this)
Round 2	Acceleration	Handout distributed at start of Round 2
Round 3	Reckoning	Handout distributed at start of Round 3
Round 4	Normalization	Handout distributed at start of Round 4

- Each phase handout describes the evolving AI landscape, market conditions, and industry dynamics for that round
- Phase materials are located in the 04_Participant_AI_Adoption_Arcs/ folder
- Read each phase handout before beginning your decision preparation for that round

I. PRACTICE MICRO-ROUND

Before Round 1, all participants complete a 10-minute practice micro-round to familiarize with the process.

Micro-Round Script

1. Facilitator reads: "It's Q2 2026. Your industry is growing 5% YoY. A new AI competitor has entered the market. A major customer is asking about your AI strategy."
2. Inject card given: "A major customer has requested a custom AI solution — high-revenue opportunity but 6-month timeline and \$2M development cost."
3. Individual preparation: 5 minutes. Analyze: Do we bid on this? What are risks? How does it align with our industry position?
4. Decision submission: 1-minute deadline. Participant submits a simple 2-sentence decision on the practice worksheet (one per assigned industry).
5. Scoring: Facilitator scores on the 3 dimensions (banded). Participant sees scoring rubric in action.
6. Debrief: Facilitator explains: "This is the pace and format for Rounds 1–4."

Goal: Demystify the process, show participants the format, and calibrate scoring expectations. Highlight that each industry decision is scored independently.

J. DECISION WORKSHEET FIELDS

Every Decision Worksheet must include these fields (completed before the 3-minute deadline). Submit **one worksheet per industry**.

CRITICAL INSTRUCTION: Use bands; avoid invented precision. The goal is to build intuition about AI adoption trajectories, not to produce spreadsheet-ready estimates.

Required Fields

#	Field	Options & Instructions
1	Industry Identification	Write industry name (e.g., Retail, Healthcare Provider, Big Tech)
2	Decision Statement	1–2 sentences: What action is this industry taking? (e.g., "Healthcare Provider: Deploy AI-assisted clinical diagnostics in 5 pilot hospitals.")

#	Field	Options & Instructions
3	Spend/Commitment Band	Select one + 1–2 sentence justification: <input type="checkbox"/> Absorbable (BAU): Can absorb cost within existing operational budget <input type="checkbox"/> Material but manageable: Requires budget reallocation but within annual capital envelope <input type="checkbox"/> Transformational bet: Significant multiyear commitment; requires board approval or capital raise <input type="checkbox"/> Potentially existential: Bet-the-company scale decision; existential upside or downside
4	Time-to-Impact Band	Select one + 1–2 sentence justification: <input type="checkbox"/> 0–3 months: Quick wins, near-term visibility <input type="checkbox"/> 3–12 months: Mid-term execution; results visible by end of year <input type="checkbox"/> 1–2 years: Long-term investment; payoff deferred <input type="checkbox"/> 2+ years: Structural transformation; multi-year runway before payoff
5	Execution Complexity Band	Select one + 1–2 sentence justification: <input type="checkbox"/> Low: Straightforward to execute; organizational capability exists; proven in pilots <input type="checkbox"/> Medium: Requires cross-functional coordination or capability building; moderate risk <input type="checkbox"/> High: Significant technical, organizational, or integration risk; unproven in your context <input type="checkbox"/> Very high: Novel approach for your organization; high capability gap; execution risk substantial

#	Field	Options & Instructions
6	Dependency Band	Select one + 1–2 sentence justification: <input type="checkbox"/> Mostly internal: Success depends primarily on internal execution <input type="checkbox"/> Requires vendor/partner: Success depends on third-party technology or services <input type="checkbox"/> Requires regulator/union/standards: Success contingent on regulatory approval, union negotiation, or standards alignment <input type="checkbox"/> Requires ecosystem shift: Success depends on industry-wide changes or network effects
7	Scale	<input type="checkbox"/> Pilot: Isolated test (1 location, 1 function, limited user base) <input type="checkbox"/> Regional: Multiple locations or functions within a region <input type="checkbox"/> National: Broad rollout; most locations/functions included <input type="checkbox"/> Global: International expansion or organization-wide
8	Success Metric	Directional or range-based, NOT single-point estimate (e.g., "10–20% margin improvement," "Enables new market entry or revenue stream")
9	Key Risk & Mitigation	Top 1–2 risks + mitigations (1–2 sentences each): What is the main downside risk? What guardrails or mitigation are in place?
10	Cross-Industry Synergies	One sentence on coordination with another industry, or "Independent decision"
11	Private Card Influence	How this round's private card(s) shaped the decision, or "None" (Rounds 1–3 only)

#	Field	Options & Instructions
12	Continuity Note	How this builds on or modifies a prior-round decision (Rounds 2–4 only). E.g., "Builds on R1 clinical AI pilot; now expanding to full hospital system after successful safety validation."

K. KEY RULES SUMMARY

Rule	Details
Participants	5–11 industry representatives + 1-2 facilitators. Ideal: 8+ industries assigned. Facilitator plays unassigned industries (up to 3).
Decisions Per Round	1 Decision Worksheet per industry per round. If you represent 2 industries, you submit 2 worksheets. Up to 11 total decisions per round across all participants.
Submission Deadline	Hard 3-minute cutoff per round. Miss the deadline for any industry → that industry auto-drops to secondary.
Default Fallback	If no worksheet submitted for an industry, that industry defaults to its pre-defined base case (scores 0 across all dimensions).
Binding Decisions	Once submitted, decisions cannot be reversed. Sunk costs are real.
Industry Scoring	Each industry decision scored independently by the facilitator on 3 dimensions {-2, 0, +2}; +/-3 only on red-flag triggers.
Private Cards	3 cards per industry per round (Rounds 1–3). Card 1 may be shared between related industries; Cards 2–3 unique per industry. Confidential to your industry. Must note card influence on each worksheet. No card in Round 4.
AI Adoption Arc	Foundation phase in pre-read. Phases 2–4 distributed as handouts at start of Rounds 2–4.

Rule	Details
Collective Bonus	Optional each round: participants may nominate an industry for +2 (strong strategy) or -2 (risky strategy). Awarded only if 3+ participants agree. Max one +2 and one -2 per round.
Facilitator Market Shock	R2 only: facilitator imposes one constraint each on 2-3 industries from constraint menu.
Industry Health	Cumulative score determines condition (Surge +15+ / Tailwind +6-14 / Steady -5+5 / Headwind -6-14 / Crisis -15-). Announced each round starting R2. Headwind = 1 constraint. Crisis = 2 constraints + defensive play required.
Continuity	Industry decisions build on prior-round decisions for that industry. Note cross-industry synergies if applicable.
Cross-Industry Discussion	27-minute dedicated period each round for all participants to discuss inter-industry dynamics, challenge reasoning, and surface spillovers.
Cross-Industry Information	Industries operate independently during decision prep. Cross-industry discussion is the forum for information sharing and peer evaluation.

End of Rules of Play — V7.4

[INDUSTRY NAME] Industry Packet

FACILITATOR NOTE: The Core Packet section should be distributed 1 week before the exercise. Private Information Cards and AI Adoption Arc updates (Rounds 2-4) are distributed as separate handouts during the exercise. See *05_Participant_Private_Cards/* and *04_Participant_AI_Adoption_Arcs/* folders.

IMPORTANT: Before the exercise begins, review the **Rules of Play** document (*03zRulesofPlay.md*) for round structure, decision submission format, card reveal mechanics, and scoring criteria.

Core Packet

Industry Role

[1 paragraph. Define the company archetype: scale, revenue, employee count, market position, geographic footprint. Written in second person ("You are..."). Must be specific enough to anchor decisions but generic enough to represent the industry, not a single real company.]

Strategic Context

[3-5 paragraphs covering:]

- [Current competitive dynamics specific to this industry]
- [AI adoption status: what's deployed, what's piloted, what's aspirational]
- [Key cross-industry impacts: how other industries' decisions affect you]
- [Core strategic tension: the fundamental trade-off this industry faces]

Objectives

Objective	Target (Banded/Directional)	Driver
[Objective 1]	[Directional target, no false precision]	[What drives this]
[Objective 2]
[Objective 3]
[Objective 4]

Constraints

Constraint	Impact	Implications
[Constraint 1]	[Quantified where possible]	[Decision implications]
[Constraint 2]
[Constraint 3]
[Constraint 4]

Resources & Levers

[Category 1] Assets:

- [Bullet list of specific resources, quantified where possible]

[Category 2] Assets:

- [Bullet list]

Potential Paths Forward:

- [Bullet list of specific levers the participant can pull, each with brief description of ROI/risk profile]
-

AI Adoption Arc — Foundation Phase

[1-2 paragraphs. ONLY the Foundation phase (Phase 1) of the 4-phase arc. Phases 2-4 are distributed as separate handouts round-by-round. See

04_Participant_AI_Adoption_Arcs/04_x_[INDUSTRY]_Arc.md.]

Foundation (2025 – Q1 2026): [Description of current AI deployment status, early pilots, organizational readiness, margin impact for this specific industry.]

Strategic Considerations

1. [Consideration 1: Frame as trade-off or tension, not directive.] [1-2 sentences exploring the decision space. Use deliberative language: "The trade-off between...", "Consider whether...", "The question is...". Avoid imperative directives like "Prioritize X" or "Invest in Y".]
2. [Consideration 2] [Same pattern]
3. [Consideration 3] [Same pattern]
4. [Consideration 4] [Same pattern]

5. [Consideration 5] [Same pattern]

04

Participant AI Adoption Arcs

AI Adoption Arc — Distribution Guide

Overview

This folder contains the **AI Adoption Arc** for each of the 11 industries in Project Threshold V7.4. Each industry's arc describes how AI adoption evolves across 4 phases aligned to the exercise's 4 rounds.

Design Rationale

Only the **Foundation phase** (Phase 1) is included in each industry's pre-read packet. Phases 2-4 are distributed round-by-round as handouts so that participants experience the arc as it unfolds — preserving surprise and forcing genuine decision-making under uncertainty.

File Format

Each file (`04_x_[Industry]_Arc.md`) contains all 4 phases, each on a new page:

- Phase 1: Foundation — Marked as "already in pre-read" (included for facilitator reference)
- Phase 2: Acceleration — Distribute at start of Round 2
- Phase 3: Reckoning — Distribute at start of Round 3
- Phase 4: Normalization — Distribute at start of Round 4

Distribution Instructions

1. Print each industry's arc file before the exercise
2. Separate at page break markers between phases
3. Set aside Phase 1 pages (already in pre-read packet)
4. Distribute at the start of each round:
 - Phase 2 → Start of Round 2
 - Phase 3 → Start of Round 3
 - Phase 4 → Start of Round 4
5. Participants keep all phases and may refer to them throughout the exercise

Files

File	Industry
04aRetail_Arc.md	Retail
04bCPG_Arc.md	CPG
04cHealthcareProviderArc.md	Healthcare Provider
04dHealthcarePayerArc.md	Healthcare Payer
04eFinance_Arc.md	Finance
04fConsulting_Arc.md	Consulting
04gLaw_Arc.md	Law
04hManufacturing_Arc.md	Manufacturing
04iLogistics_Arc.md	Logistics
04jBigTechArc.md	Big Tech
04kSaaS_Arc.md	B2B/B2C SaaS

AI Adoption Arc — Retail

FACILITATOR NOTE: Phase 1 (Foundation) is already in the participant's pre-read packet. Distribute Phase 2 at the start of Round 2, Phase 3 at the start of Round 3, and Phase 4 at the start of Round 4. Each phase starts on a new page.

Phase 1: Foundation (2025 – Q1 2026)

[Already in pre-read — included here for facilitator reference only]

Retail AI deployment is concentrated in low-risk operational areas. Demand forecasting and inventory optimization pilots are running in select distribution centers and store clusters, showing modest but validated ROI through markdown reduction and improved in-stock rates. A personalization engine was deployed to 20% of e-commerce traffic in Q4 2025, lifting conversion but also increasing return rates — net incremental profit underperformed projections. Consumer-facing AI remains limited; organizational skepticism is high among store operations leadership. The \$150M AI/ML budget for 2026 is approved but uncommitted beyond current pilots. The workforce is watching closely — any signal of large-scale automation will trigger union attention. Margin impact so far: low, but directionally positive on the operational side.

Phase 2: Acceleration (Q2 – Q3 2026)

[DISTRIBUTE AT START OF ROUND 2]

Retail AI deployment is scaling rapidly across both operational and consumer-facing domains. Inventory optimization has moved from pilot to full rollout across all 18 distribution centers, delivering validated markdown reduction and in-stock improvements. The personalization engine has been retooled based on Q4 2025 lessons and expanded to 60% of e-commerce traffic with improved recommendation targeting. Dynamic pricing is being tested in additional metro markets. Same-day delivery AI route optimization is reducing fulfillment costs meaningfully in high-density areas.

But the acceleration is generating new pressures. Consumer awareness of AI-driven pricing and personalization is rising. Social media discussions about "algorithmic pricing" at major retailers are growing. CPG vendor partners — feeling the squeeze from your data demands and private-label expansion — are accelerating DTC launch timelines. Competitors are matching your operational AI investments, compressing the advantage window. Unionized distribution center workers are raising concerns about warehouse automation timelines. The investment ramp is material: AI/ML spending is on track to exceed budget as successful pilots demand scaling capital.

What Changed Since Foundation:

- Operational AI (inventory, logistics) moved from pilot to scaled deployment with confirmed ROI
- Consumer-facing AI expanded significantly but consumer awareness and scrutiny are rising
- Vendor competitive dynamics intensified — DTC threats and data demands are accelerating

Key Tension for This Phase: *The operational AI playbook is working. The question is whether to push consumer-facing AI aggressively while the competitive window is open, or consolidate operational gains and manage the growing backlash signals.*

Phase 3: Reckoning (Q4 2026 – Early 2027)

[DISTRIBUTE AT START OF ROUND 3]

Consumer backlash against AI-driven retail experiences has intensified. Customer sentiment surveys show meaningful declines in trust scores for AI-driven shopping, concentrated among highest-value loyalty members. Specific consumer complaints — perceived price discrimination, invasive recommendations, loss of control over the shopping experience — are coalescing into a coherent narrative that media and regulators are amplifying. Private-label penetration in AI-curated product categories is outpacing non-curated categories, suggesting consumers perceive algorithmic recommendations as biased toward retailer margins rather than consumer value.

Regulatory pressure is materializing. State attorneys general are issuing formal information requests about algorithmic pricing and personalization. The FTC is drafting retail-specific guidance on deceptive AI practices. Compliance costs are rising and the regulatory trajectory points toward mandatory disclosure and transparency requirements in 2027. Meanwhile, vendor DTC channels have launched, pulling volume out of your stores and e-commerce in several high-margin categories. Competitors who moved cautiously on consumer-facing AI are positioning themselves as "trustworthy alternatives."

The market is bifurcating: aggressive AI deployers are seeing margin gains offset by customer attrition and regulatory costs, while cautious deployers are preserving trust but losing operational efficiency advantages.

What Changed Since Acceleration:

- Consumer backlash moved from early signals to measurable sentiment decline and media amplification
- Regulatory action shifted from inquiry to formal proceedings and draft guidance
- Vendor DTC channels launched, creating real revenue displacement in key categories

Key Tension for This Phase: *The aggressive deployment strategy delivered margin gains, but trust erosion and regulatory exposure are now material costs. Do you double down and absorb the turbulence, or recalibrate toward transparency and trust recovery — and can you afford the margin hit of pulling back?*

Phase 4: Normalization (2027 onwards)

[DISTRIBUTE AT START OF ROUND 4]

The retail AI landscape is stabilizing around a new equilibrium. Consumer expectations have crystallized: operational AI (fast delivery, accurate inventory, efficient supply chain) is expected and valued. Consumer-facing AI (personalization, recommendations) is accepted when transparent, opt-in, and clearly serving the customer rather than extracting margin. Dynamic pricing remains contentious but is tolerated when bounded by published fairness rules. Regulatory frameworks are taking shape — disclosure requirements, algorithmic audit standards, and pricing transparency rules are becoming the cost of doing business.

The competitive landscape has sorted itself. Retailers who invested early in operational AI and managed consumer-facing AI with discipline hold structural advantages: lower costs, better fulfillment, and preserved brand trust. Those who pushed aggressive consumer-facing AI without trust guardrails are dealing with lingering reputational damage and elevated regulatory compliance costs. Vendor DTC channels are an established feature of the market — some retailers have adapted by offering logistics-as-a-service to vendors, turning a competitive threat into a revenue stream.

The talent and workforce picture has settled. Unionized workforce negotiations have produced frameworks for automation-driven role transitions. Retail associates are increasingly using AI tools for inventory, customer service, and task management — the "AI as associate support" framing has become industry standard where it was implemented early and authentically.

What Changed Since Reckoning:

- Consumer and regulatory expectations crystallized around transparency, opt-in, and fairness
- Competitive advantage shifted to retailers with operational AI depth and preserved brand trust
- Vendor DTC became a permanent market feature; best-positioned retailers adapted their business model

Key Tension for This Phase: *The turbulence is over but the new landscape rewards different capabilities than the old one. Your long-term competitive position now depends on proprietary data assets, operational AI maturity, workforce adaptation, and brand trust — not on who was most aggressive in deployment.*

AI Adoption Arc — CPG

***FACILITATOR NOTE:** Phase 1 (Foundation) is already in the participant's pre-read packet. Distribute Phase 2 at the start of Round 2, Phase 3 at the start of Round 3, and Phase 4 at the start of Round 4. Each phase starts on a new page.*

Phase 1: Foundation (2025 – Q1 2026)

[Already in pre-read — included here for facilitator reference only]

CPG AI deployment is concentrated in back-office and operational applications. R&D teams are piloting AI-assisted formulation and consumer preference modeling with promising early results — one pilot demonstrated a 6-month reduction in development cycles and \$35M in incremental revenue from accelerated launches. Marketing is testing AI-generated content for product descriptions and campaign copy, achieving substantial cost savings versus agency work with matched conversion performance. However, an AI-generated campaign using synthetic human faces triggered meaningful negative social sentiment, underscoring brand safety risks. Demand forecasting pilots are running across select product lines with modest accuracy improvements. Consumer-facing AI is minimal; the organization remains cautious about anything that touches the brand directly. The investment pipeline for 2026 is approved but largely uncommitted beyond current pilots. Margin impact so far: low, but the R&D and marketing pilots point to material upside if scaled carefully.

Phase 2: Acceleration (Q2 – Q3 2026)

[DISTRIBUTE AT START OF ROUND 2]

CPG AI deployment is scaling across R&D, marketing, and supply chain — with DTC channel development emerging as a strategic priority. R&D acceleration has expanded from a single pilot to 12 of 45 R&D centers, with AI-assisted formulation and consumer preference modeling now integrated into the product development workflow. The validated ROI is attracting internal demand faster than the organization can absorb it. Marketing automation is scaling AI-generated content for lower-risk categories (product descriptions, technical content, social media scheduling) while holding back on lifestyle imagery and testimonials after the Q4 2025 brand safety incident.

DTC channel economics have been validated for 8 priority brands, and launch planning is underway. Retailers are aware of the DTC signals and are escalating data demands and private-label development in direct response. Demand sensing AI has expanded across the supply chain, improving forecast accuracy but exposing how few suppliers have the digital infrastructure to participate in real-time data exchange. Competitors are making similar moves — the industry is entering a land-grab phase for direct consumer relationships and AI-driven cost efficiency.

Investment is ramping significantly. AI/ML spending across R&D, marketing, supply chain, and DTC infrastructure is exceeding initial budget allocations as successful pilots demand scaling capital. The organization is stretched across multiple transformation workstreams simultaneously.

What Changed Since Foundation:

- R&D acceleration moved from single pilot to multi-center deployment with confirmed ROI pulling resources
- DTC launch planning activated for priority brands, triggering retailer countermoves
- Marketing AI scaled in low-risk categories but brand safety constraints are limiting consumer-facing content automation

Key Tension for This Phase: *The ROI from R&D and operational AI is validated, but scaling requires capital and organizational bandwidth that competes with the DTC buildout. Meanwhile, retailer relationships are deteriorating as your DTC intentions become visible — and you still depend on retailers for 55% of revenue.*

Phase 3: Reckoning (Q4 2026 – Early 2027)

[DISTRIBUTE AT START OF ROUND 3]

The consumer authenticity backlash is hitting CPG brands hard. Rising public awareness that major CPG companies use AI-generated marketing content has triggered brand trust erosion, concentrated in categories where authenticity matters most — personal care, baby products, and premium food. A viral social media investigation cataloguing AI-generated content across major CPG brands (including yours) has amplified the narrative. Consumers are equating AI-generated content with "cutting corners" and questioning whether AI-driven product claims are substantiated.

Regulatory pressure is intensifying. The FTC is drafting guidance specifically targeting AI-generated endorsements, synthetic testimonials, and non-disclosed AI content in CPG advertising and packaging. State attorneys general are issuing inquiries about content generation practices. The compliance trajectory points toward mandatory disclosure requirements and restrictions on AI-generated claims in 2027.

Meanwhile, the DTC channel launches are producing mixed results. Consumer acquisition costs are higher than modeled, and retailer retaliation is materializing as projected — shelf-space reductions and pulled promotional support are hitting retail channel revenue. The brands that launched DTC are gaining direct consumer relationships but at a higher cost than anticipated. R&D acceleration continues to deliver strong results, largely insulated from the consumer-facing backlash. The market is bifurcating between companies doubling down on AI efficiency and those pivoting to "authentically human" brand positioning.

What Changed Since Acceleration:

- Consumer authenticity backlash moved from early signals to measurable brand trust erosion and media amplification
- Regulatory action on AI-generated content shifted from general inquiry to CPG-specific draft guidance
- DTC launches produced mixed results — consumer acquisition costs higher, retailer retaliation materializing

Key Tension for This Phase: *Your consumer-facing AI strategy (marketing content, DTC personalization) is under pressure from both consumers and regulators. Your back-office AI strategy (R&D, supply chain) is delivering. Do you separate these two tracks — double down on operational AI while pulling back consumer-facing AI — or does the brand trust erosion eventually contaminate everything?*

Phase 4: Normalization (2027 onwards)

[DISTRIBUTE AT START OF ROUND 4]

The CPG AI landscape is stabilizing around a new competitive equilibrium. Consumer expectations have crystallized: AI-driven product innovation and supply chain efficiency are accepted and valued (consumers want better products faster). AI-generated marketing content is tolerated only when disclosed, high-quality, and clearly supplementary to human-created brand storytelling. Synthetic testimonials and AI-generated endorsements are effectively banned by regulatory guidance and consumer expectation alike.

The competitive landscape has sorted. CPG companies that invested early in R&D acceleration hold structural advantages — faster innovation cycles, better consumer preference prediction, and lower development costs create a compounding edge. Those that managed the marketing AI transition with discipline (transparent disclosure, maintained human creative for premium brands) preserved brand equity through the backlash. Companies that pushed undisclosed AI content aggressively are dealing with lingering brand trust deficits and elevated compliance costs.

DTC channels have found their level. They are a meaningful but not dominant revenue stream for premium and high-affinity brands. Retailer relationships have stabilized around new terms — data sharing frameworks are formalized, DTC coexistence is accepted, and private-label competition has intensified but plateaued. The companies best positioned are those with diversified channel strategies (retail + DTC + emerging platforms) rather than those that bet heavily on any single channel.

The supply chain has matured. AI-driven demand sensing is now standard across digitally ready supplier networks. Manufacturing optimization is delivering consistent efficiency gains. The CPG companies with the deepest supplier ecosystem integration hold cost advantages that are difficult for competitors to replicate.

What Changed Since Reckoning:

- Consumer expectations crystallized: operational AI valued, consumer-facing AI accepted only with transparency and disclosure
- R&D acceleration became a structural competitive advantage for early movers
- DTC and retail channels reached a stable coexistence after the turbulence of retailer retaliation

Key Tension for This Phase: *The new CPG landscape rewards R&D speed, brand authenticity, and channel diversification. Your long-term competitive position depends on whether you built the right capabilities during the turbulent years — or spent your capital and organizational bandwidth fighting fires.*

AI Adoption Arc — Healthcare Provider

FACILITATOR NOTE: Print this document. Each phase starts on a new page. Phase 1 (Foundation) is already included in the participant's pre-read packet — set aside for facilitator reference. Distribute Phase 2 at the start of Round 2, Phase 3 at the start of Round 3, and Phase 4 at the start of Round 4.

Phase 1: Foundation (2025 - Q1 2026)

Already in pre-read packet. Included here for facilitator reference.

Your AI deployment is concentrated in back-office clinical operations where regulatory risk is manageable and ROI is demonstrable. Medical coding automation has improved billing accuracy by 12% and reduced coding rework cycles. Documentation support tools are operational in targeted specialties, reducing physician administrative time. Patient risk stratification models are running across your ambulatory network, identifying high-risk patients for care coordination intervention.

Diagnostic AI pilots in radiology and pathology are running in controlled clinical settings. Accuracy results are promising — meaningful improvement over radiologist baseline in specific imaging modalities. But these remain pilots, not deployments. Regulatory uncertainty on FDA validation requirements and CMS Conditions of Participation for AI-assisted diagnostics is the primary constraint. You cannot deploy clinical AI at scale without regulatory clarity, and that clarity has not arrived.

Physician sentiment is split. Administrative AI that reduces documentation burden is welcomed — medical coding and prior auth prediction tools see reasonable adoption. Diagnostic AI that increases validation overhead faces resistance from radiologists and senior clinicians who view the additional review burden as unacceptable. Budget tensions persist: administrative AI delivers near-term savings, clinical AI requires 18-24 month regulatory and validation investment before any revenue impact.

What Changed:

- Administrative AI (coding, documentation, risk stratification) is operational and delivering ROI
- Clinical diagnostic AI pilots show accuracy gains but face physician workflow friction
- Regulatory uncertainty on FDA/CMS clinical AI requirements prevents scale deployment
- Budget tension between quick-win administrative AI and long-cycle clinical AI is unresolved

Key Tension: Accuracy improvements in clinical AI pilots are real, but physician adoption and regulatory approval are the binding constraints — not technology performance.

Phase 2: Acceleration (Q2 - Q4 2026)

FDA and CMS guidance on clinical AI requirements arrives in Q2 2026. The regulatory clarity you have been waiting for is here — but it comes with substantial compliance burden. FDA requires external clinical validation, demographic bias audits, and real-time performance monitoring for all diagnostic AI tools. CMS introduces new Conditions of Participation requiring demonstrated clinical benefit, physician oversight governance, and ongoing algorithm performance reporting.

Clinical validation costs increase sharply. Each diagnostic AI tool requires \$5-10M in external validation through academic medical center partnerships, 12-18 months of pre-market work, and ongoing monitoring infrastructure. Diagnostic AI pilots begin transitioning to limited rollout — initially constrained by geography (select facilities) and specialty (radiology first, then pathology). Prior authorization and administrative AI refinement continues, with incremental efficiency gains.

The physician adoption challenge intensifies. Regulatory clarity actually increases physician anxiety — liability frameworks remain unclear even as deployment timelines become concrete. Some physicians view the new governance requirements as validation of their concerns; others see regulatory frameworks as enabling responsible adoption. Medical staff leadership is demanding formal governance structures before broader clinical AI deployment.

What Changed:

- FDA/CMS regulatory guidance arrives — clarity on requirements but substantial compliance burden
- Clinical validation costs crystallize at \$5-10M per tool, 12-18 month timelines
- Diagnostic AI transitions from pilot to limited rollout (geographic, specialty-constrained)
- Physician adoption tension intensifies as deployment becomes real rather than theoretical
- Smaller health systems begin exiting clinical AI due to validation costs; consolidation pressure increases

Key Tension: Regulatory clarity enables action but also reveals the true cost. Early movers who began validation work 12-18 months ago have a significant head start over those who waited.

Phase 3: Reckoning (Q4 2026 - Q1 2027)

Clinical AI patient harm incidents emerge across the industry. Not at your facilities specifically (unless your deployment decisions were aggressive), but across the healthcare sector — diagnostic AI misinterpretations contributing to delayed diagnoses, missed findings in radiology, and inappropriate treatment recommendations. The incidents are rare but high-profile, generating media coverage, congressional attention, and patient advocacy group mobilization.

Physician adoption slows further. The harm incidents validate physician concerns about liability and autonomy. Malpractice insurers announce premium adjustments for physicians using AI diagnostic tools without formal governance frameworks. Your physician advisory board (if empowered) gains leverage; if overridden in earlier rounds, physician resistance hardens. Regulatory burden increases — FDA announces additional validation requirements and mandatory adverse event reporting for clinical AI. CMS tightens Conditions of Participation with new audit requirements.

Investment appetite for clinical AI contracts. Board and investor confidence in clinical AI timelines erodes. Capital becomes more expensive as the broader economic environment tightens. The temptation to retreat to administrative AI (lower risk, faster ROI) intensifies. But organizations that invested in governance, validation, and physician engagement in earlier phases are better positioned to weather the storm — their deployments are smaller, better governed, and have physician buy-in.

What Changed:

- Industry-wide clinical AI patient harm incidents generate regulatory and public backlash
- Physician adoption decelerates; malpractice insurers adjust premiums for ungoverned AI use
- FDA and CMS add validation requirements and mandatory adverse event reporting
- Investment pullback on clinical AI; capital becomes more expensive
- Organizations with strong governance and physician engagement are more resilient

Key Tension: Clinical AI promise collides with clinical AI reality. The question shifts from "can we deploy?" to "should we have deployed differently?" — and the answer depends on how much governance and physician engagement you invested in earlier.

Phase 4: Normalization (2027+)

Clinical AI pathways become well-established. The regulatory framework — FDA validation, CMS Conditions of Participation, state medical board standards — is no longer novel; it is the cost of doing business. Clinical validation costs decline as methodologies mature, academic partnerships standardize, and shared validation infrastructure emerges. What was a \$5-10M, 18-month process per tool begins compressing toward \$3-5M and 12 months as the ecosystem develops.

Early clinical AI adopters — those who invested in validation, governance, and physician engagement during the Foundation and Acceleration phases — hold a 2-3 year lead on competitors. Their clinical AI tools are deployed, generating clinical outcome improvements and operational efficiencies that late entrants cannot replicate quickly. This advantage is durable because it is rooted in organizational capability (governance, physician trust, validation infrastructure) rather than technology alone.

HIPAA compliance, clinical governance, regulatory approval, and physician oversight become standard practice requirements. Clinical AI becomes a margin improvement tool and quality differentiator, not a competitive moat — because the regulatory framework ensures all deployers meet minimum standards. The sector has consolidated: smaller health systems that could not afford validation and compliance costs have exited clinical AI or been acquired. Barriers to entry are higher but the playing field among incumbents is leveling.

What Changed:

- Regulatory frameworks are established and compliance costs are declining (but remain material)
- Early adopters hold 2-3 year capability lead on clinical AI deployment
- Clinical AI transitions from competitive advantage to operational necessity
- Sector consolidation: smaller systems exited or were acquired; barriers to entry increased
- Physician governance and oversight is standard — organizations without it cannot deploy

Key Tension: AI becomes infrastructure, not strategy. The winners are not those with the best models but those who built the organizational capability — governance, physician trust, regulatory relationships — to deploy responsibly and sustain through the reckoning.

AI Adoption Arc — Healthcare Payer

FACILITATOR NOTE: Print this document. Each phase starts on a new page. Phase 1 (Foundation) is already included in the participant's pre-read packet — set aside for facilitator reference. Distribute Phase 2 at the start of Round 2, Phase 3 at the start of Round 3, and Phase 4 at the start of Round 4.

Phase 1: Foundation (2025 - Q1 2026)

Already in pre-read packet. Included here for facilitator reference.

Your operational AI deployment is mature and delivering measurable ROI. Prior authorization AI generated \$22M in savings through optimized coverage determinations and reduced processing time. AI-driven claims processing reduced fraud loss by 15% through pattern detection on claim submissions and provider billing anomalies. Medical coding adjudication automation is reducing rework cycles and improving accuracy. These are enterprise-scale deployments, not pilots — and internal pressure to show continued near-term returns is intense.

However, the next wave of AI investment faces a different risk profile. Coverage determination algorithms are drawing regulatory attention around fairness and disparate impact. State insurance regulators are signaling increased scrutiny of algorithmic coverage decisions — particularly prior authorization denials. Your fraud detection models are showing early signs of degradation as AI-generated synthetic claims become more sophisticated. The arms race between your detection capabilities and AI-enabled fraud is already underway.

You are positioned to move faster than clinical AI providers — your regulatory cycles are shorter and your ROI is more immediate. But the speed advantage comes with growing compliance risk that your current governance infrastructure was not designed to handle at this scale. Algorithmic transparency, bias testing, and fairness auditing are not yet embedded in your operational processes.

What Changed:

- Operational AI (prior auth, claims, fraud detection) is enterprise-scale and delivering ROI
- Coverage determination algorithms face growing regulatory scrutiny on fairness and disparate impact
- Fraud detection models show early performance degradation against AI-generated synthetic claims
- Governance infrastructure lags behind deployment pace — transparency and bias testing gaps exist

Key Tension: You can move faster than clinical providers, but speed without governance creates regulatory exposure that is materializing in real time.

Phase 2: Acceleration (Q2 - Q4 2026)

Coverage determination AI becomes table stakes across the insurance industry. Every major competitor is deploying prior authorization optimization, claims automation, and member risk stratification. The competitive question is no longer whether to deploy but how to deploy responsibly — because regulatory scrutiny is catching up to deployment speed.

Fair coverage denials become a major regulatory focal point. Competitor enforcement actions emerge: CMS and state insurance boards cite major insurers for algorithmic coverage denials that exhibit patterns of inappropriate denial. These are not your enforcement actions (yet), but they dramatically increase compliance burden industry-wide. Regulators begin requesting algorithmic documentation, bias testing results, and disparate impact analyses from all major insurers. Your compliance team is scrambling to produce documentation for algorithms that were deployed for operational speed, not regulatory transparency.

Payer AI pricing and profitability models face their own regulatory questions. State insurance boards begin questioning whether AI-driven underwriting and risk adjustment models incorporate prohibited factors through proxy variables. Medical loss ratio calculations using AI-assisted claims processing attract actuarial scrutiny. The regulatory environment is not hostile, but it is no longer permissive — the era of "deploy and optimize" is transitioning to "deploy, document, and defend."

What Changed:

- Coverage determination AI is industry-wide; competitive differentiation shifts to governance quality
- Competitor enforcement actions for inappropriate prior auth denials increase compliance burden for all insurers
- Regulators demand algorithmic documentation, bias testing, and disparate impact analysis
- Payer AI pricing and profitability models face underwriting and actuarial regulatory scrutiny
- Smaller insurers begin exiting AI-heavy product lines due to compliance costs

Key Tension: Operational AI that was deployed for efficiency must now be defended for fairness. Organizations that built governance alongside deployment are positioned; those that bolted on compliance after the fact face expensive retrofit.

Phase 3: Reckoning (Q4 2026 - Q1 2027)

Prior authorization enforcement actions escalate. CMS and state insurance boards announce significant penalties against major insurers whose prior auth algorithms exhibited statistically significant disparate impact — higher denial rates correlated with member demographics, geographic proxies, and socioeconomic indicators. The enforcement actions include financial penalties, mandatory algorithm audits, required member remediation programs, and public disclosure requirements. Congressional hearings on "AI-driven healthcare denials" generate sustained media coverage.

Industry-wide prior auth audits are now certain within 12-18 months. Every major insurer will face regulatory examination. Your exposure analysis reveals vulnerabilities — denial rate disparities correlated with member zip code and age cohort that may not be intentional but meet the regulatory standard of disparate impact. Simultaneously, your fraud detection accuracy continues to degrade. The synthetic fraud arms race has intensified: AI-generated claims are more sophisticated, and your competitors who invested in next-generation detection 12-18 months ago are pulling ahead.

Regulatory pressure to slow down and increase governance is intense. Capital becomes more expensive as the broader economic environment tightens. The combined cost of prior auth remediation, fraud detection investment, and enhanced compliance infrastructure strains your technology budget. Investment trade-offs that were manageable in earlier phases become acute: you cannot fund everything simultaneously.

What Changed:

- Major enforcement actions for prior auth disparate impact with financial penalties and mandatory remediation
- Industry-wide prior auth audits certain within 12-18 months; your exposure analysis reveals vulnerabilities
- Synthetic fraud detection arms race intensifies; competitors who invested early are pulling ahead
- Combined compliance and investment costs strain technology budgets; acute capital allocation trade-offs
- Congressional and media scrutiny of AI-driven healthcare coverage denials

Key Tension: The bill comes due. Prior auth savings built on unaudited algorithms face remediation costs. Fraud detection deferred faces compounding losses. The question is no longer "how much do we invest?" but "what do we deprioritize when we cannot fund everything?"

Phase 4: Normalization (2027+)

Regulatory frameworks for payer AI mature and stabilize. Prior authorization testing, algorithmic governance, explainability requirements, and disparate impact auditing become industry standard operating procedures. State insurance boards publish model governance frameworks. CMS finalizes Conditions of Participation for AI-assisted coverage determinations. The compliance burden is permanent but predictable — and organizations that invested in governance infrastructure early bear lower marginal costs.

Payer AI deployment accelerates on a level playing field. With regulatory frameworks established, the competitive advantage shifts from deployment speed to operational excellence — model accuracy, governance quality, member satisfaction, and cost efficiency within compliance boundaries. Back-office operations are entirely transformed: claims processing, coding adjudication, and routine prior authorization are overwhelmingly automated with human oversight reserved for high-complexity and high-risk decisions.

Fraud detection capability separates winners from losers. Insurers who invested in next-generation synthetic fraud detection during the Acceleration phase have materially lower fraud loss rates. Those who delayed face compounding losses and are acquiring capability at premium prices. The fraud detection arms race is permanent — but the leaders established their position 18-24 months ago.

The sector has consolidated. Higher compliance costs and governance requirements create barriers to entry for smaller insurers and new market entrants. The largest insurers absorb compliance costs more efficiently and benefit from data scale advantages in model accuracy. AI becomes the cost of doing business in health insurance — not a source of competitive differentiation, but a prerequisite for operating at scale.

What Changed:

- Regulatory frameworks are established and compliance is standardized (permanent but predictable)
- Prior auth governance, bias testing, and explainability are industry standard
- Back-office operations are overwhelmingly automated; human oversight reserved for high-risk decisions
- Fraud detection capability is a durable competitive differentiator for early investors
- Sector consolidated; higher compliance costs create barriers to entry; AI is cost of doing business

Key Tension: AI becomes infrastructure. The winners are those who invested in governance and fraud detection capability during the turbulence — not because they had better technology, but because they built the organizational discipline to deploy responsibly and the detection capability to stay ahead of evolving threats.

Finance AI Adoption Arc

FACILITATOR NOTE: Print this document and separate at page breaks. Phase 1 (Foundation) is already included in the Finance Industry Packet pre-read — set aside for facilitator reference. Distribute Phase 2 at start of Round 2, Phase 3 at start of Round 3, Phase 4 at start of Round 4. Participants keep all phases and may refer to them throughout the exercise.

Phase 1: Foundation (2025 – Q1 2026)

Already included in Finance Industry Packet pre-read. Included here for facilitator reference.

AI deployment in the financial services industry is concentrated in back-office and internal risk management operations where regulatory requirements are well-understood and ROI has been validated over multiple budget cycles. Fraud detection models are operational across transaction monitoring, claims adjudication, and identity verification workflows. KYC/AML compliance processes are partially automated, reducing manual review volume while maintaining audit trail quality. Document review and data extraction tools are deployed in lending operations and insurance claims processing, delivering measurable labor cost reduction.

Internal risk models now incorporate machine learning for loss estimation, portfolio stress testing, and capital adequacy projections, though regulatory reporting still relies on traditional actuarial and statistical approaches. Consumer-facing AI deployment remains conservative. Loan underwriting pilots are in controlled testing with mandatory human-in-the-loop review for all credit decisions. Robo-advisory functionality is limited to account aggregation, basic portfolio rebalancing, and educational content — no autonomous investment recommendations. The institution has invested heavily in data infrastructure, model governance frameworks, and bias testing protocols, but the regulatory pathway for deploying AI in lending and underwriting decisions at scale remains unclear. Talent competition is fierce, with fintech challengers and Big Tech firms aggressively recruiting data science and ML engineering teams. Margin impact is modest but directionally positive in back-office efficiency; front-office AI revenue contribution is minimal pending regulatory clarity.

What Changed:

- Fraud detection and KYC/AML automation operational and delivering validated ROI
- Internal risk models incorporating ML for stress testing and capital adequacy
- Underwriting AI in controlled pilot with human-in-the-loop; no autonomous deployment
- Model governance frameworks and bias testing protocols established
- Consumer-facing AI limited to low-risk, non-decisional features
- Aggressive talent competition from fintech and Big Tech constraining hiring

Key Tension: Back-office AI delivers reliable returns, but the high-value frontier — AI-driven underwriting and lending — is locked behind regulatory uncertainty that no amount of internal investment can resolve unilaterally.

Phase 2: Acceleration (Q2 – Q4 2026)

Regulatory clarity begins to arrive. The OCC and Federal Reserve issue joint guidance on AI model governance expectations for supervised institutions, providing a clearer (though not complete) framework for AI deployment in lending and underwriting decisions. The guidance establishes minimum standards for model documentation, explainability, bias testing, and ongoing monitoring — but stops short of prescribing specific model architectures or explicitly approving ensemble methods. This ambiguity creates both opportunity and risk: institutions with strong governance infrastructure can move forward with deployment, while those without face longer timelines to meet the new standards.

Front-office AI deployment accelerates. Underwriting models begin production deployment for select lending products — initially personal loans and auto lending, where regulatory risk is lower and loss data is abundant. Insurance underwriting AI expands beyond pilots into production for standard P&C lines. Robo-advisory platforms add AI-driven financial planning recommendations, though fiduciary review processes remain human-supervised. Early results are strong: institutions deploying AI underwriting report measurable improvement in loss ratios and approval rates versus traditional methods. Meanwhile, fair lending enforcement actions against fintech competitors increase industry-wide compliance burden. Every major bank increases compliance spending and accelerates bias testing infrastructure buildout. Talent competition intensifies further as the entire sector scales AI teams simultaneously.

What Changed:

- OCC/Federal Reserve joint guidance on AI model governance provides partial regulatory clarity
- Underwriting AI moves from pilot to production for select lending and insurance products
- Early deployers report measurable loss ratio and approval rate improvement
- Fair lending enforcement actions against fintechs increase compliance burden across all banks
- Compliance and governance spending accelerates across the sector
- Robo-advisory expands features but maintains human supervisory review
- Talent competition intensifies as entire sector scales AI simultaneously

Key Tension: Regulatory guidance unlocks deployment but also raises the compliance bar — first movers gain competitive advantage in underwriting, but also become the first targets for examination under the new framework.

Phase 3: Reckoning (Q4 2026 – Q1 2027)

The CFPB enforcement action against a major fintech for AI-driven lending discrimination sends shockwaves through the industry. The penalties are substantial, the remediation requirements are invasive, and the compliance monitoring agreement extends for years. Within weeks, OCC and Federal Reserve examiners notify all major banks that enhanced AI model examinations will begin. Fair lending audits cascade across the sector. Every institution deploying AI in lending or underwriting must now demonstrate — under examination — that their models do not produce disparate impact against protected classes.

The examination pressure exposes a spectrum of readiness. Institutions that invested heavily in governance and bias testing infrastructure (per Card 1 and Phase 2 guidance) are better positioned but still face material examination burden. Those that deployed aggressively without proportionate governance investment face model remediation demands, potential consent orders, and customer remediation costs. Synthetic fraud detection becomes a parallel arms race — AI-generated fraud continues to escalate, and institutions that fell behind in Phase 2 now face compounding losses. Model explainability moves from regulatory preference to regulatory requirement. The era of "deploy and iterate" is over for customer-facing AI; the standard is now "demonstrate compliance before deployment." Capital markets react: bank stocks with significant AI exposure face volatility as investors price in regulatory uncertainty. Cost of capital rises modestly across the sector.

What Changed:

- CFPB enforcement action triggers industry-wide fair lending audit cascade
- Enhanced AI model examinations begin at all major supervised institutions
- Model explainability becomes regulatory requirement, not preference, for customer-facing AI
- Synthetic fraud escalation compounds losses for institutions behind on detection capability
- Governance and compliance investment becomes clearly differentiated competitive advantage
- Capital markets reprice AI risk in bank stocks; cost of capital rises modestly
- Regulatory uncertainty shifts from "what are the rules?" to "how strictly will they be enforced?"

Key Tension: Institutions that invested in governance infrastructure early are vindicated — but even they face material examination costs and deployment slowdowns. The question shifts from "can we deploy AI?" to "can we prove our AI is fair?"

Phase 4: Normalization (2027+)

Regulatory frameworks for AI in financial services mature into established practice. Fair lending testing, model governance, explainability documentation, and ongoing bias monitoring become standard components of the examination process — burdensome but predictable. The institutions that survived the Reckoning phase with their governance infrastructure intact emerge stronger: they can deploy new AI applications faster because they have the compliance machinery to support examination. Smaller institutions and fintechs that lack governance scale face higher per-model compliance costs, creating barriers to entry that favor large, well-capitalized players.

The competitive landscape has reshaped. AI-driven underwriting is standard across major banks, with interpretable model architectures as the norm for customer-facing decisions. The accuracy penalty of interpretable models is partially offset by improvements in interpretable ML techniques and richer data inputs. Fraud detection has stabilized into a managed arms race — major institutions invest continuously in synthetic identity detection as a cost of doing business, with periodic model refreshes and vendor partnerships. Back-office transformation is largely complete: claims processing, KYC/AML, document review, and trading operations are heavily automated. The frontier has shifted to AI-driven wealth management personalization, commercial lending optimization, and cross-product customer lifecycle management — areas where AI creates competitive differentiation rather than just operational efficiency. Premium AI advisory services emerge within the sector: specialized regulatory technology firms and in-house capabilities serve smaller institutions that cannot build their own governance infrastructure.

What Changed:

- Regulatory examination of AI models is standard, predictable, and ongoing
- Interpretable model architectures are the norm for customer-facing decisions
- Governance infrastructure is a competitive moat favoring large, well-capitalized institutions
- Back-office AI transformation is largely complete across the sector
- Fraud detection is a managed, continuous investment — arms race is ongoing but stabilized
- Competitive frontier shifts to wealth management personalization and commercial lending optimization
- Compliance costs are material but create barriers to entry that favor incumbents

Key Tension: AI in finance has moved from strategic opportunity to operational necessity. The winners are not the institutions that deployed fastest, but those that built the governance infrastructure to deploy sustainably under regulatory scrutiny.

Consulting AI Adoption Arc

FACILITATOR NOTE: Phase 1 (Foundation) is included in the Consulting pre-read packet (03fConsulting.md). This document contains all 4 phases for facilitator reference and for round-by-round distribution. Print and separate at page breaks. Distribute Phase 2 at the start of Round 2, Phase 3 at the start of Round 3, and Phase 4 at the start of Round 4.

Phase 1: Foundation (2025 – Q1 2026)

[Already distributed in pre-read packet — included here for facilitator reference]

AI copilot deployment across your firm is underway and producing measurable results, but the transition is messier than the pilot data suggested. Research and analysis copilots are active on approximately 80% of engagement teams, delivering 25-40% time savings on routine analytical workstreams — market sizing, competitive benchmarking, data synthesis, and first-draft deliverable production. Slide production copilots have reduced structured deliverable creation time by roughly 50%. These are real gains that have caught the attention of both your partnership and your clients.

But the gains are unevenly distributed and come with costs that were not fully anticipated. Quality varies by use case: copilots perform well on structured, data-rich tasks but produce unreliable output on ambiguous or novel analytical questions. Partners and managers are absorbing a meaningful quality assurance burden — reviewing and correcting AI output is faster than producing from scratch, but not as fast as reviewing strong junior work. Junior consultant utilization is declining as routine tasks migrate to copilots. Associates report spending more time on QA and editing than on the complex analytical work they expected when they joined the firm.

On the revenue side, early AI advisory engagements with financial services and healthcare clients are generating premium pricing and strong client feedback. But these engagements are still a small fraction of total revenue. Meanwhile, a handful of sophisticated clients are beginning to ask pointed questions about how AI is affecting your delivery costs — and whether those cost savings should be reflected in pricing. The pricing pressure is early but directional: it will intensify.

What Changed:

- AI copilots deployed across 80% of engagement teams with 25-40% productivity gains on routine work
- Junior consultant utilization declining — early signs of talent pipeline stress
- Early AI advisory engagements with financial services and healthcare clients at premium rates
- Client pricing pressure beginning but manageable
- Partner quality review burden increasing — partially offsetting efficiency gains
- MBA recruiting acceptance rates declining; junior attrition rising

Key Tension: Productivity gains are real, but they are destabilizing the leverage model and talent pipeline faster than you can adapt the business around them.

Phase 2: Acceleration (Q2 – Q4 2026)

[Distribute at start of Round 2]

The dynamics that emerged in the Foundation phase are accelerating — and new competitive pressures are compounding them. AI advisory has become a major revenue driver. Client demand for AI transformation strategy, deployment support, governance frameworks, and organizational change management has surged across every sector. Financial services clients need AI model governance and regulatory compliance. Healthcare clients need clinical AI deployment and ethics frameworks. Manufacturing clients need automation strategy and workforce planning. Energy clients need AI-driven operational optimization. Your firm's cross-industry positioning is a genuine advantage — you can staff multi-sector engagements that specialists cannot.

But the competitive landscape is intensifying faster than anticipated. Specialized AI consultancies — firms with 200-500 employees, deep technical talent, and AI-native operating models — are winning a growing share of transformation engagements. They staff faster, deliver cheaper, and position themselves as more technically credible on implementation work. The major cloud and AI platform vendors (AWS, Google, Microsoft) are expanding their professional services arms, bundling advisory with platform access at rates your firm cannot match. And your largest clients are building internal AI Centers of Excellence with explicit mandates to reduce external consulting spend.

Pricing pressure has moved from early signals to active negotiation. Multiple Fortune 100 clients have formally requested AI-adjusted rate cards. Two top-20 clients have rebid active engagements, citing AI-driven efficiency expectations. Your blended rates on AI-related work are running 15-22% below traditional rates. Time-based billing is becoming untenable on an expanding share of engagements — clients see the compressed timelines and refuse to pay for hours that were not worked.

Talent repositioning is at its most difficult and most critical. You are actively redeploying junior consultants toward client relationship management, change management support, and synthesis roles. Retraining programs are running across all offices. The investment is substantial, the results are uncertain, and the internal disruption is significant. Some junior cohorts are thriving in new roles; others are struggling with the transition. Attrition among first- and second-year consultants remains elevated.

What Changed:

- AI advisory revenue surging — becoming a top-3 revenue driver across the firm
- Specialized AI consultancies and platform vendors winning 15-20% of AI transformation deals
- Major clients building in-house AI capability, reducing external consulting demand for implementation
- Pricing pressure accelerating — blended rates on AI work 15-22% below traditional rates
- Talent repositioning underway at scale — results mixed, investment high, internal disruption significant
- Partner economics under strain — revenue per partner declining on AI-assisted engagements
- Win rate on horizontal AI strategy work declining; win rate on vertical-specialized AI work holding steady

Key Tension: AI advisory is your biggest growth opportunity, but you are fighting for share against faster, cheaper, more specialized competitors — and your own clients' internal teams — while simultaneously managing a pricing transition and talent overhaul.

Phase 3: Reckoning (Q4 2026 – Q1 2027)

[Distribute at start of Round 3]

The competitive and commercial pressures that built through Acceleration are now producing structural consequences. Disintermediation is accelerating: your clients are not just building AI capability — they are using it to reduce their dependence on consulting firms for a widening range of work. Fortune 100 companies that previously relied on your firm for market analysis, competitive intelligence, organizational diagnostics, and even strategic framing are now performing this work internally using AI tools and small internal strategy teams. The work that remains for external consultants is more complex, more specialized, and harder to staff — but there is less of it.

The pricing model is in active crisis. Time-based billing now covers less than 60% of your AI-related engagements, down from 85% eighteen months ago. The remaining 40% is a mix of fixed-fee, milestone-based, and outcome-based arrangements negotiated on a deal-by-deal basis without consistent commercial frameworks. Margin performance varies wildly: some value-based engagements are highly profitable; others have been margin-destructive due to scope creep, client expectation misalignment, and underpricing. Your CFO is flagging that the firm's overall margin trajectory is negative for the first time in a decade — not because revenue is declining, but because the revenue mix is shifting toward lower-margin engagement structures faster than costs are adjusting.

Market consolidation is beginning. Two mid-tier consulting firms have announced mergers in the past quarter, seeking scale to compete with your firm and the Big Four. A major specialized AI consultancy was acquired by a Big Four rival, giving them a capability your firm lacks. Rumors of additional M&A activity are credible. The market is bifurcating: premium strategic firms with deep vertical expertise are defending margins and growing; commodity delivery firms are facing consolidation pressure; and firms stuck in the middle — broad but not deep, efficient but not specialized — are losing positioning.

Your firm's talent model is under existential scrutiny. The partnership is divided on the path forward. One faction argues for aggressive headcount reduction in junior ranks (acknowledging that the pyramid is permanently flattened) and reinvestment in senior, specialized talent. Another argues that gutting the junior pipeline will destroy the firm's ability to develop future leaders and that the right answer is to redesign junior roles around synthesis, judgment, and client management. Both sides have merit. Neither has a clean financial model.

What Changed:

- Client disintermediation accelerating — Fortune 100 companies performing more analytical and strategic work in-house using AI
- Time-based billing now covers less than 60% of AI-related engagements; commercial frameworks for alternatives are immature
- Firm-level margin trajectory negative for first time in a decade — revenue mix shifting to lower-margin structures
- Market consolidation beginning — mergers among mid-tier firms; Big Four acquiring specialized AI consultancies
- Partnership divided on talent model — aggressive restructuring vs. redesign-and-redeploy
- Specialized AI firms winning on technical depth; premium strategic firms winning on vertical expertise; firms in the "broad but shallow" middle are losing

Key Tension: The consulting business model is being restructured by external forces faster than most firms can adapt internally. The firms that survive will be those that make hard choices now about what they are (and what they are not) — but making those choices requires consensus among a partnership that is divided, anxious, and facing personal financial consequences.

Phase 4: Normalization (2027+)

[Distribute at start of Round 4]

The shakeout from the Reckoning phase has produced a new market structure that is recognizable but fundamentally different from the consulting industry of 2024. The market has bifurcated along two axes: depth of expertise and delivery model.

The premium tier — where your firm aspires to operate — consists of firms with deep vertical expertise, trusted C-suite relationships, and the ability to integrate strategic judgment with AI-augmented delivery. These firms have successfully transitioned to value-based pricing (or hybrid models with value-based components). They employ fewer junior consultants but invest heavily in their development, treating the junior role as a 2-3 year intensive apprenticeship in synthesis, client management, and industry expertise rather than an analytical production line. Partner economics have been restructured: compensation is tied more to client impact and relationship depth than to billable hour volume. Operating margins for premium firms are holding at or slightly above pre-AI levels, driven by higher revenue per engagement and lower delivery costs — but the revenue base is more concentrated in fewer, larger, more complex engagements.

The commodity tier consists of firms (and AI-native consultancies) that compete primarily on efficiency, speed, and price for well-defined analytical and implementation work. These firms have fully embraced AI-driven delivery, operate with very small teams, and price aggressively. They serve the market that previously generated the bulk of mid-tier consulting revenue: benchmarking, market analysis, process improvement, standard organizational design. Margins in this tier are thin but delivery is fast and scalable. Several of these firms are former mid-tier consultancies that chose efficiency over specialization.

The advisory opportunities that emerged during the Acceleration and Reckoning phases have matured. AI governance, responsible AI deployment, regulatory compliance, and organizational transformation advisory are now established practice areas — not novel offerings. Clients expect their consulting partners to have this capability. The firms that built these practices early have defensible positions; latecomers are struggling to differentiate. Vertical expertise — particularly in Financial Services, Healthcare, and Manufacturing — remains the strongest predictor of premium pricing and client retention.

The talent market has stabilized around a new equilibrium. Top graduates still enter consulting, but their expectations have shifted: they expect to work with AI from day one, they expect accelerated exposure to client-facing work, and they expect clear progression toward senior roles that AI cannot replicate. Firms that offer this experience are winning the talent war. Firms that still treat junior consultants as analytical labor — whether human or AI-augmented — are losing.

What Changed:

- Market bifurcated: premium strategic firms (deep vertical expertise, value-based pricing) vs. commodity delivery firms (AI-driven, price-competitive, thin margins)
- Value-based pricing is standard for complex engagements; time-based billing survives only for well-scoped, lower-complexity work
- AI governance and transformation advisory are established practice areas, not differentiators — table stakes for premium firms
- Vertical expertise in Financial Services, Healthcare, Manufacturing is the primary differentiator for premium pricing
- Junior consultant role redesigned: fewer hires, higher quality, faster exposure to complex and client-facing work
- Partner economics restructured around client impact and relationship depth, not billable hour volume
- Market consolidation complete — fewer firms, higher barriers to entry, clearer segmentation

Key Tension: The new equilibrium rewards firms that made hard, early choices about identity and

positioning. The window for strategic repositioning is closing — firms that have not yet committed to a tier (premium or commodity) face declining competitiveness in both.

Law AI Adoption Arc

***FACILITATOR NOTE:** Phase 1 (Foundation) is already included in the Law Industry Packet pre-read. It is reproduced here for facilitator reference. Distribute Phase 2 at the start of Round 2, Phase 3 at the start of Round 3, and Phase 4 at the start of Round 4.*

Phase 1: Foundation (2025 – Q1 2026)

Already distributed in pre-read packet. Reproduced here for facilitator reference.

AI deployment in the legal industry remains confined to controlled pilots within individual practice groups. The most mature deployments are contract review copilots in corporate/M&A practices, where AI handles first-pass review of standard agreements, NDAs, and mid-complexity transactional documents. Litigation practices are using legal research AI tools — primarily LexisNexis+ and Westlaw+ — for preliminary case research, statutory analysis, and citation verification, with mandatory senior associate review before any AI-assisted research enters work product. The IP practice has begun testing AI-assisted patent landscape analysis for due diligence engagements.

Bar rules on AI use in legal practice are evolving rapidly but remain inconsistent. As of early 2026, approximately 33 state bars have issued some form of guidance — ranging from mandatory disclosure of AI use in court filings to general advisories encouraging responsible use. The remaining states are silent or have guidance in draft. Firms operating across multiple jurisdictions face a patchwork compliance challenge. Malpractice insurers are beginning to scrutinize AI deployment: the major legal malpractice carriers have requested information on AI governance policies and quality assurance protocols, though premium adjustments have not yet materialized. Client demand for AI governance and regulatory compliance advisory services is strong and growing — early engagements in this area command premium rates and face no pricing pressure.

What Changed:

- Contract review copilots deployed in pilot mode (corporate/M&A practice)
- Legal research AI tools adopted by associates for preliminary work (litigation, IP)
- Bar rules actively evolving — 33 states with some form of guidance, significant gaps remain
- Malpractice insurers requesting AI governance information (no premium changes yet)
- AI governance advisory practice showing strong early client demand
- Margin impact: minimal — pilots are small scale; quality review overhead offsets productivity gains

Key Tension: Productivity gains are real but come wrapped in liability uncertainty — every efficiency improvement must be weighed against malpractice risk in a profession where errors have legal consequences.

Phase 2: Acceleration (Q2 – Q4 2026)

Distribute at start of Round 2.

AI deployment in law firms has moved beyond pilots into mainstream practice operations. Contract review copilots are now standard tools in corporate/M&A, real estate, and financial services practices at most AmLaw 50 firms. Legal research AI is embedded in day-to-day associate workflows across litigation, regulatory compliance, and IP practices. Due diligence processes that once required associate teams for weeks are being completed in days. Document drafting copilots are producing first drafts of motions, memoranda, and transactional documents that require substantive but not foundational editing.

The competitive landscape has sharpened. Legal AI platforms — Harvey.ai, LexisNexis+, Westlaw+, and several newer entrants — have been adopted by a growing number of Fortune 500 in-house legal departments. Corporate general counsel offices are handling routine contract work, regulatory monitoring, and preliminary legal research internally, reducing their reliance on outside firms for commoditized work. Alternative legal service providers (ALSPs) have scaled AI-native delivery models and are winning mandates for due diligence, document review, and compliance monitoring at rates 30-50% below traditional firm pricing. AI-native boutique firms with focused expertise in patent prosecution, regulatory filings, and data privacy are winning mandates on speed and sophistication.

Bar rule guidance has arrived in most major jurisdictions. New York, California, Illinois, Texas, and Florida have all issued substantive rules on AI use in legal practice — covering disclosure requirements, supervision standards, and confidentiality obligations. The rules are more prescriptive than many expected: mandatory disclosure of AI use in court filings is now required in over 20 jurisdictions. Malpractice insurers have begun differentiating premiums based on AI governance maturity — firms with documented protocols receive standard rates; firms without face surcharges. Client pricing pressure is accelerating: general counsel offices are conducting their own analysis of AI-assisted delivery speed and demanding rate adjustments.

What Changed:

- AI copilots scaled to mainstream practice operations across most AmLaw 50 firms
- In-house legal departments adopting AI platforms, reducing outside counsel reliance on routine work
- ALSPs and AI-native boutiques gaining material market share in commoditized segments
- Bar rule guidance issued in all major jurisdictions — disclosure requirements now widespread
- Malpractice insurers differentiating premiums based on AI governance maturity
- Pricing pressure intensifying — clients demanding rate reductions for AI-assisted work
- AI governance advisory practice growing rapidly (15-20% revenue growth in this segment)

Key Tension: The competitive pressure to deploy AI broadly is now urgent, but the regulatory and liability frameworks are imposing real constraints — firms that move fast without governance risk sanctions, and firms that move slow risk losing clients.

Phase 3: Reckoning (Q4 2026 – Q1 2027)

Distribute at start of Round 3.

The legal industry is experiencing a simultaneous regulatory and economic reckoning. Bar rule enforcement has moved from guidance to action. The Southern District sanctions opinion — finding that an attorney violated ethical obligations by submitting an AI-generated brief with fabricated citations without disclosure — has been cited in over a dozen subsequent rulings across multiple jurisdictions. State bars have initiated disciplinary proceedings against attorneys at three firms for failure to comply with AI disclosure requirements. The New York State Bar has finalized mandatory AI ethics CLE requirements effective mid-2027; California is expected to follow. The message from regulators is unambiguous: AI in legal practice is permitted, but accountability, supervision, and disclosure are non-negotiable.

The malpractice landscape has caught up. Two significant malpractice claims involving AI-assisted legal work have been filed — one involving a missed material contract provision in an M&A transaction, and another involving incorrect regulatory analysis in a financial services compliance matter. Neither has been adjudicated, but the claims have sent a clear signal through the industry. Malpractice insurers have implemented formal AI governance requirements: firms must maintain documented AI use policies, quality assurance protocols, error tracking systems, and disclosure procedures to qualify for standard premium rates. Firms without these frameworks face 10-15% premium surcharges.

The economic reckoning is equally severe. The hourly billing model is cracking. Clients have become sophisticated about AI-assisted delivery — they know that contract review that once took 40 associate hours now takes 15, and they refuse to pay traditional hourly rates for the time saved. Effective billing rates for AI-assisted work have declined 15-22% across commodity categories. Large corporate clients are pushing aggressively for fixed-fee, outcome-based, and subscription pricing models. Several clients have shifted routine legal work entirely to AI platforms or ALSPs, maintaining outside counsel only for complex, judgment-intensive matters. Associate utilization on billable work has declined — the volume of routine work available to associates is shrinking, creating career path anxiety and recruitment challenges.

What Changed:

- Bar rule enforcement actions: sanctions, disciplinary proceedings, and mandatory ethics CLE requirements
- Malpractice claims filed for AI-assisted legal work — industry signal on liability exposure
- Malpractice insurers requiring formal AI governance frameworks for standard premium rates
- Hourly billing rates for AI-assisted work declining 15-22% in commodity categories
- Clients shifting routine work to AI platforms and ALSPs; outside counsel retained for complex matters only
- Associate utilization declining — career path disruption becoming visible
- AI governance practice remains strongest growth segment but facing competition from Big Four legal practices

Key Tension: The legal profession's economic model and regulatory obligations are colliding — bar rules demand more oversight and governance (increasing cost), while the market demands lower prices and faster delivery (compressing revenue). Firms that cannot reconcile these pressures face margin collapse.

Phase 4: Normalization (2027+)

Distribute at start of Round 4.

The legal industry has reached a new structural equilibrium — different from the pre-AI status quo, but no longer chaotic. The firms that navigated the transition early are operating from positions of strength. Those that delayed are either restructuring or being absorbed.

Bar rules on AI in legal practice are now established across all 50 states. While variation persists in specific requirements, the core framework is consistent: attorneys must supervise AI-generated work product, disclose AI use where required by jurisdiction or tribunal rules, maintain competency in AI tools used in practice, and ensure client confidentiality is protected in all AI deployments. Mandatory AI ethics CLE is the norm in major legal markets. The bar's message is clear: AI is an accepted tool of legal practice, subject to the same professional responsibility obligations that govern all attorney conduct. Malpractice frameworks have matured — carriers have developed standardized AI governance assessment criteria, and firms with strong protocols benefit from favorable premiums.

The market has bifurcated. On one side: complex, high-value legal work — major litigation, transformative M&A, novel regulatory challenges, AI governance and compliance — commands premium rates and is dominated by firms with deep expertise, strong client relationships, and credible institutional knowledge. These firms have successfully repositioned their associate development models around judgment, strategy, and client counseling. On the other side: commoditized legal work — routine contracts, standard compliance filings, basic corporate maintenance — has been permanently repriced. AI platforms, ALSPs, and in-house legal departments handle the majority of this work. Full-service firms that retain commodity practices do so through AI-assisted delivery at compressed but sustainable margins, often under alternative fee arrangements.

Associate career paths have been redesigned. The traditional model — three to five years of document production before transitioning to substantive legal work — is gone. Associates now begin with AI-augmented research and analysis, moving quickly to judgment-intensive work under partner supervision. The total number of associates at major firms has declined, but per-associate revenue and compensation have increased. The profession is smaller but more specialized. Alternative fee arrangements — fixed fee, value-based, subscription, and hybrid models — account for a growing share of revenue at major firms. Hourly billing persists for complex, unpredictable matters but is no longer the default.

What Changed:

- Bar rules established in all 50 states; AI ethics CLE mandatory in major markets
- Malpractice frameworks mature; standardized AI governance assessments in place
- Market bifurcated: premium advisory/litigation firms vs. commoditized AI-assisted delivery
- Alternative fee arrangements now standard for routine and mid-complexity work
- Associate career paths redesigned around judgment and client counsel, not document production
- Total associate headcount at major firms declined; per-associate revenue increased
- AI governance practice fully established as a major, sustainable revenue stream

Key Tension: The new equilibrium is stable but structurally different — firms must accept that a meaningful share of historical revenue has been permanently repriced, and build sustainable economics around the work that remains high-value.

Manufacturing AI Adoption Arc

FACILITATOR NOTE: The Foundation phase is included in the Industry Packet distributed before the exercise. Acceleration, Reckoning, and Normalization phases are distributed as separate handouts at the start of Rounds 2, 3, and 4 respectively. Do not distribute future phases early.

Foundation (2025 - Q1 2026)

Predictive maintenance pilots in 4 manufacturing plants validated a 19% reduction in unplanned downtime and meaningful extension of equipment life over a 12-month trial. The pilots demonstrated that AI-driven condition monitoring works in industrial settings with legacy equipment, provided sufficient sensor infrastructure is in place. Quality inspection AI was piloted in 2 high-variability plants, improving defect detection rates and reducing scrap costs. These early wins built organizational confidence that manufacturing AI delivers real operational value, not just vendor promises.

However, the pilots also exposed the scale of the challenge ahead. OT/IT integration assessments at pilot sites revealed that legacy production systems (PLCs, SCADA, proprietary protocols) require substantial custom work to connect with modern cloud and analytics infrastructure. Each plant has a different equipment mix and vintage, making standardized deployment impractical. Limited supplier digitalization partnerships were established with 5 strategic vendors to test real-time data sharing, but most suppliers showed limited interest or capability. The organization has proven that manufacturing AI works in controlled settings; the question is whether it can scale across 28 plants within capital and labor constraints.

What Changed:

- Predictive maintenance validated at pilot scale (4 plants, 19% downtime reduction)
- Quality inspection AI demonstrated in 2 high-variability plants
- OT/IT integration complexity revealed as the primary scaling barrier
- Supplier digitalization initiated with 5 vendors; limited traction with broader base
- Organizational confidence built; internal skeptics partially converted

Key Tension: Pilot success creates pressure to scale fast, but infrastructure reality demands patience.

Acceleration (2026-2027, Exercise Timeline)

The organization moves from pilot validation to scaling deployment. Predictive maintenance expands to 8 highest-value plants, targeting \$45-55M in annual savings. Quality inspection AI rolls out across additional plants as the technology matures and integration playbooks from pilots are applied. OT/IT integration begins in earnest, with the first phase targeting 4-6 plants for full digital connectivity between legacy production systems and cloud analytics platforms. The capital commitment is significant: sensor infrastructure, edge computing, network upgrades, and systems integration work consume a growing share of the AI budget.

Warehouse automation enters active deployment in 2 pilot facilities, delivering material labor productivity and throughput improvements. But this is where the labor transition becomes real. Automation-driven headcount reductions in warehouse and material handling roles trigger union attention. "No-layoff" retraining agreements are negotiated at material cost (\$35-40M), preserving labor relations but reducing net automation ROI by 20-25%. Retraining programs require 12-18 months before redeployed workers reach full effectiveness. The organization begins to confront the central tension of manufacturing AI: operational gains are real, but the human transition is slower, more expensive, and more politically sensitive than technology deployment plans assumed.

What Changed:

- Predictive maintenance scales to 8 highest-value plants; ROI tracking begins at scale
- Quality inspection AI deployed across multiple plants
- OT/IT integration program launched (Phase 1: 4-6 plants)
- Warehouse automation pilots deliver results; labor displacement becomes real
- Union negotiations produce "no-layoff" retraining agreements (\$35-40M cost)
- Capital allocation tensions emerge between plant retrofits and other AI demands

Key Tension: Scaling AI is a capital allocation and labor relations challenge, not a technology problem.

Reckoning (2027-2028)

The scaling program hits structural friction. OT/IT integration reveals cost overruns and timeline delays as the heterogeneity of legacy systems across 28 plants defeats standardized approaches. Each plant requires custom integration work, and the program is consuming more of the AI capital budget than planned. Some plants that were scheduled for Phase 2 integration are pushed to Phase 3 or beyond.

Supplier digitalization dependencies become fully apparent. Production scheduling AI, designed to optimize across the full supply chain, is limited by the analog reality of most suppliers. Only a small number of strategic suppliers have come online with real-time data sharing; the rest remain on manual processes, email-based ordering, and phone-call coordination. End-to-end optimization remains aspirational.

Union contract renewal negotiations occur in this period, and labor displacement concerns peak. The combination of warehouse automation, predictive maintenance (which reduces maintenance headcount), and OT/IT modernization (which changes skill requirements) creates a cumulative labor transition story that unions frame as systematic workforce reduction. Retraining programs are underway but have not yet produced enough success stories to counter the narrative. Strike risk during peak production periods is elevated.

A geopolitical supply disruption tests the AI forecasting and scheduling systems. Models trained on historical data fail to anticipate the disruption, and production schedules built on AI recommendations require rapid manual revision. The disruption exposes model brittleness and forces a reassessment of how much operational decision-making should be delegated to AI systems.

What Changed:

- OT/IT integration cost overruns; timeline extended; some plants deferred
- Supplier digitalization stalls; end-to-end optimization remains out of reach
- Union contract renewal negotiations amid peak labor displacement concerns
- Geopolitical disruption exposes AI forecasting model brittleness
- Confidence in "AI transforms everything" narrative shaken; pragmatism returns

Key Tension: The organization discovers that scaling AI in manufacturing is constrained by ecosystem readiness and labor dynamics, not just internal capability.

Normalization (2028-2030)

Predictive maintenance becomes table-stakes across the manufacturing sector. ROI stabilizes at 12-15% improvement in maintenance cost and equipment availability, a meaningful but no longer differentiating capability. Every major manufacturer has deployed some version of condition monitoring AI; the competitive advantage has shifted from "having predictive maintenance" to "executing maintenance operations faster and more reliably."

Quality inspection AI is standard practice in high-variability manufacturing. Defect rates have improved materially across the industry, and the technology is increasingly embedded in production equipment by OEMs rather than deployed as a standalone initiative. The early-mover advantage has been absorbed into the industry baseline.

OT/IT integration continues but at a sustainable pace. Full connectivity across all 28 plants remains years away, but the highest-value plants are connected and generating real optimization value. The organization has accepted that "fully smart factories" is a decade-long journey, not a 3-year program. Autonomous warehouse systems are nascent but not yet at scale; the labor transition from the Acceleration and Reckoning phases has stabilized, with retraining programs producing redeployed workers who are now productive in new roles.

The competitive landscape has normalized. Differentiation shifts from data and models to execution speed, manufacturing resilience, and the ability to operate effectively in mixed digital+analog supply chain environments. The organizations that managed the labor transition well emerge with stronger workforce relationships and institutional knowledge intact.

What Changed:

- Predictive maintenance is industry standard; ROI stabilized at 12-15%
- Quality inspection AI embedded in standard manufacturing practice
- OT/IT integration ongoing but at sustainable pace; full connectivity still years away
- Labor transition stabilized; retraining programs producing results
- Competitive differentiation shifts from technology to execution and resilience

Key Tension: Manufacturing AI is no longer a source of advantage but a cost of doing business; the winners are those who managed the transition without destroying their workforce relationships or overcommitting capital.

Logistics AI Adoption Arc

FACILITATOR NOTE: The Foundation phase is included in the Industry Packet distributed before the exercise. Acceleration, Reckoning, and Normalization phases are distributed as separate handouts at the start of Rounds 2, 3, and 4 respectively. Do not distribute future phases early.

Foundation (2025 - Q1 2026)

Route optimization pilots in 3 logistics zones covering 800 trucks achieved 14% fuel savings and 12% on-time delivery improvement over a 12-month trial. The pilots demonstrated that AI-driven routing delivers measurable value when combined with driver acceptance: trucks that followed AI-recommended routes consistently outperformed those using traditional dispatch methods. Predictive vehicle maintenance pilots showed meaningful reductions in unplanned breakdowns and extended equipment service life. Driver assistance and safety systems (collision avoidance, fatigue monitoring) were tested across a subset of the fleet with positive safety outcomes and generally favorable driver reception.

The organization also initiated monitoring of the autonomous vehicle landscape, establishing observation partnerships with AV specialists to track technology and regulatory developments. No operational AV pilots were launched due to regulatory uncertainty, but intelligence gathering began. Limited demand forecasting AI was piloted in partnership with a small number of digitally capable shippers, testing real-time data sharing for dynamic routing. Results were promising but constrained by the small scale and the limited digital readiness of most logistics partners. The organization validated that logistics AI works at regional scale; the challenge now is full-fleet deployment across 5,000+ vehicles while managing driver acceptance, union dynamics, and the autonomous vehicle question.

What Changed:

- Route optimization validated at regional scale (3 zones, 800 trucks, 14% fuel savings)
- Predictive vehicle maintenance piloted with positive results
- Driver safety systems tested with favorable driver reception
- AV landscape monitoring initiated; no operational pilots due to regulatory uncertainty
- Demand forecasting AI tested in limited partnership; partner readiness a constraint

Key Tension: Pilot results are strong, but scaling from 800 trucks to 5,000+ is a driver acceptance and capital allocation problem, not a technology problem.

Acceleration (2026-2027, Exercise Timeline)

The organization shifts from regional pilots to fleet-wide ambition. Route optimization expands beyond the initial 3 zones, targeting full fleet coverage of 5,000+ trucks with potential annual savings of \$180-200M. Predictive vehicle maintenance scales alongside route optimization, with telematics hardware deployed to a growing portion of the fleet. The capital commitment is significant: hardware retrofits, data infrastructure, and systems integration for thousands of vehicles consume a large share of the AI budget.

Driver acceptance becomes the central operational challenge. Younger drivers (<35) embrace routing AI at high rates (~85% voluntary adoption), but older drivers (>50) resist (~55% adoption). The generational divide creates a split fleet: AI-optimized routes on some trucks, traditional dispatch on others. Union representatives raise formal concerns about "algorithmic management," requesting contractual protections around routing AI override capabilities, driver data privacy, and the right to reject AI-recommended routes. Change management programs are launched, including transparency initiatives, override capabilities, fuel efficiency bonuses, and targeted engagement with older driver cohorts.

Autonomous vehicle partnerships with AV specialists (Waymo Via, Aurora Innovation) advance to limited corridor pilot planning, but federal regulatory timelines remain uncertain (2028-2030 at earliest). The organization faces a strategic allocation question: invest management attention and capital in AV pilots (uncertain timeline, high strategic option value) or redirect all resources to proven route optimization and fleet efficiency (near-term ROI, lower risk).

What Changed:

- Route optimization expanding toward full fleet; \$180-200M savings potential at scale
- Predictive vehicle maintenance scaling with telematics hardware deployment
- Driver acceptance emerges as the critical bottleneck; sharp generational divide
- Union raises "algorithmic management" concerns; contractual negotiations begin
- AV partnership pilots in planning; regulatory timeline remains uncertain
- Capital allocation tension between AV positioning and proven fleet optimization

Key Tension: The technology scales; driver acceptance and union dynamics do not scale at the same pace.

Reckoning (2027-2028)

Last-mile profitability analysis delivers an uncomfortable truth. Detailed route-level economics confirm that AI route optimization delivers meaningful value on high-density urban and regional corridors but cannot fix the structural economics of low-density last-mile delivery. Rural routes with low stop density are unprofitable regardless of optimization. The organization must decide whether to continue subsidizing unprofitable routes (network completeness, customer retention) or rationalize the network (margin improvement, service risk). This is a business model question, not a technology question, and AI cannot resolve it.

Autonomous vehicle partnerships fail to materialize into operational deployments. Federal DOT/NHTSA regulatory timelines continue to slip, and AV specialists shift their focus to passenger vehicles and controlled environments rather than commercial long-haul trucking. Internal AV advocates who pushed for early investment face scrutiny. The capital allocated to AV partnership preparation could have been deployed to proven fleet optimization with clear near-term returns.

Driver acceptance plateaus. Change management programs have moved the older driver cohort from 55% to the mid-60s in adoption rate, but further gains prove difficult. The remaining holdouts have deep resistance rooted in professional identity and autonomy concerns that incentive programs and transparency initiatives cannot fully address. Productivity gains from route optimization stall at a level below the full-fleet potential, constrained by the adoption ceiling.

A geopolitical disruption tests demand forecasting models. AI systems trained on historical shipping patterns fail to anticipate the disruption's impact on freight volumes and routing needs. The system overcommits capacity to pre-disruption patterns and is slow to reallocate. The event exposes model brittleness and forces the organization to invest in robustness, manual override capabilities, and scenario-based forecasting.

What Changed:

- Last-mile profitability limits confirmed; business model question surfaces
- AV partnerships stall; regulatory timelines slip; AV capital questioned
- Driver acceptance plateaus in mid-60s%; full-fleet optimization ceiling hit
- Demand forecasting models fail disruption test; model brittleness exposed
- Confidence in "AI solves logistics" narrative tempered by structural realities

Key Tension: AI optimization has real but bounded value; the organization must distinguish between problems technology can solve and problems that require business model or structural changes.

Normalization (2028-2030)

Route optimization becomes ubiquitous across the logistics industry. Every major fleet operator has deployed AI-driven routing; the competitive advantage of early adoption has eroded as the technology commoditizes. Fuel efficiency and on-time delivery improvements are now industry-standard expectations, not differentiators. The 14% fuel savings that felt transformative in the Foundation phase is the new baseline.

Autonomous vehicles remain limited to specific controlled corridors or structured partnership arrangements. Long-haul autonomous trucking is still years away from commercial-scale deployment, constrained by regulatory frameworks that have not kept pace with technology development. The organizations that over-invested in AV positioning during the Acceleration phase have written down some of those investments. The organizations that maintained discipline and focused capital on proven fleet optimization are in stronger financial positions.

Driver acceptance has reached a stable equilibrium. Routing AI is standard practice for the majority of the fleet, with override capabilities and transparency features that address earlier union concerns. The generational transition is underway as older drivers retire and younger, AI-native drivers replace them. The "algorithmic management" debate has been resolved through contractual frameworks, not technology changes.

Real-time demand forecasting and dynamic routing are standard capabilities, but differentiation has shifted from algorithm quality to network density, execution speed, and supply chain resilience. The organizations that maintained strong driver relationships and workforce culture through the transition emerge with operational advantages that are difficult for competitors to replicate.

What Changed:

- Route optimization is industry standard; competitive advantage eroded
- Autonomous vehicles remain limited to specific corridors; long-haul autonomous years away
- Driver acceptance stabilized; generational transition underway
- Demand forecasting and dynamic routing are table-stakes capabilities
- Competitive differentiation shifts from technology to network, execution, and resilience

Key Tension: Logistics AI is no longer a source of advantage but a cost of doing business; the winners are those who maintained workforce relationships, managed capital discipline, and built operational resilience during the transition.

Big Tech AI Adoption Arc

FACILITATOR NOTE: Phase 1 is included in the pre-read packet. Phases 2-4 are distributed as separate handouts at the start of each corresponding round.

Phase 1: Foundation (2025 - Q1 2026)

[Already included in pre-read packet — provided here for facilitator reference]

Your AI product integration effort is in early execution. AI-powered features have launched across core products — search, email, productivity suites, cloud services, advertising — with mixed adoption. Enterprise cloud customers are piloting AI workloads on your platform, but conversion from pilot to production is slower than projected. Your data centers are operational and scaling, but GPU/TPU utilization rates are below target as inference demand has not matched infrastructure buildout. Advertising AI (generative creative, AI-powered targeting) is in limited deployment with promising early signals but no material revenue impact yet. Internally, engineering teams are stretched across too many AI product initiatives simultaneously; prioritization and focus are emerging as the primary execution risks. Capex is committed and accelerating — infrastructure spending is outpacing the revenue it was designed to support.

What Changed:

- AI features launched in core products; adoption is real but below projections
- Enterprise cloud AI pilots are active; conversion to production workloads is slow
- Infrastructure utilization rates are below target — capex outpacing demand
- Engineering bandwidth is the binding constraint; too many parallel initiatives
- Margin compression from infrastructure costs, partially offset by cloud revenue growth

Key Tension: You are building infrastructure for demand that has not yet arrived — and you cannot stop building without ceding ground to competitors who are still building.

Phase 2: Acceleration (Q2 - Q3 2026)

AI features are becoming table-stakes across every major technology product. Every productivity suite has a copilot. Every cloud platform offers managed AI services. Every search engine is integrating generative responses. The competitive landscape has shifted from "who has AI features" to "whose AI features are better, faster, and more deeply integrated." Differentiation is narrowing as all major players converge on similar capabilities.

Enterprise cloud is where the real battle is unfolding. Customers are migrating AI workloads to production at increasing velocity, but they are ruthlessly price-sensitive. Open-source models are setting the pricing floor — customers benchmark your managed AI services against the cost of running open-source alternatives on bare metal. Your cloud pricing is under pressure from both competition (cloud provider price wars) and substitution (customers bringing their own models). The companies winning cloud workloads are those offering the best managed services (fine-tuning, monitoring, compliance, SLAs) — not the cheapest raw compute.

Talent competition has peaked. AI engineering, product, and data science compensation is at historic highs. Startups, finance, and healthcare are competing for the same talent pool you need. Mid-career engineer attrition is accelerating. The engineers who remain are overloaded; quality issues and launch delays are becoming visible.

What Changed:

- AI features are universal — differentiation now depends on quality and integration depth
- Enterprise cloud AI workloads are scaling to production; pricing pressure is intense
- Open-source models have set a commodity pricing floor for AI services
- Talent competition has peaked; attrition and compensation costs are at historic highs
- Margin compression continues — cloud revenue growth partially offsets but does not cover capex burden

Key Tension: You are in a margin squeeze — cloud customers demand lower prices while your infrastructure costs remain high and talent costs are rising.

Phase 3: Reckoning (Q4 2026 - Early 2027)

AI capex ROI is under direct scrutiny from investors, boards, and analysts. The infrastructure you built in the Foundation and Acceleration phases is generating revenue — but not enough to justify the capex at current utilization rates. Investor patience is thinning. Earnings calls are dominated by questions about capex payback periods and infrastructure utilization metrics. Some competitors have already announced capex moderation programs, signaling that the "build first, monetize later" phase is ending.

Regulatory enforcement is arriving. Antitrust actions targeting platform practices, data governance, and AI service bundling are moving from investigation to enforcement. Data separation mandates, API access requirements, and platform conduct restrictions are being formalized. Compliance will consume significant engineering and legal resources for 12-18 months — resources that would otherwise be building AI products. The regulatory environment is not just a cost; it is a strategic constraint that limits how aggressively you can leverage your platform advantages.

Enterprise customers are consolidating vendors. The proliferation of AI tools and services from the Acceleration phase has created "AI tool fatigue" — customers want fewer vendors, more integrated solutions, and simpler procurement. This benefits your platform scale but only if your integrated offering is genuinely better than specialized alternatives. Customers who consolidated around your platform expect premium service; those who did not are now harder to win back.

What Changed:

- Investor scrutiny on AI capex ROI is intense; capex discipline is now mandatory
- Regulatory enforcement is active — data separation, API access, and platform conduct restrictions are being implemented
- Engineering and legal resources are diverted to compliance; AI product velocity slows
- Enterprise customers are consolidating vendors; platform integration is the competitive advantage
- Profitability pressure is acute; cost-cutting and portfolio rationalization are underway

Key Tension: You must simultaneously demonstrate capex discipline to investors, comply with regulatory mandates that constrain your competitive advantages, and defend market position against competitors who face less regulatory scrutiny.

Phase 4: Normalization (2027 Onwards)

AI is embedded in every major product and service you operate. The question is no longer whether AI is part of the product but how effectively it is integrated and how much value it delivers to customers. Differentiation has shifted from "having AI" to the quality of AI-powered workflows, the depth of ecosystem integration, and the reliability of managed AI services. The companies that balanced infrastructure investment with product execution and customer experience are the market leaders. Those that over-invested in infrastructure without commensurate demand, or under-invested in product quality, are restructuring.

Cloud infrastructure has matured into a stable, high-margin business for the survivors. Enterprise customers have completed their initial AI workload migrations and are now optimizing for cost, performance, and vendor consolidation. Pricing has stabilized — the commodity pricing floor set by open-source models is accepted, and premium pricing is defensible only for genuinely differentiated managed services. Cloud revenue growth has moderated from the explosive rates of 2025-2026 but remains healthy. The capex cycle is moderating — new infrastructure buildout is targeted and demand-driven rather than speculative.

Regulatory frameworks are established. Compliance is operationalized. The initial disruption of enforcement has passed; data governance and API access requirements are built into product architecture and business processes. Companies that proactively adapted are now advantaged — they have cleaner data practices, more transparent pricing, and stronger customer trust than those that resisted and were forced to comply reactively.

What Changed:

- AI is fully embedded; differentiation is depth of integration and quality of experience
- Cloud infrastructure is stable and profitable; capex cycle is moderating
- Regulatory compliance is operationalized; proactive adapters are advantaged
- Talent market has stabilized; AI engineering is a mature discipline with established career paths
- Margin expansion is underway as capex moderates and AI revenue mixes improve

Key Tension: The winners are defined — the question now is whether you can sustain and extend your position, or whether new disruptions (next-generation AI architectures, geopolitical shifts, new regulatory regimes) reset the competitive landscape again.

B2B/B2C SaaS AI Adoption Arc

FACILITATOR NOTE: Phase 1 is included in the pre-read packet. Phases 2-4 are distributed as separate handouts at the start of each corresponding round.

Phase 1: Foundation (2025 - Q1 2026)

[Already included in pre-read packet — provided here for facilitator reference]

Your AI feature integration is in early deployment. Copilot features and predictive analytics are live in core product modules, with adoption concentrated among your most sophisticated enterprise customers — roughly 15-20% of your installed base. Customer feedback is positive on capability but sharply negative on pricing; most customers expect AI features bundled into their existing subscription. Third-party API costs (OpenAI, Anthropic, Cohere) are running above budget, and gross margin on AI features is 10-15 points below legacy features. Three AI-native startups have emerged as serious competitive threats in your core vertical, winning proof-of-concept engagements with your mid-market customers. Your AI engineering team is deployed but stretched across too many parallel initiatives. Two senior AI engineers departed for startups in Q4 2025. The overall margin impact is low but directionally negative from infrastructure costs, while revenue impact is directionally positive from early upsell but below plan.

What Changed:

- AI features are live in core product; adoption is real but concentrated in top-tier customers
- Pricing tension is immediate — customers expect AI features bundled at no additional cost
- Gross margin on AI features is 10-15 points below legacy product margin
- AI-native startups are winning competitive engagements in the mid-market
- AI engineering team is stretched; early talent attrition signals are emerging

Key Tension: You are investing heavily in AI features that customers want but refuse to pay extra for — and the cost of delivering those features is higher than your business model assumed.

Phase 2: Acceleration (Q2 - Q3 2026)

AI features have become table-stakes in enterprise software. Every major SaaS platform now offers copilots, predictive analytics, and automation workflows. The window for AI-driven differentiation is closing — what was novel six months ago is now expected. Customer conversations have shifted from "do you have AI?" to "why isn't your AI as good as what I've seen from your competitors?" The competitive pressure is relentless and comes from both directions: Big Tech platforms are embedding AI features that overlap with your core product, and AI-native startups are offering focused, lightweight solutions that outperform you in specific use cases.

Pricing pressure has intensified dramatically. Your enterprise customers are consolidating vendor relationships and using AI feature parity across platforms as leverage in contract negotiations. The mid-market — where most of your new customer growth has historically come from — is increasingly choosing AI-native alternatives over your full platform. Your win rate in competitive deals has declined meaningfully. Sales cycles are longer. Deal sizes are flat or declining as customers resist price increases.

The cost side is not improving. Third-party API costs remain high at scale. Internal AI development is consuming more engineering capacity than planned, crowding out maintenance, security updates, and non-AI feature requests from customers. Your product organization is stretched thin, and the quality of both AI and non-AI features is suffering from the resource crunch.

What Changed:

- AI features are universal across SaaS; differentiation window is closing
- Big Tech and AI-native startups are both competing directly with your product
- Pricing pressure has intensified; customers expect AI bundled and are leveraging competition
- Win rate in competitive deals has declined; sales cycles are longer
- Engineering capacity is overextended; quality issues are emerging across the product

Key Tension: You are caught in a three-way squeeze — Big Tech from above, AI-native startups from below, and customers demanding more for less in the middle.

Phase 3: Reckoning (Q4 2026 - Early 2027)

The market is separating winners from losers. SaaS companies that executed their AI integration well — with clear pricing, high-quality features, and disciplined cost management — are gaining market share and deepening customer relationships. Those that over-promised and under-delivered, or that failed to manage the cost-margin equation, are losing customers and facing investor pressure to restructure.

Customer churn is materializing. The top-account churn signals from the Acceleration phase are converting into actual losses. Several mid-market customers who evaluated AI-native alternatives have switched. Net revenue retention has declined below the company's historical range. The customer success team is in firefighting mode, focused on retention rather than expansion. Upsell and cross-sell activities have stalled as customer conversations center on contract terms, AI feature gaps, and competitive alternatives.

Talent constraints have become the binding constraint on execution. AI engineering attrition has accelerated as startups offer equity packages and Big Tech offers compensation and scale that you cannot match. Your remaining AI engineers are overloaded, and product quality is the casualty. The engineers you need to build the features that would win back customers are the ones you are losing. Meanwhile, your customer success and professional services teams are seeing AI-driven role automation; voluntary attrition in these teams is rising as employees sense their roles are at risk. You face a dual workforce challenge that demands simultaneous investment in acquisition and transition.

What Changed:

- Market bifurcation: well-executed SaaS companies are pulling ahead; laggards are falling behind
- Customer churn is materializing; net revenue retention has declined
- Customer success is in retention mode; upsell and expansion have stalled
- AI talent attrition has accelerated; roadmap execution is constrained by people, not capital
- Services workforce is under automation pressure; dual workforce challenge is acute

Key Tension: You need to invest in AI to retain customers, but the talent to build AI features is leaving — and the customers who are leaving are taking the revenue you need to fund talent investment.

Phase 4: Normalization (2027 Onwards)

The SaaS market has restructured. The companies that survived the Reckoning phase are stronger, leaner, and more focused. AI is fully embedded in enterprise software — it is no longer a feature but an expected capability layer across every product module. Customer expectations have normalized: AI features are bundled into base subscriptions, and the pricing model has adjusted accordingly. Companies that made the bundling decision early and managed the margin transition are in the strongest position. Those that delayed are still catching up on pricing perception and customer trust.

Competitive dynamics have stabilized. The AI-native startups that gained traction in 2026 have either been acquired, scaled into serious competitors, or flamed out. The survivors are now established players with real customer bases and viable business models. Big Tech has largely retreated from competing directly with specialized SaaS platforms in vertical-specific use cases, focusing instead on horizontal AI infrastructure and productivity tools. The competitive landscape is clearer: SaaS companies compete on domain depth, workflow integration, customer relationships, and AI feature quality — not on raw AI capability.

The talent market has reached a new equilibrium. AI engineering compensation has stabilized (high, but no longer spiraling). The workforce transition in customer success and professional services is underway — roles are evolving toward higher-value consulting, solution architecture, and customer strategy, supported by AI tools that handle routine tasks. Companies that invested in retraining and internal mobility programs are seeing returns; those that relied on layoffs are rebuilding capability from scratch.

What Changed:

- AI is embedded and expected; pricing has normalized around bundled models
- Competitive landscape has stabilized; domain depth and integration are the differentiators
- AI-native startups have been absorbed, scaled, or exited; the field is clearer
- Talent market has stabilized; workforce transition from services to higher-value roles is underway
- Margin recovery is underway as AI infrastructure costs moderate and revenue mixes improve

Key Tension: Survival is secured, but sustainable growth requires continued investment in domain-specific AI capabilities, customer relationships, and workforce evolution — the companies that coast on their current position will be vulnerable to the next cycle of disruption.

05

Participant Private Cards

Private Information Cards — Distribution Guide

Overview

This folder contains **Private Information Cards** for each of the 11 industries in Project Threshold V7.4. Each industry has 3 cards (one for each of Rounds 1-3). Round 4 has no private cards.

File Format

Each file (05_x_[Industry]_PC.md) contains 3 cards separated by page breaks. Each card starts on a new page so the facilitator can print the document and hand out clean, individual pages.

Distribution Instructions

1. Print each industry's card file before the exercise
2. Cut or separate at the page break markers between cards
3. Distribute face-down at the start of each round's decision preparation phase:
 - Card 1 → Round 1
 - Card 2 → Round 2
 - Card 3 → Round 3
4. Cards are confidential to the industry — participants should not share card content with other industries unless the facilitator permits during cross-industry discussion
5. Cards accumulate — participants keep all cards and may refer to them in later rounds
6. Participants must note card influence on their Decision Worksheet (one-sentence summary)

Card Design: Shared vs. Unique

- Card 1 (Round 1): Shared intelligence across related industries (e.g., Retail and CPG both receive cards with the same underlying intelligence, but framed with industry-specific decision tensions)
- Cards 2 & 3 (Rounds 2-3): Unique per industry (industry-specific intelligence that other industries do not receive)

Files

File	Industry	Cards
05aRetail_PC.md	Retail	3
05bCPG_PC.md	CPG	3
05cHealthcareProviderPC.md	Healthcare Provider	3
05dHealthcarePayerPC.md	Healthcare Payer	3
05eFinance_PC.md	Finance	3
05fConsulting_PC.md	Consulting	3
05gLaw_PC.md	Law	3
05hManufacturing_PC.md	Manufacturing	3
05iLogistics_PC.md	Logistics	3
05jBigTechPC.md	Big Tech	3
05kSaaS_PC.md	B2B/B2C SaaS	3

Retail Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Retail participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: Personalization ROI Reality Check

Card Type: Operational Intelligence

Reveal Timing: Round 1 situation update

Classification: Restricted

Source: Internal A/B testing results from Q4 2025 e-commerce personalization pilot

The Intelligence:

Your AI personalization engine, deployed to 20% of e-commerce traffic in Q4 2025, lifted conversion rates but increased return rates. Root cause analysis reveals that AI recommendations are misaligned with actual customer expectations — the model optimizes for purchase probability, not satisfaction. Net incremental profit is roughly one-sixth of the original projection. Customers receiving AI-driven recommendations show measurably lower repeat purchase rates and lower lifetime value scores, indicating emerging negative sentiment toward the experience.

Dynamic pricing pilots in three metro test markets show a similar pattern: AI-optimized pricing extracted short-term margin improvement but triggered a spike in price-comparison app usage and negative reviews citing "unfair pricing." Social listening detects early signals of a consumer narrative forming around algorithmic price discrimination by major retailers.

Decision Tension:

Do you accelerate consumer-facing AI (personalization, dynamic pricing, recommendation algorithms) to maximize short-term margin expansion, knowing the early data shows customer satisfaction erosion and emerging backlash? Or do you pivot investment toward proven operational AI (inventory optimization, demand forecasting, supply chain automation) and sacrifice the consumer-facing margin upside?

Questions to Consider:

- What is the monetizable value of brand trust in lifetime customer value? Is it higher than the margin gains from aggressive personalization?
 - How would you redesign the personalization engine to prioritize customer satisfaction over conversion lift? What margin are you willing to trade for better recommendations?
 - Your dynamic pricing pilots show short-term margin gain but long-term trust erosion. At what threshold of negative sentiment do you implement price consistency rules?
 - Amazon already does aggressive personalization and dynamic pricing. Can you afford not to match them, or is "trustworthy pricing" a viable competitive differentiation?
-

Card 2 — Round 2

Title: Vendor DTC Launch Intelligence & Channel Threat Assessment

Card Type: Competitive Intelligence

Reveal Timing: Round 2 situation update

Classification: Confidential

Source: Confidential vendor conversations and competitive intelligence

The Intelligence:

Multiple top CPG vendor partners have informed you (confidentially) that they are planning direct-to-consumer launches in Q3 2026, using AI-driven marketing, demand sensing, and personalized pricing. If successful, these DTC channels could disintermediate you on significant SKU volumes — representing material at-risk revenue across your highest-margin product categories. Vendors are motivated in part by your own push for data transparency and AI-driven merchandising; they perceive your private-label expansion and algorithmic shelf-space allocation as commoditizing their brands.

In parallel, two of your largest vendors are demanding real-time access to your AI demand forecasts, customer purchase-intent data, and category-level consumer insights. They frame this as a "partnership requirement" for continued promotional support and innovation investment. Confidential intelligence suggests that at least one competitor retailer has already shared comparable data with these same vendors, creating competitive asymmetry.

Your analysis indicates that sharing this data would improve joint demand forecasting accuracy by 15-20% but would also give vendors the precise consumer intelligence they need to optimize their DTC channels against you.

Decision Tension:

Do you share proprietary demand forecasts and consumer insights with your largest vendors to preserve relationships, maintain shelf economics, and prevent accelerated DTC launches — knowing it gives them the data to bypass you more effectively? Or do you protect your data advantage, accelerate private-label as a defensive hedge, and accept the risk of vendor retaliation through DTC and reduced promotional support?

Questions to Consider:

- What data is truly proprietary and defensible versus what is table-stakes to share? Can you segment by strategic value — share aggregated insights, delay tactical data, exclude high-margin categories?
 - If vendors launch DTC successfully, what is your defensive strategy? Accelerated private-label? Exclusive product arrangements? Logistics-as-a-service to capture vendor DTC fulfillment?
 - What contractual terms (exclusivity, non-competition, revenue share) would make data sharing acceptable while limiting competitive exposure?
 - How quickly do you need to decide? The Q3 2026 vendor launch timeline creates urgency — waiting means facing DTC competition without having shaped the terms.
-

Card 3 — Round 3

Title: Consumer Sentiment Shift & Regulatory Pressure on Retail

Card Type: Market Intelligence

Reveal Timing: Round 3 situation update (post-inject)

Classification: Restricted

Source: Q2 2026 consumer perception study and regulatory intelligence

The Intelligence:

Your Q2 2026 consumer perception study reveals a concerning sentiment shift. Among consumers aware of AI-driven shopping experiences, sentiment has declined meaningfully from prior measurement periods. Specific concerns cited by respondents: "pricing feels unfair because I know I'm tracked," "recommendations feel invasive," "I don't want a robot deciding what I should buy." Sentiment is notably worse among your highest-value loyalty program members, who perceive AI as enabling price discrimination against loyal customers.

Private-label penetration in AI-curated product categories is rising faster than in non-curated categories. Your data team's analysis suggests customers perceive AI recommendations as biased toward higher-margin products and private-label, undermining trust in the recommendation engine as a consumer tool versus a margin extraction tool.

Regulatory pressure is accelerating. State attorneys general have issued information requests to major retailers regarding algorithmic pricing transparency, personalization data practices, and labor displacement from automation. The FTC is drafting guidance on "deceptive AI" in retail, specifically targeting dynamic pricing, non-disclosed AI recommendations, and AI-generated product reviews. Compliance and disclosure requirements will increase materially in 2027.

Decision Tension:

Do you continue scaling AI-driven personalization and dynamic pricing to maintain margin trajectory, accepting the reputational erosion and regulatory risk as manageable costs of doing business? Or do you recalibrate toward transparency, opt-in personalization, and price consistency — sacrificing conversion lift and margin expansion in exchange for trust recovery and regulatory positioning?

Questions to Consider:

- What specific transparency and fairness commitments would credibly restore customer trust? Opt-in personalization? Recommendation explanations? Posted price-consistency guarantees?
- Your highest-value customers are the most skeptical. How do you design a loyalty program experience that leverages AI for convenience without triggering surveillance concerns?
- Can you get ahead of the FTC guidance by proactively implementing disclosure and fairness guardrails? Is there a first-mover advantage in "trustworthy AI retailing"?
- If you pull back on personalization intensity, what is the margin impact over 6-12 months? Can operational AI (inventory, supply chain) compensate for lost consumer-facing margin?
- How do you communicate this shift to your workforce, especially store associates who have been told AI is their co-pilot? Messaging consistency matters for morale and retention.

CPG Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the CPG participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: AI Content & R&D Acceleration — ROI Reality Check

Card Type: Operational Intelligence

Reveal Timing: Round 1 situation update

Classification: Restricted

Source: Internal campaign performance analysis and R&D pilot results

The Intelligence:

Your AI-generated marketing campaign pilot (household products division) achieved significant cost efficiency — substantially lower production cost than agency-created work — and matched conversion performance in direct A/B testing. The business case for scaling AI content across your \$2.1B marketing spend is compelling on pure economics. However, the pilot also generated a meaningful brand safety incident: AI-generated human faces in lifestyle imagery triggered significant negative social sentiment (#FakeAds, #CPGDontLieToUs). PR contained the episode, but the signal is clear — consumers are developing sensitivity to AI-generated brand content, and your iconic brands are disproportionately exposed to authenticity scrutiny.

In parallel, the R&D acceleration pilot has delivered strong results. AI-assisted formulation and consumer preference modeling reduced a product development cycle by 6 months and generated \$35M in incremental revenue from an accelerated launch. The validated ROI is material and the reputational risk is low (consumers don't see the R&D process). The question is how aggressively to scale across your 45 R&D centers and 35 brand portfolios.

Decision Tension:

Do you aggressively scale AI-generated marketing content across your brand portfolio to capture the cost efficiency, knowing that brand safety incidents are likely and consumer sensitivity to synthetic content is rising? Or do you constrain AI content to low-risk applications (product descriptions, technical specs, data visualization) and protect brand authenticity as your competitive moat — accepting higher marketing costs? Separately, how aggressively do you scale the R&D acceleration pilot given the validated ROI?

Questions to Consider:

- Which content categories are safe to automate (product descriptions, technical specifications, data visualizations) versus which require human creation (brand storytelling, lifestyle imagery, testimonials, any content featuring people)?
- Is the \$35M incremental revenue from R&D acceleration strong enough to justify scaling across all 45 R&D centers? What is the organizational readiness to absorb AI tools across diverse brand teams?

- How do you handle disclosure? Should you proactively label AI-generated content, or does disclosure itself undermine brand perception? What is the regulatory trajectory on this?
 - Your competitors are scaling AI content faster. If you constrain AI marketing to protect authenticity, do you lose cost competitiveness — or does "human-made brand storytelling" become a premium differentiator?
-

Card 2 — Round 2

Title: DTC Launch Decision & Retailer Power Dynamics

Card Type: Competitive Intelligence

Reveal Timing: Round 2 situation update

Classification: Confidential

Source: Retailer partnership negotiations and competitive intelligence

The Intelligence:

Your top retail partners are demanding real-time access to your AI demand forecasts, supply chain visibility, and consumer insights to improve their own AI models. They frame this as a "partnership requirement" and are subtly threatening shelf-space reduction and accelerated private-label development in categories where you withhold data. Confidential intelligence confirms that at least two competitor CPG companies have already shared comparable data with these same retailers, creating competitive asymmetry — retailers are using shared data to optimize private-label products that directly compete with your brands.

In parallel, your DTC readiness assessment is complete. AI-driven personalization, dynamic pricing, and customer service automation make the unit economics of direct brand channels viable for 8 of your 35 brands (those with strong direct consumer affinity and premium positioning). DTC margins would be materially higher than retail channel margins. However, your modeling projects that an aggressive DTC launch would trigger retailer retaliation: estimated shelf-space reduction across your portfolio and potential delisting of promotional support, which would reduce retail channel revenue significantly in the near term.

The timing pressure is real. Competitors are moving toward DTC. Retailers are building private-label alternatives. Every quarter you delay DTC launch, competitors establish direct consumer relationships that become harder to displace.

Decision Tension:

Do you share proprietary demand data and consumer insights with retail partners to preserve shelf space, maintain promotional support, and delay retailer retaliation — knowing it accelerates their private-label capabilities against you? Or do you protect your data, launch DTC on select brands, and accept the near-term revenue hit from retailer retaliation in exchange for building direct consumer relationships and higher-margin channels?

Questions to Consider:

- Can you pursue a middle path — share aggregated, delayed data with retailers while launching DTC on brands where retail overlap is minimal? Or will retailers see through the segmentation?
- What is the breakeven timeline for DTC revenue to offset retailer retaliation? Can your balance sheet absorb 12-18 months of revenue pressure?
- Which 8 brands should lead the DTC launch? Should you prioritize brands with the strongest direct consumer affinity or brands where retailer private-label pressure is already eroding your position?
- How do you message the DTC launch to retailers? "Brand building" and "consumer engagement" framing — will retailers accept it, or does any DTC signal trigger immediate retaliation?
- What logistics and fulfillment infrastructure do you need? Can you partner with a retailer's fulfillment network for DTC (turning a rival into a logistics partner)?

Card 3 — Round 3

Title: Consumer Authenticity Backlash & Regulatory Pressure on CPG

Card Type: Market Intelligence

Reveal Timing: Round 3 situation update (post-inject)

Classification: Restricted

Source: Q2 2026 brand perception tracking and regulatory intelligence

The Intelligence:

Your Q2 2026 brand health tracking reveals a concerning shift in consumer sentiment. Awareness that CPG companies use AI-generated marketing content has increased substantially over the past year, and among aware consumers, brand trust scores have declined meaningfully. The damage is concentrated in categories where authenticity matters most — personal care, baby products, and premium food brands — where consumers explicitly associate AI-generated content with "cutting corners" and "not caring about quality."

Two specific developments are escalating the risk. First, a viral social media investigation ("Are Your Favorite Brands Lying to You?") has catalogued AI-generated content across major CPG brands, including yours. The investigation found AI-generated testimonials, synthetic lifestyle imagery, and AI-written product claims — some of which make factual assertions about product efficacy that cannot be substantiated. Second, the FTC is drafting guidance on "deceptive AI in consumer products," specifically targeting AI-generated endorsements, synthetic testimonials, non-disclosed AI content in packaging and advertising, and AI-derived product claims. State attorneys general are also issuing inquiries to CPG companies about content generation practices.

Your competitor intelligence shows the market is bifurcating: some CPG companies are doubling down on AI content to maintain cost advantage, while others are pivoting to "authentically human" brand positioning. Neither strategy has a proven track record yet.

Decision Tension:

Do you continue scaling AI-generated marketing content to maintain cost efficiency, accepting brand trust erosion and regulatory risk as manageable costs? Or do you pivot to a "brand authenticity" positioning — human-created content, transparent disclosure, real consumer stories — sacrificing cost efficiency but potentially differentiating against competitors who are perceived as inauthentic?

Questions to Consider:

- What is the cost of a full pivot to human-created content for your highest-trust brands (personal care, baby, premium food) while maintaining AI efficiency for lower-stakes categories (household cleaning, commodity products)?
- Can you get ahead of FTC guidance by proactively implementing content disclosure and eliminating AI-generated testimonials? Is there a first-mover advantage in "authentically human" CPG branding?
- How do you handle the viral investigation? Proactive disclosure and policy change, or wait for the news cycle to pass? What is the brand equity cost of each approach?
- Your R&D acceleration investments (back-office AI) are not affected by the consumer authenticity backlash. Should you double down on R&D AI while pulling back on consumer-facing AI content — making the R&D advantage your competitive moat instead of marketing efficiency?
- What are the regulatory compliance priorities for 2027? Can you turn compliance into competitive advantage by being the first CPG company with transparent AI content policies?

Healthcare Provider Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Healthcare Provider participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: Clinical Diagnostic AI Workflow Friction + Prior Authorization Liability Exposure

Card Type: Operational Intelligence

Classification: Operational Intelligence

Source: Clinical Operations, Radiology Department, Risk Management

Reveal Timing: Round 1 Decision Preparation

The Intelligence:

Your diagnostic radiology AI pilot shows meaningful improvement in diagnostic accuracy compared to radiologist baseline. However, documentation analysis reveals radiologists are spending significantly more time validating AI-generated reports than reviewing standard imaging studies. Radiologists report high cognitive load despite high accuracy — the AI system improves diagnostics but worsens radiologist workflow. Early warning: accuracy does not equal clinical adoption.

Separately, your prior authorization prediction tool (designed to help physicians anticipate insurer denials) surfaced an internal finding: the tool's underlying model identified cases where the affiliated health plan's prior auth algorithm recommended denial of clinically appropriate care. The cases are rare but material. The clinical risk is real — patient harm from inappropriate denial creates liability exposure for you as the treating provider (physician practice liability, malpractice risk) and for the payer (coverage appropriateness liability). Implementing mandatory human review for high-risk denial predictions would add physician burden but near-eliminate false negatives.

Decision Tension:

Your diagnostic AI improves accuracy but increases physician validation burden. Do you pause the radiology pilot and redesign to minimize radiologist workflow friction? Push through adoption friction and accept higher physician burden during transition? Or pivot to different clinical domains (pathology, EKG interpretation) with different workflow profiles?

On prior auth, your prediction tool has revealed liability risk in the coverage determination process. Do you escalate to the payer side and demand algorithm changes? Implement your own human review overlay for flagged cases (adding physician burden)? Or accept the risk and focus your limited capital on clinical AI instead?

Questions to Consider:

- What is the minimum workflow burden reduction required for physician adoption of diagnostic AI? What design changes reduce validation overhead without sacrificing accuracy?
 - For the prior auth finding, what is your actual liability exposure if a patient is harmed by an inappropriate denial you identified but did not escalate?
 - Can you design tiered review for prior auth predictions (high-confidence predictions require no action; low-confidence predictions trigger physician review) to manage burden?
 - How do you operationalize continuous monitoring of diagnostic AI performance and physician workflow impact? What metrics trigger escalation or redesign?
-

Card 2 — Round 2

Title: Physician Advisory Board Resistance to Clinical AI Deployment

Card Type: Risk Event Intelligence

Classification: Risk Event Intelligence

Source: Physician Advisory Board, Medical Staff Leadership, Chief Medical Officer

Reveal Timing: Round 2 Decision Preparation

The Intelligence:

Your physician advisory board is expressing significant and organized resistance to clinical AI deployment. The resistance is not fringe — it represents senior physicians and academic faculty who carry substantial influence over clinical practice standards at your facilities. Key concerns:

1. **Autonomy and clinical judgment:** Diagnostic AI is perceived as reducing physician autonomy and substituting algorithmic recommendations for clinical expertise. Senior physicians view this as fundamentally threatening to the physician-patient relationship and professional identity.
2. **Liability exposure:** Physicians are personally liable for clinical decisions. If an AI-recommended diagnosis contributes to patient harm, liability cascades to the treating physician and their malpractice insurer. Physicians are asking: who bears liability when they follow AI recommendations that turn out to be wrong? Their malpractice insurers are asking the same question.
3. **Administrative burden perception:** Prior authorization prediction tools and documentation AI are perceived by some physicians as "more technology to manage" rather than burden reduction. Physician leaders report that early-career physicians are more receptive, but senior physicians and department heads are actively discouraging adoption.

The advisory board has formally requested a 6-month moratorium on new clinical AI deployments pending development of a physician-led governance framework, liability clarity, and workflow burden assessment.

Decision Tension:

Do you accept the moratorium and invest 6 months in physician-led governance development — potentially losing competitive ground but building sustainable adoption? Do you reject the moratorium and accelerate deployment, risking physician attrition and organized resistance? Or do you negotiate a middle path — continue existing pilots under enhanced physician oversight while developing the governance framework in parallel?

Questions to Consider:

- What would credibly address physician concerns about autonomy and liability? Transparent algorithm design? Explicit liability frameworks? Limited scope of deployment to advisory/support roles only?

- How much physician consensus is required for sustainable clinical AI adoption? Can you deploy with partial buy-in, or does resistance from senior physicians poison adoption across the organization?
 - What role should physician leadership play in AI governance? Advisory only, or veto authority over clinical AI deployments?
 - How do you sequence physician engagement — do you invest in early advisory input now, or push through and address concerns later? What are the costs of each approach?
-

Card 3 — Round 3

Title: FDA Clinical AI Validation Requirements + CMS Conditions of Participation

Card Type: Regulatory Development

Classification: Regulatory Intelligence

Source: Regulatory Affairs, FDA Pre-Submission Contacts, CMS Policy Team

Reveal Timing: Round 3 Decision Preparation

The Intelligence:

Through FDA and CMS contacts, you have received advance notice of upcoming regulatory guidance that will materially affect your clinical AI strategy:

FDA (Q2 2026): New clinical validation requirements for diagnostic AI tools are imminent. External validation will be required (internal validation insufficient). Bias audits across demographic groups will be mandatory. Real-time performance monitoring with reporting obligations will be required post-deployment. Implementation estimate: 12-18 months of pre-market work per diagnostic AI tool. Estimated compliance cost: \$5-10M per tool.

CMS (Q2 2026 draft, finalizing 2027): New Conditions of Participation for AI-assisted diagnostics will require demonstrated clinical benefit (not just accuracy), governance protocols with physician oversight, and ongoing algorithm performance monitoring with CMS reporting. Hospitals that deploy clinical AI without meeting these conditions risk losing CMS certification — which would be catastrophic for Medicare/Medicaid reimbursement.

The regulatory clarity you have been waiting for is arriving, but it comes with substantial compliance burden. Early movers who filed pre-submission meetings and began validation work 12-18 months ago will have a significant head start. Those who waited for final guidance face 18-24 month delays before any clinical AI can reach deployment.

Decision Tension:

Do you accelerate clinical AI investment now to position ahead of regulatory requirements — accepting the \$50M+ multi-year compliance cost for your clinical AI portfolio? Or do you slow clinical AI and redirect capital to administrative AI (coding, documentation, prior auth) where regulatory burden is lower and ROI is faster? Do you file with FDA immediately (accepting risk that guidance may shift) or wait for final guidance (slower but more certain)?

Questions to Consider:

- If you invested in early FDA pre-submission engagement and academic validation partnerships in earlier rounds, you are ahead. If not, you face 18-24 month delays. How does your prior positioning affect this decision?
- What is the competitive value of being an early clinical AI deployer vs. a fast follower? In healthcare, is first-mover advantage real or illusory?
- How do you manage investor and board expectations for AI ROI given longer regulatory timelines for clinical AI?

- Do you prioritize low-risk administrative AI (medical coding, documentation) for near-term ROI, or invest in clinical diagnostic AI for long-term differentiation?
- How do you sequence regulatory filings across your clinical AI portfolio? Which tools file first?

Healthcare Payer Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Healthcare Payer participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: Prior Authorization False Negative Liability + Clinical Diagnostic AI Friction

Card Type: Operational Intelligence

Classification: Operational Intelligence

Source: Claims Audit & Quality Team, Clinical Review, Risk Management

Reveal Timing: Round 1 Decision Preparation

The Intelligence:

Your prior authorization AI achieved strong operational results — \$22M in cost savings and meaningful reduction in processing time. However, internal audit has identified cases where the algorithm recommended denial of clinically appropriate care. The cases are rare but material. Clinical review confirms: these are genuine false negatives where the AI's cost-optimization logic overrode clinical appropriateness signals in the training data.

The clinical risk is real. Patient harm from an inappropriate denial creates liability exposure for you as the coverage decision-maker (coverage appropriateness liability, regulatory enforcement risk) and for the treating provider (physician practice liability). Implementing mandatory human review for high-risk denials would reduce net savings by 20-30% but near-eliminate false negatives. A tiered approach (human review only for high-cost or high-risk denial categories) could preserve 80-85% of savings while catching the most dangerous false negatives.

Separately, intelligence from your provider network indicates that clinical diagnostic AI pilots are showing accuracy improvements but creating physician workflow friction. This matters to you: if providers' clinical AI deployments stall due to physician resistance, the downstream clinical data quality improvements you need for your own risk adjustment and member health management models will be delayed.

Decision Tension:

Do you accept the \$22M savings and manage the litigation and regulatory risk of occasional inappropriate denials? Implement mandatory human review across all denials and reduce savings to \$16-18M? Or design a tiered review system — human review for high-risk denial categories, algorithmic decisioning for low-risk approvals — to balance savings and safety?

The provider-side diagnostic AI friction is outside your direct control, but it affects your data quality roadmap. Do you invest in incentivizing provider AI adoption (data sharing agreements, joint governance) or plan around continued data quality limitations?

Questions to Consider:

- What is your actual liability exposure per inappropriate denial? Does your corporate liability insurance cover AI-driven coverage denials that lead to patient harm?
 - Can you implement tiered human review (high-confidence AI decisions pass through; medium-confidence sampled; low-confidence mandatory review) to reduce costs while managing liability?
 - How do you operationalize continuous monitoring of prior auth accuracy? What false negative rate triggers escalation or algorithm redesign?
 - If regulators begin auditing your prior auth algorithms for inappropriate denials, is your current documentation and governance sufficient to withstand scrutiny?
-

Card 2 — Round 2

Title: Synthetic Fraud Detection Arms Race

Card Type: Risk Event Intelligence

Classification: Risk Event Intelligence

Source: Fraud Detection Team, Special Investigations Unit, Cybersecurity

Reveal Timing: Round 2 Decision Preparation

The Intelligence:

Your transaction monitoring system has flagged a meaningful increase in synthetic claim submissions over the past quarter. AI-generated provider credentials, member data, and claim narratives are becoming increasingly difficult to distinguish from legitimate claims. Your fraud detection model performance has degraded relative to prior periods — the false negative rate (fraudulent claims passing undetected) has increased, while the false positive rate (legitimate claims flagged as fraud) has also risen, creating provider relations friction.

You are losing the AI arms race. Offenders are using generative AI to forge claims faster and more convincingly than your detection models can adapt. Competitor intelligence suggests that at least two rival insurers have deployed more advanced synthetic claim detection capabilities, putting you at a competitive disadvantage — sophisticated fraud rings will preferentially target insurers with weaker detection.

Developing next-generation synthetic fraud detection requires material investment and a 12-18 month development timeline. Options include: (1) build internally with your existing data science team (slower, retains IP), (2) acquire a specialized fraud-detection startup (faster, expensive, integration risk), or (3) partner with a cybersecurity vendor (moderate speed, shared IP, ongoing cost).

Decision Tension:

Do you accelerate investment in next-generation fraud detection now, accepting the material cost and 12-18 month timeline? Do you acquire a specialized startup for speed at the cost of integration risk and premium pricing? Or do you accept higher fraud costs as an ongoing cost of business and reallocate capital to other priorities?

Questions to Consider:

- What is the ROI threshold for next-generation fraud detection investment? At what fraud loss rate does the arms race become existential rather than incremental?
- Your fraud detection model degradation is accelerating. What is your exposure if you delay investment by 6-12 months? Can you quantify the cost of inaction?
- Build vs. buy vs. partner: what are the trade-offs for each approach? Which gives you the fastest path to competitive parity?
- How does synthetic fraud interact with your claims data quality constraints? If fraudulent claims are contaminating your training data, does your model degradation compound over time?

- Can you coordinate with industry consortia or law enforcement to share synthetic fraud intelligence, or does competitive dynamics prevent it?
-

Card 3 — Round 3

Title: Prior Authorization Enforcement Cascade + Disparate Impact Exposure

Card Type: Regulatory Development

Classification: Regulatory Intelligence

Source: Regulatory Affairs, State Insurance Board Contacts, CMS Policy Intelligence

Reveal Timing: Round 3 Decision Preparation

The Intelligence:

CMS and state insurance boards are preparing to announce enforcement actions against major insurers for inappropriate prior authorization denials. The target companies' prior auth algorithms exhibited statistically significant disparate impact — higher denial rates for certain member demographics (age, race, geographic location, socioeconomic proxies). The enforcement actions will include material financial penalties, mandatory algorithm audits, and required member remediation programs.

This enforcement action will trigger industry-wide prior auth audits within 12-18 months. Every major insurer will face regulatory examination of their prior authorization algorithms for fairness, accuracy, and disparate impact. Your internal exposure analysis suggests vulnerabilities in your current prior auth algorithm — specifically, denial rate disparities correlated with member zip code (a known socioeconomic proxy) and age cohort. These disparities may not be intentional, but the regulatory standard is disparate impact, not intent.

Proactive remediation requires material investment: comprehensive bias auditing of all coverage determination algorithms, algorithm redesign where disparities are found, member remediation for historically affected populations, and enhanced governance and reporting infrastructure. Reactive posture (waiting for your audit) risks larger penalties, mandatory corrective action under regulatory supervision, and reputational damage.

Decision Tension:

Do you immediately launch proactive bias auditing and algorithm remediation across all coverage determination AI — accepting the material cost and operational disruption — to position ahead of the regulatory wave? Or do you monitor the enforcement actions against competitors and react when your audit arrives, preserving capital but accepting higher regulatory and reputational risk?

Questions to Consider:

- Is early bias remediation a strategic advantage (early mover credibility with regulators, reduced penalty exposure) or just an expense with no competitive benefit?
- What is your actual disparate impact exposure? Can you quantify the scope of affected member populations and potential remediation costs?
- How do you communicate proactive remediation to regulators, members, and investors? Does disclosure create legal exposure, or does it demonstrate good faith?
- If your prior auth AI requires fundamental redesign (not just parameter tuning), what is the impact on your \$22M savings and operational efficiency gains?

- How do you balance prior auth compliance costs against fraud detection investment needs? Both require material capital in the same timeframe. What gets funded first?

Finance Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Finance participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: Underwriting Explainability Audit — Proxy Variable Risk & Regulatory Exposure

Card Type: Regulatory Development

Classification: Regulatory Intelligence

Source: Internal Model Risk Management Team; Regulatory Affairs

Reveal Timing: Round 1, start of Decision Preparation phase

The Intelligence:

Your AI-driven insurance underwriting model outperforms traditional actuarial methods and is your strongest near-term business case for AI deployment. However, an internal explainability audit has surfaced material concerns. The model relies on complex ensemble methods that regulators may challenge under current examination standards. More critically, the model uses proxy variables — combinations of geographic, behavioral, and demographic features — that correlate with protected characteristics. Your state insurance regulatory contacts have informally flagged that models using similar proxy structures at competitor institutions are drawing scrutiny.

Your regulatory affairs team assesses material risk of enforcement action if you deploy this model in its current form without further refinement. The accuracy advantage of the ensemble model over an interpretable alternative is 6-8 percentage points — meaningful in underwriting economics. An interpretable model would be regulator-friendly but would sacrifice competitive edge.

Your Model Risk Management team has proposed a dual-architecture approach: deploy the interpretable model for regulatory-facing underwriting decisions, and use the high-accuracy ensemble model internally for risk management and portfolio analytics where explainability requirements are lower. This is technically feasible but operationally complex — maintaining two parallel model stacks increases cost, governance burden, and the risk of model drift between the two.

Decision Tension:

Regulatory compliance versus competitive accuracy. Your interpretable models sacrifice meaningful accuracy. Do you deploy the high-accuracy ensemble for internal risk management only and present the interpretable model to regulators? How do you manage the dual-model complexity, cost, and governance risk? Or do you accept the accuracy penalty across the board and deploy a single interpretable model, simplifying governance at the cost of competitive disadvantage?

Questions to Consider:

- What is your regulatory tolerance? Can you defend a dual-tier underwriting architecture (high-accuracy internal, compliant-but-less-accurate external) if regulators challenge why you use different models for different purposes?
- What is the business cost of 6-8 points of accuracy loss in underwriting? How does that translate to loss ratios, approval rates, and ROE impact?
- How do you sequence the regulatory engagement? Do you proactively brief OCC and state insurance regulators on your model governance framework, or wait for them to examine?
- If a competitor deploys a high-accuracy model and faces no immediate regulatory action, does that change your calculus?

***SHARED INTELLIGENCE:** Consulting and Law industry participants have received related intelligence about AI governance and regulatory compliance challenges across professional services. Their perspectives on regulatory strategy and governance frameworks may surface during cross-industry discussion.*

Card 2 — Round 2

Title: Synthetic Fraud Detection Arms Race — Model Degradation & Competitive Intelligence

Card Type: Risk Event Intelligence

Classification: Risk Event Intelligence / Competitive Intelligence

Source: Fraud Detection Team; Cybersecurity Operations; Competitive Intelligence Unit

Reveal Timing: Round 2, start of Decision Preparation phase

The Intelligence:

Your transaction monitoring system has flagged a meaningful increase in synthetic identity applications over the past quarter. AI-generated work histories, income documentation, and identity papers are becoming increasingly difficult to distinguish from legitimate documents. Your fraud detection model — which was best-in-class 18 months ago — has experienced measurable performance degradation. False negative rates have increased as adversarial AI generates more sophisticated forgeries.

You are losing the AI arms race. Fraudsters are using generative AI to create synthetic identities and forge documentation faster than your detection models can adapt. The estimated annual fraud loss exposure if current trends continue is material to your P&L. Your fraud team estimates that developing next-generation synthetic identity detection capability requires significant investment and a 12-18 month development timeline.

Competitive intelligence adds urgency. Two of your top-5 banking peers have deployed more advanced synthetic identity detection systems, reportedly achieving meaningfully better detection rates. One acquired a specialized fraud-detection startup; the other partnered with a cybersecurity vendor. Your current detection rate is falling behind competitive benchmarks, and the gap is widening each quarter as adversarial AI improves.

Decision Tension:

Investment urgency versus capital discipline. Developing next-generation fraud detection requires material capital commitment on a 12-18 month timeline with uncertain ROI. Do you accelerate investment now to close the competitive gap, or accept elevated fraud losses as an ongoing cost of business while deploying capital elsewhere? Build versus buy: do you invest in internal R&D, acquire a specialized startup, or partner with a cybersecurity vendor? Each path has different cost, timeline, and capability trade-offs.

Questions to Consider:

- What is the ROI threshold for next-generation fraud detection investment? At what false negative rate does synthetic fraud become an existential cost rather than a manageable loss?
- Build, buy, or partner? Internal R&D preserves IP control but is slowest. Acquisition is fastest but most expensive and creates integration risk. Partnership is flexible but creates vendor dependency. What fits your risk profile?
- How do you prioritize fraud detection investment against other AI initiatives (underwriting, wealth management, compliance infrastructure) competing for the same capital and talent?

- If competitors have already deployed superior detection, what is the cost of waiting another 6-12 months? Does competitive gap create customer attrition risk or regulatory exposure?

Card 3 — Round 3

Title: CFPB Fair Lending Enforcement Action — Audit Cascade & Model Vulnerability Exposure

Card Type: Regulatory Development

Classification: Regulatory Intelligence

Source: Regulatory Affairs; Board Contacts (OCC, Federal Reserve); General Counsel

Reveal Timing: Round 3, start of Decision Preparation phase

The Intelligence:

The CFPB is preparing to announce an enforcement action against a major fintech competitor for AI-driven lending discrimination. The company's underwriting model exhibited statistically significant disparate impact against protected classes in mortgage and auto lending. The enforcement action will include substantial penalties, mandatory model remediation, and a multi-year compliance monitoring agreement.

This enforcement action will trigger fair lending audits across all major banks within 12-18 months. Your regulatory contacts confirm that OCC and Federal Reserve examiners are already developing enhanced AI model examination procedures. Every institution deploying AI in lending or underwriting decisions should expect heightened scrutiny.

Your internal exposure analysis — conducted by your Model Risk Management team in response to Card 1 intelligence — has identified vulnerabilities in your underwriting systems. The proxy variable concerns flagged in Round 1 remain partially unresolved. Your remediation options and their estimated costs:

- Full model refactoring (replace ensemble with fully interpretable architecture): Significant investment; 12-18 month timeline; resolves regulatory concern but sacrifices accuracy.
 - Proxy variable remediation (remove or constrain problematic features): Moderate investment; 6-9 month timeline; reduces but does not eliminate disparate impact risk.
 - Customer remediation program (proactive outreach to potentially affected borrowers): Material cost; demonstrates good faith; reduces penalty severity if enforcement action occurs.
 - Enhanced compliance infrastructure (independent model audit, ongoing bias monitoring, regulatory reporting): Substantial annual cost; necessary regardless of which remediation path you choose.
-

Decision Tension:

Proactive remediation versus reactive monitoring. Do you immediately launch comprehensive model remediation and customer outreach — expensive but demonstrating regulatory good faith and potentially reducing enforcement severity? Or do you monitor the enforcement action against your competitor, learn from their experience, and remediate only what examiners specifically require — cheaper in the near term but risking harsher treatment if regulators view your response as inadequate? The first-mover question cuts both ways: early remediation is expensive and may fix problems examiners never would have flagged, but late remediation after an examination finding carries consent order risk and reputational damage.

Questions to Consider:

- What is your proactive versus reactive stance on fair lending compliance? Is early, comprehensive remediation a strategic advantage (demonstrating industry leadership and regulatory good faith) or just premature expense?
- How do you sequence remediation? Full model refactoring takes 12-18 months. Proxy variable remediation is faster but incomplete. Do you do both? In what order?
- What is the cost of a consent order versus the cost of proactive remediation? Consider not just financial penalties but reputational damage, customer attrition, and impact on deposit growth and wealth management retention.
- How does this enforcement landscape affect your AI deployment pace for other initiatives? Does regulatory uncertainty cause you to slow down across the board, or do you accelerate governance investment to enable continued deployment?
- How do you communicate this regulatory risk to your board, investors, and employees without triggering panic or premature market reaction?

Consulting Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Consulting participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: Copilot Productivity Gains and Junior Talent Pipeline Stress

Card Type: Workforce Intelligence

Classification: Internal Operations Intelligence / Workforce Intelligence

Source: Internal AI Deployment Team; Human Capital Analytics; Recruiting Operations

Reveal Timing: Round 1 — Decision Preparation Phase

Shared Intelligence: This card shares a common intelligence base with Finance and Law (regulatory/business pressure intelligence). Each industry receives the same underlying signal framed with industry-specific decision tensions.

The Intelligence:

Your internal AI copilot deployment is producing stronger productivity gains than initially projected — but the second-order effects are more disruptive than anticipated.

Productivity Data:

- Research and analysis copilots are reducing time-to-deliverable by 30-40% on standard analytical workstreams (market sizing, competitive benchmarking, industry analysis, financial modeling support)
- First-draft slide production time has dropped by approximately 50% for structured deliverables (status updates, market overviews, benchmarking decks)
- Data synthesis tasks that previously required 2-3 analysts working 4-5 days are now completed by one analyst with copilot support in 1-2 days

The Hidden Problem: Junior consultant utilization is declining. Associates and analysts are spending less time on the routine work that historically filled their days — and more time validating and editing AI output than doing genuinely complex, judgment-intensive work. The fundamental skill mix is shifting: junior talent is being displaced from production work but not yet absorbed into higher-value activities at sufficient scale.

Early Warning Signs:

- Junior consultant utilization has dropped from 66% to 58% over the past two quarters — below the threshold that supports current staffing levels

- Your most recent MBA recruiting cycle saw a 12% decline in acceptance rates at target schools; candidates cite "uncertainty about the consulting career path in an AI world"
- First-year analyst attrition is running 8 percentage points above the prior year; exit interviews cite "not learning enough" and "doing QA on AI output instead of real consulting work"
- Partners are spending 15-20% more time reviewing AI-augmented deliverables than they spent reviewing traditional analyst work — quality assurance overhead is real

Partner Compensation Pressure: The leverage model that drives partner economics is under strain. Fewer junior billable hours per engagement means lower revenue per partner — unless pricing shifts to compensate. Early modeling suggests that if current utilization trends continue without pricing adaptation, partner compensation could face meaningful downward pressure within 18-24 months.

Decision Tension:

Productivity gains vs. talent development pipeline. Your copilots are making delivery faster and cheaper, but they are hollowing out the entry-level experience that trains future partners and senior leaders. Do you invest heavily in redefining the junior role (retraining, new career paths, mentorship redesign) — accepting near-term cost and complexity? Or do you lean into the efficiency gains, accept reduced junior headcount, and restructure the economics around a smaller, more senior workforce?

Questions to Consider:

- What is the minimum junior consultant utilization rate you can sustain before headcount restructuring becomes necessary? What is the cost of restructuring vs. redeployment?
- How do you redefine the junior consulting role so that top graduates still see it as a compelling career launchpad? What does "day one" look like for a new analyst if AI handles the traditional entry-level work?
- How much partner time for AI quality review is acceptable before it erodes the senior capacity that generates the highest-margin client work?
- How do you communicate internally about the changing talent model without triggering a retention crisis among junior and mid-level staff?
- What is the cross-sector signal? If your clients in Finance, Healthcare, and Manufacturing are experiencing similar junior talent displacement, does that create an advisory opportunity (workforce transformation consulting) or a cautionary tale?

Card 2 — Round 2

Title: Competitive Disintermediation — Losing Deals to Specialists and In-House Teams

Card Type: Competitive Intelligence

Classification: Competitive Intelligence / Market Intelligence

Source: Sales & Pipeline Analytics; Win/Loss Review Board; Partner Client Feedback

Reveal Timing: Round 2 — Decision Preparation Phase

The Intelligence:

Your quarterly win/loss analysis reveals a pattern of competitive losses that has accelerated over the past two quarters. The losses are concentrated in AI-related engagements and follow a consistent pattern: you are losing to more specialized competitors on technical depth, and to client in-house teams on cost and speed.

Win/Loss Summary:

Segment	Reason for Loss	% of Lost Deals (AI-Related)	Competitor Type
AI Transformation Strategy	Deeper technical expertise, faster staffing	15%	AI-Native Consultancies (e.g., Palantir Advisory, specialized boutiques)
AI Deployment & Implementation	Lower cost, embedded in client tech stack	12%	Cloud/Platform Vendor Advisory (e.g., AWS, Google, Microsoft professional services)
Strategic Organizational Change	In-house team preferred for cultural sensitivity	8%	Client Internal AI/Transformation Teams
Data & Analytics Strategy	Technical depth perceived as stronger	10%	Specialized Data/Analytics Firms

Pattern Analysis:

- Your win rate on horizontal AI transformation engagements (generic "help us build an AI strategy") has declined from 45% to 31% over the past 12 months
- Your win rate on deep vertical AI engagements (e.g., "help us deploy AI in clinical trial operations" or "redesign our lending model governance") remains strong at 62% — essentially unchanged
- Clients that previously would have hired you for end-to-end AI transformation are now splitting the work: hiring specialists for technical implementation and keeping you (if at all) for strategic framing and change management
- Three of your top-20 clients have formally established internal AI Centers of Excellence in the past 6 months, with explicit mandates to reduce external consulting spend on AI-related work

Revenue Impact:

- AI transformation advisory was projected to be your fastest-growing segment (25%+ growth). Actual growth is running closer to 12%, with the gap attributable to competitive losses and deal deferrals
- Average deal size for AI engagements is declining — clients are scoping work more narrowly and retaining more execution in-house
- Your pipeline of AI-related opportunities remains strong in dollar terms, but conversion rates are weakening

Bright Spot: Engagements where you combine deep industry expertise with AI advisory are outperforming. A Financial Services AI governance engagement (regulatory compliance + organizational design + AI deployment) closed at 2.3x the average AI engagement size and at premium rates. A Healthcare AI ethics and deployment engagement (clinical workflow redesign + regulatory navigation) closed at 1.8x average. Vertical depth is the differentiator.

Decision Tension:

Breadth vs. depth. You are losing horizontal AI engagements to specialists who are faster, cheaper, and more technically credible. You are winning vertical AI engagements where your industry expertise creates differentiated value. Do you double down on 3-4 verticals and accept that you will cede the horizontal AI transformation market to specialists? Or do you try to compete on both fronts — investing in technical depth to match specialists while maintaining your vertical advantages? The first path is focused but narrowing. The second is expensive and may spread you too thin against better-resourced competitors.

Questions to Consider:

- Can you credibly compete with AI-native consultancies on technical depth, or is that a losing battle? Where exactly is the boundary between "we can win this" and "we should partner or refer"?
- What is the cost and timeline to build genuinely differentiated vertical AI practices in 3-4 industries? How many specialized hires (PhD-level AI researchers, industry regulatory experts, domain engineers) do you need, and what does that cost?
- Should you acquire a specialized AI consultancy to accelerate capability building? What are the cultural integration risks?
- How do you manage partners who lead horizontal AI practices and face declining revenue? Internal politics are real — restructuring practice areas around verticals means winners and losers among the partnership.
- What does your competitive positioning statement become? "We do AI strategy" is commoditized. What is the 15-word version of your differentiated value proposition?

Card 3 — Round 3

Title: Pricing Model Collapse — The Billing Rate Crisis

Card Type: Operational Intelligence

Classification: Financial Intelligence / Commercial Strategy

Source: CFO; Pricing & Commercial Excellence Team; Client Contract Analytics

Reveal Timing: Round 3 — Decision Preparation Phase

The Intelligence:

Your financial analysis of AI-enabled service delivery has surfaced a structural pricing problem that threatens your revenue model more acutely than the competitive dynamics in Card 2.

The Core Problem: Associates using AI copilots are dramatically more productive — but clients will not pay traditional hourly rates for AI-assisted work. The efficiency gain that improves your delivery cost simultaneously erodes your pricing power.

Pricing Impact Data:

Service Type	Traditional Billing Rate	AI-Assisted Rate (Client Demand)	Discount	Productivity Gain	Net Margin Impact
Strategic Analysis & Research	\$450-550/hr	\$320-400/hr	25-30%	35-40% faster	Modestly positive
Financial Modeling & Valuation	\$400-500/hr	\$300-375/hr	20-25%	30-35% faster	Roughly neutral
Benchmarking & Market Sizing	\$350-450/hr	\$250-300/hr	25-35%	40-50% faster	Negative
Deliverable Production (Decks)	\$300-400/hr	\$200-275/hr	30-35%	45-55% faster	Negative
Organizational Design & Change Mgmt	\$500-650/hr	\$475-600/hr	5-10%	10-15% faster	Positive

The Structural Problem:

- Time-based billing rewards inefficiency. AI increases efficiency, which mechanically reduces billable hours. A two-week analysis compressed to three days generates 70% fewer billable hours — even at full rates, total revenue per engagement drops

- Clients are becoming sophisticated about AI productivity. Three of your top-10 clients have formally requested AI-assisted delivery pricing tiers. Two have threatened to rebid engagements if rate structures do not reflect AI efficiency
- Your average blended billing rate across AI-assisted engagements is running at a 22% discount to traditional engagements — materially above the 10-12% discount you modeled when deploying copilots
- Large enterprise clients are pushing for outcome-based pricing: fixed fees for defined deliverables ("complete this market entry analysis," not "provide 400 hours of consultant time"). This transfers execution risk to you

Revenue Model Projections:

- If current pricing trends continue without structural adaptation, your firm faces 8-12% revenue compression over the next 18-24 months on AI-assisted engagements (which are becoming the majority of engagements)
- Margin impact is partially offset by lower delivery costs — but the offset is incomplete because quality assurance costs, AI infrastructure investment, and talent redeployment expenses absorb much of the efficiency gain
- The services with the highest productivity gains (benchmarking, market sizing, deliverable production) are also the services with the steepest pricing declines — your most efficient work is becoming your least profitable

The Value-Based Pricing Challenge: Six value-based pricing pilots are underway with trusted clients. Early results are mixed:

- Two pilots (Financial Services AI governance, Healthcare operational transformation) are generating above-traditional margins — clients are willing to pay a premium for assured quality and defined outcomes in high-stakes domains
- Two pilots (general strategy projects) are generating below-traditional margins — scope creep and client expectation management are proving difficult without hourly billing as a natural boundary
- Two pilots are too early to evaluate

Decision Tension:

Fight the pricing collapse or adapt to it. You can try to hold traditional rate structures by emphasizing quality, brand, and the irreplaceable value of human judgment — but clients are becoming more resistant, and competitors who accept lower rates will take share. Or you can accelerate the transition to value-based pricing — but the commercial infrastructure, risk management frameworks, and partner skill sets required for value-based pricing are immature, and the transition involves a period of margin compression while you learn. Neither path is painless. The question is which pain is more survivable.

Questions to Consider:

- What is your 18-month pricing strategy? Can you maintain a dual-pricing model (traditional rates for non-AI work, discounted rates for AI-assisted work) or does that create adverse selection where clients push all work into the discounted tier?
- How do you price value-based engagements to protect margin? What scoping, risk management, and client management disciplines are required?

- Which services should you stop offering if they cannot be profitably delivered at AI-assisted rates? Is there work you should walk away from?
- How do you restructure partner compensation when the revenue-per-partner metric is structurally declining? What replaces "book of business" as the primary partner performance measure?
- How do you communicate the pricing transition to clients, investors, and employees without signaling weakness? Is there a way to frame value-based pricing as a premium offering rather than a concession?
- What is the competitive implication? If you move to value-based pricing and competitors do not, do you gain or lose? If competitors move first, are you forced to follow on their terms?

Law Private Information Cards

***FACILITATOR NOTE:** Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Law participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.*

Card 1 — Round 1

Title: Contract Review Productivity vs. Quality Risk — Internal Pilot Results

Card Type: Operational Intelligence

Classification: Operational Intelligence / Regulatory Risk

Source: AI Deployment Steering Committee; Practice Group Leaders (Corporate/M&A, Litigation); General Counsel's Office

Reveal Timing: Round 1 Decision Preparation

Shared Intelligence: This card shares an underlying intelligence base with Finance and Consulting. All three industries are facing parallel AI deployment pressures — regulatory/business model disruption from early AI adoption — but each experiences the pressure through a different lens.

The Intelligence:

Your internal AI deployment data from the corporate/M&A and litigation practice pilots tells a nuanced story.

Productivity gains are real and significant. Associates using the contract review copilot are completing first-pass reviews 40-55% faster than manual workflows. Legal research AI tools are cutting preliminary research time by approximately 35%. Document drafting copilots produce serviceable first drafts in a fraction of the time associates previously required. Across pilot practice groups, effective associate productivity has increased materially.

But the quality picture is more complicated. Associates are spending 15-25% of their total work time validating and correcting AI output — checking citations for accuracy, verifying jurisdictional applicability, confirming that contract language matches client-specific requirements, and ensuring no hallucinated provisions or precedents have been introduced. Malpractice review protocols have caught substantive AI errors on approximately 3-5% of AI-generated work product — errors that, if undetected, could have created direct liability exposure (fabricated case citations in two instances, incorrect regulatory references in three, and a missed change-of-control provision in a mid-market M&A contract).

Bar rule compliance remains uncertain and resource-intensive. Your regulatory tracking team monitors AI guidance across 50 state bars. As of Q1 2026, 33 states have issued some form of AI guidance for legal practice — but the requirements differ materially. Some require disclosure of AI use in court filings; others require only internal documentation. Several states have not addressed AI at all. Your firm is applying the most conservative interpretation firm-wide, which adds administrative overhead but reduces compliance risk.

Malpractice insurers are paying attention. Your primary malpractice insurer has requested a formal briefing on AI deployment scope, quality assurance protocols, and error rates. They have not yet adjusted premiums but have signaled that firms without documented AI governance frameworks may

face surcharges beginning mid-2026.

Decision Tension:

The productivity gains are compelling — but they come with a quality and liability tail that is difficult to quantify. The 3-5% error rate on AI-generated work product is manageable with current review protocols, but those protocols consume a significant portion of the productivity gain. If you scale AI deployment beyond the current pilots to mainstream practice, the volume of AI-generated work requiring review will increase dramatically. Can your quality assurance infrastructure scale with it? And bar rule uncertainty in 15+ states means every new matter requires jurisdictional compliance analysis that adds friction and cost.

The core trade-off: Do you scale AI deployment aggressively to capture the productivity gains and competitive positioning (accepting higher quality assurance costs and residual liability risk)? Or do you hold at current pilot scale until bar rules stabilize and malpractice frameworks mature (accepting that competitors who move faster may capture market position)?

Questions to Consider:

- At what error rate does AI-assisted legal work become an unacceptable malpractice risk? Is 3-5% tolerable at scale, or only at pilot volumes?
- How much partner and senior associate time are you willing to allocate to AI quality review? At what point does review overhead eliminate the productivity gain?
- Should you adopt a uniform firm-wide AI policy (most conservative jurisdiction's rules) or a jurisdiction-by-jurisdiction approach (more efficient, higher compliance risk)?
- What is your malpractice insurer engagement strategy — proactive disclosure and governance demonstration, or wait for them to set requirements?
- How do you message AI deployment to clients? Transparency builds trust but may invite pricing pressure. Silence risks discovery and credibility damage.

Card 2 — Round 2

Title: Competitive Disintermediation — Legal AI Platforms and Specialist Firms Eroding Market Share

Card Type: Competitive Intelligence

Classification: Competitive Intelligence

Source: Client Development & Market Intelligence Team; Practice Group Revenue Analytics; Partner Retreat Working Group on Competitive Strategy

Reveal Timing: Round 2 Decision Preparation

Unique to Law. This card contains intelligence specific to the legal industry that other industries do not receive.

The Intelligence:

Your market intelligence and revenue analytics teams have completed a comprehensive competitive loss analysis for the past 12 months. The findings confirm what partners have been reporting anecdotally: you are losing work across multiple segments to AI-enabled competitors.

Competitive Loss Summary:

Segment	Work Type Lost	Estimated Share Lost	Primary Competitor Type
Routine Contract Work	Contract review, NDA drafting, standard agreements	~12% of prior-year volume	Legal AI platforms (Harvey.ai, LexisNexis+, Westlaw+) adopted by in-house legal departments
Focused Expertise	Patent prosecution, regulatory filings, specialized compliance	~7% of prior-year volume	AI-native boutique law firms with deep practice focus
AI Advisory & Governance	AI regulatory compliance, liability assessment, governance frameworks	~5% of prior-year volume	Consulting firms and Big Four legal practices with integrated AI advisory
Due Diligence	M&A due diligence, portfolio reviews, loan documentation	~4% of prior-year volume	ALSPs (alternative legal service providers) with AI-native delivery

The pattern is clear: commoditized, high-volume work is migrating fastest. Your largest corporate clients are piloting direct AI legal tools — three of your top-ten clients by revenue have deployed Harvey.ai or comparable platforms within their in-house legal departments in the past six months. They are not yet replacing your firm on complex matters, but they are pulling back on the routine work that generates reliable associate billings.

The specialist threat is different but equally concerning. AI-native boutique firms are winning mandates not on price alone, but on speed and technical sophistication. A patent prosecution boutique using AI-assisted prior art analysis is completing landscape reviews in days rather than weeks, with comparable accuracy. A regulatory compliance specialist is offering real-time regulatory monitoring that your practice groups cannot match with current tools.

Revenue impact to date: Your total revenue has not yet declined — growth in complex advisory work (particularly AI governance and litigation) has partially offset losses in routine segments. But the mix is shifting. Revenue from commoditized work categories declined 8-10% year-over-year. Revenue from complex, judgment-intensive work grew 12-15%. The net effect is modest overall growth, but the trend line is unmistakable.

Decision Tension:

Breadth vs. depth. You are a full-service firm. That positioning has been your competitive identity for decades — clients come to you because you can handle everything from a billion-dollar M&A transaction to a routine employment agreement. But AI-enabled competitors are unbundling that value proposition. Specialists are winning focused mandates. Platforms are winning commodity work. Your full-service model is being attacked from both ends.

Do you defend breadth — invest in AI-enabling every practice group to compete on efficiency across the full service spectrum, accepting that you will lose some commodity work but maintaining the one-stop-shop relationship? Or do you pivot toward depth — concentrate investment in complex, judgment-intensive practice areas (litigation, M&A advisory, regulatory strategy, AI governance) where human expertise is most defensible, and accept permanent share loss in commoditized segments?

Questions to Consider:

- What percentage of your current revenue comes from work that AI platforms could handle directly within 2-3 years? What is the replacement revenue strategy?
- Can you credibly compete with AI-native specialists on speed and cost in focused practice areas, or should you cede those segments and invest the savings in defensible expertise?
- Your top clients are building in-house AI legal capability. How do you maintain relevance to their general counsel — as a high-stakes advisor, or as a full-service provider? Can you be both?
- If you pivot toward specialization, how do you manage the cultural and economic disruption to practice groups that currently generate reliable (but increasingly threatened) commodity revenue?
- What is your competitive response timeline? Boutique firms and AI platforms are gaining market share now. How quickly can you reposition?

Card 3 — Round 3

Title: Bar Rule Enforcement Arrives and Hourly Billing Model Under Siege

Card Type: Regulatory Development

Classification: Regulatory Intelligence / Financial Intelligence

Source: Professional Responsibility Committee; Chief Financial Officer; Client Billing Analytics; State Bar Regulatory Affairs Liaison

Reveal Timing: Round 3 Decision Preparation

Unique to Law. This card contains intelligence specific to the legal industry that other industries do not receive.

The Intelligence:

Two simultaneous developments are creating acute strategic pressure.

Bar Rule Enforcement Is No Longer Theoretical.

State bars are moving from guidance to enforcement. Three significant developments in the past quarter:

1. **Mandatory AI disclosure:** Eight states have now adopted mandatory disclosure requirements for AI use in court filings and client-facing legal work product. Four additional states have proposed similar rules with effective dates in 2026-2027. Non-disclosure is sanctionable.
2. **AI ethics certification:** Two states (New York and California) have proposed mandatory AI ethics continuing legal education (CLE) requirements for all practicing attorneys. If adopted, every attorney in your firm practicing in those jurisdictions will need AI-specific ethics certification within 18 months.
3. **Enforcement precedent:** A mid-size firm in the Southern District received sanctions for submitting an AI-generated brief containing fabricated citations without disclosure. The court's opinion explicitly stated that attorneys have an affirmative duty to verify AI-generated legal work and disclose AI assistance. This is now citable precedent.

Your malpractice insurer has responded. Firms without documented AI governance frameworks — formal policies on AI use, quality assurance protocols, error tracking, and disclosure procedures — will face premium surcharges of 10-15% beginning Q3 2026. Your firm has governance frameworks in place, but several practice groups have been slow to implement them consistently.

The Hourly Billing Model Is Under Direct Attack.

Your CFO's billing analysis reveals what the market has been signaling:

Service Category	Avg. Rate Change (AI-Assisted vs. Traditional)	Client Demand for Alt. Fee	Volume Trend
Contract Review	-22% effective rate	65% of clients requesting fixed fee	Declining (AI platform substitution)
Due Diligence	-18% effective rate	55% of clients requesting fixed fee	Declining (ALSP competition)
Legal Research & Memo Drafting	-20% effective rate	50% of clients requesting capped fee	Stable (but rate-compressed)
Regulatory Compliance Advisory	-8% effective rate	30% of clients requesting value-based	Growing (AI governance demand)
Complex Litigation	+3% effective rate	15% of clients requesting alt. fee	Growing (judgment premium)
M&A Advisory (Complex)	+5% effective rate	10% of clients requesting alt. fee	Stable to growing

The pattern is stark. AI-assisted work is commanding materially lower effective rates. Clients can see that AI accelerates delivery and are demanding proportional pricing adjustments. For commoditized work categories, clients are pushing aggressively for fixed-fee or outcome-based arrangements — and they have competitive alternatives if you refuse. Meanwhile, complex and judgment-intensive work is holding or increasing in rate, but this category represents a smaller share of total billable hours.

The math is unforgiving. If the current rate erosion trend continues, your blended effective billing rate will decline 8-12% over the next 18 months. At current volume, that translates to a \$400-650M annual revenue reduction. Margin compression follows unless you reduce cost-to-serve proportionally — which means either fewer associates, lower associate compensation, or dramatically higher utilization on the work that remains.

Decision Tension:

Regulatory compliance and pricing pressure are hitting simultaneously, and they compound each other.

On the regulatory side: proactive bar rule compliance is expensive (firm-wide AI governance implementation, ethics CLE, disclosure protocols, malpractice insurer engagement) but positions you as a trusted, compliant firm in a market where trust is becoming a competitive differentiator. Waiting for clarity is cheaper in the short term but risks enforcement actions, sanctions, and reputational damage that could be catastrophic.

On the pricing side: the hourly billing model is eroding for AI-assisted work, and the erosion is accelerating. You can fight it (maintain rate cards, lose price-sensitive clients to competitors) or adapt (develop alternative fee models, accept lower revenue per matter but potentially higher volume and margin through AI-assisted delivery). Neither option is comfortable.

The compound problem: Proactive compliance investment increases your cost base at exactly the moment pricing pressure is reducing your revenue per matter. The firms that emerge strongest will be those that invest in compliance and governance (building trust and market credibility) while simultaneously restructuring their pricing and delivery models (maintaining margin despite rate erosion). But the capital and leadership bandwidth required to do both at once is substantial.

Questions to Consider:

- Do you implement firm-wide AI governance and disclosure protocols now (ahead of mandate in most jurisdictions) or wait for clarity? What is the cost of each approach? What is the risk?
- How do you handle the malpractice insurer requirements — proactive demonstration of governance (reduced premiums, trust signal) or minimum compliance (lower cost, higher premium risk)?
- What is your pricing strategy for AI-assisted work? Accept rate erosion and compete on volume/efficiency? Or establish premium pricing for AI-assisted work with guaranteed quality assurance (positioning quality review as value-add, not cost)?
- At what point do you restructure associate headcount? The revenue math suggests the current model is unsustainable if rate erosion continues. What is your workforce transition plan?
- How do you communicate to partners that the economics of law firm practice are permanently changing? Partner compensation will be affected. What is the narrative?

Manufacturing Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Manufacturing participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: Predictive Maintenance ROI Scaling Costs & Plant Retrofit Reality

Card Type: Operational Intelligence

Reveal Timing: Round 1 Situation Update

Classification: Internal Operations & Financial Analysis

Source: Manufacturing Operations Review + Finance (Q4 2025 / Q1 2026)

Shared Intelligence: *This card shares a common base with the Logistics industry participant. Both industries received predictive maintenance scaling cost data from the same operations review. Manufacturing and Logistics participants may discover overlapping information during cross-industry discussion.*

The Intelligence:

Your predictive maintenance AI, deployed in 4 pilot plants, reduced unplanned downtime and extended equipment life in a 12-month trial. Technical validation is strong. However, scaling to remaining plants requires material sensor infrastructure, edge computing, and OT/IT integration investment per plant. Total scaling cost across all 28 plants is significant and constrained by your annual capital budget.

Equipment intelligence data reveals that your 8 highest-maintenance plants account for a disproportionate share of annual maintenance costs. Prioritizing these plants first could generate meaningful annual savings with a substantially shorter payback timeline than full deployment.

At current capital pace, full deployment across all 28 plants takes multiple years. Equipment leasing and financing arrangements could accelerate the timeline but increase total lifetime cost. Every dollar allocated to plant retrofits competes with other AI investment demands across the organization.

Decision Tension:

Do you prioritize rapid full-deployment across all 28 plants (accelerate via leasing/financing, maximize total revenue potential) or selective high-value deployment (slower, proven ROI, lower total cost)? Capital constraint forces a direct trade-off between speed and fiscal discipline. Your choice also signals to equipment vendors, unions, and investors what kind of AI transition you are running.

Questions to Consider:

- Which 8-12 plants should be prioritized for first-wave deployment, and what criteria drive that selection?
- What is an acceptable payback period for manufacturing AI capex? (3 years? 4 years? 5 years?)
- Does leasing/financing make sense to accelerate deployment, or is conservative capital allocation preferable given long asset lifecycles?

- How does your plant retrofit pace affect the broader organization's AI capital budget (including fleet telematics and other demands)?
- Should you signal large-scale investment commitment to suppliers and customers to build competitive advantage, or keep deployment plans quiet until ROI is proven at scale?

Card 2 — Round 2

Title: Warehouse Automation Results & Labor Displacement Signal

Card Type: Workforce Intelligence

Reveal Timing: Round 2 Situation Update

Classification: Technology Validation + Labor Relations Intelligence

Source: Manufacturing Automation Trials + Union Monitoring (Q2-Q3 2026)

The Intelligence:

Your warehouse automation pilots in 2 manufacturing facilities have delivered results:

Warehouse Automation Performance: Pick-and-place robots and sorting systems improved labor productivity and throughput materially in the pilot facilities. Quality consistency also improved, with fewer handling-related defects. The technology works. The question is no longer technical validation but organizational readiness for scale deployment.

Labor Impact & Union Response: Your largest manufacturing facility is planning major automation upgrades that will reduce headcount in warehouse and material handling roles. Union leadership is monitoring closely and has signaled that any further automation expansion without corresponding retraining commitments will trigger formal grievance proceedings. The political sensitivity is high: your plant-level automation decisions are being watched as a signal for the entire manufacturing sector's labor transition approach.

Retraining Economics: You have negotiated "no-layoff" agreements committing to retraining and redeployment of affected workers. This is a material cost (\$35-40M across affected facilities) but preserves labor relations, safety culture, and institutional knowledge. Retraining costs reduce net automation savings by 20-25%, but the reputational and operational benefits of maintaining union cooperation are substantial. Retraining requires a 12-18 month productivity ramp before redeployed workers reach full effectiveness.

OT/IT Integration Update: Early OT/IT integration work at pilot plants has revealed that legacy production systems are more heterogeneous than expected. Integration cost per plant is running above initial estimates. The phased approach (4-6 plants first, then scale) appears validated, but the timeline for full integration across all plants has extended.

Decision Tension:

Do you accelerate warehouse automation rollout across more facilities, knowing that labor transition costs are high and union sensitivity is peaked? Or do you pause to consolidate labor relations and complete retraining before expanding? How do you balance the operational case for automation against the reputational and labor relations cost of moving too fast?

Questions to Consider:

- Should you accelerate warehouse automation to additional facilities, or pause until retraining programs are further along?
- Is the \$35-40M retraining/redeployment cost justified by the labor relations and organizational stability benefit?
- How do you communicate automation plans to the workforce without triggering a broader labor relations crisis?
- Does the OT/IT integration cost overrun change your deployment sequencing or capital allocation?
- Does a public commitment to "no-layoff" retraining become a competitive disadvantage (higher costs than rivals) or advantage (talent retention, reputation, union cooperation)?

Card 3 — Round 3

Title: Supplier Digitalization Gap & OT/IT Integration Cost Overruns

Card Type: Market Intelligence

Reveal Timing: Round 3, Post-Inject

Classification: Strategic Intelligence + Supply Chain Risk Assessment

Source: Supply Chain Intelligence Team + Manufacturing Technology Office (Q4 2026 / Q1 2027)

The Intelligence:

Your push for end-to-end manufacturing AI optimization has revealed critical ecosystem and infrastructure constraints:

Supplier Digitalization Gap (Procurement Perspective): Your top suppliers (representing the majority of procurement spend) are not digitally ready to share real-time production data, quality metrics, or demand signals. Most operate legacy systems and manual processes. Integrating supplier data into your production scheduling and demand forecasting AI requires material investment per supplier for digital interfaces, data standardization, and connectivity. Total aggregate cost to digitalize your top supplier base is significant. Most suppliers are resistant, citing implementation costs, data security concerns, and uncertain ROI. Without supplier data integration, your production scheduling AI operates on incomplete information, limiting optimization potential.

Realistic Digitalization Pathway: Your analysis suggests a small number of strategic suppliers (those with existing digital capability or strong commercial incentive) can be brought online through long-term contracts, volume commitments, and technical support. However, the majority of your supplier base will remain analog for the medium term. Full end-to-end manufacturing optimization is constrained by supplier readiness, not your internal technical capability.

OT/IT Integration Cost Overruns: The OT/IT convergence program, now underway in pilot plants, is running materially above initial cost estimates. Legacy production systems (PLCs, SCADA, proprietary protocols) require custom integration work that does not scale easily across plants. Each plant has a different equipment mix and vintage, making standardized deployment impractical. The timeline for full OT/IT integration has extended, and the program is consuming more of your AI capital budget than planned.

Supply Chain Disruption Test: A recent geopolitical supply disruption tested your demand forecasting and production scheduling AI. Models trained on historical data failed to anticipate the disruption. Suppliers without real-time visibility could not react quickly. Your production scheduling system underestimated input volatility and overcommitted to production plans that required rapid revision.

Decision Tension:

Do you invest heavily in supplier digitalization (ecosystem transformation, long timeline, distributed ROI) or accept a hybrid digital+analog supply chain and optimize within those constraints? How do you manage OT/IT integration cost overruns without derailing your broader manufacturing AI program? And how do you build production resilience against disruptions that AI models cannot predict?

Questions to Consider:

- Which 10-20 suppliers should be prioritized for digitalization partnership, and what incentives will actually work?
- Should you invest in supplier technical enablement, or expect suppliers to self-fund their digital transformation?
- How do you design production scheduling AI that works effectively with a mixed digital+analog supplier base?

- Does the OT/IT cost overrun require you to slow down the integration program, or can you absorb the additional cost?
- What production resilience measures (safety stock, supplier diversification, demand buffering) should complement AI forecasting?
- How should the disruption test inform your AI model design going forward? (Stress-test forecasts, add volatility margins, build manual override capabilities?)

Logistics Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Logistics participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: Predictive Maintenance ROI Scaling Costs & Fleet Telematics Reality

Card Type: Operational Intelligence

Reveal Timing: Round 1 Situation Update

Classification: Internal Operations & Financial Analysis

Source: Fleet Operations Review + Finance (Q4 2025 / Q1 2026)

Shared Intelligence: *This card shares a common base with the Manufacturing industry participant. Both industries received predictive maintenance scaling cost data from the same operations review. Manufacturing and Logistics participants may discover overlapping information during cross-industry discussion.*

The Intelligence:

Your predictive maintenance and telematics AI, deployed across 800 trucks in 3 pilot logistics zones, achieved meaningful fuel savings and on-time delivery improvement in a 12-month trial. Vehicle condition monitoring reduced unplanned maintenance events and extended equipment life. Technical validation is strong. However, scaling to the full fleet of 5,000+ vehicles requires material telematics hardware, data infrastructure, and integration investment per vehicle. Total scaling cost is significant and constrained by your annual capital budget.

Fleet analytics reveal that your highest-utilization routes and vehicle segments account for a disproportionate share of fuel and maintenance costs. Prioritizing these segments for full telematics deployment first could generate meaningful annual savings with a substantially shorter payback timeline than fleet-wide rollout.

At current capital pace, full fleet deployment takes multiple years. Vehicle lease cycles (5-8 years) create natural upgrade windows, but waiting for lease turnover delays the efficiency gains. Accelerating through upfront hardware retrofits increases near-term cost but captures savings sooner.

Decision Tension:

Do you prioritize rapid fleet-wide telematics deployment (accelerate via hardware retrofits, maximize total savings potential) or selective high-utilization deployment (slower, proven ROI, lower total cost)? Capital constraint forces a direct trade-off between fleet coverage and fiscal discipline. Your route optimization AI is only as good as the data it receives, meaning partial fleet coverage creates optimization gaps on unequipped routes.

Questions to Consider:

- Which vehicle segments and routes should be prioritized for first-wave telematics deployment?
- What is an acceptable payback period for fleet AI investment? (2 years? 3 years? 4 years?)

- Should you accelerate through hardware retrofits or align with natural vehicle lease replacement cycles?
- How does your fleet telematics pace affect the broader organization's AI capital budget (including plant retrofits and other demands)?
- Does partial fleet coverage meaningfully limit the value of route optimization AI, or can you optimize around data gaps?

Card 2 — Round 2

Title: Route Optimization Labor Impact & Autonomous Vehicle Pilot Intelligence

Card Type: Workforce Intelligence

Reveal Timing: Round 2 Situation Update

Classification: Technology Validation + Labor Relations Intelligence + Regulatory Monitoring

Source: Fleet Operations Testing + AV Partnership Intelligence + Union Monitoring (Q2-Q3 2026)

The Intelligence:

Your AI systems pilots across fleet operations reveal a mixed picture of technological promise and human complexity:

Route Optimization at Scale: Your AI-driven route optimization, now expanded beyond the initial 3-zone pilot, is achieving material fuel reduction and on-time delivery improvement in equipped zones. Scaling to the full fleet could generate significant annual savings (\$180-200M potential at full deployment). The algorithm works. But fleet-wide deployment depends entirely on driver acceptance.

Driver Acceptance Reality: Driver response to routing AI is sharply generational. Younger drivers (<35) embrace the system at high rates (~85% voluntary adoption), citing fuel bonuses and reduced planning burden. Older drivers (>50) resist at significantly lower rates (~55% adoption), citing concerns about GPS tracking, algorithmic control over routing decisions, loss of professional autonomy, and perceived surveillance. Union representatives have raised formal concerns about "algorithmic management" and are requesting contractual protections around routing AI override capabilities and driver data privacy.

Autonomous Vehicle Partnership Intelligence: Your monitoring of the autonomous vehicle landscape reveals that federal DOT/NHTSA rules on commercial autonomous long-haul trucks remain years away (2028-2030 at earliest). Competitor intelligence confirms that AV specialists (Waymo Via, Aurora Innovation) are materially ahead of any internal fleet operator capability. Partnering with an AV specialist for limited corridor pilots is faster and cheaper than internal development, but the regulatory and commercialization timeline remains deeply uncertain. Internal AV development would require material capex with multi-year timelines and no guaranteed regulatory pathway.

Decision Tension:

How do you scale route optimization across a workforce with sharply divided acceptance? Do you mandate adoption (risking union conflict and driver attrition) or invest in change management (slower, more expensive, but preserves workforce relationships)? Separately, how do you position for autonomous vehicles: invest in partnership pilots now (strategic option value, uncertain payoff) or redirect all capital to proven route optimization and fleet efficiency (near-term ROI, lower strategic risk)?

Questions to Consider:

- How do you improve older driver adoption of routing AI? What change management programs, incentives, or override capabilities would move the needle?
- Should routing AI adoption be mandatory or voluntary? What are the labor relations consequences of each approach?
- Is a Waymo/Aurora partnership pilot worth the capital and management attention given regulatory uncertainty, or should you redirect those resources to proven fleet optimization?
- How do you manage union concerns about "algorithmic management" and driver data privacy?
- Does a public commitment to driver autonomy and "AI as assistant" framing help or hurt your competitive positioning?

Card 3 — Round 3

Title: Supplier Ecosystem Readiness Gap & Last-Mile Profitability Limits

Card Type: Market Intelligence

Reveal Timing: Round 3, Post-Inject

Classification: Strategic Intelligence + Network Economics Assessment

Source: Supply Chain Intelligence Team + Network Economics Analysis + Demand Forecasting Review (Q4 2026 / Q1 2027)

The Intelligence:

Your push for end-to-end logistics AI optimization has revealed critical ecosystem and economic constraints:

Supplier & Shipper Ecosystem Readiness Gap: Your logistics AI optimization depends on data from shippers, receivers, and supply chain partners. Most of your shipper base lacks the digital capability to share real-time shipment data, demand forecasts, or inventory positions. Integrating partner data into your route optimization and demand forecasting systems requires material investment per partner for API development, data standardization, and connectivity. The aggregate cost to digitalize your top partners is significant. Most shippers are resistant, viewing the investment as benefiting your operations more than theirs. Without partner data integration, your demand forecasting and dynamic routing AI operates on incomplete information.

Last-Mile Profitability Analysis: Your detailed route-level profitability analysis confirms what pilots suggested: AI route optimization delivers meaningful value on high-density urban and regional corridors but cannot fix the structural economics of low-density last-mile delivery. Rural routes with fewer than a threshold number of stops per route are unprofitable regardless of optimization. AI improves margins on profitable routes but does not convert unprofitable routes to profitable ones. A material portion of your last-mile network operates below breakeven, and no amount of algorithmic improvement will change the underlying stop density economics.

Demand Forecasting Model Brittleness: A recent geopolitical supply disruption tested your AI demand forecasting accuracy. Models trained on historical shipping patterns failed to anticipate the disruption's impact on freight volumes and routing needs. The system overcommitted capacity to pre-disruption patterns and was slow to reallocate. Post-mortem analysis revealed that your forecasting models lack robustness to regime changes and tail events, performing well in stable conditions but poorly when the underlying demand distribution shifts.

Competitive Differentiation Erosion: As competitors also deploy route optimization and demand forecasting AI, the efficiency advantage of early adoption is diminishing. Differentiation is shifting from algorithm quality to execution speed, network density, and supply chain resilience.

Decision Tension:

Do you invest heavily in shipper/partner digitalization (ecosystem transformation, long timeline, uncertain partner cooperation) or accept a hybrid data environment and optimize within those constraints? How do you handle the last-mile profitability problem: continue to subsidize unprofitable rural routes (customer retention, network completeness) or rationalize the network and outsource to regional carriers (margin improvement, service risk)? And how do you build demand forecasting resilience against disruptions that historical models cannot predict?

Questions to Consider:

- Which shippers and supply chain partners should be prioritized for data integration, and what incentives will motivate them?
- Should you invest in partner digital enablement, or expect partners to self-fund integration?

- How do you handle unprofitable last-mile routes: subsidize, outsource, or exit? What are the customer retention and network completeness implications?
- How do you design demand forecasting AI that works with incomplete partner data and remains robust to regime changes?
- What supply chain resilience measures (capacity buffers, flexible routing, demand smoothing) should complement AI forecasting?
- Does your logistics optimization advantage persist as competitors catch up, or is it temporary? What is the next source of differentiation?

Big Tech Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the Big Tech participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: AI Infrastructure Costs Exceed Projections & Capex Escalation

Card Type: Operational Intelligence

Reveal Timing: Round 1 — distribute face-down at start of decision preparation phase

Classification: Restricted

Source: Internal capex tracking, cloud infrastructure utilization dashboards, and competitive intelligence

Shared Intelligence (also received by SaaS — framed differently):

The Intelligence:

Your internal capex models for AI infrastructure are tracking significantly above projections. Data center buildout, custom silicon procurement, and inference infrastructure costs are escalating faster than anticipated. Your GPU/TPU utilization rates are lower than expected — inference demand from enterprise cloud customers and internal AI product workloads has not scaled as quickly as projected. Many clusters are idle between peak usage windows. You are on track to spend \$18B on AI-related infrastructure capex in 2026, up from \$12B projected six months ago. The payback period on this infrastructure investment may exceed 3-5 years.

Competitive intelligence suggests all major cloud providers are experiencing similar cost pressures and utilization shortfalls. A market bifurcation is emerging: only well-capitalized players can sustain this level of infrastructure investment. Mid-market cloud providers and smaller AI companies are being forced to rely on open-source models and shared infrastructure, accelerating the commoditization of AI compute.

The critical question for you is not whether to build AI infrastructure — that ship has sailed. The question is how to drive utilization rates up, how to convert enterprise AI pilots into production workloads that generate sustained cloud revenue, and how to manage the gap between committed capex and realized demand. Your cloud pricing strategy, managed services offering, and enterprise sales execution all directly affect whether this infrastructure investment pays off or becomes stranded capital.

Decision Tension:

Do you accelerate infrastructure buildout further to capture enterprise AI workload migration (betting that demand will catch up to supply)? Or do you moderate capex growth and focus on driving utilization of existing infrastructure through pricing incentives, managed services, and customer success investment?

Questions to Consider:

- What is the utilization rate threshold at which your current infrastructure investment becomes profitable? How far are you from that threshold?
- Can you accelerate enterprise cloud customer conversion from AI pilot to production workload? What are the bottlenecks?

- How does your cloud pricing strategy compare to competitors? Is aggressive pricing required to capture workload share, or does it destroy margin?
- What is the competitive consequence if you moderate capex growth while competitors continue to build? How quickly does market share shift?
- How long can you sustain 30%+ capex growth before investor pressure forces a reset?

Card 2 — Round 2

Title: Enterprise Cloud Pricing Collapse & Open-Source Commoditization Pressure

Card Type: Market Intelligence

Reveal Timing: Round 2 — distribute face-down at start of decision preparation phase

Classification: Confidential

Source: Enterprise cloud sales team intelligence, customer contract negotiations, and competitive pricing analysis

The Intelligence:

Your enterprise cloud customers are demanding access to AI services at commodity pricing. The pricing floor is being set by open-source models: customers are benchmarking your managed AI services against the cost of running LLaMA, Mistral, or Phi on their own infrastructure or through discount cloud providers. Your premium pricing for proprietary AI services is under direct pressure.

Specific signals from the field:

- Your top 50 enterprise cloud accounts (representing 25% of cloud revenue) are renegotiating contracts. They want 20-30% volume discounts on AI compute and managed services, citing competitive offers from other cloud providers and internal open-source alternatives.
- Three Fortune 100 customers have announced internal "open-source first" AI strategies, explicitly reducing dependence on proprietary cloud AI services. They are building internal capabilities to run open-source models on bare-metal infrastructure, bypassing your managed services entirely.
- Your cloud sales team reports that new AI workload deals are taking 40% longer to close than 12 months ago. Customers are conducting extended proof-of-concept evaluations, comparing your platform against competitors on price, performance, and flexibility.
- Open-source model quality continues to improve. For 70%+ of enterprise use cases (classification, summarization, basic generation, structured data extraction), open-source models deliver acceptable performance at a fraction of proprietary costs. Your premium is justified only for complex, high-stakes applications (advanced reasoning, multimodal, domain-specific fine-tuning).

Your cloud margin is compressing. The gap between infrastructure cost (which you bear) and the price customers are willing to pay (which is falling) is narrowing. If you match competitive pricing to retain volume, margins erode. If you hold premium pricing, you lose workload share to competitors and open-source alternatives.

Decision Tension:

Do you defend premium pricing for proprietary AI cloud services (protecting margin but risking workload share loss)? Or do you match commodity pricing to retain volume and market position (preserving share but compressing margins and potentially triggering a race to the bottom)?

Questions to Consider:

- What percentage of your cloud AI revenue comes from use cases where open-source models are viable substitutes? How fast is that percentage growing?
- Can you differentiate managed services (security, compliance, fine-tuning, SLA guarantees) enough to justify a premium over open-source alternatives?
- What is the revenue impact of losing your top 50 enterprise accounts to competitors or open-source alternatives? What retention investment is justified?

- If you cut pricing aggressively, what is the margin impact? Can you offset it with volume growth?
- What does a "two-tier" pricing strategy look like — commodity pricing for basic AI compute, premium pricing for advanced managed services?

Card 3 — Round 3

Title: Antitrust Enforcement Targets Platform Practices & Data Governance

Card Type: Regulatory Development

Reveal Timing: Round 3 — distribute face-down at start of decision preparation phase

Classification: Restricted

Source: Regulatory intelligence, legal counsel briefing, and government affairs team analysis

The Intelligence:

Federal antitrust enforcement is moving from investigation to action. The enforcement focus is on your platform practices, data governance, and the degree to which your market position creates barriers to competition in AI services. An enforcement action is being drafted that would likely include:

- **Data separation mandates:** You cannot use enterprise cloud customer data for improving your own AI products or services without explicit, granular, per-use-case consent. Cross-product data flows (e.g., using search data to improve cloud AI services, or using cloud customer usage patterns to optimize advertising) would require separate opt-in agreements. Compliance cost is estimated at \$500M-\$1B over 18 months.
- **API access mandates:** You must provide non-discriminatory API access to your cloud AI infrastructure and managed services. Competitors and third-party developers must be able to access the same services, at the same pricing, as your own internal product teams. This reduces the integration advantage your internal product teams currently enjoy.
- **Platform conduct restrictions:** Bundling restrictions on how you package AI features with existing products (e.g., cannot require cloud subscription to access AI features; cannot preferentially surface your own AI services over competitors in your marketplace or app store). This directly affects your ecosystem lock-in strategy.

Parallel intelligence: European regulators are advancing similar enforcement actions under the Digital Markets Act, with potentially stricter requirements including forced interoperability and data portability mandates.

Compliance will require significant organizational restructuring of data governance, pricing, and product architecture. Legal and engineering teams will be consumed by compliance work for 12-18 months. The distraction and resource diversion are as costly as the direct financial impact.

Decision Tension:

Do you adopt a proactive compliance posture — voluntarily implementing data separation, open API access, and fair pricing before enforcement arrives (reducing regulatory risk but also reducing competitive advantage)? Or do you contest enforcement through litigation and lobbying while maintaining current practices (preserving competitive advantage near-term but risking more severe enforcement outcomes and reputational damage)?

Questions to Consider:

- What is the competitive impact of data separation mandates? How much does your AI product quality degrade if cross-product data flows are restricted?
- Can you turn API access mandates into a revenue opportunity (charging competitors for access to your infrastructure)?
- What is the reputational cost of aggressive litigation vs. proactive compliance? How do enterprise customers react to each approach?

- How do you prioritize engineering resources between compliance implementation and AI product roadmap execution? What slips?
- If European enforcement is stricter, does it make sense to implement a global compliance standard (one set of practices) or region-specific approaches?

B2B/B2C SaaS Private Information Cards

FACILITATOR NOTE: Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the SaaS participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.

Card 1 — Round 1

Title: AI Feature Costs Erode Margins & Third-Party Infrastructure Expense Escalates

Card Type: Operational Intelligence

Reveal Timing: Round 1 — distribute face-down at start of decision preparation phase

Classification: Restricted

Source: Internal finance analysis, cloud infrastructure billing data, and engineering cost modeling

Shared Intelligence (also received by Big Tech — framed differently):

The Intelligence:

Your internal cost models for AI feature delivery are tracking significantly above projections. The AI features you have deployed (copilots, predictive analytics, intelligent search, automation workflows) are consuming more cloud compute, more third-party API calls, and more engineering maintenance than anticipated. Gross margin on AI-powered features is running 10-15 points below your legacy product margin — and the gap is widening as adoption scales and inference volume grows.

Specific cost pressures:

- Third-party API costs (OpenAI, Anthropic, Cohere) are the largest line item. Your negotiated volume pricing is better than retail, but at scale the cost is material — approximately \$2-3 per active user per month for AI features, against a blended subscription price of \$25-30 per user per month. As AI feature adoption increases across your customer base, this cost grows linearly while subscription revenue grows only if you can upsell.
- Internal fine-tuning and model optimization work is consuming 30% of your AI engineering capacity, leaving less bandwidth for new feature development. The "build once, deploy everywhere" assumption has not held — each vertical and customer segment requires model customization.
- Vector database and embedding storage costs are a new and growing expense category that was not in your original infrastructure budget.

Competitive intelligence suggests all major SaaS companies are experiencing similar margin pressure from AI features. The companies managing costs most effectively are those with disciplined make-vs-buy strategies — leveraging open-source models for commodity tasks and reserving proprietary development for high-differentiation use cases.

Decision Tension:

Do you absorb the margin compression as an investment in competitive positioning (betting that AI feature adoption will drive retention and upsell that offsets the cost)? Or do you aggressively optimize AI infrastructure costs — switching to open-source models, reducing third-party API dependence, and

limiting AI feature scope — to protect margins now?

Questions to Consider:

- What is the true gross margin on your AI features at current scale? At projected scale in 12 months?
- Which AI features can be delivered using open-source models without meaningful quality degradation? Which require proprietary APIs?
- Can you pass any AI infrastructure costs to customers through pricing changes, or will the market not bear it?
- What is the competitive consequence of slowing AI feature development to protect margins? How quickly do AI-native competitors gain ground?
- What is your break-even point — the adoption level at which AI feature revenue (upsell, retention) exceeds AI infrastructure cost?

Card 2 — Round 2

Title: Enterprise Customers Demand Free AI & Competitive Displacement Accelerates

Card Type: Competitive Intelligence

Reveal Timing: Round 2 — distribute face-down at start of decision preparation phase

Classification: Confidential

Source: Enterprise sales team intelligence, customer renewal negotiations, and competitive win/loss analysis

The Intelligence:

Your top 20 enterprise customers have made their position explicit in contract renewal discussions: they expect AI features to be bundled into their existing subscriptions at no additional cost. The framing from customers is consistent — "AI is table-stakes, like search was a decade ago. We shouldn't pay extra for it." Willingness to pay a premium for AI capabilities is minimal across your customer base: only 5% of accounts are willing to pay 10%+ more for AI-powered features.

The situation is compounding:

- Deal slippage: Your enterprise sales cycle has extended by 30-40%. Prospects are evaluating AI-native alternatives in parallel with your platform. Three deals representing \$45M in annual recurring revenue are stalled as customers compare your AI capabilities against specialized competitors.
- Churn signals: Three of your top 20 accounts (representing 8-10% of total revenue) have initiated vendor evaluation processes. Two are evaluating AI-native CRM/ERP competitors; one is building internal AI tools on open-source models. Your customer success team assesses these as "high risk" renewals.
- Mid-market displacement: Below your top accounts, AI-native startups are winning new customer acquisition in your mid-market segment. Five accounts that would typically have been your wins in the past quarter chose AI-native alternatives, citing faster AI capabilities, lower cost, and simpler implementation. Your win rate in competitive deals has dropped from 45% to 32% over the past two quarters.
- Pricing contagion: Customers who have seen competitors offer AI features at no additional cost are now demanding the same from you. Your pricing team estimates that bundling AI features into base subscriptions would reduce blended revenue per user by 8-12%.

Decision Tension:

Do you bundle AI features into base subscriptions to defend market position and reduce churn risk (accepting margin compression)? Or do you hold premium pricing for AI features and invest in differentiation to justify the premium (risking accelerated churn and competitive displacement)?

Questions to Consider:

- What is the revenue impact of bundling AI into base subscriptions vs. the revenue impact of losing 8-10% of customers to competitors?
- Can you create a tiered AI offering — basic AI bundled, advanced AI as premium — that satisfies most customers while preserving some pricing power?
- How do you respond to the three high-risk top-20 account renewals? What retention investment is justified?

- What is your competitive win rate trend, and at what point does declining win rate become an existential threat?
- If you bundle AI, how do you communicate this to investors without triggering concerns about margin erosion?

Card 3 — Round 3

Title: AI Talent Cost Spiral & Roadmap Execution Constraints

Card Type: Workforce Intelligence

Reveal Timing: Round 3 — distribute face-down at start of decision preparation phase

Classification: Restricted

Source: HR department talent market analysis, engineering leadership assessment, and compensation benchmarking

The Intelligence:

AI engineering talent costs are spiraling beyond your compensation model's capacity. The situation has moved from "competitive pressure" to "structural constraint on execution."

Specific signals:

- Salary inflation: AI engineering compensation (base + equity) has increased 25-30% year-over-year. Your budget assumed 10-15% increases. Senior AI engineers (5+ years experience in ML/AI) are commanding \$400K-\$600K total compensation packages. Your compensation structure — designed for enterprise software engineers — is not competitive without case-by-case exceptions that create internal equity problems.
- Attrition acceleration: You have lost 8 AI engineers in the past two quarters (out of a team of 45). Four went to well-funded AI startups; two went to Big Tech; two went to financial services firms. Each departure creates a 3-6 month productivity gap and a 6-12 month replacement timeline. Your AI feature roadmap is now 2-3 months behind plan, with the gap widening.
- Offer acceptance decline: Your offer acceptance rate for senior AI engineering roles has fallen from 65% to 40% over the past year. Candidates are receiving multiple competing offers and choosing based on equity upside (startups), problem quality (Big Tech/research labs), or compensation (financial services). Your enterprise software brand is less attractive to AI-native talent than consumer tech or pure-AI companies.
- Downstream effects: The engineering capacity constraint is rippling through your product organization. Feature launches are being delayed. Quality issues are emerging as remaining engineers are stretched across too many projects. Customer-facing bugs in AI features are increasing, damaging the perception of AI capability that you are trying to build.
- Services workforce transition: Separately, your customer success and professional services teams (1,500+ employees) are seeing role automation from AI tools. Voluntary attrition in these teams is rising as employees sense their roles are at risk. You face a dual challenge: acquiring expensive AI talent you cannot easily afford while managing the transition of a large existing workforce whose roles are changing.

Decision Tension:

Do you dramatically increase AI talent compensation (breaking your compensation model and creating internal equity issues) to retain and recruit competitively? Or do you accept slower AI roadmap execution and focus on retaining your current team through non-compensation levers (interesting problems, career development, stability) while selectively using M&A and contracting to fill gaps?

Questions to Consider:

- What is the revenue impact of a 3-6 month AI roadmap delay? How does it affect competitive positioning and customer retention?
- Can you restructure compensation specifically for AI engineering roles without creating morale problems across the rest of the engineering organization?
- What is the true cost of replacing an experienced AI engineer (recruitment, ramp-up, lost productivity)? Is retention investment cheaper?
- Can M&A (acqui-hiring AI startups) solve the talent problem faster than organic recruitment? What are the integration risks?
- How do you manage the services workforce transition? What is the cost of proactive retraining vs. reactive restructuring?
- Is there a talent strategy that addresses both sides — acquiring AI talent and retaining/transitioning existing talent — without breaking the budget?

[INDUSTRY NAME] Private Information Cards

FACILITATOR NOTE: *Print this document and separate at page breaks. Distribute one card per round, face-down, at the start of each round's decision preparation phase. Cards are confidential to the [Industry] participant. Cards accumulate — the participant keeps all cards and may refer to them in later rounds.*

Card 1 — Round 1

Title: [Descriptive title — concise, captures the core tension]

Card Type: [Market Intelligence | Competitive Intelligence | Regulatory Development | Operational Intelligence | Risk Event Intelligence | Workforce Intelligence]

Reveal Timing: Round 1 — Decision Preparation Phase

Classification: [Private | Shared with [Industry Name(s)]]

Source: [Internal or external source(s) — e.g., Internal Analytics Team; Competitive Intelligence Unit; Regulatory Affairs]

Shared Intelligence: [OPTIONAL — Include only when this card shares an underlying intelligence base with another industry. Format: "This card shares a common intelligence base with [Industry]. Both industries received [brief description]. [Industry] participants may discover overlapping information during cross-industry discussion."]

The Intelligence:

[Core content. 2-5 paragraphs describing the situation, data, and context. Include quantitative data where relevant but do not force financial precision — directional ranges and qualitative indicators are appropriate. Use bullet lists, bold text, and sub-headers for readability.]

Decision Tension:

[1-2 paragraphs framing the trade-off this intelligence creates. What opposing choices does this information surface? What makes the decision difficult? Written in second person ("Do you...")]

Questions to Consider:

- [Question 1 — specific, actionable, forces the participant to weigh a trade-off]
 - [Question 2]
 - [Question 3]
 - [Question 4]
-

Template Notes (remove in final cards):

- Each industry receives 3 cards (Rounds 1-3). Round 4 has no new private card.
- Cards should create genuine decision tensions — not obvious choices.
- Quantitative data is included where it sharpens the decision but is not required for every card.
- Shared Intelligence notes are included only when cross-industry information overlap exists. Not every card is shared.
- Card Type helps the facilitator and participant categorize the nature of the intelligence. Use the most descriptive type from the list above.

- Page breaks (`<div style="page-break-before: always;"></div>`) separate cards for printing.

Card 2 — Round 2

[Same structure as Card 1]

Card 3 — Round 3

[Same structure as Card 1]

06

Rounds

Round Materials Index

This directory contains all situation updates, injects, decisions, and integration guidance for the four core rounds of the Project Threshold Tabletop Exercise V7.4.

Files

Core Round Materials

File	Time Period	Theme	Decision Type
06Round1_Foundation.md	0–6 Months (Mar–Aug 2026)	Initial AI adoption, early bifurcation, emerging labor concerns	1 industry decision
06Round2_Acceleration.md	6–18 Months (Aug 2026 – Aug 2027)	Dramatic divergence, competitive pressure, visible labor displacement	1 industry decision
06Round3_Reckoning.md	18–30 Months (Aug 2027 – Aug 2028)	Productivity plateau, regulatory tightening, market consolidation	1 industry decision
06Round4_Normalization.md	30–48 Months (Aug 2028 – Early 2030)	AI embedded as infrastructure, valuations normalize, labor stabilizes	1 industry decision
06CrossRound_Integration.md	All Rounds	Spillovers, cascades, cross-industry dynamics, valuation integration, participant quality rubric	Facilitator reference

Round-by-Round Overview

Round 1: Foundation (0–6 Months)

What Happens:

- Foundation models (GPT-5 equivalent) are commercially available

- Early adopters (large, well-capitalized firms) are deploying AI copilots
- Market bifurcation is emerging: winners moving fast, laggards cautious
- Labor displacement is starting to become visible in early-adopting industries
- Regulatory environment is still forming (guidance, not binding rules)

Industry Participant Decision Type (V7.4 Individual Model):

- 1 industry decision per participant
- Theme: "To AI or not to AI? And how fast?"
- Use banded framework: avoid invented precision

Facilitator Focus:

- This is participants' first decision. Scoring feedback is critical for learning the framework
- Time pressure is moderate (not extreme)
- Use this round to surface industry risk tolerance and strategic orientation
- Emphasize: decisions in one industry will cascade to others

Duration: 65 minutes (5 min update + 3 min inject + 15 min individual prep + 3 min submission + 12 min scoring + 27 min cross-industry discussion)

Round 2: Acceleration (6–18 Months)

What Happens:

- Early Round 1 adopters report strong financial results (if execution went well)
- Competitive pressure accelerates. Laggards now face "catch-up or die" pressure
- Labor displacement is more visible and politically sensitive
- Regulatory clarity increases (formal guidance on transparency, audit rights, etc.)
- Market consolidation begins via M&A (larger firms acquiring smaller ones for AI talent/IP)
- AI talent market is now tight; wages inflating

Industry Participant Decision Type (V7.4 Individual Model):

- 1 industry decision per participant
- Theme: "How to respond to divergence? Accelerate, adapt, or defend?"
- Use banded framework; avoid false precision

Facilitator Focus:

- Time pressure increases (competitive dynamics accelerating)
- This is where strategy differences become financially visible
- Emphasize spillovers: decisions in one industry affect others (via labor market, consolidation, regulatory pressure, customer demand)
- Call out cross-industry impacts explicitly in discussion

Duration: 65 minutes

Round 3: Reckoning (18–30 Months)

What Happens:

- First-wave productivity gains from Round 1-2 deployments are now visible but plateau
- Regulatory environment tightens. Binding rules are in place
- Labor displacement is severe and politically visible (millions of jobs)
- Market consolidation is extreme. Mid-market players are being acquired or exiting
- Consumer trust varies widely by industry (low in Healthcare Provider, Healthcare Payer, Finance if not managed; higher in Consulting)
- Tail risk injects hit hard (major AI system failures, market disruptions)

Industry Participant Decision Type (V7.4 Individual Model):

- 1 industry decision per participant
- Theme: "Pivot or persist? Consolidate or stand alone? Manage labor transition or fight?"
- Use banded framework

Facilitator Focus:

- Time pressure is now extreme
- Participants that made good Round 1-2 decisions are winning; those that didn't are desperate
- Tail risk scenarios are now probable, not theoretical
- Collective Bonus accountability is critical (regulatory direction is now set by policy)
- Cross-industry spillovers are now severe (labor market very tight, regulatory pressure cascading, market consolidation extreme)

Duration: 65 minutes

Round 4: Normalization (30–48 Months)

What Happens:

- AI is now embedded infrastructure in all industries
- Market structure is consolidated: winners are clearly defined, losers have exited
- Labor market includes new pools from AI displacement (retraining programs create new talent)
- Valuations normalize: AI productivity gains are reflected in baseline multiples, no longer a surprise
- Regulatory environment is mature: compliance is routine, innovation continues within bounds
- Tail risks from earlier rounds either materialized or were hedged

Industry Participant Decision Type (V7.4 Individual Model):

- 1 industry decision per participant
- Theme: "Final positioning. Exit? Consolidate? Double down? Prepare for next disruption?"
- IMPORTANT: Round 4 has NO PRIVATE CARDS (all information is transparent)
- Use banded framework

Facilitator Focus:

- This is the final strategic bet. Participants are positioning for the 2030+ environment
- Debrief preparation: start noting themes for final debrief discussion on cross-industry dynamics
- Score final decisions and announce final results

- Prepare for 60-minute debrief on cross-industry dynamics, spillovers, and no-regrets actions

Duration: 65 minutes

How to Use Round Materials

Before Exercise

1. Facilitator reads all four round files (06Round1Foundation, 06Round2Acceleration, 06Round3Reckoning, 06Round4Normalization)
 - Understand situation progression across all four rounds
 - Identify key injects for each round
 - Note the escalating complexity and time pressure
2. Facilitator reads 06CrossRound_Integration.md
 - Understand how decisions cascade across rounds
 - Understand spillovers between 11 industries
 - Prepare to emphasize integration throughout exercise and in debrief
3. Facilitator prepares situation update slides using content from round files
 - One slide per round with situation statement
 - Organize inject cards by round
 - Prepare speaking notes for smooth delivery
 - Prepare spillover prompts for cross-industry discussion

During Exercise

1. Practice Round: Use simple, low-stakes situation (typically not in round files; facilitator creates it)
 - Purpose: Participants learn decision format and scoring
 - Participants see banded scoring rubric in action
 - Confidence-building before high-stakes rounds
2. Rounds 1-4: Execute in sequence following 8-hour timeline
 - Read situation update aloud (5 min)
 - Deploy 2-3 injects per round (3 min)
 - Individual decision preparation (15 min; shorter than team deliberation)
 - Decision submission (3 min)
 - Facilitator scoring + Industry Health update (12 min)
 - Cross-industry discussion & results (27 min) — this is the primary discussion period

During Cross-Industry Discussion

Use 06CrossRound_Integration.md rubric to assess participant quality:

- Do participants acknowledge spillovers from other industries?

- Do participants recognize cascades from prior rounds?
- Do participants adapt strategy based on new information?
- Do participants understand how their industry's decision affects others?

In Debrief

Use 06CrossRound_Integration.md prompts:

- "What surprised you about how your industry's Round 1 decision cascaded?"
 - "Which participants/industries recognized spillovers best?"
 - "How did cross-industry dynamics shape your strategy?"
 - "What no-regrets actions would your industry pursue in reality?"
-

Decision Framework Across Rounds

Round 1: Foundation (Establish Position)

- Decision Focus: Strategic orientation (aggressive vs. defensive)
- Constraint: Time (lots available), talent (expensive, available), regulation (ambiguous)
- Scoring Emphasis: Strategic Fit (does this align with your industry?)

Round 2: Acceleration (Respond to Divergence)

- Decision Focus: Pace (accelerate, maintain, slow down)
- Constraint: Time (pressuring), talent (scarce, expensive), regulation (clarifying)
- Scoring Emphasis: Execution Risk (can you actually do this?)

Round 3: Reckoning (Adapt or Defend)

- Decision Focus: Strategic pivot (consolidate, partner, stand alone, exit)
- Constraint: Time (extreme), talent (scarce, very expensive), regulation (binding)
- Scoring Emphasis: Tail Risk (does this hedge downside? create new downside?)

Round 4: Normalization (Final Positioning)

- Decision Focus: Long-term position (2030+ strategy)
 - Constraint: Time (less pressure, but stakes higher), talent (available from retraining), regulation (mature)
 - Scoring Emphasis: Tail Risk + Strategic Fit (is this sustainable? aligned with market structure?)
-

Private Information Cards (Rounds 1-3 Only)

Round 1: Private Card 1 distributed to each participant from `05_Participant_Private_Cards/`

- Card 1 is shared across related industries (e.g., Retail and CPG receive the same card)
- Examples: Market research data, competitive intelligence, regulatory guidance

Round 2: Private Card 2 distributed to each participant from `05_Participant_Private_Cards/`

- Card 2 is unique per industry
- Examples: Earnings data from competitors, customer sentiment, labor market trends

Round 3: Private Card 3 distributed to each participant from `05_Participant_Private_Cards/`

- Card 3 is unique per industry
- Examples: Regulatory enforcement actions, major competitor moves, supply chain disruptions

Round 4: NO PRIVATE CARDS

- All information is transparent
- This creates a level playing field for final decisions

Facilitator Tip: Distribute private cards DURING INJECT PROCESSING TIME (not during deliberation). This gives participants time to absorb before deciding.

Cross-Round Integration

See **06CrossRound_Integration.md** for detailed guidance on:

- How decisions cascade across rounds
- Spillovers between industries
- Valuation integration mechanics
- Participant response quality rubric
- Facilitator checklist for emphasizing integration

Key Concepts:

- Earnings Cascades: R1-R2 decisions → R2-R3 financial results → R3-R4 competitive pressure
 - Regulatory Cascades: R1-R2 ambiguity → R2-R3 clarity → R3-R4 binding enforcement
 - Labor Market Dynamics: R1-R2 talent scarcity → R3 wage inflation + cross-industry talent competition → R4 displacement pool
 - M&A Consolidation: R2 early acquisitions → R3 market consolidation → R4 market structure set
 - Collective Bonus Dynamics: R1-R4 optional peer nomination → accountability pressure → strategy adaptation across all industries
 - Cross-Industry Spillovers: Retail and CPG decisions affect labor/talent → Healthcare Provider and Healthcare Payer affected by wage pressure; Big Tech and B2B/B2C SaaS face engineer shortages; all industries feel ripple effects
-

Timeline Reference

For detailed timeline including practice round and all 4 rounds, see:

- 01_Exercise_Overview/01_Timeline.md

For 8-hour minute-by-minute schedule, see:

- 07_Supporting_Materials/07_Lead_Facilitator_Briefing.md

File Structure

```
06_Rounds/
├── 06_Round_1_Foundation.md           (Situation, injects, decisions, scoring)
├── 06_Round_2_Acceleration.md        (Situation, injects, decisions, scoring)
├── 06_Round_3_Reckoning.md           (Situation, injects, decisions, scoring)
├── 06_Round_4_Normalization.md       (Situation, injects, decisions, scoring)
├── 06_Cross_Round_Integration.md     (Spillovers, cascades, facilitator guidance)
└── 06_README.md                      (This file)
```

Key Facilitator Tips

1. Emphasize Cross-Industry Spillovers: During each cross-industry discussion, explicitly call out how decisions in one industry affect others (labor costs, customer demand, regulatory cascades, talent competition)
2. Track Valuation: Note which industries are outperforming/underperforming; use that in Round X+1 injects
3. Highlight Adaptation: Reward participants that pivot strategy based on new information and recognize spillovers
4. Manage Pressure: Rounds 1-2 are moderate pressure; Rounds 3-4 are high pressure with complex cascades
5. Extend Discussion Time: Round 4 is NOT a truncated round. Use full 27-min cross-industry discussion for deep exploration of no-regrets actions and lessons learned
6. Prepare Debrief: Throughout exercise, note industry strengths and blind spots for debrief discussion, especially regarding cross-industry awareness

Document Version: Project Threshold V7.4 — Round Materials Index **Last Updated:** March 2026
Format: V7.4 single 8-hour format, 11 industries (Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS), individual decision-maker model, Collective Bonus + Facilitator Market Shock, emphasis on cross-industry spillovers

Cross-Round Integration Guide

Part of the Project Threshold Tabletop Exercise V7.4 — Round Materials

Cross-Round Spillovers & Cascades

Throughout Rounds 1-4, the following causal flows are critical for industry participants to track and for facilitators to emphasize:

1. Earnings & Valuation Cascades (All Rounds)

Real earnings data is the most powerful signal. Participants must calibrate their own productivity assumptions against what the market is seeing in reported results. Early-movers who reported strong results in R1-R2 set benchmarks that laggards must match or beat in R3-R4.

What happens:

- Round 1: Participants make aggressive/defensive AI deployment decisions
- Round 2: Early adopters report earnings gains (or disappointments)
- Round 3: Market prices in early performance. Laggards face competitive pressure to match
- Round 4: Valuation multiples reflect realized vs. promised AI productivity

Facilitator move: In Round 2+ share-outs, explicitly call out earnings surprises: "Retail reported 8% margin expansion. Finance reported only 2%. Notice the divergence? That reflects execution quality, not just deployment pace."

2. Regulatory Decisions → Industry-Specific Constraints (R1-R3)

Regulatory clarity (or uncertainty) cascades downstream. Healthcare regulation tightens adoption timelines; finance regulation creates compliance friction; labor regulation constrains workforce optimization. Participants must anticipate and plan for regulatory pressure in their industry.

What happens:

- Round 1: Regulatory landscape is ambiguous. Participants choose fast or cautious
- Round 2: Facilitator Market Shocks and injects clarify regulatory direction
- Round 3: Binding regulatory requirements are in place. Participants pivot strategy
- Round 4: Regulation is part of competitive landscape. Winners are those who built governance early

Facilitator move: When regulatory injects hit, explicitly connect to participant decisions: "In Round 1, you chose aggressive deployment. Regulation has now shifted the landscape. Your strategy must adapt. This is real."

3. Labor Market Dynamics → Talent Scarcity & Wage Inflation (R1-R4)

As AI adoption accelerates, talent becomes scarce and wages rise. Organizations that invested early in AI talent acquisition (R1) have an advantage in R2-R3. Those that delayed face higher costs and lower

skill availability. This is a self-reinforcing dynamic.

What happens:

- Round 1: AI talent is available but expensive. Early investors gain skill advantage
- Round 2: Talent scarcity increases. Wage inflation accelerates
- Round 3: Talent market is tight. Organizations without strong teams are at disadvantage
- Round 4: Labor displacement from AI is visible. Retraining candidates from displaced workers become available (too late for market leaders)

Facilitator move: In Round 2+ injects, emphasize talent constraints: "You want to expand AI deployment, but talent market is tight. You can hire at higher wages, or you can slow deployment and invest in reskilling internal teams. Cross-industry competition for talent is fierce."

4. M&A Consolidation → Market Structure (R2-R4)

Larger firms acquire smaller competitors for talent and IP. This consolidates the market, reduces competitive intensity, and increases barriers to entry for mid-market players. Consolidation trends in Round 2 set up the market structure in Rounds 3-4.

What happens:

- Round 1: Market structure is stable. Participants make organic growth choices
- Round 2: Consolidation begins. Larger companies acquire smaller ones for talent and capability
- Round 3: Market is more concentrated. Mid-market players are consolidation targets or acquired
- Round 4: Market structure is highly concentrated. Winners are the survivors; losers have exited or been acquired

Facilitator move: When M&A happens, emphasize the strategic consequences: "Big Tech just acquired an AI startup for their talent. Now they have enhanced platform capabilities. But they also have integration risk and antitrust scrutiny. The regulatory environment is watching."

5. Consumer Trust & Transparency → Speed of Adoption (R1-R4)

Low consumer trust in high-stakes domains (Healthcare Provider, Healthcare Payer, Finance, Retail) slows adoption and raises compliance costs. Participants that invest early in transparency and governance (R1) build trust and move faster in R2-R4. Those that prioritize speed over governance face regulatory and reputational friction.

What happens:

- Round 1: Participants choose aggressive or transparent approach to AI deployment
- Round 2: Consumer sentiment data emerges. Transparent participants have higher trust scores
- Round 3: Regulatory and consumer pressure targets low-trust approaches. Governance becomes mandatory
- Round 4: Transparency is a competitive advantage. Participants that built trust early are favored by regulation and customers

Facilitator move: In Round 2+ injects, emphasize trust dynamics: "Consumer trust in AI recommendations is at 45%. Industries that disclosed AI usage have higher trust (55%). Those that didn't disclose have lower trust (35%) and now face regulatory pressure."

6. Peer Accountability → Strategic Discipline (R1-R4)

The Collective Bonus creates a lightweight but recurring accountability loop. Every round, participants

may publicly recognize strong strategies or flag risky ones. This shapes behavior: participants know their peers are watching and may nominate.

What happens:

- Round 1: Initial nominations reflect first impressions and declared intent
- Round 2: Nominations begin reflecting execution against Round 1 commitments
- Round 3: Nominations reflect demonstrated track record; cumulative score impact becomes material
- Round 4: Final nominations assess full strategic arc across all rounds

Facilitator move: After each Collective Bonus, connect the results to strategy quality: "Three participants nominated Finance as a risky strategy this round. Why? They cited the mismatch between aggressive automation claims and the regulatory enforcement reality. That's peer accountability in action — and it costs Finance -2 on their cumulative score."

Industry Health Assessment Across Rounds

At the end of each round, facilitators should:

1. Collect participant decision data from decision worksheets
2. Assess each industry's strategic position using the banded scoring framework:
 - Strategic Fit: {-2, 0, +2} — Does decision align with industry fundamentals?
 - Execution Risk: {-2, 0, +2} — Can this be executed?
 - Tail Risk: {-2, 0, +2} — Does this create or hedge risk?
3. Assess how each industry's health outlook is evolving heading into the next round
4. Communicate industry health trajectory to participants so subsequent decisions are made with awareness of how prior choices are shaping their strategic position

Example: Retail participant decides to aggressively automate customer service. Facilitator scores:

- Strategic Fit: 0 (neutral; aligns with margin pressure but creates brand risk)
- Execution Risk: +1 (feasible but requires talent and change management)
- Tail Risk: -1 (creates customer trust risk if not managed well)
- Total: 0 points this round

Then facilitator signals to Retail participant: "Your aggressive automation decision is strategically sound, but risky on execution and customer trust. In Round 2, we'll see if you execute well. Your industry health outlook will weaken if customer trust drops."

Participant Response Quality Rubric

For facilitators evaluating industry participant responses (and communicating feedback):

Strong Response Characteristics:

- Acknowledges both opportunity and risk/execution challenge
- References specific data from injects or prior rounds

- Considers second-order effects (e.g., labor relations, regulatory response, market consolidation)
- Articulates clear decision rationale tied to financial/strategic goals
- Anticipates cross-industry spillovers and adjusts strategy accordingly
- Updates prior assumptions when new data contradicts them
- Recognizes how decisions in one industry create constraints/opportunities for other industries

Example (Strong): "In Round 1, we chose aggressive automation to match competitors' margin expansion. In Round 2, we see labor resistance and customer trust drops. We're now investing in transparency and reskilling programs, accepting slower margin gains but protecting brand and regulatory relationships. This decision also affects Healthcare Provider's talent costs and Finance's headcount reduction strategies. We're coordinating with them on transition timing."

Weak Response Characteristics:

- Overconfident in ability to execute (ignores complexity or risk)
- Dismissive of contrary data ("This won't happen in our industry")
- Ignores labor, regulatory, or consumer sentiment factors
- Makes decisions in isolation without considering cross-industry spillovers
- Repeats R1 strategy in R2-R3 without adjusting for new circumstances
- Confuses correlation with causation (assumes past trend will continue)
- Misses how their decisions cascade to other industries

Example (Weak): "We're doubling down on automation. We don't care about labor sentiment or consumer trust. Our competitors are moving fast, so we move faster. What other industries are doing doesn't affect us."

Individual Industry Decisions (V7.4 Model)

In V7.4, each industry participant submits ONE decision per round. This reflects the individual decision-maker model where each industry is represented by a single decision-maker.

What is an Industry Decision?

- One strategic choice per industry per round
- Decision reflects the industry's response to injects and market conditions
- Use banded framework to avoid false precision
- Decision must address industry-specific challenges and opportunities

Cross-Industry Spillovers: Spillovers propagate across all eleven industries. Example:

- Retail participant makes aggressive automation decision
- This creates labor market spillover affecting all industries' talent costs
- Healthcare Provider and Healthcare Payer face wage pressure as labor becomes scarcer
- Big Tech and B2B/B2C SaaS must compete for engineering talent at higher wages
- Finance faces pressure to defend junior talent
- Consulting sees demand surge for workforce transition advisory
- Facilitator notes in share-out: "Retail's aggressive automation has cascaded into the talent market. Notice how all industries are now facing wage pressure?"

Facilitator Guidance:

- Explicitly announce industry decisions and their spillovers after each round
 - Explain the spillover's competitive implications across all industries
 - Use spillovers as learning moments: "Why did Big Tech face talent constraints after Retail's automation decision? What's the connection?"
-

Spillover Tracking for Eleven Industries

Retail Spillovers — Receives From:

- Finance: Capital availability for expansion; payment processing dynamics
- Consulting: AI deployment strategy advisory
- Manufacturing and Logistics: Supply chain optimization affects customer fulfillment speed
- Healthcare Provider and Healthcare Payer: Trust/regulatory dynamics in healthcare set consumer expectations for data security
- Big Tech: AI capabilities and pricing available for deployment; competitive platform offerings
- B2B/B2C SaaS: AI-embedded tools available for retail operations
- CPG: Demand signals; shared supply chain dynamics

CPG Spillovers — Receives From:

- Retail: Demand signals; channel dynamics; shared supply chain infrastructure
- Manufacturing and Logistics: Supply chain and distribution optimization
- Consulting: AI deployment strategy advisory
- Big Tech and B2B/B2C SaaS: AI capabilities for demand forecasting and marketing

Healthcare Provider Spillovers — Receives From:

- Healthcare Payer: Reimbursement dynamics; shared regulatory environment
- Finance: Capital availability for clinical AI investment
- Consulting: Compliance consulting; AI deployment advisory
- Law: Regulatory advisory; malpractice and liability guidance
- Big Tech and B2B/B2C SaaS: AI capabilities for clinical decision support; talent competition for data scientists
- Retail: Consumer trust in AI affects patient willingness to accept AI-assisted care
- Manufacturing and Logistics: Medical device and pharmaceutical supply chain optimization

Healthcare Payer Spillovers — Receives From:

- Healthcare Provider: Clinical outcomes data; shared regulatory environment

- Finance: Capital availability; shared regulatory dynamics in financial services
- Consulting: Compliance consulting; claims automation advisory
- Law: Regulatory advisory; compliance guidance
- Big Tech and B2B/B2C SaaS: AI capabilities for claims processing; talent competition

Finance Spillovers — Receives From:

- Consulting: Compliance advisory; AI deployment strategy
- Law: Regulatory and litigation advisory
- Retail: Customer demand for AI-driven financial services
- Healthcare Provider and Healthcare Payer: Regulatory constraints on healthcare create consulting demand
- Manufacturing and Logistics: Supply chain disruption affects financial operations
- Big Tech and B2B/B2C SaaS: AI capabilities for trading, underwriting, and service delivery; talent competition

Consulting Spillovers — Receives From:

- All Other Industries: Demand for AI advisory, implementation, and compliance consulting
- Finance: Demand for AI-driven financial advisory and compliance
- Healthcare Provider and Healthcare Payer: Demand for compliance and governance consulting
- Law: Shared professional services dynamics; regulatory advisory partnerships
- Big Tech and B2B/B2C SaaS: AI tools for delivery augmentation; talent competition

Law Spillovers — Receives From:

- All Other Industries: Demand for regulatory advisory, litigation defense, and compliance review
- Finance: Demand for financial regulatory advisory
- Healthcare Provider and Healthcare Payer: Demand for healthcare regulatory and malpractice advisory
- Consulting: Shared professional services dynamics; partnership on compliance engagements
- Big Tech and B2B/B2C SaaS: AI tools for document review and legal research; talent competition

Manufacturing Spillovers — Receives From:

- Logistics: Shared supply chain dynamics; distribution optimization
- Retail and CPG: Demand volatility and product mix changes affect production planning
- Finance: Capital for infrastructure investment
- Consulting: AI deployment advisory for manufacturing operations
- Big Tech and B2B/B2C SaaS: AI capabilities for production optimization; talent competition

Logistics Spillovers — Receives From:

- Manufacturing: Production scheduling and output affect logistics planning
- Retail and CPG: Demand volatility and fulfillment requirements
- Finance: Capital for fleet and infrastructure investment
- Consulting: AI deployment advisory for logistics operations
- Big Tech and B2B/B2C SaaS: AI capabilities for route optimization and demand forecasting; talent competition

Big Tech Spillovers — Receives From:

- All Other Industries: Demand for cloud, enterprise AI platforms, ads services, and device ecosystem capabilities
- Retail and CPG: Consumer demand for personalization and optimization capabilities
- Healthcare Provider and Healthcare Payer: Regulatory requirements for clinical AI; talent competition
- Finance: Demand for AI in trading, underwriting, and compliance workflows
- Consulting and Law: Demand for AI in professional service delivery
- Manufacturing and Logistics: Demand for logistics and production optimization AI
- B2B/B2C SaaS: Competitive dynamics in enterprise software; shared talent pool

B2B/B2C SaaS Spillovers — Receives From:

- All Other Industries: Demand for AI-embedded software products
- Retail and CPG: Consumer demand for AI-powered tools
- Healthcare Provider and Healthcare Payer: Demand for clinical and claims management AI; compliance requirements
- Finance: Demand for AI in financial workflows
- Consulting and Law: Demand for AI in professional service delivery
- Manufacturing and Logistics: Demand for production and logistics optimization tools
- Big Tech: Competitive dynamics in enterprise software; platform dependency; shared talent pool

Round-by-Round Integration Checklist

After Each Round, Facilitator Should:

During Share-Out Phase:

- Listen for first, second, and third-order effects participants mention (or miss)
- Note which participants recognize spillovers from other industries
- Note which participants ignore Collective Bonus feedback (warning sign of weaker strategy)
- Track cross-industry dynamics: How does one industry's decision affect others?

During Scoring Phase:

- Assign higher scores to participants that anticipate spillovers
- Assign lower scores to participants that ignore cross-industry dynamics
- Explicitly call out in feedback: "You scored well on Strategy, but missed the labor market consequence of your automation decision. That impacts Execution Risk across all industries."

Between Rounds:

- Assess industry health trajectory based on all Round X decisions
- Note which industries are outperforming/underperforming expectations
- Prepare injects for Round X+1 that emphasize spillovers from Round X decisions
- Track how one industry's decision creates constraints or opportunities for others

In Debrief (After Round 4):

- Ask participants: "What surprised you about how your industry's Round 1 decision cascaded through Rounds 2-4?"
- Ask: "How did cross-industry spillovers change your strategy?"
- Highlight participants that adapted well (recognized spillovers, pivoted strategy)
- Highlight participants that missed spillovers (got blindsided by consequences)

Key Themes to Emphasize in Debrief

1. Cross-industry spillovers are real: In the exercise, as in reality, decisions don't stay siloed. Your industry's AI deployment strategy affects talent market, regulatory landscape, consumer trust, market consolidation, and all other industries.
2. Adaptation is required: Participants that made one decision and stuck with it underperformed. Participants that adapted strategy based on new information (injects, competitor moves, Collective Bonus feedback, cross-industry impacts) outperformed.
3. Governance pays off: Participants that invested in transparency, governance, and stakeholder management in early rounds had easier time adapting in later rounds and faced less regulatory/labor friction.
4. Execution beats strategy: A good strategy executed poorly scored lower than a mediocre strategy executed well. The gap was Execution Risk scoring.
5. Peer accountability matters: Participants that ignored Collective Bonus feedback were blindsided when cumulative score penalties compounded. Participants that listened to peer nominations and adapted their strategies built competitive advantages.
6. Cross-industry awareness: Participants that recognized how their decisions cascaded to other industries (labor markets, talent competition, customer demand, regulatory pressure) and coordinated or adapted accordingly outperformed those that made isolated decisions.

Document Version: Project Threshold V7.4 — Cross-Round Integration Guide **Last Updated:** March 2026 **Format:** Operational guidance for facilitators; emphasizes 11-industry dynamics, Collective Bonus accountability, and cross-industry spillovers

Round 1: Foundation

Part of the Project Threshold Tabletop Exercise V7.4 — Consolidated Round Materials

Round 1: FOUNDATION

0–6 Months from Exercise Start (~March–August 2026)

A. ROUND OVERVIEW

Five to eleven industry participants converge for Round 1. The AI landscape has matured significantly—foundation models are reaching 70-80% performance on standardized exams, early automation is affecting 40% of white-collar task volume, and equity markets have rallied (S&P +12% YTD). Yet labor concerns are emerging visibly. Each industry participant must decide how to navigate this inflection point: accelerate adoption, pilot cautiously, or wait for regulatory clarity.

Time Period: 0–6 months (March–August 2026) **Key Theme:** Foundation and early deployment **What Participants Face:** Capability announcements, industry-specific productivity data, nascent regulatory signals, and early visible labor displacement

B. SITUATION UPDATE

(To be read aloud by facilitator at round open.)

Welcome to Spring 2026. The AI landscape has matured significantly, and economic implications are becoming impossible to ignore.

Model and Capability Announcements: The leading frontier labs have released increasingly capable copilot systems. We're seeing reliable multi-step task automation in customer service, content creation, and routine business process management. Foundation models are reaching 70-80% performance on standardized exams; competitor models are narrowing the capability gap. However, no "artificial general intelligence" milestone has been achieved—we remain in the "very capable narrow AI" era. Latency and cost have both improved, making real-time deployment feasible.

Enterprise Adoption: Pilots from 2024-2025 are graduating to production. Companies with dedicated AI teams are moving past proof-of-concept and seeing measurable productivity gains (15-25% in software and finance; more modest in traditional industries). Smaller and mid-market firms are moving cautiously, struggling with integration costs. Early instances of "AI buyer's remorse" are appearing—enterprises that rolled out copilots without proper change management are seeing limited adoption or high switching costs.

Labor Market Signals: Hiring growth in routine cognitive roles (data entry, junior analysis, claims processing) has visibly slowed. Entry-level hiring is down 40% in some industries. Conversely, demand for "AI-adjacent" roles (prompt engineers, AI training specialists, change management consultants) is booming. A few high-profile companies have announced modest headcount reductions (3-8% in specific divisions) with AI productivity explicitly cited. Unemployment remains at 4%, but job displacement is concentrated in specific geographies and demographics. Union activity around AI is emerging, particularly in consulting, law, and logistics.

Valuation and Capital Markets: AI-beneficiary stocks have rallied strongly (S&P +12% YTD). The "Magnificent AI" subset has outperformed. Industries with high displacement potential have seen selective weakness. VC deal flow for AI startups remains robust, but focus has shifted from model companies to "implementation and integration" plays. Enterprise software multiples have not collapsed but have stabilized. Regulatory conversation is nascent but clarifying: transparency, auditability, and disclosure will be demanded.

Regulatory Developments: Congress has held hearings on AI in financial services and healthcare. The SEC, FDA, and CFTC all have active AI task forces. No major binding regulation has passed, but the direction is clear: "responsible AI" frameworks are coming. State regulations (California, New York) are advancing. Regulators are concerned about concentration risk, labor displacement, and systemic financial risk.

Consumer Sentiment: Polling shows stark divergence: 68% trust AI for retail personalization; only 41% trust it for healthcare; only 35% for financial decisions. Across all industries, 72% demand transparency ("Tell me if I'm talking to an AI"). Primary concern is "job displacement"—50% of respondents worry AI will reduce job availability. When companies disclose AI use transparently and describe human oversight, trust increases 15-25 points.

C. CORE INJECTS (2-3 Maximum)

R1-01: Copilot Displacement Spike

Classification: Market Signal / Capability Announcement **Time:** Opening (start of round)

Narrative: A major frontier AI lab announces a new copilot generation capable of reliably executing multi-step customer service interactions with minimal human oversight. Early customer case studies show 60-70% of tier-1 and tier-2 support tickets resolved autonomously. The system handles refunds, returns, billing disputes, and escalation routing. Enterprise availability is immediate; pricing is per-interaction. A Fortune 500 retail customer is prominently featured. Stock markets react positively to beneficiary companies; some BPO and call-center stocks dip 3-5%. Within weeks, entry-level hiring is down 40% across industries—the visible labor displacement signal arrives.

Ambiguity/Unknowns:

- Real-world performance on complex edge cases remains untested at scale.
- True cost of ownership is non-obvious.

Industry Impact:

Industry	Direct Impact	Implied Action
Retail	Immediate cost pressure in customer service; displacement in entry-level roles	Accelerate deployment or face competitive disadvantage
CPG	Downstream demand signal effects; marketing and consumer engagement automation	Evaluate AI-driven demand forecasting and personalization
Healthcare Provider	Affects clinical workflow automation and patient scheduling	Opportunity to optimize operations while regulators are watching
Healthcare Payer	Back-office claims processing and compliance automation	Accelerate claims automation to maintain margins
Finance	Knowledge work automation; junior analyst and trading roles under pressure	Shift business model or rationalize headcount
Consulting	Client advisory delivery automation; junior consultant roles under pressure	Accelerate AI-augmented delivery or risk margin erosion
Law	Document review, due diligence, and contract automation	Deploy AI tools or face cost disadvantage vs. competitors
Manufacturing	Quality control and production scheduling automation; entry-level hiring pressure	Deploy AI to maintain margins and quality
Logistics	Customer service and logistics routing automation; entry-level hiring pressure	Deploy AI to maintain margins
Big Tech	Immediate pressure to integrate AI features into cloud, enterprise software, ads, and device platforms; competition from AI-native startups	Accelerate feature integration or lose competitive position
B2B/B2C SaaS	Pressure to embed AI into product offerings; competition from AI-native startups	Accelerate AI feature development or face churn

Facilitator Guidance:

- Early deployments show ~15-20% silent failure rate on edge cases.
- Pricing is lower than traditional BPO but requires upfront infrastructure; total cost of ownership is non-obvious.
- The flagship customer experienced 30% headcount reduction in pilot divisions but also faced 5% increase in escalations due to customer frustration with AI limitations.

- Within 8-10 weeks, competitors will announce similar capabilities. First-mover window is short.

R1-02: Regulatory Signal

Classification: Regulatory Signal / Policy Coordination **Time:** Mid-round (Day 1 afternoon or Day 2 morning)

Narrative: Three regulatory bodies issue concurrent signals. The SEC proposes guidance on AI-related disclosure and risk management for public companies. A Senate Finance Committee hearing features testimony from both banking advocates (citing efficiency gains) and consumer advocates (citing algorithmic discrimination and labor displacement). The FDA publishes draft guidance on AI-assisted clinical decision support, requiring manufacturers to demonstrate safety, transparency, and clinician oversight mechanisms. The message is clear: transparency, auditability, and disclosure will be demanded—but the timeline and enforcement intensity remain uncertain. Healthcare Provider, Healthcare Payer, and Finance participants face immediate pressure to assess compliance gaps.

Ambiguity/Unknowns:

- Timeline for finalization: months or years?
- Enforcement intensity: advisory or binding?

Industry Impact:

Industry	Direct Impact	Implied Action
Healthcare Provider	Direct regulatory scrutiny; need for governance structures and disclosures	Assess compliance gaps; proactively build governance
Healthcare Payer	Direct regulatory scrutiny; claims processing AI under review	Assess compliance gaps; prepare governance frameworks
Finance	Direct regulatory scrutiny; compliance risk	Assess compliance gaps; prepare governance frameworks
Consulting	Secondary impact via client demand for compliance advisory	Prepare compliance consulting offerings
Law	Secondary impact via demand for regulatory advisory and compliance review	Prepare regulatory advisory offerings
Retail	Secondary impact via customer trust and transparency demands	Prepare disclosure mechanisms
CPG	Minimal direct impact in this round	Monitor for indirect effects via retail channel
Manufacturing	Minimal direct impact in this round	Monitor for indirect effects
Logistics	Minimal direct impact in this round	Monitor for indirect effects

Industry	Direct Impact	Implied Action
Big Tech	Secondary impact via government scrutiny of cloud and enterprise AI platforms	Prepare for transparency and auditability requirements
B2B/B2C SaaS	Secondary impact via customer compliance requirements for AI-embedded products	Prepare for transparency and auditability requirements

Facilitator Guidance:

- SEC guidance will likely require disclosure of model performance, training data provenance, and bias risks. Enforcement bar has risen.
- Senate hearing will attract media. Two senators will call for a "financial services AI transparency act." Industry will testify against "innovation-killing regulation."
- FDA guidance is non-binding but is expected to become de facto standard. Any manufacturer deviating will face skepticism.
- One large healthcare system is lobbying FDA behind the scenes for looser requirements; this will leak within 2-3 months.

R1-03: Market Signal

Classification: Market Signal / Competitive Dynamics **Time:** Mid-round (concurrent with or slightly after R1-02)

Narrative: Q1 2026 earnings season and venture funding data converge to send a clear signal. Large, well-capitalized enterprises that deployed AI in late 2025 are reporting visible productivity gains (15-25% in software and finance). Cloud providers cite 3-4 percentage points of margin expansion from AI infrastructure optimization. Consulting firms report 22% increase in billable hours per consultant year-over-year. Financial institutions cite 150 bps of operating margin improvement. However, one major software company misses guidance—attributed to slower enterprise AI adoption than expected; the stock drops 15%. The divergence is acute: proven AI adopters are beating expectations; others are disappointing. Wall Street begins explicitly ranking companies by "AI execution capability" and incorporating AI adoption maturity into valuation models. Venture capital is consolidating into mega-funds; smaller VC firms are struggling to raise capital.

Ambiguity/Unknowns:

- Is the productivity improvement sustainable, or does it plateau as low-hanging fruit is exhausted?
- Are valuation premiums for "AI winners" justified by fundamentals, or is a bubble forming?

Industry Impact:

Industry	Direct Impact	Implied Action
Retail	Market is measuring productivity gains; laggards face valuation compression	Demonstrate AI-driven margin improvement or face investor scrutiny
CPG	Supply chain and demand forecasting gains being measured	Publicize AI-driven efficiency to avoid valuation compression
Healthcare Provider	Compliance costs and governance investments may dampen near-term margins; slower adoption	Balance regulatory caution with investor expectations for AI gains
Healthcare Payer	Claims processing automation gains attracting attention	Demonstrate AI-driven cost reduction to maintain multiples
Finance	Market comparing financial services efficiency; aggressive AI transformation rewarded	Demonstrate AI-driven margin expansion
Consulting	Billable hour gains attracting attention; delivery model transformation	Demonstrate AI-augmented delivery efficiency
Law	Slower adoption but efficiency gains in document review being tracked	Show productivity improvements in key practice areas
Manufacturing	Capital efficiency and production optimization attracting attention	Publicize AI-driven efficiency to avoid valuation compression
Logistics	Logistics optimization gains being measured	Publicize AI-driven efficiency to avoid valuation compression
Big Tech	Market expects rapid AI feature integration into cloud, ads, devices, and enterprise platforms; laggards face valuation compression	Demonstrate AI-driven growth and margin expansion
B2B/B2C SaaS	Market expects AI-embedded products; laggards face churn and valuation compression	Demonstrate AI-driven growth and retention

Facilitator Guidance:

- This inject validates the scenario's "S&P +12%" figure.
- Participants that chose aggressive R1 strategies should prepare to show earnings aligned with announced productivity in Round 2.
- Participants that chose conservative approaches should show minimal AI impact in Round 2 earnings (neutral or slightly negative).
- The "AI execution capability" ranking becomes a real analyst talking point. Market rotates toward winners.

- One company that "missed guidance" signals that not all organizations can execute transformation equally. Market rewards execution; punishes stumbles.

D. OPTIONAL INJECTS (Use Only If Table Finishes Early)

R1-OPT-01: Insider Trading Indictment

Classification: Legal / Regulatory / Reputational **Time:** Late round (Day 2 afternoon)

Narrative: The SEC announces an indictment against a former executive of a major AI company for insider trading. The executive, based on non-public knowledge of model capability delays, sold shares before a public announcement that disappointed investors. The indictment is widely reported and raises questions about information asymmetry in AI development and the reliability of company guidance. The incident reflects growing investor attention to AI capability announcements and the risk of overhyped claims. Market reaction is modest, but the reputational damage to the company is material. Insurance companies and governance advisors begin offering "AI disclosure risk" insurance.

Ambiguity/Unknowns:

- Will this become a precedent for other prosecutions?
- How will it affect investor confidence in AI company guidance?

Industry Impact:

Industry	Direct Impact	Implied Action
All Industries	Investor confidence in AI ROI claims is shaken; disclosure liability increases	Tighten internal governance on AI capability claims; avoid overhyped guidance

E. COLLECTIVE BONUS (Optional, within Cross-Industry Discussion)

Facilitator Script (Read Aloud):

"Does anyone want to recognize an especially strong strategy this round, or flag one that seems particularly risky?"

Procedure:

Step	Action	Time
1	Participation is optional — no one is required to speak	—
2	Participants who wish to respond name one industry as "strong strategy" and/or one as "risky strategy" (cannot nominate own industry)	3 min
3	If 3+ participants agree on the same industry, facilitator applies adjustment	1 min

Scoring:

- If 3+ participants agree an industry has a strong strategy: +2 cumulative score bonus
- If 3+ participants agree an industry has a risky strategy: -2 cumulative score penalty
- Maximum one +2 and one -2 per round
- If no consensus or no one speaks up, no bonus applied

Facilitator Guidance:

- The Collective Bonus creates lightweight peer accountability without forcing participation.
- Encourage participants to cite specific reasoning if they do nominate.
- Announce results openly — nominations are public, not anonymous.
- In Round 1, nominations often reflect initial impressions. Later rounds will reflect track record.

F. PRIVATE CARD DISTRIBUTION

Facilitator Note: Distribute the appropriate Private Information Card from [05 Participant Private Cards](#) to each participant face-down. Round 1 uses Private Card 1. Card 1 is shared across related industries (e.g., Retail and CPG receive the same card; Healthcare Provider and Healthcare Payer receive the same card; Finance, Consulting, and Law receive the same card; Manufacturing and Logistics receive the same card; Big Tech and B2B/B2C SaaS receive the same card). Private cards contain industry-specific data, internal metrics, or confidential intelligence relevant to each participant's decisions in this round. These cards are intentionally asymmetric—participants have incomplete information about other industries, forcing reliance on public inferences and cross-industry inference.

Private cards are used in Rounds 1-3 only. Round 4 operates on symmetric information.

G. INDIVIDUAL INDUSTRY DECISION SUBMISSION REQUIREMENTS — V7.4

Each industry participant submits:

1. Industry Decision (1 REQUIRED):
 - Each participant must submit an explicit decision for their industry per round.
 - Use banded framework: avoid invented precision, use directional language.
 - Participant states their choice, rationale using bands, and expected outcome.
 - The decision reflects the industry's strategic response to the injects and market conditions.

Submission Format:

[INDUSTRY NAME] – ROUND 1 DECISION

DECISION TITLE: [Title]

RATIONALE:

[Why this choice? How does it respond to injects and your private card?]

BANDED ASSESSMENT:

- Spend/Commitment: [Band]
- Time-to-Impact: [Band]
- Execution Complexity: [Band]
- Dependency: [Band]
- Scale: [Band]

EXPECTED OUTCOME:

[What do you expect in Round 2 earnings and industry conditions?]

INDUSTRY REPRESENTATIVE: [Name]

Time to Submit: 18 minutes after injects conclude.

Key Notes:

- One decision per participant per round.
- Use banded framework to avoid false precision.
- The decision should reflect the industry's response to the injects and cross-industry dynamics you anticipate.

H. INDUSTRY HEALTH BASELINE — ROUND 1

All industries start at cumulative score 0 (Steady condition). Industry Health conditions will be announced for the first time at the start of Round 2, based on Round 1 scoring outcomes.

End of Round 1 Materials

Round 2: Acceleration

Part of the Project Threshold Tabletop Exercise V7.4 — Consolidated Round Materials

Round 2: ACCELERATION

6–18 Months (~August 2026 – August 2027)

A. ROUND OVERVIEW

Twelve months have passed. AI momentum is visible now in earnings, employment data, and market valuations. The bifurcation between early adopters and laggards has accelerated sharply. Agentic systems are moving from lab demo to limited production. Autonomous trading, AI-assisted hiring, and autonomous logistics are beginning to show visible economic impact. The regulatory environment is hardening. Labor unrest is mounting. This round tests whether each industry participant's Round 1 decisions deliver returns or backfire.

Time Period: 6–18 months (August 2026 – August 2027) **Key Theme:** Acceleration and visible differentiation between winners and losers **Macro Scenario:** Agentic AI at 60% autonomous trading; AI recruitment filtering 80% of applications; autonomous logistics at 30% of long-haul routes; back-office unemployment at 8%; AI Accountability Act advancing; deepfake fraud \$3B; Nasdaq down 18% from highs; market volatility increasing.

B. SITUATION UPDATE

(To be read aloud by facilitator at round open.)

Twelve months have passed. The AI momentum is unmistakable in earnings, employment data, and market valuations. The bifurcation that emerged in Round 1 has accelerated sharply.

Model Capabilities & Competitive Landscape: Frontier labs have released new model generations with meaningfully improved reasoning, multi-modal capabilities, and latency suitable for real-time interaction. Code generation is now reliable enough that junior developer productivity has visibly increased. Agentic systems (AI that can autonomously execute multi-step tasks with periodic human validation) are moving from lab demo to limited production in specialized domains: data analysis, customer service, supply chain optimization, autonomous trading at 60% capability. A notable open-source release has narrowed the capability gap with frontier models for many applications, triggering competitive price pressure and commoditization. The conversation has shifted from "Will AI be impactful?" to "Who will capture the value?"

Enterprise Adoption & Financial Results: Earnings season (Q2-Q3 2027) reveals stark divergence.

Large, well-capitalized enterprises that deployed AI in Round 1 are reporting measurable financial benefits: productivity gains of 15-25%, margin expansion of 100-200 bps in some cases. Software companies and financial institutions are leading; they are guiding higher FY2027 earnings driven explicitly by AI productivity. However, a subset of enterprises—those that deployed AI without adequate change management or governance—are reporting integration failures, productivity gains that failed to materialize, and writedowns of AI-related investments. The divergence is driving a "winner-take-most" dynamic: capital and talent are flowing to proven AI adopters, while laggards are struggling to catch up.

Labor Market Transformation: The labor market is showing visible stress. Certain occupational categories—data entry, junior analysis, routine coding, claims processing, customer service agents—are declining in headcount. Unemployment remains around 4% in aggregate, but job displacement is concentrated in specific geographies (back-office unemployment now at 8%) and demographics (less-educated workers, those in routine cognitive roles). Union activity has intensified: strikes and organizing campaigns have occurred at two major tech companies and one large consulting firm, explicitly around "AI and jobs." Calls for government support (retraining, wage insurance) are increasing from labor groups and progressive lawmakers. Three states (California, New York, Washington) have begun AI-era workforce transition planning. Consumer advocacy groups are launching campaigns against "AI layoffs."

Valuation & Capital Markets: AI-beneficiary stocks have continued to rally, but more selectively. Investors are increasingly discriminating between proven winners (cloud providers, large enterprise software, financial services companies with AI deployment success) and speculative plays. Mega-cap tech companies have hit valuation multiples unseen since the dot-com era. Venture capital is consolidating: mega-funds are leading large Series B/C rounds, while smaller VC firms are struggling to raise capital. The "hot" part of the venture market is now "AI adoption enablers" (consulting, integration, training, change management), not pure model or data companies. Nasdaq is down 18% from highs due to profit-taking and growing concerns about profitability timelines.

Regulatory Landscape Crystallizing: Regulatory certainty is increasing. The SEC has finalized guidance on AI disclosure and risk management. The FDA has finalized clinical AI guidance; several healthcare systems have deployed diagnostics under these guidelines. The "AI Accountability Act" is advancing in Congress—it would require explainability, bias auditing, and third-party audit trails for AI systems used in high-stakes decisions. The EU's AI Act is now in enforcement phase, creating compliance costs and precedent-setting litigation. State regulations are proliferating. Notably, no major federal labor-specific AI legislation has passed, but several proposals are in Congressional committees with bipartisan interest.

Fraud & Security Threats: Deepfake fraud has emerged as a significant threat. The first coordinated deepfake fraud ring was detected, targeting financial institutions and retail companies. Losses to deepfake fraud are estimated at \$3B globally in 2027. Regulatory and law enforcement response is intensifying. Authentication technologies (digital signatures, blockchain, biometric verification) are becoming a growth market, but adoption lags behind the problem. Media coverage is sensational, amplifying public anxiety about AI-generated disinformation.

B1. AI ADOPTION ARC DISTRIBUTION

Facilitator Note: Distribute Phase 2 (Acceleration) from [04_Participant_AI_Adoption_Arcs/](#) to each participant. This provides each industry's updated AI adoption trajectory based on Round 1 decisions and macro developments.

C. CORE INJECTS (2-3 Maximum)

R2-01: Regulatory Backlash

Classification: Regulatory Signal / Enforcement Action **Time:** Opening (start of round)

Narrative: Three concurrent regulatory actions land on participants simultaneously. The SEC launches trading audits on AI-driven autonomous trading systems, questioning risk controls and backtesting methodologies. The FDA delays approvals for two AI diagnostic systems, citing inadequate real-world validation and bias testing. A coalition of state attorneys general announce coordinated enforcement actions against financial institutions using AI credit scoring systems without adequate bias auditing. The message is unmistakable: regulators are shifting from "guidance" to "enforcement." Healthcare Provider, Healthcare Payer, and Finance participants face immediate cost pressure and execution delays. Compliance budgets double or triple. The window for aggressive, loosely-governed AI deployment has closed.

Ambiguity/Unknowns:

- How severe will enforcement actions be? Fines? Forced system shutdowns?
- Will enforcement slow or accelerate AI innovation in regulated industries?

Industry Impact:

Industry	Direct Impact	Implied Action
Healthcare Provider	Direct regulatory scrutiny; enforcement actions; compliance costs spike	Pause new deployments; conduct emergency audits; implement governance
Healthcare Payer	Direct regulatory scrutiny on claims processing AI; compliance costs spike	Pause new deployments; conduct emergency audits; implement governance
Finance	Direct regulatory scrutiny; enforcement actions; compliance costs spike	Pause new deployments; conduct emergency audits; implement governance
Consulting	Surge in demand for compliance advisory services	Capitalize on compliance consulting demand
Law	Surge in demand for regulatory advisory and litigation defense	Capitalize on regulatory advisory demand
Retail	Secondary impact via cost of compliance passed to consumers	Pricing may increase for AI-driven services
CPG	Minimal direct impact	Monitor for contagion effect
Manufacturing	Minimal direct impact	Monitor for contagion effect in other regulated domains
Logistics	Minimal direct impact	Monitor for contagion effect in other regulated domains

Industry	Direct Impact	Implied Action
Big Tech	Secondary impact via customer compliance requirements for cloud and enterprise AI platforms	Prepare for customer demand for governance frameworks
B2B/B2C SaaS	Secondary impact via customer compliance requirements for AI-embedded products	Prepare for customer demand for governance frameworks

Facilitator Guidance:

- SEC trading audits will force firms to document model performance, backtesting, and stress-testing methodologies. Several firms will be forced to reduce AI trading exposure.
- FDA delays will set back clinical AI deployments by 6-12 months in some cases.
- State enforcement actions will trigger a wave of bias audits and remediation efforts across financial services.
- Compliance costs will increase 200-300% for firms that aggressively deployed AI without governance. This is a major penalty for Round 1 aggressive strategies.
- Firms that built proactive governance in Round 1 will face minimal friction; they will be positioned as "responsible actors" and gain competitive advantage over penalized peers.

R2-02: Market Stress Event

Classification: Market Signal / Macro Shock **Time:** Mid-round (Day 1 afternoon or Day 2 morning)

Narrative: A "flash crash" event occurs in early August 2027, triggered by a cascade of AI trading systems executing simultaneously in response to a volatility spike. The crash itself is brief (30-45 minutes), but the reputational damage is significant. Nasdaq drops 4% intraday; circuit breakers halt trading twice. Regulators issue emergency statements questioning the safety of autonomous trading systems. Several pension funds and asset managers announce they are reviewing their AI trading exposure. The Nasdaq finishes the trading season down 18% from July highs. Investor confidence in AI-driven market infrastructure is shaken. Market volatility increases significantly. Credit markets tighten slightly. This is the first visible "black swan" event attributed to AI systems, and it triggers a broad reassessment of AI risk across the board.

Ambiguity/Unknowns:

- How quickly will market confidence in AI systems recover?
- Will this trigger broader regulatory constraints on autonomous trading?

Industry Impact:

Industry	Direct Impact	Implied Action
Finance	Direct impact; trading losses; reputational damage; new scrutiny	Pause autonomous trading expansion; increase human oversight
Consulting	Surge in demand for AI risk management advisory	Capitalize on risk consulting demand
Law	Surge in demand for regulatory and litigation advisory	Capitalize on litigation and advisory demand
Retail	Secondary impact via market volatility; consumer confidence may decline	Reassure customers about safety and oversight
CPG	Secondary impact via market volatility	Monitor consumer spending effects
Manufacturing	Secondary impact via credit market tightening; capex may become more expensive	Monitor financing availability for R2 expansion plans
Logistics	Secondary impact via credit market tightening	Monitor financing availability for expansion plans
Healthcare Provider	Secondary impact via credit market tightening	Monitor financing availability for expansion plans
Healthcare Payer	Secondary impact via credit market tightening	Monitor financing availability for expansion plans
Big Tech	Secondary impact via investor sentiment on AI valuations for cloud and enterprise platforms	Monitor for valuation compression
B2B/B2C SaaS	Secondary impact via investor sentiment on AI valuations	Monitor for valuation compression in AI-heavy segments

Facilitator Guidance:

- The flash crash is not a catastrophic failure—AI systems recovered gracefully and circuit breakers worked as designed. But the psychological impact is large.
- Finance participants that have AI trading systems will see immediate margin pressure and forced risk reductions.
- Nasdaq down 18% from highs signals broader profit-taking and growth concerns. Valuations may compress across the board, not just AI-heavy industries.
- This is a "black swan" that was theoretically predictable but practically ignored by risk managers. It raises questions about whether other AI systems have similar hidden risks.
- Consulting participants will see demand surge for "AI risk management consulting." This is a new revenue stream.

R2-03: Labor and Social Backlash

Classification: Labor/Social Signal / Union Action **Time:** Mid-round (concurrent with or just before R2-02)

Narrative: Workers at a large technology company strike over AI-related grievances regarding invasive monitoring, speed-up requirements, and job security threats. The strike settles after 3 weeks with: (1) Transparency on AI-assisted tasks; (2) Human oversight requirements for employment decisions; (3) \$10M retraining fund. The settlement becomes a precedent. Three other companies face union organizing campaigns explicitly around AI. Progressive politicians call for federal AI labor standards; business groups counter that overregulation will stifle productivity gains. Consumer advocacy groups launch campaigns against "AI layoffs." Multiple newspapers run investigative stories on AI-driven job displacement in specific communities. Public sentiment on AI shifts noticeably. Retail workers, call-center workers, and back-office workers are organizing.

Ambiguity/Unknowns:

- Will the settlement become a template for other negotiations?
- Will political pressure for federal AI labor standards build into binding law?

Industry Impact:

Industry	Direct Impact	Implied Action
Retail	High risk: customer service and back-office workforce organizing; retraining costs spike	Proactively engage unions; announce transparent workforce transition plans
CPG	Moderate risk: manufacturing and distribution workforce organizing	Prepare for labor negotiations; budget for retraining
Manufacturing	High risk: production and back-office workers organizing; automation may face labor opposition	Prepare for labor negotiations; budget for retraining
Logistics	High risk: warehouse and logistics workers organizing; automation may face labor opposition	Prepare for labor negotiations; budget for retraining
Healthcare Provider	Secondary risk: back-office and clinical staff organizing	Monitor labor market; prepare for wage pressures
Healthcare Payer	Secondary risk: claims processing staff organizing	Monitor labor market; prepare for wage pressures
Finance	High risk: back-office workers organizing; junior roles under pressure	Proactively manage talent retention; announce reskilling programs
Consulting	High risk: junior consultant roles under pressure; organizing around AI displacement	Proactively manage talent retention; announce reskilling programs

Industry	Direct Impact	Implied Action
Law	Moderate risk: associate and paralegal roles under pressure	Proactively manage talent retention; announce reskilling programs
Big Tech	High risk: engineering and technical workers organizing around AI labor displacement	Proactively manage talent retention; announce reskilling programs
B2B/B2C SaaS	Moderate risk: engineering workforce organizing	Proactively manage talent retention; announce reskilling programs

Facilitator Guidance:

- The struck company sees modest productivity gains (5-8%) vs. projected 20%+; transparency requirements add friction.
- Worker retention improves (20-30% higher in retraining cohorts vs. layoff scenarios).
- Other companies are quietly executing AI headcount reductions via "voluntary severance" before labor standards solidify.
- Firms that transparently manage workforce transitions will avoid union action and maintain brand reputation; firms that execute silent headcount reductions will face backlash.
- This is a major penalty for Round 1 aggressive headcount reduction strategies. Firms that chose aggressive optimization will now face union action and retraining costs.
- Firms that chose transparent transition or "AI + Human" will be positioned as responsible actors and gain reputational advantage.

D. OPTIONAL INJECTS (Use Only If Table Finishes Early)

R2-OPT-01: Deepfake Fraud Ring

Classification: Security / Fraud / Regulatory **Time:** Mid-round to late round (Day 2 morning or afternoon)

Narrative: A sophisticated deepfake fraud ring is detected. The ring used AI-generated video and audio to impersonate C-suite executives and initiate fraudulent wire transfers totaling \$85M across multiple financial institutions and retail companies. One victim company—a major financial institution—suffered a \$20M loss before fraud was detected. The incident makes international headlines: "AI Deepfakes Used in \$85M Fraud Ring." Law enforcement launches investigation. Multiple companies announce emergency reviews of authentication and transaction controls. Regulatory bodies issue urgent guidance on deepfake detection and authentication. Media speculation runs wild: "Is video evidence unreliable now?" "Can we trust any digital communication?"

Ambiguity/Unknowns:

- How widespread is deepfake fraud risk?

- What authentication technologies will be effective?

Industry Impact:

Industry	Direct Impact	Implied Action
Finance	High risk: transaction authentication; fraud detection; reputational damage if fraud goes undetected	Invest in authentication technology and fraud detection
Consulting	Opportunity: AI risk and fraud advisory demand	Develop fraud and authentication advisory offerings
Law	Opportunity: litigation and regulatory advisory demand	Develop fraud litigation and advisory offerings
Retail	Secondary impact via customer concern about security of video-based interactions	Communicate security measures to customers
CPG	Minimal direct impact	Monitor for brand impersonation risks
Manufacturing	Secondary impact via supplier authentication concerns	Ensure supplier authentication mechanisms are robust
Logistics	Secondary impact via supplier and carrier authentication concerns	Ensure authentication mechanisms are robust
Healthcare Provider	Secondary impact via patient and provider authentication	Implement multi-factor authentication for clinical systems
Healthcare Payer	Secondary impact via claims authentication	Implement multi-factor authentication for claims systems
Big Tech	Opportunity: develop authentication and fraud detection capabilities for cloud and enterprise customers	Accelerate product development in security domain
B2B/B2C SaaS	Opportunity: develop authentication capabilities for customers	Accelerate product development in security domain

Facilitator Guidance:

- This incident is not an existential threat to AI or digital communications, but it is a major psychological shock.
- Companies with robust transaction controls and human oversight will survive with minimal damage.
- Companies that relied on AI authentication alone will face significant reputational damage and fraud losses.
- Authentication technology and fraud detection consulting become high-growth markets.

E. FACILITATOR MARKET SHOCK (3 minutes)

The facilitator imposes external Market Shocks on 2-3 industries. This replaces the former IC Constraint Imposition — the facilitator acts as the market, not a participant role.

Facilitator Script (Read Aloud):

"The market doesn't wait for consensus. Based on macro developments — regulatory actions, competitive moves, and market dynamics — the following industries are facing an external shock this round."

Procedure:

Step	Action	Time
1	Facilitator selects 2-3 industries to receive one constraint each	Pre-round prep
2	Facilitator announces constraints to the table	2 min
3	Affected participants note constraint on their decision worksheet	1 min

Constraint Menu (select one per targeted industry; cannot impose same type on multiple industries):

1. Regulatory Halt: Regulators have announced a review of AI deployments in this industry. All new AI projects are frozen for this round — no new pilots, no expansion to new sites, no go-live on tools still in testing. The participant's strategy must work within their current AI footprint.
2. Labor Cost Surge: Union negotiations, wage competition, or mandatory retraining obligations have driven labor costs sharply higher — enough to compress margins and force trade-offs. Any strategy that depends on labor-intensive rollout now costs significantly more. The participant must narrow scope, slow timelines, or offset the pressure.
3. Capital Tightening: Market conditions have significantly reduced available capital. The board has pulled back discretionary spending authorization. Capital-intensive strategies (large infrastructure investments, acquisitions, multi-site expansions) are constrained. The participant must go capital-light or justify a smaller, more focused investment.
4. Reputational Pressure: Public backlash — workforce displacement, algorithmic fairness, or customer trust erosion — has forced a defensive posture. Stakeholders expect visible action. Aggressive AI expansion without a credible "responsible transition" component will be scored harshly. The participant must address the backlash in their strategy.
5. Competitive Response: A major competitor has used AI to aggressively undercut pricing, poach customers, or leapfrog product capability. The threat is immediate and material. Purely internal-facing or long-horizon strategies that ignore the near-term competitive threat are risky. The participant must defend their position.
6. Litigation Risk: The industry faces active or imminent AI-related litigation — algorithmic bias, labor displacement, IP disputes, or product liability. Legal exposure is material and uncertain. Any strategy that expands litigation surface area carries additional risk. The participant should reflect the legal overhang in their reasoning.

Key difference from Industry Health Signal constraints: Market Shocks can target ANY industry —

including well-performing ones. Health Signal constraints are mechanical consequences of poor cumulative scores. Market Shocks reflect external forces that hit regardless of performance.

Facilitator Guidance:

- Prepare constraint selections before the round begins — do not deliberate in front of participants.
- Target industries where constraints create the most interesting strategic tension, not just the weakest performers.
- Constraints interact with Round 1 decisions: an industry that chose "aggressive" strategy and then receives a constraint will face severe pressure. An industry that chose "prudent" strategy will be better positioned to absorb it.
- Announce quickly and move on. This is a 3-minute mechanic, not a deliberation.

E1. INDUSTRY HEALTH SIGNALS — ROUND 2

Facilitator Note: After scoring Round 1 decisions and applying fallbacks, announce Industry Health conditions for each industry. This is the first time conditions are announced.

Procedure:

1. Calculate cumulative aggregate score for each industry (Round 1 scores only at this point)
2. Look up condition in the table: Surge (+15+), Tailwind (+6 to +14), Steady (-5 to +5), Headwind (-6 to -14), Crisis (-15 or worse)
3. Announce condition for each industry (~2 minutes)
4. For Headwind industries: impose one additional constraint from the constraint menu
5. For Crisis industries: impose two constraints; inform participant their primary decision must include a defensive component

Note: Facilitator Market Shock constraints from Section E are separate from Health Signal constraints. An industry could receive both a Market Shock constraint AND a Health Signal constraint if their Round 1 performance was poor AND the facilitator also targeted them.

E2. COLLECTIVE BONUS (Optional, within Cross-Industry Discussion)

Facilitator Script (Read Aloud):

"Does anyone want to recognize an especially strong strategy this round, or flag one that seems particularly risky?"

Procedure:

Step	Action	Time
1	Participation is optional — no one is required to speak	—
2	Participants who wish to respond name one industry as "strong strategy" and/or one as "risky strategy" (cannot nominate own industry)	3 min
3	If 3+ participants agree on the same industry, facilitator applies adjustment	1 min

Scoring:

- If 3+ participants agree an industry has a strong strategy: +2 cumulative score bonus
- If 3+ participants agree an industry has a risky strategy: -2 cumulative score penalty
- Maximum one +2 and one -2 per round
- If no consensus or no one speaks up, no bonus applied

Facilitator Guidance:

- The Collective Bonus creates lightweight peer accountability without forcing participation.
- Encourage participants to cite specific reasoning if they do nominate.
- Announce results openly — nominations are public, not anonymous.
- In Round 2, nominations should reflect how well participants adapted to new information from Round 1 results and this round's injects.

F. PRIVATE CARD DISTRIBUTION

Facilitator Note: Distribute the appropriate Private Information Card from [05_Participant_Private_Cards/](#) to each participant face-down. Round 2 uses Private Card 2. Card 2 is unique per industry—each participant receives information specific to their industry. Private cards contain updated industry-specific data, financial metrics, competitive intelligence, or internal developments relevant to each participant's decisions in this round. These cards are intentionally asymmetric—participants have incomplete information about other industries' actual performance from Round 1, forcing reliance on public injects and cross-industry inference.

Private cards are used in Rounds 1-3 only. Round 4 operates on symmetric information.

G. INDIVIDUAL INDUSTRY DECISION SUBMISSION REQUIREMENTS — V7.4

Each industry participant submits:

1. Industry Decision (1 REQUIRED):

- Each participant must submit an explicit decision for their industry per round.
- Use banded framework: avoid invented precision, use directional language.
- Address how you're responding to any Market Shock constraints or Industry Health Signal constraints imposed in this round.

Submission Format:

[INDUSTRY NAME] – ROUND 2 DECISION

CONSTRAINT IMPOSED: [Name of constraint, if any – Market Shock or Health Signal]

Response to Constraint:

[How are you adapting strategy in response?]

DECISION TITLE: [Title]

RATIONALE:

[Why this choice? How does it respond to R2 injects and constraints?]

BANDED ASSESSMENT:

- Spend/Commitment: [Band]
- Time-to-Impact: [Band]
- Execution Complexity: [Band]
- Dependency: [Band]
- Scale: [Band]

EXPECTED OUTCOME:

[What do you expect in Round 3?]

INDUSTRY REPRESENTATIVE: [Name]

Time to Submit: 18 minutes after injects and Market Shock constraints are announced.

Key Notes:

- One decision per participant per round.
- Banded framework prevents false precision and focuses discussion on real trade-offs.
- Address constraint response explicitly if one is imposed on your industry.

H. INDUSTRY HEALTH UPDATE — ROUND 2

Round 2 Industry Health conditions will be updated at the start of Round 3 based on cumulative scores (R1 + R2), including any Collective Bonus adjustments (+2/-2 per round). Strong Round 2 performance can shift an industry from Steady to Tailwind; poor performance can drop an industry into Headwind.

End of Round 2 Materials

Round 3: Reckoning (18–30 Months)

AI First Wave: Early Gains Plateau, Regulatory Tightens, Winners Emerging

A. Round Overview

By August 2027, the AI hype cycle has collided with operational reality. First-wave productivity gains (code generation, routine automation, customer service triage) are plateauing. Early adopters in software, finance, and logistics captured significant advantages; competitors now catching up, eroding cost advantages. Regulatory frameworks from Round 2 are being actively enforced, raising compliance costs. Capital markets are repricing AI bets downward as productivity data disappoints vs. 2026 expectations. Market concentration accelerates: top quartile firms pulling away from laggards. Labor displacement visible and accelerating in routine cognitive and administrative roles. This round is the "reckoning"—where hype meets reality, and strategic divergence becomes stark.

B. Situation Update

The Productivity Plateau & Reality Reckoning

First comprehensive data on AI productivity outcomes through August 2027 is in. Aggregate productivity gain across eleven industries: +1.8% annualized—at the lower end of 2026 forecaster expectations and well below the +3-5% hype cycle projections. Distribution is deeply skewed.

- Retail and CPG: AI-driven personalization, logistics optimization, and demand forecasting delivering 2.8-3.1% productivity gains for early adopters (Amazon, Walmart); mid-market retail struggling at +0.7-1.2%.
- Healthcare Provider and Healthcare Payer: Healthcare systems seeing modest gains (+1.0-1.8%) due to regulatory friction and mandatory human review requirements. Clinical AI deployment slowed by FDA guidance; payer-side claims processing benefits at +2.0-2.5%.
- Manufacturing and Logistics: Logistics and manufacturing leaders (autonomous scheduling, quality control, demand planning) achieving 2.5-3.0%; smaller manufacturers at +0.3-0.8%.
- Finance, Consulting, and Law: Code generation, trading automation, and document automation enabling 2.2-3.0% gains for large firms; mid-market services at +0.6-1.5%.
- Big Tech and B2B/B2C SaaS: AI-native feature integration driving 3.5-4.5% productivity gains in SaaS; Big Tech cloud, ads, devices, and enterprise platform capex efficiency gains 4.0-5.0%.

Critical finding: Most low-hanging fruit is picked. Organizations explicitly report that "first-wave gains are exhausted." Second-wave gains (deep process redesign, organizational restructuring, end-to-end workflow transformation) are slowing and showing declining ROI. Sustaining productivity growth requires expensive, slow organizational change.

Enterprise AI Economics: ROI Deteriorating

Average payback period for AI projects: 24-36 months (vs. 18-24 months estimated in early 2026). Total cost of ownership 40-60% higher than initially projected. Mid-market and small firms increasingly unable

to achieve comparable ROI to mega-caps due to talent/data/capital constraints. **Capability inequality accelerating:** gap between Fortune 500 and mid-market AI ROI widened from 1.5x in 2026 to 3-4x by end of 2027.

Labor Market Shift: Barbell Emerging

Unemployment risen to 4.3%. Job losses in routine cognitive and administrative roles accelerating: data entry (-35% YoY), customer service agents (-15%), junior analysts (-18%), routine content creation (-22%). Wage growth for AI-adjacent technical roles strong (+8-12% YoY), but wage stagnation for routine roles widening income inequality.

- Retail and CPG: Retail/e-commerce customer service and warehouse roles hardest hit. CPG distribution and marketing roles affected.
- Manufacturing and Logistics: Logistics and manufacturing back-office reductions acute. Production and warehouse roles under pressure.
- Healthcare Provider and Healthcare Payer: Healthcare administrative roles; clinical support roles facing automation. Claims processing staff reductions.
- Finance, Consulting, and Law: Finance back-office consolidation; junior analyst/associate roles compressed; senior expert roles strong. Consulting junior roles under pressure. Law associate and paralegal roles facing automation.
- Big Tech and B2B/B2C SaaS: Engineering wage pressure from global talent competition; junior developer roles increasingly challenged by AI code generation.

Regional disparities pronounced. Manufacturing Midwest and financial back-office centers (Kansas City, Omaha) hardest hit. Tech hubs (Seattle, San Francisco, Boston, Austin) seeing labor market tightness for AI talent. Union organizing intensified; strikes in logistics, customer service, and tech explicitly citing AI-driven job losses. Two state legislatures (California, New York) passed "AI Worker Protection" laws.

Regulatory Tightening: Compliance Now Structural

- SEC enforcement: Three high-profile actions against companies misrepresenting AI capabilities in investor disclosures. Total fines >\$200M. Signals aggressive enforcement.
- FDA clinical AI: Final binding rule mandating human review, explicit safety testing, transparent failure reporting. Compliance costs material; healthcare AI deployment timelines extended.
- Financial services: Major bank fined \$150M for algorithmic bias (non-AI rule-based system, but regulators explicitly flagging AI will face stricter scrutiny). Compliance infrastructure costs rising sharply.
- EU AI Act: Full enforcement; non-compliance costs mounting. State-level fragmentation creating compliance burden for national companies.

Regulatory burden now real and slowing adoption, but no innovation-killing regulations have emerged.

Market Structure Shift: Winners Pulling Away

Research from major consulting firms documents sharp bifurcation. Top 3 firms in each industry capturing 70-85% of AI-driven productivity gains. Large firms' advantages in data quality, integration capability, and capital multiplying. M&A increasingly driven by "acquirer wants target's AI talent and IP." Multiple major AI startups have missed growth targets; valuations compressed; acquisitions at steep discounts.

B1. AI ADOPTION ARC DISTRIBUTION

Facilitator Note: Distribute Phase 3 (Reckoning) from 04_Participant_AI_Adoption_Arcs/ to each participant. This provides each industry's updated AI adoption trajectory based on Round 1-2 decisions and macro developments.

C. Core Injects (2-3 Maximum)

R3-01: Regulatory Enforcement Wave

Title: SEC & FDA Coordinate AI Enforcement Actions **Classification:** Regulatory **Time:** Opening (start of round)

Narrative: SEC announces three high-profile enforcement actions against public companies for misrepresenting AI capabilities or downplaying algorithmic risks in investor disclosures. Total fines exceed \$200M. FDA simultaneously issues final rule on AI-assisted clinical decision support: mandatory human review, explicit safety testing, transparent failure reporting. Major national bank fined \$150M by Federal Reserve and OCC for algorithmic bias in lending (rule-based, but precedent-setting for AI oversight). Combined effect signals regulators will actively enforce AI transparency and safety requirements; non-compliance carries material financial consequences. Compliance costs rise sharply; insurance companies launch "AI compliance liability" coverage.

Industry Impact:

Industry	Impact	Constraint	Implication
Healthcare Provider	Very High	Mandatory compliance infrastructure (audit, risk mgmt, documentation); deployment timelines extended 6-12 months	Clinical AI deployment delays; must retrofit with human review layers
Healthcare Payer	Very High	Mandatory compliance infrastructure; claims processing AI under review	Claims AI deployment delays; must retrofit with audit and review layers
Finance	High	Mandatory compliance; liability insurance; trading governance	Must retrofit autonomous trading with human checkpoints
Consulting	Moderate-High	Compliance consulting demand surges; own AI delivery under scrutiny	Capitalize on compliance demand; audit own AI-augmented delivery
Law	Moderate-High	Regulatory advisory demand surges; own AI tools under scrutiny	Capitalize on regulatory demand; audit own AI tools

Industry	Impact	Constraint	Implication
Retail	Moderate	Disclosure audit; liability insurance costs; transparency documentation	Personalization less affected; customer-facing AI requires documentation
CPG	Low-Moderate	Compliance if consumer-facing AI deployed	Monitor for downstream effects
Manufacturing	Low-Moderate	Compliance if autonomous systems deployed	Autonomous deployment governance scrutiny
Logistics	Low-Moderate	Compliance if autonomous systems deployed	Autonomous deployment governance scrutiny
Big Tech	Moderate	Product liability for AI features in cloud, enterprise platforms, ads, and devices; compliance documentation	Liability/compliance responsibilities for platform AI increase
B2B/B2C SaaS	Moderate	Product liability for AI-embedded products; compliance documentation	Liability/compliance responsibilities for SaaS vendors increase

Facilitator Guidance: SEC planning second wave of enforcement against 15-20 companies. FDA rule more prescriptive than anticipated: human review mandatory; documentation burden material. Bank fine was rule-based but memo explicitly noted "AI systems will face higher scrutiny." Compliance is now structural cost, not temporary friction.

R3-02: Market Concentration Evidence

Title: Winner-Take-Most Dynamics Accelerating (Consulting Research) **Classification:** Market Research **Time:** Early-to-mid round

Narrative: Major consulting firm publishes "AI and Market Concentration: Winner-Take-Most in Industry 4.0." Top 3 firms in each industry captured 70-85% of AI-driven productivity gains (2025-2027). Amazon's AI-driven supply chain and personalization driving retail consolidation. Leading cloud platforms consolidating adoption. Mega-cap banks' agent-based M&A processing providing material cost advantage. Large manufacturers' AI-driven quality and scheduling optimization accelerating; smaller manufacturers falling behind. Research concludes: "AI exhibits strong network effects and economies of scale. Top quartile pulls away from bottom three quartiles. Mid-market faces existential competitive pressure unless they find defensible niches." Triggers immediate strategic question: compete or consolidate via M&A?

Industry Impact:

Industry	Impact	Constraint	Implication
Retail	High	Market share concentration evident; mid-market must choose: aggressive transformation, niche positioning, or M&A	Large retailers can raise prices; smaller retailers under margin pressure
CPG	High	Brand consolidation accelerating; mid-market CPG squeezed	Large CPG firms gaining; mid-market must find defensible niches
Manufacturing	High	Scale advantages in automation evident	Small/mid manufacturers must find defensible niches or seek acquisition
Logistics	High	Scale advantages in route optimization and automation evident	Smaller logistics firms must consolidate or find niches
Healthcare Provider	Moderate	Regulatory fragmentation allows mid-market players to compete in niches	System consolidation in regulated states; mid-market viable in fragmented regions
Healthcare Payer	Moderate	Claims processing scale advantages emerging	Mid-market payers under pressure; niche specialists may survive
Finance	High	Finance consolidation evident	Consolidation accelerating; boutique specialists may survive; mid-market squeezed
Consulting	High	Large-firm advantage in AI-augmented delivery	Large consulting firms pulling away; boutique specialists in AI niches viable
Law	Moderate-High	Large-firm advantage in AI-powered practice areas	Large law firms gaining; mid-market squeezed in commoditized practice areas
Big Tech	High	Platform lock-in and network effects favor incumbents in cloud, ads, enterprise platforms, and devices	AI-native startups under pressure; platform consolidation accelerating

Industry	Impact	Constraint	Implication
B2B/B2C SaaS	High	Platform lock-in and network effects favor incumbents	SaaS consolidation; smaller SaaS firms under acquisition pressure

Facilitator Guidance: Three mega-cap retailers increased market share 400 bps in past 18 months. Antitrust regulators (FTC, EU) beginning to examine AI-driven consolidation; early opinions: superior execution legal; proprietary data/closed algorithms might be challenged. M&A multiples elevated (6-8x revenue for high-talent targets). Some mid-market firms exploring "AI co-ops" (shared data platforms).

R3-03: System Failure Incident (Choose ONE)

Title: [A] Clinical AI Misdiagnosis / [B] Autonomous Vehicle Accident / [C] Algorithmic Trading Flash Crash
Classification: Incident/Crisis **Time:** Mid-to-late round

OPTION A: Clinical AI Misdiagnosis

Healthcare system's AI diagnostic assistant (pneumonia detection in radiology) makes systematic errors on edge-case presentations, leading to 12 delayed diagnoses and 3 patient deaths over 4 months. Root cause: training data underrepresented atypical presentations from specific patient populations. Healthcare system faces \$500M+ malpractice liability, FDA investigation, state medical board inquiry. Regulatory response: FDA accelerates clinical AI guidance; investigation of other AI diagnostic systems. Industry reckoning: trust in AI-assisted diagnosis erodes; deployment slows sharply. Healthcare systems retrofit AI systems with stricter validation and human review.

OPTION B: Autonomous Vehicle Accident

Autonomous vehicle operated by major logistics company encounters novel scenario (construction-zone condition underrepresented in training data) and makes unsafe maneuver, resulting in multi-vehicle accident with 2 fatalities, 8 injuries. Manufacturer and operator face criminal negligence investigations, civil liability claims, NHTSA inquiry. Regulatory response: NHTSA announces new safety certification requirements; autonomous deployment paused pending investigation. Industry reckoning: AV deployment timelines extended 18-24 months; liability framework unclear, deterring investment.

OPTION C: Algorithmic Trading Flash Crash

Autonomous trading agent at major financial institution triggers market microstructure anomaly during stressed conditions, causing 15-minute flash crash in major index. Losses ~\$200M; systemic risk concerns triggered. SEC investigation, congressional hearing. Regulatory response: SEC announces "Algorithmic Trading Governance Rule" requiring mandatory circuit breakers, human oversight checkpoints, pre-trade risk limits on AI trading systems. Finance industry reckoning: autonomous trading deployment paused; compliance burden material.

Industry Impact (using Option A as example):

Industry	Impact	Constraint	Implication
Healthcare Provider	Very High	Clinical AI faces stricter validation, longer approval timelines, higher liability	Healthcare AI productivity gains paused; compliance costs spike
Healthcare Payer	High	Claims AI under additional scrutiny; liability exposure	Claims processing AI deployment slowed; audit requirements increase
Finance	Moderate	Liability/consulting implications	"AI risk management" consulting demand spikes
Consulting	Moderate	Demand for AI risk advisory surges	Capitalize on risk consulting demand
Law	Moderate	Demand for litigation and regulatory defense	Capitalize on litigation demand
Big Tech	Moderate	Product liability for clinical AI platform vendors	Clinical AI platform vendors face litigation risk; liability insurance costs
B2B/B2C SaaS	Moderate	Product liability for clinical AI SaaS vendors	Clinical AI SaaS vendors face litigation risk
Retail	Low	Secondary effects	Health data sensitivity elevated
CPG	Low	Minimal direct impact	Monitor for consumer trust spillover
Manufacturing	Low	Spillover if Option B	Autonomous manufacturing deployment deferred
Logistics	Low-Moderate	Spillover if Option B	Autonomous logistics deferred

Facilitator Guidance: Incident is credible but isolated to specific vendor/deployment. Root cause attributable to training data bias, not fundamental unsuitability. Regulatory response measured: oversight, not prohibition. Erodes confidence temporarily; second-wave deployments slower and more expensive. Exercise probes participants' risk tolerance for AI under uncertainty.

D. Optional Injects

R3-OPT-01: Deepfake Disinformation Event

Title: Deepfake Political Video Triggers Regulatory & Trust Crisis **Classification:** Incident/Policy **Time:**

Early or mid-round

Sophisticated deepfake of prominent US Senator surfaces, purporting inflammatory statements on AI regulation. Spreads to 5M social media impressions before debunking (6-8 hour lag). White House issues Executive Order: "Deepfakes and Synthetic Media: Detection, Disclosure, and Accountability." Consumer trust in digital media drops sharply (polling: only 35% trust video/image online; 42% reduced trust overall). **Impact:** Retail and Finance face customer trust challenges; compliance burden for synthetic media disclosure; consumer sentiment on AI deteriorates.

E. COLLECTIVE BONUS (Optional, within Cross-Industry Discussion)

Facilitator Script (Read Aloud):

"Does anyone want to recognize an especially strong strategy this round, or flag one that seems particularly risky? By Round 3, you have real track records to evaluate. Reference what worked and what didn't."

Procedure:

Step	Action	Time
1	Participation is optional — no one is required to speak	—
2	Participants who wish to respond name one industry as "strong strategy" and/or one as "risky strategy" (cannot nominate own industry)	3 min
3	If 3+ participants agree on the same industry, facilitator applies adjustment	1 min

Scoring:

- If 3+ participants agree an industry has a strong strategy: +2 cumulative score bonus
- If 3+ participants agree an industry has a risky strategy: -2 cumulative score penalty
- Maximum one +2 and one -2 per round
- If no consensus or no one speaks up, no bonus applied

Facilitator Guidance:

- By Round 3, participants have seen two rounds of results. Nominations should reflect demonstrated execution, not just declared intent.
 - The Collective Bonus creates real consequences: cumulative score shifts can push industries into or out of Headwind/Crisis territory.
 - Encourage participants to cite specific evidence from prior rounds' outcomes if they do nominate.
 - Announce results openly — nominations are public, not anonymous.
-

E1. INDUSTRY HEALTH SIGNALS — ROUND 3

Facilitator Note: After scoring Round 2 decisions (including fallbacks) and applying any Collective Bonus adjustments from Rounds 1-2, update and announce Industry Health conditions.

Procedure:

1. Calculate cumulative aggregate score for each industry (R1 + R2 scores, including any Collective Bonus adjustments)
 2. Look up condition for each industry
 3. Announce conditions (~2 minutes)
 4. Apply Headwind/Crisis constraints as applicable
-

F. Private Card Distribution — Round 3

Facilitator Announcement: "Distribute the appropriate Private Information Card from 05_Participant_Private_Cards/ to each participant face-down. Round 3 uses Private Card 3. Card 3 is unique per industry—each participant receives information specific to their industry. This is the LAST round with private cards. Round 4 will have no private information—all information is public."

G. INDIVIDUAL INDUSTRY DECISION SUBMISSION REQUIREMENTS — V7.4

Each industry participant submits:

1. Industry Decision (REQUIRED): One strategic choice tied to core Round 3 scenario (workforce restructuring, M&A positioning, regulatory compliance, product strategy).
 - Use banded framework.
 - 1-2 sentences explaining rationale.

Submission Format:

[INDUSTRY NAME] – ROUND 3 DECISION

DECISION TITLE: [Title]

RATIONALE:

[Why this choice? How does it respond to R3 injects and market dynamics?]

BANDED ASSESSMENT:

- Spend/Commitment: [Band]
- Time-to-Impact: [Band]
- Execution Complexity: [Band]
- Dependency: [Band]
- Scale: [Band]

EXPECTED OUTCOME:

[What do you expect in Round 4?]

INDUSTRY REPRESENTATIVE: [Name]

Submission Window: 15 minutes.

Key Notes:

- One decision per participant per round.
 - Banded framework for all submitted decisions.
 - This is the final round with private information.
-

End of Round 3: Reckoning

Round 4: Normalization (30–48 Months)

AI Embedded Infrastructure: Valuations Stabilize, Winners Clear, Labor Market Adjusts

A. Round Overview

By early 2030, AI has become embedded in normal business operations for leading companies. The "AI transformation" narrative has matured into a more nuanced story: productivity gains are real but uneven; labor displacement visible but manageable; regulatory frameworks in place; valuations normalized. AI is now table-stakes infrastructure, not a standalone bet. Winners and losers are clear. This is the final competitive positioning round before debrief—each industry consolidates long-term competitive advantage in a mature AI environment.

B. Situation Update

AI as Ordinary Infrastructure (Not Miracle Cure)

Frontier models now at "GPT-6" equivalence or beyond; reasoning, multimodal, and agentic capabilities broadly commoditized. Value of AI no longer in the model itself but in application, data, and integration. Companies with high-quality data, clear business problems, and strong execution creating value; companies without these capabilities seeing diminishing returns. Open-source models fully closed capability gap with frontier models for many applications. Enterprise AI market dominated by implementation and integration services, not model companies. AI startup ecosystem consolidating around few surviving players and much larger long tail of niche specialists.

Productivity—Final Tally Through 2029

Average productivity gain across eleven industries: +1.8% annualized (realized outcome matches forecast). Distribution remains skewed: top quartile +3.5% annualized; bottom quartile +0.4% annualized. Organizations sustaining productivity growth have moved past obvious automation plays and are redesigning processes, roles, and business models around AI. This is slow, expensive, requires deep organizational change. Many companies have hit plateau and are unsure how to move forward. Low-hanging fruit fully harvested; further gains require genuine innovation, not just tool deployment.

Labor Market—Stabilization After Transition

Unemployment settled at 4.35-4.55% (0.4pp higher than 2026 baseline). Total AI-driven job displacement (690K) largely realized; pace of new displacement slowed. Routine cognitive and administrative roles still declining, but slower. AI-adjacent roles (trainers, governance, compliance, integration consultants) grown to ~350-400K and now becoming routine skills (not scarce). "Barbell" labor market persists but stabilizing. Regional inequality durable: Midwest and back-office centers seen persistent job losses; tech hubs thrived. Wage growth for high-skill workers strong; routine cognitive workers flat to negative. Retraining programs trained ~200-250K workers; success rates mixed (60-70% find comparable/better roles; 30-40% do not). Union activity decreased as novelty wore off. Few "AI-era workforce" agreements struck, typically with productivity sharing, notice requirements, retraining funds. Federal labor policy unsettled; state-level policies matured and becoming standard.

Valuation & Capital Markets—Normalization

AI-beneficiary stock multiples normalized from peaks but remain elevated vs. historical averages. Software: 6.0x-7.0x EV/Revenue (down from 7.5x-8.5x Round 2 peak; still above pre-AI 4.5x-5.5x). Finance: 0.95x-1.1x Price/Book (stable). Mega-cap tech multiples compressed from peak levels. VC valuations compressed; Series B/C rounds 30-40% lower than 2027 peaks. M&A normalized; acquisitions still driven by "talent and IP" but at rational prices. Venture market for "pure AI" matured; growth area is "AI adoption services" (consulting, integration, change management). Multiple VC-backed categories shut down: AI hardware (except winners); narrow-purpose AI software (displaced by commodity models); general-purpose autonomous agents (timelines extended).

Regulatory Frameworks—Mature & Enforced

SEC, FDA, industry-specific regulations now well-established and routinely enforced. Compliance is material cost but predictable and manageable. EU AI Act in mature enforcement; US regulations industry-by-industry but increasingly harmonized. State-level regulations persist; fragmentation manageable. No major new regulatory shocks expected; framework is stable. Regulation has not stifled innovation; it has channeled innovation toward "responsible AI" and "AI for compliance." Innovation in healthcare, finance, regulated industries continues, constrained but not halted.

Macroeconomic Picture

GDP growth back to trend (2.2-2.5%); productivity improvements from AI absorbed into normal economic growth. "AI productivity dividend" real but modest—roughly 0.3-0.5pp of annual productivity growth (vs. 2-4pp claimed in 2026-2027). Inflation stable; interest rates normalized. Asset valuations broadly in line with historical ranges. Credit markets functioning normally. US economy successfully absorbed AI-driven labor displacement without major downturn. Income inequality increased slightly (high-skill workers pulling away; routine workers stagnant). Triggered modest policy attention but not major political upheaval.

Societal & Competitive Dynamics

Consumer attitudes toward AI normalized. AI now routine technology; consumers understand and interact regularly. No longer novel or controversial (except high-stakes domains like healthcare, where skepticism persists). Regulatory mandates on transparency being met. Job displacement concerns persist among affected workers but have not triggered political crisis. Small "AI skepticism" community remains; trust in AI for high-stakes decisions (medical, financial, legal) lower than for routine tasks. No major backlash has emerged; AI integrated into normal business and life.

Internationally, US, EU, and China developed divergent but stable regulatory regimes. Competition in AI global; US companies lead, but European and Chinese competitors matured. Standards-setting bodies (NIST, ISO, OECD) increasingly influential; interoperability improving. AI talent and compute capacity concentrated in few jurisdictions; creates international tensions but no major conflict.

B1. AI ADOPTION ARC DISTRIBUTION

Facilitator Note: Distribute Phase 4 (Normalization) from 04 Participant AI Adoption Arcs/ to each participant. This provides each industry's final AI adoption trajectory based on Round 1-3 decisions and macro developments.

C. Core Injects (2-3 Maximum)

R4-01: Policy Finalization

Title: Federal AI Regulation Framework Enacted **Classification:** Policy **Time:** Opening (start of round)

Narrative: Federal AI Regulation and Responsibility Act of 2029 passes both chambers with bipartisan support. Key provisions: (1) Mandatory AI impact assessments for systems affecting >100K users in regulated industries (finance, healthcare, employment); (2) Federal "AI Compliance Fund" providing grants and tax credits to small/mid-market firms for governance infrastructure; (3) Establishment of National AI Safety Board with authority to audit high-risk systems, impose fines for non-compliance; (4) 24-month regulatory ramp period, with incremental enforcement. Act is moderate: establishes clear rules but avoids banning or severely restricting AI deployment. Market reaction positive: clarity reduces compliance uncertainty. Small/mid-market firms see funding support. Large firms face compliance costs but predictable. Venture funding for "AI governance and compliance" services accelerates.

Industry Impact:

Industry	Impact	Constraint	Implication
Healthcare Provider	Moderate (Framework clarifies prior uncertainty)	Compliance standardization; mandatory impact assessments; federal fund support available	Healthcare providers can now deploy with clear rules; uncertainty premium lifts
Healthcare Payer	Moderate (Framework clarifies prior uncertainty)	Compliance standardization; mandatory impact assessments; federal fund support available	Healthcare payers can now deploy with clear rules; uncertainty premium lifts
Finance	Moderate (Framework clarifies prior uncertainty)	Compliance standardization; mandatory impact assessments; federal fund support available	Finance can now deploy with clear rules; uncertainty premium lifts
Consulting	Moderate-High (Strategic opportunity)	Compliance consulting demand sustained; own AI delivery benefits from clarity	"AI governance" advisory becomes sustained service line
Law	Moderate (Strategic opportunity)	Regulatory advisory demand sustained	Regulatory advisory becomes sustained practice area
Retail	Low-Moderate	Transparency requirements for personalization; impact assessment if system >100K users	Large retailers trigger assessment requirements
CPG	Low	Minimal direct impact; compliance if consumer-facing AI scales beyond thresholds	CPG companies largely unaffected

Industry	Impact	Constraint	Implication
Manufacturing	Low	Compliance if autonomous systems scale beyond thresholds	Manufacturing companies largely unaffected
Logistics	Low	Compliance if autonomous systems scale beyond thresholds	Logistics companies largely unaffected
Big Tech	Moderate	Sustained consulting opportunity for customer compliance via cloud and enterprise platforms; product liability insurance normalizes	"AI governance" services for platform customers becomes key offering
B2B/B2C SaaS	Moderate	Consulting opportunity for customer compliance; product liability insurance normalizes	"AI governance" services for SaaS platform becomes key offering

Facilitator Guidance: Federal framework reduces patchwork of state regulations; national firms benefit from clarity. Compliance Fund is material assistance to mid-market (estimated \$5-15B over 10 years). National AI Safety Board is independent but politically insulated. Ramp period gives companies time for implementation without punitive enforcement.

R4-02: Market Consolidation Outcome

Title: AI-Leading Tech Company Consolidates Competitors; Valuations Reset **Classification:** Market Signal **Time:** Mid-round

Narrative: Major AI-leading tech company (Prometheus AI equivalent) announces acquisition of struggling second-tier AI competitor at \$5-7 billion, representing 40-50% discount to 2027 valuations. Acquirer explicitly states deal is about "acquiring talent and IP for frontier layer consolidation." 30-40% of target's engineering team to be integrated into acquirer's core research; remainder may be redundant. First major "down round" acquisition of once-hyped AI startup signals venture ecosystem's previous valuations were inflated. News triggers sharp sell-off in other mid-tier AI startups; broader revaluation of VC-backed AI companies. Multiple venture firms that bet heavily on AI startups face marked-to-market losses. Concurrently, large consulting firm acquires smaller AI-for-enterprise-services company at 2.5x revenue multiple (vs. 4.0x+ in 2027), reflecting compressed VC-exit valuations. Data shows: AI startups that pivoted away from "pure AI" to "AI + vertical applications" valued higher than pure-play models.

Industry Impact:

Industry	Impact	Constraint	Implication
Retail	Moderate	Consolidation may trigger vendor changes for smaller retailers relying on AI startups	Large retailers benefit from vendor consolidation/discipline
CPG	Low-Moderate	Consolidation may affect AI vendor landscape for CPG firms	Monitor for vendor stability
Manufacturing	Low	Limited vendor exposure	Minimal disruption
Logistics	Low	Limited vendor exposure	Minimal disruption
Healthcare Provider	Low	Established vendors remain; startups already constrained by compliance	Limited impact
Healthcare Payer	Low	Established vendors remain; startups already constrained by compliance	Limited impact
Finance	Moderate (Strategic)	Finance firms evaluating AI vendor stability	Monitor vendor consolidation; secure long-term contracts
Consulting	High (Strategic)	Consulting firms acquiring AI companies signal shift to "adoption services" model	"AI for business transformation" becomes core consulting service line
Law	Moderate	Law firms evaluating AI vendor stability; M&A advisory demand	Capitalize on M&A advisory demand
Big Tech	High (Strategic)	Platform companies acquiring AI startups for talent and IP; consolidation around core cloud, enterprise, ads, and device platforms	Platform companies consolidating AI innovation around core products
B2B/B2C SaaS	High (Strategic)	SaaS companies acquiring AI startups for talent and IP; consolidation among AI-first software companies	SaaS consolidation accelerating

Facilitator Guidance: Acquisition price defensible on "talent and IP" basis but signals end of "standalone AI unicorn" era. Future winners will be integrated into existing platforms or build vertical-specific solutions. 2-3 additional similar acquisitions will occur in R4 as consolidation continues.

Venture returns in AI will trail broader market for 2027-2029 vintages. LPs in AI-focused funds will demand higher scrutiny on startups; next funding round will be much more disciplined on fundamentals (revenue, path to profitability, defensible IP).

R4-03: Labor Market Normalization Signal

Title: Federal Workforce Transition Legislation & New Job Categories Emerge **Classification:** Policy
Time: Mid-round

Narrative: AI Workforce Development and Transition Act of 2029 passes both chambers with bipartisan support and is signed into law. Act creates: (1) "AI Transition Tax Credits" for companies retraining displaced workers; firms receive 30-40% tax credit on retraining expenses (up to \$50K per worker) if demonstrating >70% placement rates into comparable/better roles; (2) "AI Workforce Transition Bonds" funding state-level retraining and extended unemployment benefits (up to 24 months at 80% prior income) for AI-displaced workers. Act also mandates 60 days' notice to workers and union consultation for AI-driven layoffs. Act does not restrict AI deployment but significantly raises cost of labor displacement and incentivizes transparent, managed transitions. Market reaction mixed: tech/services stocks dip on compliance costs; manufacturing/logistics stocks rise (retraining credits offset labor savings). Concurrently, data released showing emergence of new job categories (AI trainers, AI governance specialists, AI safety auditors) and wage stabilization in routine cognitive roles (previous decline arrested). Labor market stabilizing; displacement crisis narrative softening.

Industry Impact:

Industry	Impact	Constraint	Implication
Retail	High (Labor-intensive)	Retraining credits available for retail workforce optimization	Large retailers can offset labor costs with federal credits; smaller retailers face higher net costs
CPG	Moderate (Labor-intensive)	Retraining credits available for distribution and manufacturing workforce	CPG firms can utilize tax credits for workforce transitions
Manufacturing	High (Labor-intensive)	Manufacturing can utilize tax credits; 60-day notice requirement	Workforce reductions now require managed transitions, costly but tax-supported
Logistics	High (Labor-intensive)	Logistics can utilize tax credits; 60-day notice requirement	Workforce reductions now require managed transitions, costly but tax-supported
Healthcare Provider	Moderate	Healthcare systems face labor costs; clinical staff reductions modest due to regulatory needs	Healthcare providers benefit from clinical talent stability

Industry	Impact	Constraint	Implication
Healthcare Payer	Moderate	Claims processing staff transitions; retraining credits available	Payers can utilize credits for claims staff transitions
Finance	High (Labor-intensive)	Finance can utilize tax credits; 60-day notice requirement	Workforce transitions in back-office and junior roles now managed and tax-supported
Consulting	High (Labor-intensive)	Consulting can utilize tax credits; 60-day notice requirement; advisory demand for transition planning	Workforce transitions managed; capitalize on transition advisory demand
Law	Moderate	Associate and paralegal transitions; retraining credits available	Workforce transitions in commoditized practice areas managed and tax-supported
Big Tech	Moderate	Engineering workforce shifts; new job categories (AI trainers, governance specialists) favorable for talent supply in cloud, enterprise, and platform operations	Tech firms see new hiring/training opportunities and talent mobility
B2B/B2C SaaS	Moderate	Engineering workforce shifts; new job categories favorable for talent supply	SaaS firms see new hiring/training opportunities and talent mobility

Facilitator Guidance: Tax credits real and material; detailed IRS guidance specifies "comparable or better" as wage-equivalent within 10% (adjusted by industry/region). At least two major companies will design "AI Transition Programs" positioning as "responsible AI" leaders; marketing value exceeds tax credit value. State implementation of transition bonds will be uneven; some states (California, Massachusetts, Texas) proactive; others lag. Early estimates: 40-60% of companies undertaking AI-driven reductions will pursue tax credits; rest will use traditional severance. New job categories are emerging (AI trainers, auditors, governance specialists); labor market adapting faster than expected.

D. Optional Injects

R4-OPT-01: Black Swan Event

Title: [A] Unexpected AI Capability Breakthrough / [B] Catastrophic AI System Failure **Classification:** Incident/Disruptive Event **Time:** Late round (for dramatic effect)

OPTION A: Unexpected Capability Breakthrough

Major frontier lab announces breakthrough in "long-horizon AI planning": agents can now reliably plan and execute multi-quarter strategic plans ("grow revenue in region X by 20% over three quarters given constraints Y and Z"). System decomposes complex goals into sub-tasks, adapts to mid-course corrections, coordinates across teams. Early case studies show mixed results: some firms report highly effective strategic execution; others require heavy human oversight. Technology positioned as "AI for enterprise strategy and operations"; available for licensing by large enterprises. Announcement met with excitement (strategic AI applications) and skepticism (can AI really do strategy?). Technology credible enough to trigger immediate interest from CFOs and COOs at large enterprises. Market reacts: "strategic AI" category becomes new investment thesis; consulting firms launch "Strategic AI Implementation" practices within weeks.

OPTION B: Catastrophic System Failure

Series of autonomous vehicle accidents across multiple companies over 4-week period causes 25+ fatalities. Preliminary investigation reveals systemic issue affecting multiple vendors' AI perception systems under specific weather/lighting conditions that weren't adequately tested. NHTSA halts autonomous vehicle testing and deployment nationwide (60-day moratorium, likely extended). Congressional hearings triggered. Public trust in autonomous systems collapses (only 18% of public willing to ride in autonomous vehicle, down from 35%). Industry reckoning: AV timeline extended 3+ years; liability framework overhaul required; certification standards fundamentally reworked. Side effect: broader skepticism about autonomous systems in other domains (logistics, manufacturing); autonomous deployment paused across industries pending NHTSA guidance.

Facilitator Guidance (Option A): Capability real but overstated in marketing; actual system requires heavy human input (40-60% of decisions rejected/modified). Customers report 1-2 quarter lead time advantage on tactical moves; meaningful but not decisive. Licensing model expensive (\$500K-\$2M annually); only large enterprises justify cost. Within 18 months, major consulting firm launches "Strategic AI Implementation" practice, positioning as expert integrator.

E. COLLECTIVE BONUS (Optional, within Cross-Industry Discussion)

Facilitator Script (Read Aloud):

"Final round. Does anyone want to recognize an especially strong strategy across all four rounds, or flag one that seems particularly risky? This is your last chance to weigh in on each other's strategies before the debrief."

Procedure:

Step	Action	Time
1	Participation is optional — no one is required to speak	—
2	Participants who wish to respond name one industry as "strong strategy" and/or one as "risky strategy" (cannot nominate own industry)	3 min

Step	Action	Time
3	If 3+ participants agree on the same industry, facilitator applies adjustment	1 min

Scoring:

- If 3+ participants agree an industry has a strong strategy: +2 cumulative score bonus
- If 3+ participants agree an industry has a risky strategy: -2 cumulative score penalty
- Maximum one +2 and one -2 per round
- If no consensus or no one speaks up, no bonus applied

Facilitator Guidance:

- The final Collective Bonus is the culmination of four rounds of observation. Participants now have full track records to evaluate.
- Encourage participants to reference the full arc — not just Round 4 decisions, but how strategies evolved across all rounds.
- Announce results openly — nominations are public, not anonymous.
- The final Collective Bonus results feed directly into debrief discussion: "Why did peers recognize your strategy — or flag it as risky?"

E1. FINAL INDUSTRY HEALTH SIGNALS — ROUND 4

Facilitator Note: After scoring Round 3 decisions (including fallbacks), announce final Industry Health conditions. These are the definitive conditions for the debrief.

Procedure:

1. Calculate cumulative aggregate score for each industry (R1 + R2 + R3 scores, including any Collective Bonus adjustments)
2. Look up final condition for each industry
3. Announce final conditions (~2 minutes)
4. Apply Headwind/Crisis constraints for Round 4 decisions

Debrief reference: Final Industry Health conditions provide the starting point for the debrief discussion. Which industries ended in Surge vs. Crisis? What decisions drove those trajectories?

F. No Private Cards — Round 4

Facilitator Announcement: "No private cards this round. All information is public. Everyone has complete visibility into market conditions and industry performance."

G. INDIVIDUAL INDUSTRY DECISION SUBMISSION REQUIREMENTS — V7.4

Each industry participant submits:

1. Industry Decision (REQUIRED): One strategic choice for long-term competitive positioning (M&A strategy, market positioning, capability buildout, etc.).
 - Use banded framework.
 - 1-2 sentences explaining rationale and expected long-term outcome.

Submission Format:

[INDUSTRY NAME] – ROUND 4 DECISION

DECISION TITLE: [Title]

RATIONALE:

[Why this choice? How does it position your industry for long-term competitive advantage?]

BANDED ASSESSMENT:

- Spend/Commitment: [Band]
- Time-to-Impact: [Band]
- Execution Complexity: [Band]
- Dependency: [Band]
- Scale: [Band]

EXPECTED OUTCOME:

[What long-term competitive position does this establish?]

INDUSTRY REPRESENTATIVE: [Name]

Submission Window: 15 minutes.

Key Notes:

- Round 4 has NO PRIVATE CARDS (all information is transparent).
- One decision per participant.
- Banded framework for all decisions.
- Focus on long-term competitive positioning, not short-term gains.

H. Transition to Debrief

After Round 4 Decision Submissions are collected and the final Collective Bonus is resolved:

Facilitator Transition Script: "Round 4 concludes the Project Threshold simulation. You have now navigated the full arc of AI adoption across four years (2026-2030): Foundation, Acceleration, Reckoning, and Normalization. Market winners and losers are clear. Regulatory frameworks are in place. Labor markets have adjusted. AI is now infrastructure.

We will now move into a 60-minute facilitated debrief where we will discuss:

- What strategic decisions paid off? What did you learn?
- How did cross-industry dynamics and spillover effects shape outcomes?
- What surprised you about the labor market, regulatory environment, or competitive dynamics?
- What would you do differently if you could replay the simulation?
- What are the no-regrets actions your industry will pursue?

This debrief is your opportunity to extract lessons for your organization's actual AI strategy. Let's begin."

End of Round 4: Normalization

07

Supporting Materials

Supporting Materials Index

Overview

This directory contains comprehensive supporting materials for running the Project Threshold Tabletop Exercise in the V7.4 single 8-hour format with up to 11 individual industry representatives + 1-2 facilitators. All files are organized for easy reference during preparation and execution.

File Structure

Core Documents

File	Title	Purpose	Audience
07_Lead_Facilitator_Briefing.md	Lead Facilitator Briefing	Comprehensive guidance for facilitating the exercise, banded scoring framework, debrief prompts, individual model, 11 industries, 5–11 participants	Lead Facilitator
07_Post_Exercise_Synthesis_Memo.md	Post-Exercise Synthesis Memo Template	One-page memo template for industry participants to externalize learning; includes industry-specific examples for V7.4	Participants (after exercise)
07_Room_Setup_Logistics.md	Room Setup & Logistics Guide	Detailed physical setup for individual seating (up to 11 positions), tech requirements, materials checklist, staffing model (1 facilitator + helper)	Logistics Coordinator, Helper

File	Title	Purpose	Audience
07_Diagnostic_Survey.md	Pre-Exercise Diagnostic Survey	8-question survey to capture baseline assumptions from individual participants; analysis and usage guidance for V7.4 industries	Participants (before exercise)
07_Alternative_Scenario_Overrides.md	Alternative Scenario Overrides (Optional)	Non-default scenario configurations for research or follow-up exercises	Facilitators (if needed)
07_Historical_Reference_Data.md	Historical Reference Data	Consolidated historical benchmarks: industry-specific operational data, valuation multiples (EV/EBITDA, P/E), S&P 500 sector returns, Treasury yields	Facilitators (reference)

Related Support Documents (Outside 07 Directory)

File	Title	Purpose	
02_Base_Case_Fallback_Bank.md	Base Case Fallback Bank	Pre-defined default actions for industries without explicit decisions	Facilitators (reference)
05_Participant_Private_Cards/	Private Information Cards	Confidential per-industry cards distributed in Rounds 1–3	Facilitators (distribute during exercise)
04_Participant_AI_Adoption_Arcs/	AI Adoption Arc	Round-by-round AI adoption progression per industry	Participants (reference during rounds)

How to Use These Files

Pre-Exercise (1-2 weeks before)

1. Facilitator reads 07_Lead_Facilitator_Briefing.md (30 min)

- Understand scoring framework, common facilitator mistakes, debrief structure
 - Brief yourself on facilitation scenarios and moves for the 11-industry model
 - Review the 8-hour timing and pacing guidance
 - Plan industry selection if running fewer than 11 participants
2. Logistics coordinator reads 07_Room_Setup_Logistics.md (30 min)
 - Confirm room size and layout (1,000–1,500 sq ft depending on participant count)
 - Create materials checklist and print management plan
 - Arrange tech (projector, timer display, WiFi backup)
 - Plan seating for the specific number of industry positions needed
 3. Optional: Administer 07_Diagnostic_Survey.md to all participants (1-2 weeks before)
 - Send as Google Form or paper survey
 - Analyze results within 1 week
 - Use aggregate results in opening remarks
 - Flag outliers for debrief discussion
 4. All staff review 07_Room_Setup_Logistics.md setup checklists (1 hour before exercise start)
 - 1:00–1:30 before: room layout, tech testing, materials organization
 - 30-45 min before: final walkthrough, participant arrival
 - At start time: welcome, industry card distribution, exercise launch

During Exercise

1. Facilitator uses 07_Lead_Facilitator_Briefing.md as reference
 - Scoring framework for real-time adjudication
 - Debrief prompts for final 60-minute session
 - Common scenarios for participant dynamics management
2. Helper follows 07_Room_Setup_Logistics.md timing and materials distribution
 - Execute minute-by-minute schedule
 - Distribute private information cards from 05_Participant_Private_Cards/ at correct times (Rounds 1-3 only)
 - Post industry scores visibly after each round

Post-Exercise

1. Participants complete 07_Post_Exercise_Synthesis_Memo.md (15 min)
 - Distribute template to each industry participant at end of exercise
 - Participants complete before departure or submit within 24 hours
 - Use memos to document learning and strategic implications
2. Facilitator debriefs using prompts from 07_Lead_Facilitator_Briefing.md (60 min)
 - Ask the core debrief questions
 - Highlight divergence in industry strategies
 - Connect to organizational reality
3. Collect feedback and close
 - Distribute post-exercise feedback form (optional)
 - Thank participants and staff
 - Archive all materials

File Descriptions

07LeadFacilitator_Briefing.md

Scope: 400+ lines

Key Sections:

- Purpose and exercise hypotheses (why we run Project Threshold)
- Facilitator best practices (tone, decision-making style, mistakes to avoid)
- When to increase/decrease pressure
- Practice micro-round guidance
- Collective Bonus mechanics and private card management (referencing 05_Participant_Private_Cards/)
- Simplified staffing model (1 facilitator + 1 helper)
- 8-hour timing with minute-by-minute schedule
- Debrief facilitation (core questions + moves for up to 11 industry participants)
- Common facilitation scenarios and facilitator moves
- Banded scoring framework ({-2, 0, +2} with +/-3 red flags)

Use When: Preparing to facilitate the exercise; during exercise for real-time adjudication and participant management; during debrief

07PostExerciseSynthesisMemo.md

Scope: 400+ lines

Key Sections:

- One-page memo template with 5 key sections:
 - Key Forecast Surprises
 - What We Got Wrong (Starting Assumptions)
 - Strategic Implications (Real-World Industry)
 - Surprises About Decision Quality
 - No-Regrets Actions (3 universal hedges)
- Industry-specific examples (Consumer/Retail, Finance, Healthcare, Consulting, Law)

Use When: End of exercise; distribute template to each industry participant; participants complete in 15 minutes

07RoomSetup_Logistics.md

Scope: 500+ lines

Key Sections:

- Space requirements (1,000–1,500 sq ft depending on participant count)
- Floor plan with flexible seating for 5–11 industry positions
- Detailed position specs for each seat, facilitator station, breakout area, observer seating
- Technology setup checklist (pre-exercise + day-of)
- Participant materials checklist (industry cards, worksheets, private cards from 05_Participant_Private_Cards/)
- 8-hour timing with minute-by-minute schedule
- Resource allocation table (staff, budget, printing volume)
- Contingency planning (tech failures, participant issues)
- Pre-exercise and post-exercise logistics checklists

Use When: Planning logistics 1-2 weeks before; 1:30 before exercise start (setup); during exercise (follow timing)

07DiagnosticSurvey.md

Scope: 350+ lines

Key Sections:

- 8 questions (Q1-Q7 multiple choice; Q8 open-ended)
 - Q1: Speed of AI adoption
 - Q2: Labor displacement magnitude
 - Q3: Regulatory posture
 - Q4: Value capture by industry (ranking from 11 industries)
 - Q5: Regulatory risk in your industry
 - Q6: Market concentration trend
 - Q7: Tail risk / black swan event
 - Q8: Most important strategic question (open-ended)
- Scoring rubric for each question
- Facilitator guidance for analysis and usage
- How to use results before, during, and after exercise

Use When: Administer 1-2 weeks before exercise; analyze results 1 week before; reference during exercise for outlier prompting; use in debrief for assumption testing

07AlternativeScenario_Overrides.md

Scope: 200+ lines (much shorter than V3)

Key Sections:

- Clear note that alternatives are OPTIONAL extensions (not recommended for first-time exercises)
- When to use alternatives (only for research or follow-up exercises)

- Four alternative configurations:
 - Slow Deployment + Type A (longer competitive window)
 - Baseline + Type B (more autonomous AI, more disruption)
 - Fast Deployment + Type A (extreme time pressure)
 - Fast Deployment + Type B (most transformative scenario)
- For each: key assumption change, round-by-round narrative changes, impact on industry decisions, use cases
- Industry-specific override notes for Consulting and Law
- Implementation guidance and recommended pairings for multiple exercises

Use When: You are running multiple Project Threshold exercises and want to vary the scenario; you have specific research questions; you have advanced participants

Quick Reference Checklist

1 Week Before Exercise:

- Read 07LeadFacilitator_Briefing — note banded framework and individual industry model sections
- Read 07RoomSetup_Logistics — note individual seating arrangement for target participant count
- Review 02BaseCaseFallbackBank (Base Case Fallback Bank) — understand default fallback options for inactive industries
- Administer 07DiagnosticSurvey if desired
- Select which industries are in play (if fewer than 11 participants)
- Book room (1,000–1,500 sq ft depending on participant count)
- Arrange catering (coffee, lunch, snacks)
- Test tech (projector, WiFi, timer)
- Print materials (industry cards, worksheets, injects, private cards from 05_Participant_Private_Cards/)

1 Day Before Exercise:

- Final room walkthrough
- Confirm all tech working
- Brief facilitator and helper on 07LeadFacilitator_Briefing content
- Organize all materials by industry

Day of Exercise (1:30 before start):

- Follow setup checklist in 07RoomSetup_Logistics (1:00–1:30 before, 0:30–0:45 before, 0:15–0:30 before)
- Facilitate using timing and scenarios in 07LeadFacilitator_Briefing
- Track decisions and scores using 07LeadFacilitator_Briefing framework

End of Exercise:

- Distribute 07PostExerciseSynthesisMemo to industry participants
- Debrief using questions from 07LeadFacilitator_Briefing

- [] Collect feedback
- [] Pack and archive all materials

Document Versions & Updates

File	Version	Date	Status
07LeadFacilitator_Briefing.md	V7.4	Mar 2026	Final
07PostExerciseSynthesis Memo.md	V7.4	Mar 2026	Final
07RoomSetup_Logistics.md	V7.4	Mar 2026	Final
07DiagnosticSurvey.md	V7.4	Mar 2026	Final
07AlternativeScenario_Overrides.md	V7.4	Mar 2026	Final
07HistoricalReference_Data.md	V7.4	Mar 2026	Final

Support & Troubleshooting

For facilitation questions: See *07LeadFacilitator_Briefing.md* sections on "Common Facilitator Mistakes" and "Facilitation Scenarios & Moves"

For logistics issues: See *07RoomSetup_Logistics.md* section on "Contingency Planning"

For survey analysis: See *07DiagnosticSurvey.md* section on "Survey Administration Guidance"

For scoring questions: See *07LeadFacilitator_Briefing.md* section on "Banded Scoring Framework"

Directory Status: Complete **Format:** Markdown (.md) files, UTF-8 encoding **Total Lines:** ~3,000 lines across 6 files **Last Updated:** March 2026 **Version:** V7.4 (individual participant model, 11 industries, Collective Bonus + Facilitator Market Shock) **Prepared by:** Project Threshold Development Team

Alternative Scenario Overrides (Optional Extensions)

NOTE: V7.4 is optimized for single 8-hour format with Baseline Speed + Type A (Copilots Everywhere) configuration. Alternative scenarios below are optional extensions only.

Overview

Project Threshold V7.4 **default and recommended configuration** is **Baseline Speed + Type A (Copilots Everywhere)**. This document provides optional narrative and mechanic overrides for 4 alternative configurations if you want to extend or vary the exercise beyond the standard format.

Alternative Configurations (Optional):

1. Slow Deployment + Type A (Copilots Everywhere) — Same capability type, slower adoption rate
2. Baseline + Type B (Agentic Workflows) — Same speed, different AI type (more autonomous)
3. Fast Deployment + Type A (Copilots Everywhere) — Faster adoption of copilots
4. Fast Deployment + Type B (Agentic Workflows) — Faster adoption of agentic systems

Important: These alternatives are provided for specialized use cases (e.g., follow-up workshops, deep-dive analysis for specific industries, or research purposes). They are NOT recommended for first-run Project Threshold exercises.

When to Use Alternative Configurations

Use Baseline + Type A (Default) if:

- This is your first Project Threshold exercise
- You have participants unfamiliar with AI tabletop exercises
- You want to complete exercise in single 8-hour session
- You want broad applicability across industries
- You are not researching a specific AI scenario

Use Alternative Configurations only if:

- You are running multiple Project Threshold sessions and want to contrast different AI futures
 - You have a specific research question (e.g., "What if agentic workflows mature faster than copilots?")
 - You have individual participants with advanced strategic expertise comfortable with complexity
 - You have extended time (2-day format recommended)
-

Configuration 1: Slow Deployment + Type A (Copilots Everywhere)

Key Assumption Change

Foundation models advance on schedule, but enterprise adoption is slowed by organizational friction, regulatory caution, and integration costs. AI tools exist and are capable, but deployment is careful and phased. By 2030, copilot adoption in large enterprises is 40-50% (vs. 70-80% in Baseline).

Round-by-Round Narrative Changes

Practice Round & Round 1: (Same as Baseline)

Round 2: Emphasis slower enterprise adoption than Baseline. "AI copilots are now available, but adoption rates lag expectations. Enterprise integration challenges are higher than anticipated. Smaller firms especially struggle with change management."

Round 3: "Adoption accelerates slightly but remains below Baseline pace. Regulatory clarity helps somewhat. Winners are firms that invested early in training and change management. Laggards face higher deployment costs in later rounds."

Round 4: "By 2030, copilot adoption has reached 45-50% of large enterprises. Market consolidation is real but less concentrated than Baseline due to longer competitive window. Mid-market players have slightly better survival odds."

Impact on Industry Decisions

- Participants have more time to make strategic pivots before competitive pressure forces action
- Regulatory constraints feel more binding (regulatory injects feel more realistic)
- Labor displacement is slower and more dispersed
- M&A activity is lower (less pressure to consolidate)
- Margin expansion is gradual, not rapid
- Consulting/Law specific: Slower disruption gives professional services industries more time to redesign business models, but also reduces urgency and may create complacency

Use This If:

- You want to test strategies in a more cautious/regulatory-constrained scenario
- You want to see how participants respond to slower-moving competition
- You have Healthcare Provider, Law, or Manufacturing participants who want to model realistic deployment timelines

Configuration 2: Baseline + Type B (Agentic Workflows)

Key Assumption Change

Instead of copilots (AI assists humans in specific tasks), the scenario evolves toward agentic workflows (AI systems that autonomously execute sequences of tasks with human oversight). This is more disruptive to organizational structures and labor, but offers higher productivity gains.

Round-by-Round Narrative Changes

Practice Round & Round 1: (Same as Baseline setup, but inject language emphasizes autonomous capability)

Round 2: "AI is now capable of executing full workflows autonomously (e.g., autonomous trading, autonomous customer service, autonomous supply chain optimization). This creates both opportunity and risk. Organizations must decide: deploy with oversight (safer), or deploy autonomously (faster)."

Round 3: "Agentic deployment is accelerating in competitive industries. Organizations that deployed with strong oversight (human-in-loop) are capturing benefits safely. Organizations that deployed fully autonomously are seeing both wins and failures. Regulatory scrutiny is intense."

Round 4: "Agentic workflows are now standard in Finance, Big Tech, and Logistics. Labor displacement is visible and politically sensitive. Healthcare Provider and Law are lagging due to regulatory constraints on autonomous systems."

Impact on Industry Decisions

- Higher upside if participants bet on aggressive autonomous deployment (higher productivity, more risk)
- Labor displacement is more severe and concentrated (triggers stronger labor/regulatory injects)
- Governance and oversight become critical risk factors
- Tail risk injects (major AI system failure) hit harder
- Consulting specific: Agentic workflows could automate entire engagement workstreams, making the talent pipeline disruption even more acute
- Law specific: Autonomous legal analysis raises unauthorized practice of law concerns in a way copilot-assisted work does not; malpractice liability is more severe
- Healthcare Provider specific: Autonomous clinical decision-making faces much stricter FDA scrutiny than physician-assisted AI

Use This If:

- You want to stress-test governance and risk management capabilities
- You want to emphasize labor displacement and policy responses
- You want to explore more radical AI scenarios with longer-term implications

Configuration 3: Fast Deployment + Type A (Copilots Everywhere)

Key Assumption Change

Foundation models advance rapidly, and enterprise adoption is faster than Baseline. By 2028, copilot adoption in large enterprises reaches 85%+. Competitive pressure is extreme and relentless.

Round-by-Round Narrative Changes

Practice Round & Round 1: (Same as Baseline)

Round 2: "AI copilot adoption is outpacing expectations. 60%+ of large enterprises now have AI copilots in production. Competitive pressure is extreme. Laggards are falling behind rapidly."

Round 3: "Adoption has reached 75%+ of large enterprises. Market consolidation is accelerating. Mid-market players that have not adopted are under existential pressure."

Round 4: "By 2030, 85%+ of large enterprises have deployed AI copilots. Market structure has consolidated dramatically. Winners and losers are clearly defined."

Impact on Industry Decisions

- Time pressure is intense from the start
- Industries that don't move fast early are essentially eliminated by Round 3
- M&A consolidation happens earlier and at larger scale
- Labor displacement is more severe and more rapid
- Regulatory backlash injects hit harder (political pressure visible earlier)
- Consulting specific: Client demand for AI advisory surges but so does pricing pressure; first-mover consulting firms capture outsized share
- Law specific: Firms that adopt early gain massive efficiency advantage in contract review and due diligence; laggards lose competitive bids

Use This If:

- You want to emphasize speed-to-market and execution risk
- You want to create high-pressure competitive dynamics
- You want to show how first-mover advantage compounds over time

Configuration 4: Fast Deployment + Type B (Agentic Workflows)

Key Assumption Change

Agentic workflows mature rapidly and are adopted aggressively across industries. This is the most disruptive scenario: autonomous AI systems executing business processes with minimal human oversight.

Round-by-Round Narrative Changes

Practice Round & Round 1: (Same as Baseline setup)

Round 2: "Agentic workflows are emerging faster than expected. Early adopters (Finance, Big Tech) are seeing dramatic productivity gains (50%+ in some functions). But failures are also visible. Regulatory uncertainty is high."

Round 3: "Agentic deployment is accelerating in competitive industries. Market consolidation is extreme. Organizations that navigated governance successfully are pulling away. Organizations that over-estimated safety are facing crises."

Round 4: "By 2030, agentic workflows are standard in unregulated industries but heavily restricted in regulated industries. Market structure is highly bifurcated: winners have 5-10x productivity; laggards have exited or been acquired."

Impact on Industry Decisions

- Governance and risk management are make-or-break decisions
- Tail risk injects are frequent and severe
- Labor displacement is dramatic and politically visible
- Regulatory bifurcation between industries is extreme
- Healthcare Provider and Law face strategic paralysis: Can't move fast without regulatory approval; can't get approval without moving fast
- Consulting faces existential question: If agentic workflows can execute full consulting engagements, what is the firm's value proposition?
- Finance may lead: Trading and risk management agentic systems create massive competitive advantage for early movers

Use This If:

- You want to explore the most transformative AI scenario
- You want to emphasize governance, tail risk, and regulatory bifurcation
- You want to test leadership's appetite for real transformational risk

Industry-Specific Override Notes (V7.4 Additions)

Consulting Overrides

Across all alternative configurations, Consulting faces amplified versions of the baseline dynamics:

- Slow Deployment: More time to redesign business model, but risk of complacency; clients may reduce advisory spend if AI adoption is slow
- Agentic Workflows (Type B): Existential threat to leverage model; entire engagement workstreams could be automated; need to articulate post-agentic value proposition

- Fast Deployment: Advisory demand surges but pricing collapses; first movers capture outsized share; talent pipeline disruption accelerates

Collective Bonus Override for Consulting: In alternative configurations, expect Collective Bonus dynamics to shift for Consulting. In Fast + Type B, peers may consistently nominate Consulting as a "risky strategy" (business model under existential threat). In Slow + Type A, Consulting may receive "strong strategy" nominations (stable advisory demand with time to transform).

Law Overrides

Across all alternative configurations, Law faces unique regulatory complexity:

- Slow Deployment: Bar associations have more time to issue guidance; malpractice case law may begin to develop, reducing uncertainty
- Agentic Workflows (Type B): Autonomous legal analysis raises unauthorized practice of law concerns; malpractice liability for AI-generated legal advice becomes acute; some jurisdictions may ban agentic legal AI entirely
- Fast Deployment: Firms that adopt early gain massive efficiency advantage; bar rule arbitrage (deploying in permissive jurisdictions first) becomes viable strategy

Facilitator Note for Law: Law industry dynamics are primarily partnership-based (not public equity). Facilitator Market Shock constraints and Collective Bonus patterns for Law should reflect the industry's unique economic trajectory. Consider whether the alternative scenario benefits or threatens the billable hour model when selecting Market Shock targets in Round 2.

Implementation Guidance for Alternative Configurations

If using an alternative configuration:

1. Modify Round situation slides — Update narrative language in Rounds 2-4 to match configuration
2. Adjust inject timing and intensity — Some injects become more relevant, others less
3. Emphasize different scoring dimensions — Slow configs emphasize Strategic Fit; Fast configs emphasize Execution Risk; Agentic configs emphasize Tail Risk
4. Brief participants explicitly — Make clear which configuration you're running. Example: "In this version, we're modeling Fast Deployment + Agentic Workflows. This is a more transformative scenario."
5. Adjust Facilitator Market Shock intensity — In more extreme scenarios, the facilitator may impose additional constraints or target more industries in Round 2

Do NOT:

- Mix multiple alternative configurations in a single exercise (this creates confusion)
- Use alternative configurations without clear participant briefing (surprises undermine learning)
- Expect alternative configurations to finish in standard 8-hour time (some may run long)

Recommended Pairings (if running multiple exercises)

Exercise 1: Baseline + Type A (Default)

- Best for first-time facilitators and participants
- Broadest applicability across industries
- Completes in 8 hours comfortably

Exercise 2 (Follow-up): Baseline + Type B (Agentic Workflows)

- Assumes participants already familiar with Project Threshold format
- Explores more transformative scenario
- Emphasizes governance and tail risk
- Good for C-suite or risk committees
- Particularly illuminating for Consulting and Law participants

Exercise 3 (Research): Fast Deployment + Type A or B

- For strategy teams wanting to explore high-pressure competitive dynamics
- Can be condensed to 6 hours or extended to 10 hours depending on emphasis

Document Version: Project Threshold V7.4 — Alternative Scenario Overrides (Optional) **Last Updated:** March 2026 **Note:** Alternatives are provided for reference only. Default Baseline + Type A is strongly recommended for all standard Project Threshold exercises.

Pre-Exercise Diagnostic Survey

Administer 1-2 Weeks Before Exercise

Purpose: Capture baseline assumptions from participants. Use results to (a) identify outliers and prepare for debate, (b) inform debrief by contrasting initial assumptions with post-exercise thinking, (c) calibrate individual participants during exercise based on aggregate results.

Format: 8-question survey, 15-20 min completion time **Scoring:** Multiple choice; each question has a rubric **Participant Model:** Individual industry representatives respond individually (not as teams)

Question 1: Speed of Enterprise AI Adoption

Question Text: "By 2028 (2 years from now), what percentage of large enterprises (>\$1B revenue) will have deployed at least one AI copilot tool into their primary business processes?"

Answer Options:

- A) Less than 25% (slow adoption)
- B) 25-50% (moderate adoption)
- C) 50-75% (rapid adoption)
- D) More than 75% (ubiquitous adoption)

Scoring Rubric:

- A = Pessimistic on speed: Participant expects significant adoption friction, regulatory constraints, or organizational resistance.
- B = Moderate: Steady but not explosive adoption curve.
- C = Rapid: Majority of large enterprises move quickly.
- D = Ubiquitous: AI copilots become standard in most enterprises within 2 years.

Facilitator Guidance:

- If >40% select A/B, expect skeptical participant behavior. Prepare for debate on plausibility.
 - If >40% select D, expect overconfidence. Use injects to surface execution barriers.
 - Outliers (A vs D) should be flagged for debrief discussion.
-

Question 2: Labor Displacement Magnitude (By 2029)

Question Text: "By 2029 (3 years from now), approximately how many jobs in the US will have been structurally displaced by AI (meaning role eliminated, not retrained)?"

Answer Options:

- A) Less than 500,000 jobs (minimal displacement)
- B) 500,000 - 2 million jobs (moderate displacement)
- C) 2-5 million jobs (significant displacement)

- D) More than 5 million jobs (major structural displacement)

Scoring Rubric:

- A = Optimistic on labor: Most workers will retrain into new roles.
- B = Moderate: Some displacement but manageable transition.
- C = Pessimistic: Significant net job loss. Retraining feasibility questioned.
- D = Very pessimistic: Structural unemployment or major workforce dislocation.

Facilitator Guidance:

- If cohort splits between A/B and C/D, expect tension in discussions about labor policy.
 - If cohort heavily selects C/D, expect aggressive labor protection proposals.
 - This question signals how participants will approach workforce transition strategies.
-

Question 3: Regulatory Posture (By 2028)

Question Text: "By 2028, which best describes the regulatory environment for AI in the US?"

Answer Options:

- A) Light-touch regulation: Industry self-regulation dominates; minimal binding government rules.
- B) Sector-specific light regulation: Industry regulators (FDA, SEC, EEOC, state bar associations) issue guidance; enforcement is light.
- C) Moderate federal regulation: Congress passes binding rules on transparency, explainability, audit rights; enforcement is increasing.
- D) Strict regulation: Congress passes restrictive rules (liability caps, deployment restrictions, mandatory human oversight); enforcement is aggressive.

Scoring Rubric:

- A = Pro-innovation: Low regulatory barriers; firms can move fast.
- B = Cautious: Piecemeal guidance but no major binding rules yet.
- C = Balanced: Rules will be established and enforced moderately.
- D = Restrictive: Significant regulatory constraint on deployment.

Facilitator Guidance:

- Project Threshold default assumes B->C transition (light guidance in early rounds, binding rules in later rounds).
 - If cohort clusters on A, they may be surprised by regulatory injects.
 - If cohort clusters on D, regulatory injects will feel realistic to them.
-

Question 4: Value Capture by Industry (Rank Top 3)

Question Text: "Which three industries will capture the most economic value (earnings uplift, margin expansion, market share gain) from AI between 2026-2030? Rank them 1 (most), 2, 3."

Answer Options: (Select 3 from this list and rank)

- Retail
- CPG
- Healthcare Provider
- Healthcare Payer
- Finance
- Consulting
- Law
- Manufacturing
- Logistics
- Big Tech
- B2B/B2C SaaS

Scoring Rubric (Aggregate):

- Industries consistently ranked top 3: Finance, Big Tech, SaaS (usually expected winners)
- Industries sometimes ranked: Consulting, Healthcare Payer, Logistics (emerging value capture stories)
- Industries rarely ranked: Healthcare Provider, Manufacturing, CPG, Law, Retail (usually expected laggards due to regulatory/complexity/data constraints)
- Outliers: Anyone ranking Healthcare Provider, Law, or Manufacturing as top 3 has a contrarian thesis

Facilitator Guidance:

- If your industry representative is from a lower-ranked industry, they may rank their own industry higher. That's fine; creates ownership during exercise.
- If cohort overwhelmingly selects same top 3, you have consensus on winners. Use that in debrief: "Everyone thinks Finance & Big Tech will win. What happens to the others?"
- Outliers with idiosyncratic rankings should be probed: "You rank Law as #2. Why? What's your thesis?"
- Compare rankings between related industries: Did anyone rank Consulting higher than Finance? Healthcare Payer higher than Provider?

Question 5: Regulatory Risk in Your Industry (Self-Assessed)

Question Text (Asked by Industry): "In your assigned industry [INDUSTRY], how much regulatory risk do you foresee from AI deployment between 2026-2030?"

Options:

- A) Very low risk: Regulation will not meaningfully constrain our AI strategy.
- B) Low-to-moderate risk: Regulation may require compliance investments, but won't change our core strategy.
- C) Moderate-to-high risk: Regulation could force us to pivot our strategy or slow deployment.

- D) Very high risk: Regulatory restriction could prevent our AI strategy from succeeding.

Scoring Rubric:

- A = Confident industry: Participant believes their industry has regulatory tailwinds.
- B = Cautious: Compliance costs but manageable.
- C = Wary: Real regulatory risk; strategy pivots likely.
- D = Defensive: Regulation is a fundamental threat.

Typical Risk Profiles by Industry:

- Healthcare Provider: 70-80% select C/D (clinical AI, FDA approval, physician liability)
- Healthcare Payer: 60-70% select C/D (algorithmic coverage decisions, state insurance oversight)
- Law: 60-70% select C/D (bar rules, malpractice liability, unauthorized practice concerns)
- Finance: 50-60% select C/D (regulatory bifurcation, algorithmic fairness audits)
- Consulting: 30-40% select C/D (client data handling, professional licensing)
- Big Tech: 40-50% select C/D (antitrust, data privacy, content moderation)
- B2B/B2C SaaS: 30-40% select C/D (data privacy, contractual liability)
- Retail: 20-30% select C/D (labor risk more than regulatory)
- CPG: 20-30% select C/D (consumer data privacy)
- Manufacturing: 25-35% select C/D (safety regulations, OT system compliance)
- Logistics: 20-30% select C/D (safety, autonomous vehicle regulation)

Facilitator Guidance:

- Healthcare Provider and Law representatives will likely select C/D. That's appropriate; use it to frame decisions.
- If a Healthcare Provider representative selects A/B, they're either optimistic or not thinking deeply. Prompt them: "FDA approval timelines are 12+ months. How does that affect your deployment timeline?"
- Industry-specific outliers are interesting. If a Big Tech representative selects B while others in cohort select C/D, that person has a specific view worth probing.

Question 6: Market Concentration Trend

Question Text: "By 2030, will AI lead to greater market concentration (fewer, larger players) or market decentralization (more competition from new AI-native entrants)?"

Options:

- A) Significant decentralization: AI tools will be so widely available that new entrants can compete with incumbents. Market share will disperse.
- B) Slight decentralization: New entrants will emerge, but incumbents will retain most share.
- C) Stable concentration: Market structure will remain roughly as it is today.
- D) Increased concentration: Large, well-capitalized incumbents will pull away. Mid-market players will consolidate or exit.

Scoring Rubric:

- A = Optimistic on disruption: Participant believes AI democratizes competition. Incumbents are vulnerable.
- B = Slightly optimistic: Some new entrants but incumbents retain advantage.
- C = Neutral: No major shift in market structure.
- D = Pessimistic on disruption: Participant believes AI reinforces incumbent advantage.

Distribution Benchmarks:

- 10% select A
- 25% select B
- 20% select C
- 45% select D

Facilitator Guidance:

- Majority will select D (consolidation). This is the Project Threshold default.
- Anyone selecting A/B is betting on disruption. This person may be a VC investor or entrepreneur skeptical of incumbents.
- Use this question to prime M&A discussions in the exercise. Expect participants to consider acquisitions by Round 3-4 if they're trailing.

Question 7: Tail Risk / Black Swan Event

Question Text: "What's the event you worry about most between 2026-2030? (Select one)"

Options:

- A) Major AI system failure (misdiagnosis, trading error, autonomous vehicle crash) causing significant harm and regulatory backlash.
- B) Synthetic media / deepfake incident undermining trust in information and institutions.
- C) AI-driven market manipulation or systemic financial risk.
- D) Severe labor dislocation and political backlash forcing rapid regulatory reversal.
- E) None of the above; I think tail risks are overstated.

Scoring Rubric (Aggregate):

- Count distribution of tail risk concerns.
- A/B/C/D: Participant has identified a specific systemic risk.
- E: Participant is risk-neutral or dismissive of tail risks.

Distribution Benchmarks:

- 25% select A (major incident)
- 20% select B (deepfakes)
- 20% select C (financial manipulation)
- 25% select D (labor dislocation)
- 10% select E (overstated)

Facilitator Guidance:

- If cohort heavily selects A/B/C, they're primed for crisis injects. Use injects to surface their concerns realistically.
 - If cohort heavily selects D (labor), frame debrief around policy responses and stakeholder management.
 - Anyone selecting E is an optimist. They may be surprised by injects. Use that surprise as a teaching moment.
 - Use this question to foreshadow injects: "Several of you are worried about deepfakes. That will come up in Round 3."
-

Question 8: Most Important Strategic Question (Open-Ended)

Question Text: "What's the single most important strategic question your organization is grappling with in AI? (100 words max)"

Examples of Good Answers:

- "How do we move fast enough to keep up with larger competitors while managing execution and governance risks?"
- "How do we balance AI-driven efficiency with labor protection and brand reputation?"
- "How do we navigate regulatory uncertainty while committing capital to long-term AI investments?"
- "How do we compete with AI-native startups that have no legacy systems or embedded organizational constraints?"
- "How do we preserve our junior talent pipeline when AI can perform 60-70% of entry-level work?" (Consulting/Law specific)
- "How do we transition our billable hour model without destroying revenue during the transition?" (Law specific)

Facilitator Guidance:

- Read all open-ended responses before the exercise.
 - Identify common themes. If multiple participants mention "execution risk," that's a signal the cohort is worried about capability constraints.
 - Look for outlier concerns. If one person mentions "AI-driven market manipulation as systemic risk" and no one else does, that person is flagged for systemic risk concern.
 - Use these responses to personalize injects. If cohort is worried about execution, create injects that test execution capability.
 - In debrief, reference these responses: "Earlier, you said your biggest concern was execution risk. Did the exercise validate that? Did you find your execution assumptions too optimistic?"
-

Survey Administration Guidance

When to Administer:

- 1-2 weeks before exercise
- Digital form preferred (Google Forms, Qualtrics) for easy aggregation
- Paper forms acceptable if digital infrastructure unavailable
- Completion time: 15-20 minutes
- Request brief written responses for Q8 (open-ended)

How to Analyze Results:

1. Aggregate Results (5 min): Create summary showing distribution of answers to Q1-Q7. Example output:

Q1 (Adoption Speed): A: 10%, B: 30%, C: 40%, D: 20%

Q3 (Regulatory Posture): A: 8%, B: 32%, C: 42%, D: 18%

Q6 (Concentration): A: 12%, B: 28%, C: 18%, D: 42%
2. Identify Outliers (10 min): Note individuals with idiosyncratic responses:
 - "John (Finance) selected very pessimistic on adoption (A) but optimistic on labor (A). That's coherent: slow adoption = less displacement."
 - "Sarah (Consulting) selected pro-disruption (A) on Q6 but very pessimistic on regulation (D). She thinks regulation will be so restrictive that new entrants can't compete anyway."
3. Thematic Analysis of Q8 (10 min): Group open-ended responses by theme:
 - "Execution & Capability Risk" (mentioned by 5 participants)
 - "Labor & Workforce Transition" (mentioned by 6 participants)
 - "Regulatory Uncertainty" (mentioned by 4 participants)
 - "Business Model Disruption" (mentioned by 3 participants — especially Consulting, Law)
 - "Market Consolidation Pressure" (mentioned by 3 participants)

How to Use Results During Exercise:

1. Before Exercise Starts:
 - Present 1-minute summary of aggregate results: "Most of you expect rapid adoption and increased concentration. A few of you are more pessimistic on regulation."
 - Call out themes from Q8: "Several of you flagged execution risk as your biggest concern. We'll test that in the exercise."
2. During Exercise:
 - Reference baseline assumptions when adjudicating: "Your decision assumes light-touch regulation, which is consistent with your Q3 answer (A). But the scenario is modeling binding regulation by Round 3. That assumption is being challenged."
 - Use outliers to fuel debates: "John selected very pessimistic adoption. When competitors announce rapid deployments, are you still skeptical? What would change your view?"
3. During Debrief:
 - Compare initial assumptions to post-exercise views: "Before the exercise, 40% of you selected C (rapid adoption). After seeing injects and competitive pressure, do you still think adoption will be that fast? How did the scenario challenge your assumptions?"
 - Highlight learning: "You were right to worry about labor dislocation (Q2). The scenario bore that out. Your concern was prescient. Now, what do you actually do about it in your real organization?"

Document Version: Project Threshold V7.4 — Pre-Exercise Diagnostic Survey **Last Updated:** March 2026 **Format:** Single 8-question survey; V7.4 eleven-industry configuration (5-11 participants + 1-2 facilitators), individual participant model

Historical Reference Data

Usage note: This document consolidates historical market and financial data for all 11 industries. It serves as a reference for grounding decisions in real market dynamics. Data covers 2015-2025 where available.

Retail

The following tables provide actual historical data (2015-2025) for retail-focused metrics referenced in this industry packet. These benchmarks help contextualize the performance targets and competitive dynamics outlined in the exercise scenario.

US Retail Sales Growth (Annual %, 2015-2025)

Year	Total Retail Sales Growth
2015	2.0%
2016	3.2%
2017	4.0%
2018	2.9%
2019	3.5%
2020	5.2%
2021	10.5%
2022	4.3%
2023	2.5%
2024	3.0%
2025	2.8% (approx.)

Retail Sector Operating Margins (% , 2015-2025)

Year	Margin
2015	4.2%
2016	4.0%
2017	4.1%
2018	4.3%
2019	4.2%
2020	4.8%
2021	5.1%
2022	4.7%
2023	4.4%
2024	4.3%
2025	4.2%

Consumer Confidence Index (Annual Average, 2015-2025)

Year	CCI
2015	98.0
2016	98.6
2017	100.7
2018	99.9
2019	103.3
2020	92.6
2021	109.7
2022	99.5
2023	102.6
2024	104.7
2025	103.2 (approx.)

S&P 500 Consumer Discretionary Annual Returns (% , 2015-2025)

Year	Return
2015	-6.6%
2016	6.3%
2017	8.6%
2018	-12.3%
2019	26.4%
2020	36.5%
2021	27.3%
2022	-23.5%
2023	22.5%
2024	18.8%
2025	12.0% (approx.)

Sources: U.S. Census Bureau (Retail Sales), Conference Board (Consumer Confidence Index), S&P Global Market Intelligence, MSCI. Some 2024-2025 figures are approximations based on available quarterly data.

CPG

The following tables provide actual historical data (2015-2025) for CPG-focused metrics referenced in this industry packet. These benchmarks help contextualize the performance targets and competitive dynamics outlined in the exercise scenario.

US CPG Industry Revenue Growth (Annual %, 2015-2025)

Year	Growth Rate
2015	0.8%
2016	1.1%
2017	1.5%
2018	2.1%
2019	1.9%
2020	5.5%
2021	5.2%
2022	3.8%
2023	1.2%
2024	1.5%
2025	1.8% (approx.)

CPG Sector Operating Margins (% , 2015-2025)

Year	Margin
2015	13.5%
2016	13.2%
2017	13.4%
2018	13.8%
2019	13.6%
2020	14.1%
2021	14.5%
2022	14.2%
2023	13.9%
2024	13.8%
2025	13.7%

Consumer Confidence Index (Annual Average, 2015-2025)

Year	CCI
2015	98.0
2016	98.6
2017	100.7
2018	99.9
2019	103.3
2020	92.6
2021	109.7
2022	99.5
2023	102.6
2024	104.7
2025	103.2 (approx.)

S&P 500 Consumer Staples Annual Returns (% , 2015-2025)

Year	Return
2015	7.5%
2016	12.0%
2017	7.1%
2018	-4.8%
2019	16.0%
2020	12.2%
2021	11.8%
2022	-4.9%
2023	4.2%
2024	8.1%
2025	6.5% (approx.)

Sources: Bureau of Economic Analysis (CPG/Consumer spending), Conference Board (Consumer Confidence Index), S&P Global Market Intelligence, MSCI. Some 2024-2025 figures are approximations based on available quarterly data.

Healthcare Provider

US National Health Expenditure Growth (Annual %, 2015-2025)

Year	Growth Rate
2015	4.7%
2016	4.9%
2017	4.8%
2018	4.5%
2019	4.6%
2020	9.7%
2021	2.3%
2022	4.1%
2023	4.2%
2024	4.5%
2025	4.8% (approx.)

Hospital Operating Margins (% , 2015-2025)

Year	Margin
2015	2.8%
2016	3.1%
2017	3.4%
2018	3.2%
2019	3.5%
2020	4.2%
2021	4.8%
2022	4.1%
2023	3.9%
2024	3.6%
2025	3.4% (approx.)

S&P 500 Healthcare Annual Returns (% , 2015-2025)

Year	Return
2015	4.3%
2016	20.6%
2017	20.6%
2018	-5.3%
2019	22.5%
2020	13.3%
2021	16.8%
2022	-4.2%
2023	7.5%
2024	19.3%
2025	14.8% (approx.)

Medicare Spending Growth (Annual %, 2015-2025)

Year	Growth Rate
2015	5.5%
2016	5.1%
2017	4.9%
2018	5.2%
2019	5.4%
2020	9.8%
2021	2.8%
2022	5.6%
2023	5.1%
2024	5.3%
2025	5.5% (approx.)

Sources: Centers for Medicare & Medicaid Services (CMS), National Health Expenditure accounts, American Hospital Association (AHA), Centers for Medicare & Medicaid Services (CMMI), S&P Global Market Intelligence. Some 2024-2025 figures are approximations based on available quarterly data.

Healthcare Payer

Health Insurance Sector Combined Ratios (% , 2015-2025)

Year	Combined Ratio
2015	96.2%
2016	96.8%
2017	97.1%
2018	98.3%
2019	97.8%
2020	95.2%
2021	93.8%
2022	99.1%
2023	98.6%
2024	97.2%
2025	96.8% (approx.)

Medicaid Spending Growth (Annual %, 2015-2025)

Year	Growth Rate
2015	5.9%
2016	5.7%
2017	5.3%
2018	4.8%
2019	4.5%
2020	14.1%
2021	2.1%
2022	4.8%
2023	4.9%
2024	5.0%
2025	5.2% (approx.)

S&P 500 Healthcare Annual Returns (% , 2015-2025)

Year	Return
2015	4.3%
2016	20.6%
2017	20.6%
2018	-5.3%
2019	22.5%
2020	13.3%
2021	16.8%
2022	-4.2%
2023	7.5%
2024	19.3%
2025	14.8% (approx.)

Federal Funds Rate (Year-End % , 2015-2025)

Year	Rate
2015	0.50%
2016	0.75%
2017	1.50%
2018	2.25%
2019	1.75%
2020	0.08%
2021	0.08%
2022	4.25%
2023	5.33%
2024	4.25%
2025	3.75% (approx.)

Sources: National Association of Insurance Commissioners (NAIC), Centers for Medicare & Medicaid Services (CMS), National Health Expenditure accounts, S&P Global Market Intelligence, Federal Reserve Board. Some 2024-2025 figures are approximations based on available quarterly data.

Finance

The following tables provide actual historical data (2015-2025) for financial services metrics referenced in this industry packet. These benchmarks help contextualize the performance targets and competitive dynamics outlined in the exercise scenario.

US Bank ROE (Return on Equity, %, 2015-2025)

Year	ROE
2015	9.3%
2016	10.1%
2017	10.8%
2018	11.3%
2019	10.7%
2020	7.4%
2021	10.2%
2022	10.9%
2023	11.8%
2024	12.4%
2025	12.0% (approx.)

P&C Insurance Combined Ratios (% , 2015-2025)

Year	Combined Ratio
2015	99.2%
2016	99.8%
2017	101.8%
2018	100.5%
2019	99.3%
2020	98.2%
2021	96.8%
2022	101.2%
2023	99.6%
2024	98.8%
2025	97.5% (approx.)

S&P 500 Financials Annual Returns (% , 2015-2025)

Year	Return
2015	-3.6%
2016	16.3%
2017	17.5%
2018	-17.3%
2019	28.6%
2020	-2.1%
2021	27.9%
2022	-16.8%
2023	18.1%
2024	22.4%
2025	16.5% (approx.)

Federal Funds Rate (Year-End %, 2015-2025)

Year	Rate
2015	0.50%
2016	0.75%
2017	1.50%
2018	2.25%
2019	1.75%
2020	0.08%
2021	0.08%
2022	4.25%
2023	5.33%
2024	4.25%
2025	3.75% (approx.)

Sources: Federal Reserve (FRED), National Association of Insurance Commissioners (NAIC), S&P Global Market Intelligence. Some 2024-2025 figures are approximations based on available quarterly data and industry trends.

Consulting

The following tables provide actual historical data (2015-2025) for management consulting industry metrics referenced in this packet. These benchmarks help contextualize the performance targets and competitive dynamics outlined in the exercise scenario.

Global Management Consulting Revenue Growth (Annual %, 2015-2025)

Year	Growth Rate
2015	3.2%
2016	2.8%
2017	4.5%
2018	4.1%
2019	3.9%
2020	1.2%
2021	8.5%
2022	5.8%
2023	4.2%
2024	4.8%
2025	5.1% (approx.)

Top-Tier Consulting Firm Operating Margins (Approximate %, 2015-2025)

Year	Operating Margin
2015	16.5%
2016	16.2%
2017	17.0%
2018	17.3%
2019	16.8%
2020	13.5%
2021	18.2%
2022	17.8%
2023	16.5%
2024	16.9%
2025	17.1% (approx.)

Consulting Utilization Rates — Top-Tier Firms (% , 2015-2025)

Year	Utilization Rate
2015	64%
2016	63%
2017	66%
2018	67%
2019	65%
2020	55%
2021	71%
2022	69%
2023	65%
2024	66%
2025	64% (approx.)

Global AI/Digital Transformation Consulting Market Size (\$B, 2015-2025)

Year	Market Size
2015	18
2016	22
2017	28
2018	35
2019	42
2020	38
2021	52
2022	68
2023	82
2024	98
2025	115 (approx.)

Sources: Kennedy Research & Consulting (formerly ALM Intelligence), Source Global Research, Statista (Management Consulting Market), IDC (Digital Transformation Spending), Vault Consulting Industry Reports, company filings and investor presentations. Some 2024-2025 figures are

approximations based on available quarterly data and industry analyst estimates.

Law

The following tables provide actual historical data (2015-2025) for legal industry metrics referenced in this packet.

US Legal Services Revenue Growth (Annual %, 2015-2025)

Year	Growth Rate
2015	2.1%
2016	2.3%
2017	2.8%
2018	3.2%
2019	2.9%
2020	1.1%
2021	4.5%
2022	3.8%
2023	2.4%
2024	3.1%
2025	3.3% (approx.)

AmLaw 100 Revenue Per Lawyer (Approximate \$K, 2015-2025)

Year	RPL (\$K)
2015	755
2016	785
2017	825
2018	870
2019	895
2020	875
2021	950
2022	1,020
2023	1,055
2024	1,095
2025	1,125 (approx.)

Average Partner Billing Rate (Approximate \$/hr, 2015-2025)

Year	Partner Rate
2015	475
2016	495
2017	525
2018	565
2019	595
2020	600
2021	650
2022	720
2023	780
2024	835
2025	875 (approx.)

Sources: American Bar Association (Legal Services), Thomson Reuters (AmLaw rankings), NALP, Georgetown Law Center for the Study of the Legal Profession / Thomson Reuters Institute (Annual Report on the State of the Legal Market). Some 2024-2025 figures are approximations based on available quarterly data and industry trends.

Manufacturing

The following tables provide actual historical data (2015-2025) for manufacturing and industrial metrics referenced in this packet. These benchmarks help contextualize the operational targets and capital dynamics outlined in the scenario.

US Industrial Production Growth (Annual %, 2015-2025)

Year	Growth Rate
2015	0.5%
2016	-0.5%
2017	1.8%
2018	3.8%
2019	-0.8%
2020	-5.2%
2021	4.7%
2022	-0.5%
2023	-0.3%
2024	0.8%
2025	1.2% (approx.)

ISM Manufacturing PMI (Annual Average, 2015-2025)

Year	PMI
2015	51.5
2016	50.0
2017	52.5
2018	54.7
2019	51.2
2020	50.4
2021	58.5
2022	52.8
2023	47.2
2024	48.5
2025	49.2 (approx.)

US Manufacturing Employment (Millions, 2015-2025)

Year	Employment
2015	12.33
2016	12.35
2017	12.49
2018	12.74
2019	12.80
2020	12.29
2021	12.47
2022	12.87
2023	12.90
2024	12.92
2025	12.88 (approx.)

S&P 500 Industrials Annual Returns (% , 2015-2025)

Year	Return
2015	-4.4%
2016	5.1%
2017	13.4%
2018	-8.3%
2019	25.1%
2020	8.3%
2021	24.2%
2022	-19.2%
2023	10.7%
2024	19.5%
2025	14.2% (approx.)

Sources: Federal Reserve Economic Data (FRED), Institute for Supply Management (ISM), Bureau of Labor Statistics (BLS), S&P Global Market Intelligence. Some 2024-2025 figures are approximations based on available monthly data.

Logistics

The following tables provide actual historical data (2015-2025) for logistics and industrial metrics referenced in this packet. These benchmarks help contextualize the operational targets and market dynamics outlined in the scenario.

US Freight Volume Index (Indexed to 2015 = 100, 2015-2025)

Year	Index
2015	100.0
2016	99.5
2017	101.8
2018	103.5
2019	104.2
2020	99.8
2021	109.5
2022	111.2
2023	108.5
2024	109.8
2025	110.5 (approx.)

S&P 500 Industrials Annual Returns (% , 2015-2025)

Year	Return
2015	-4.4%
2016	5.1%
2017	13.4%
2018	-8.3%
2019	25.1%
2020	8.3%
2021	24.2%
2022	-19.2%
2023	10.7%
2024	19.5%
2025	14.2% (approx.)

US Manufacturing Employment (Millions, 2015-2025)

Included as proxy for industrial labor market dynamics that directly affect logistics workforce availability and wage pressure.

Year	Employment
2015	12.33
2016	12.35
2017	12.49
2018	12.74
2019	12.80
2020	12.29
2021	12.47
2022	12.87
2023	12.90
2024	12.92
2025	12.88 (approx.)

ISM Manufacturing PMI (Annual Average, 2015-2025)

Included as leading indicator for freight demand: manufacturing expansion/contraction directly drives logistics volume.

Year	PMI
2015	51.5
2016	50.0
2017	52.5
2018	54.7
2019	51.2
2020	50.4
2021	58.5
2022	52.8
2023	47.2
2024	48.5
2025	49.2 (approx.)

Sources: American Trucking Associations (Freight Volume Index), S&P Global Market Intelligence, Bureau of Labor Statistics (BLS), Institute for Supply Management (ISM), Federal Reserve Economic Data (FRED). Some 2024-2025 figures are approximations based on available monthly data.

Big Tech

The following tables provide actual historical data (2015-2025) for Big Tech metrics referenced in this industry packet. These benchmarks help contextualize the performance targets and competitive dynamics outlined in the exercise scenario.

Big Tech Revenue Growth (Annual %, 2015-2025)

Company	2015	2018	2021	2024	2025 (Est.)
Google (Alphabet)	12.5%	22.9%	41.4%	13.1%	12.8%
Meta	39.0%	26.0%	23.2%	16.1%	15.3%
Microsoft	14.0%	14.3%	16.4%	16.0%	15.2%
Amazon	20.5%	31.0%	21.7%	10.9%	11.2%
Apple	1.3%	9.2%	33.3%	0.9%	4.2%

Big Tech Cloud Revenue & Growth (Annual, 2015-2025)

Year	AWS Revenue (\$B)	Azure Revenue (\$B)	GCP Revenue (\$B)	Combined Growth (%)
2015	7.9	0.6	0.1	-
2016	12.2	1.2	0.2	50%
2017	17.5	2.5	0.5	55%
2018	25.7	4.2	1.0	45%
2019	35.0	6.8	1.8	42%
2020	45.4	10.2	3.0	35%
2021	62.2	15.0	5.0	38%
2022	80.1	20.0	7.0	28%
2023	90.8	25.0	10.0	23%
2024	104.0	32.0	14.0	18%
2025	118.0 (est.)	40.0 (est.)	18.0 (est.)	16% (est.)

S&P 500 Information Technology Annual Returns (% , 2015-2025)

Year	Return
2015	4.4%
2016	14.6%
2017	36.2%
2018	-9.7%
2019	48.2%
2020	33.4%
2021	28.7%
2022	-32.4%
2023	53.3%
2024	25.0%
2025	18.5% (approx.)

Software & Tech Sector Capital Expenditure as % of Revenue (2015-2025)

Year	Big Tech Cloud	Industry Average
2015	5.0%	3.0%
2016	6.0%	3.5%
2017	7.0%	4.0%
2018	8.0%	4.5%
2019	9.0%	5.0%
2020	10.0%	5.5%
2021	12.0%	6.5%
2022	12.5%	6.8%
2023	14.0%	7.5%
2024	15.5%	8.2%
2025	17.0% (approx.)	9.0% (approx.)

Sources: S&P Global Market Intelligence, MSCI, company earnings reports (Alphabet, Meta, Microsoft, Amazon, Apple), Gartner, IDC, Bureau of Economic Analysis. Some 2024-2025 figures are

approximations based on available quarterly data and industry trends.

SaaS

The following tables provide actual historical data (2015-2025) for SaaS industry metrics referenced in this industry packet. These benchmarks help contextualize the performance targets and competitive dynamics outlined in the exercise scenario.

SaaS Industry Gross Margin & Net Retention Rate (% , 2015-2025)

Year	Gross Margin (%)	Net Revenue Retention (%)
2015	78.0%	110.0%
2016	79.0%	112.0%
2017	80.0%	115.0%
2018	80.5%	118.0%
2019	81.0%	120.0%
2020	81.5%	125.0%
2021	82.0%	128.0%
2022	81.5%	125.0%
2023	80.8%	122.0%
2024	80.5%	120.0%
2025	79.8% (approx.)	118.0% (approx.)

Software Industry Operating Margins (% , 2015-2025)

Year	Cloud/SaaS Avg	Enterprise Software Avg
2015	18.5%	20.0%
2016	19.2%	20.8%
2017	20.1%	21.5%
2018	21.0%	22.2%
2019	22.3%	23.0%
2020	23.5%	23.8%
2021	24.1%	24.5%
2022	22.8%	23.2%
2023	23.5%	24.1%
2024	24.2%	25.0%
2025	23.8% (approx.)	24.8% (approx.)

Enterprise Software Spending as % of IT Budget (2015-2025)

Year	% of IT Budget
2015	18.0%
2016	18.5%
2017	19.0%
2018	19.5%
2019	20.0%
2020	21.0%
2021	22.0%
2022	21.5%
2023	21.0%
2024	20.5%
2025	20.0% (approx.)

S&P 500 Software Index Annual Returns (% , 2015-2025)

Year	Return
2015	0.8%
2016	20.0%
2017	36.5%
2018	-11.0%
2019	42.0%
2020	35.5%
2021	27.0%
2022	-35.0%
2023	55.0%
2024	28.0%
2025	20.0% (approx.)

Sources: S&P Global Market Intelligence, MSCI, company earnings reports (Salesforce, SAP, Workday, Adobe, ServiceNow), Gartner, IDC, Bureau of Economic Analysis. Some 2024-2025 figures are approximations based on available quarterly data and industry trends.

Valuation Multiples by Industry

The following tables provide actual historical EV/EBITDA and P/E multiples by industry for 2020-2025. These benchmarks help facilitators calibrate industry health assessments and add real-market color to Industry Health announcements.

Historical EV/EBITDA Ranges by Industry (2020-2025)

Retail

Period	Range	Typical
2020	12-18x	15.2x
2021	14-21x	17.8x
2022	10-16x	13.2x
2023	12-18x	15.1x
2024	13-19x	16.0x
2025	13-20x	16.5x (approx.)

Financials

Period	Range	Typical
2020	0.8-1.2x	1.0x
2021	1.0-1.4x	1.2x
2022	0.9-1.3x	1.1x
2023	1.0-1.4x	1.2x
2024	1.1-1.5x	1.3x
2025	1.1-1.6x	1.35x (approx.)

Healthcare

Period	Range	Typical
2020	11-15x	13.2x
2021	12-17x	14.8x
2022	11-15x	12.9x
2023	12-16x	14.1x
2024	13-17x	15.0x
2025	13-18x	15.5x (approx.)

Industrials (Manufacturing + Logistics)

Period	Range	Typical
2020	8-12x	10.2x
2021	10-15x	12.5x
2022	8-13x	10.8x
2023	9-14x	11.8x
2024	10-15x	12.5x
2025	10-16x	13.0x (approx.)

Consulting (Professional Services)

Period	Range	Typical
2020	15-22x	18.5x
2021	18-26x	22.0x
2022	14-20x	17.2x
2023	15-22x	18.5x
2024	16-24x	20.0x
2025	16-25x	20.5x (approx.)

Law (Legal Services)

Period	Range	Typical
2020	12-18x	15.0x
2021	14-20x	17.5x
2022	11-17x	14.2x
2023	13-19x	16.0x
2024	14-20x	17.0x
2025	14-21x	17.5x (approx.)

Historical P/E Ranges by Industry (2020-2025)

Financials P/E

Period	Range	Typical
2020	9-13x	11.0x
2021	11-16x	13.5x
2022	9-14x	11.8x
2023	10-15x	12.3x
2024	11-16x	13.5x
2025	11-17x	13.8x (approx.)

Healthcare P/E

Period	Range	Typical
2020	18-24x	20.8x
2021	20-26x	22.5x
2022	16-22x	19.2x
2023	18-24x	20.8x
2024	19-26x	22.2x
2025	19-27x	22.8x (approx.)

Consumer (Retail + CPG) P/E

Period	Range	Typical
2020	14-20x	17.2x
2021	18-26x	20.8x
2022	12-18x	15.5x
2023	15-22x	18.5x
2024	17-24x	20.1x
2025	16-24x	19.8x (approx.)

Industrials P/E

Period	Range	Typical
2020	13-18x	15.8x
2021	16-22x	18.8x
2022	12-17x	14.5x
2023	14-19x	16.2x
2024	15-21x	17.5x
2025	15-22x	18.2x (approx.)

Consulting (Professional Services) P/E

Period	Range	Typical
2020	18-25x	21.5x
2021	21-30x	25.2x
2022	16-23x	19.8x
2023	18-26x	21.5x
2024	20-28x	23.8x
2025	19-27x	23.2x (approx.)

Law (Legal Services) P/E

Period	Range	Typical
2020	14-20x	17.0x
2021	16-24x	20.0x
2022	13-19x	16.2x
2023	15-21x	18.0x
2024	16-23x	19.5x
2025	16-24x	20.0x (approx.)

Sources: S&P Global Market Intelligence, FactSet, Bloomberg, Refinitiv. Multiple ranges reflect high/low observed multiples during each year; typical represents median/consensus. 2024-2025 figures are approximations based on available quarterly data and industry trends. Law firm multiples are approximations based on limited public market data (most AmLaw firms are private partnerships); included for directional calibration only.

Broader Market Reference Data

The following tables provide actual historical data (2020-2025) for S&P 500 sector returns, forward P/E multiples, and Treasury yields. These benchmarks provide facilitators with real market context for calibrating Industry Health announcements and market narrative.

S&P 500 Sector Annual Returns (% , 2020-2025)

Year	Cons. Discretionary	Cons. Staples	Financials	Healthcare	Industrials
2020	36.5%	12.2%	-2.1%	13.3%	8.3%
2021	27.3%	11.8%	27.9%	16.8%	24.2%
2022	-23.5%	-4.9%	-16.8%	-4.2%	-19.2%
2023	22.5%	4.2%	18.1%	7.5%	10.7%
2024	18.8%	8.1%	22.4%	19.3%	19.5%
2025	12.0%	6.5%	16.5%	14.8%	14.2% (approx.)

S&P 500 Sector Forward P/E Multiples (2020-2025)

Year	Cons. Discretionary	Cons. Staples	Financials	Healthcare	Industrials
2020	17.2	21.5	11.8	20.5	15.2
2021	20.8	23.1	13.2	22.8	17.5
2022	15.5	20.2	10.5	19.8	13.8
2023	18.5	21.8	12.1	21.2	15.5
2024	20.1	22.5	13.5	23.1	16.8
2025	19.8	21.9	13.0	22.5	16.5 (approx.)

S&P 500 Total Return (Annual %, 2015-2025)

Year	Total Return
2015	1.4%
2016	12.0%
2017	21.8%
2018	-6.2%
2019	31.5%
2020	18.4%
2021	28.7%
2022	-18.1%
2023	24.2%
2024	24.0%
2025	15.0% (approx.)

US 10-Year Treasury Yield (Year-End %, 2015-2025)

Year	Yield
2015	2.24%
2016	2.45%
2017	2.40%
2018	2.69%
2019	1.92%
2020	0.92%
2021	1.52%
2022	3.88%
2023	3.88%
2024	4.24%
2025	3.95% (approx.)

Sources: S&P Global Market Intelligence, Federal Reserve Economic Data (FRED), FactSet, Bloomberg. Forward P/E and Treasury yields are point-in-time (year-end) estimates. Some 2024-2025 figures are approximations based on available quarterly data through Q1 2025.

End of Historical Reference Data — V7.4

Lead Facilitator Briefing

Purpose and Exercise Hypotheses

Project Threshold is a strategic stress-test designed to surface near-term (2026-2030) economic implications of different AI adoption timelines and capability types. The exercise is **not** a prediction; it is a structured scenario exploration that forces decision-makers to grapple with real trade-offs, uncertainty, and strategic consequences.

Core Hypotheses Being Tested:

1. Industry divergence is inevitable. AI benefits will accrue unevenly across the eleven industries based on technical feasibility, regulatory tolerance, and business model fit. Finance, Consulting, Big Tech, and SaaS will move faster than Healthcare Provider, Manufacturing, and Law.
2. Market concentration may accelerate. Large, well-capitalized players with data, talent, and capital will capture disproportionate value from AI. Mid-market players face strategic risk if they cannot execute speed-to-market.
3. Labor displacement is structural, not cyclical. Unlike previous technology cycles where displaced workers could transition to new roles over a decade, AI displacement is occurring faster (1-3 years) and in roles that do not have obvious retraining paths. Junior roles in Consulting, Law, and Finance are particularly exposed.
4. Regulatory bifurcation is likely. Policy will not be uniform. Some jurisdictions will embrace aggressive AI deployment; others will impose strict safeguards. Firms must operate across this patchwork. Law faces unique complexity with 50+ state bar jurisdictions.
5. Plausible catastrophic tail risks exist. A major incident (AI system causing patient harm, financial system manipulation, or market disruption) could trigger severe regulatory backlash and broader market disruption.
6. Organizational execution matters more than technology. Firms that combine AI deployment with disciplined change management, workforce transition planning, and customer communication will succeed. Those that treat AI as a "tool drop" will underperform.
7. Professional services business models face existential pressure. Consulting and Law face unique disruption: AI compresses the work that junior talent performs, undermining the leverage model that drives profitability. Pricing models (billable hours, per-consultant fees) are under direct threat.

What Makes a Good V7.4 Facilitator?

Tone and Energy

High-Energy Engagement Without Hype: Maintain enthusiasm for the scenario while signaling that this is serious strategic work. Avoid breathless "AI will change everything" language; instead emphasize "AI will force hard choices."

Intellectual Humility: You are not predicting the future. You are exploring one plausible future. When participants challenge assumptions, acknowledge the validity. Say: "That's a fair point. Regulation could move slower than we've modeled, which would shift outcomes. That's exactly the kind of assumption we should examine."

Authentic Interest in Participant Reasoning: When a participant presents a decision, ask follow-up questions that probe their logic. "Why did you choose aggressive automation over gradual deployment?" Listen to their reasoning. Score accordingly, but also show that you understand the tension.

Controlled Urgency: Participants should feel time pressure, but not panic. Create artificial urgency through inject timing and decision deadlines, but don't make it feel arbitrary. Say: "You have 12 minutes to decide. The market is moving fast in this scenario. What's your call?"

Decision-Making Style

Principle-Based, Not Arbitrary: Your role is to adjudicate based on transparent criteria (Strategic Fit, Execution Risk, Tail Risk). Explain your scoring publicly so participants understand the logic. If a participant challenges a score, acknowledge the challenge and explain your reasoning.

Firm but Fair on Plausibility: Some participants will propose ideas that are technologically implausible or organizationally infeasible. Your job is to enforce the bounds of the scenario without being pedantic. If a participant says "We will deploy fully autonomous systems with zero human oversight in 6 months," push back: "Regulators require human oversight. That's non-negotiable in this scenario. But you could deploy a system where AI makes 80% of routine decisions autonomously and flags exceptions for human review. That's plausible." Don't shut down creativity; redirect it.

Transparency on Tradeoffs: When adjudicating, make clear that every decision involves tradeoffs. Aggressive AI deployment = speed and efficiency, but also labor risk and regulatory risk. Defensive moves = safety, but also competitive disadvantage. Help participants see the tradeoff matrix, even if they choose poorly.

Common Facilitator Mistakes to Avoid

1. Over-Explaining

- Mistake: Spending 10 minutes explaining round setup when 3 minutes suffices.
- Fix: Read the situation scenario aloud concisely. Trust the materials. If participants have questions, answer them, but don't pre-answer every possible question.

2. Participants Don't Fully Engage

- Mistake: An industry representative doesn't think critically about their own position. Decisions feel flat or defaulted.
- Fix: Ask probing questions one-on-one during processing time. "What assumption would need to be true for your industry to move faster?" Draw out genuine thinking.

3. Not Enforcing Time Discipline

- Mistake: Industry representatives get 15 minutes to decide, but you don't stop them at 15 minutes. Exercises run long and lose impact.

- Fix: Use a visible timer. Announce "2 minutes remaining" at 13 minutes. Call time at 15 minutes. Say: "Pens down. What's your decision?" If they're still deliberating, say: "You have 60 seconds to decide: aggressive or conservative?" Force a choice.

4. Being Too Lenient on Implausible Strategies

- Mistake: A participant proposes something that violates the scenario bounds, and you allow it because "it's interesting."
- Fix: Maintain scenario boundaries. If a participant proposes AGI-level capabilities in 2026, say: "That's outside our scenario. We're modeling incremental AI progress, not discontinuous breakthroughs. What's a plausible version of your strategy?"

5. Not Creating Urgency

- Mistake: Participants feel they have unlimited time to think. Discussions become academic and lose energy.
- Fix: Announce timing upfront. Use injects to force decisions. Say: "Markets are moving. Your competitor just announced a major move. Do you respond or stick with your plan? You have 5 minutes to decide."

6. Treating Debrief as Afterthought

- Mistake: Rushing through the final debrief because time is running out.
- Fix: Protect debrief time. The debrief is where participants make sense of what they learned. Allocate 60 minutes minimum (non-negotiable). Ask: "What surprised you?" "What assumptions did we test?" "What policy questions remain unresolved?"

7. Not Managing Cross-Industry Discussion with Larger Groups

- Mistake: With 11 industry participants, discussion becomes unwieldy or dominated by 2–3 voices.
- Fix: Use structured facilitation. Call on specific industries: "Healthcare Provider, how does Finance's decision affect you?" Rotate who speaks first each round. Use the "popcorn" method — each participant responds briefly, then deeper discussion follows.

Industry Selection Facilitation (When Running Fewer Than 11)

When participant count is below 11, the facilitator must select which industries are in play. Guidance:

5 participants (minimum viable): Select one industry per industry cluster to maximize cross-industry diversity:

- Retail, Healthcare Provider, Finance, Manufacturing, Big Tech

7 participants: Add Consulting and SaaS for professional services and software depth. These two industries create rich cross-industry dynamics (Consulting advises all others; SaaS sells to all others).

9 participants: Add CPG, Healthcare Payer, and Logistics for broader coverage and intra-group tension (Provider vs. Payer, Retail vs. CPG).

11 participants: Full roster. Maximum cross-industry richness.

Collective Bonus scales naturally: The Collective Bonus is available every round regardless of participant count. With 5 participants, consensus is easier to reach; with 11, agreement among 3+ participants is more common and multiple industries may be nominated.

When to Increase Pressure vs. Give Breathing Room

Signal	Context	Facilitator Move
Participants are disengaged or passive	Industry rep is waiting for facilitator guidance instead of making proactive decisions	Increase pressure: Deploy injects. "Your competitor just announced a major move. Do you match their move or differentiate?"
Participants are pursuing obviously implausible strategies	A participant wants to "ignore AI and focus on human excellence" despite aggressive competitor moves in the scenario	Increase pressure: Push back directly. "That's a defensible posture, but the scenario assumes aggressive competitor moves. How do you compete if your rival has 30% cost advantage from automation?"
Participants are over-analyzing and under-deciding	Analysis paralysis — participant says "I need more data before deciding"	Increase pressure: Force a call. "The market doesn't wait for perfect information. Make a decision with the data you have. What's your call?"
Late rounds and endgame positioning	Round 3-4; participants should feel the weight of their cumulative decisions	Increase pressure: Force final strategic bets. Tighten timelines, deploy high-stakes injects.
A participant is working through a genuinely hard tradeoff	Participant is wrestling with labor vs. efficiency or risk vs. return — real tensions	Give breathing room: Let them deliberate. These tensions are the heart of the exercise.
A major inject has just hit and participants are reacting	Disruptive inject landed; participants need time to absorb	Give breathing room: Allow 2-3 extra minutes to process before requiring a decision.
Early rounds (Practice Round, R1) — participants are learning the format	Participants are still settling into the rhythm	Give breathing room: Let them learn the mechanics before ramping pressure.
A participant has made a bold, well-reasoned move	Participant took a risk but backed it with strong reasoning	Give breathing room: Acknowledge good reasoning even if you disagree with the strategic choice. Reward careful thinking.

The Practice Micro-Round (9:00–9:10, Pre-Round 1)

V7.4 includes a **brief practice round** (9:00–9:10) before Round 1 proper. This is crucial for:

- Familiarizing individual industry representatives with decision format and scoring
- Testing tech and materials
- Building confidence before high-stakes rounds
- Surfacing questions about rules and procedures

Facilitator Guidance:

- Keep it short (10 minutes total including decision + feedback)
 - Make decisions low-stakes (small point values)
 - Provide immediate scoring feedback so participants understand the rubric
 - Use this round to correct misunderstandings about process
 - Tell participants explicitly: "This is practice. Use it to learn the format."
-

Running Collective Bonus (Optional, Every Round, R1-R4)

Each round includes an optional **Collective Bonus** opportunity during the cross-industry discussion period (~4 minutes within the 27-minute discussion block). Participation is not mandatory; it represents an organic peer signal on strategy quality.

Collective Bonus Mechanics:

1. Open nomination. Facilitator opens the floor: "Does anyone want to recognize an especially strong strategy this round, or flag one that seems particularly risky?" Participants who wish to respond name one industry as a strong strategy and/or one as a risky strategy. Participants cannot nominate their own industry.
2. Consensus threshold. If 3+ participants agree on the same industry: +2 (strong strategy) or -2 (risky strategy) applied to cumulative score. Maximum one +2 and one -2 per round. If no consensus or no one speaks up, no bonus applied.
3. Use Collective Bonus to create cascades. Collective Bonus results have ripple effects: a +2 bonus signals peer confidence and may embolden that industry's next move; a -2 penalty signals peer concern and may prompt a strategic pivot. Call these out: "Three participants flagged Manufacturing as risky this round. Manufacturing, does that change your approach next round? Logistics, does your supply chain partner's low credibility affect your plans?"
4. Debrief Collective Bonus patterns. In final debrief, ask: "Who consistently received peer recognition? Who was consistently flagged as risky? Did the group's collective judgment track with scoring outcomes, or diverge?"

Running Facilitator Market Shock (Round 2 Only)

In Round 2, the facilitator imposes a **Market Shock** — selecting 2-3 industries and imposing one constraint each from the standard constraint menu. This replaces the former IC Constraint Imposition

action.

Facilitator Responsibilities:

1. Select 2-3 target industries. Choose based on Round 1 outcomes, strategic positioning, or to create interesting cross-industry dynamics. Any industry can be hit.
 2. Impose one constraint per target industry from the standard constraint menu: Regulatory Halt, Labor Cost Surge, Capital Tightening, Reputational Pressure, Competitive Response, Litigation Risk.
 3. Announce publicly. Read the constraints aloud: "Market Shock: Finance faces Capital Tightening. Healthcare Provider faces Regulatory Halt. Consulting faces Competitive Response." Allow ~3 minutes for this action.
 4. Create cascades. Market Shock constraints force strategy pivots and create cross-industry ripple effects. Call these out: "Finance just got hit with Capital Tightening. Consulting, does that affect your clients' advisory budgets? SaaS, does reduced financial sector spending hit your pipeline?"
-

Private Cards (Rounds 1-3 Only)

In V7.4, **private cards are distributed only in Rounds 1-3** from the separate `05_Participant_Private_Cards/` folder. Round 4 has **no private information**.

Facilitator Guidance:

1. Protect information: When distributing a private card to a Retail participant, ensure it's not visible to other industries. Read it aloud (so facilitator can control framing) or hand directly to participant.
 2. Timing matters: Distribute private cards during inject processing time, not during decision preparation. This gives participants time to absorb before deciding.
 3. Selective revelation: If a private card is shared across multiple industries, coordinate revelation timing. Example: "Finance gets this market data card 2 minutes before Consulting does, so Finance has a brief information advantage."
 4. Round 4 transparency: Remind participants in Round 4 that there are no private cards. All information is shared. This creates a different dynamic: participants are operating on equal information.
-

Simplified Staffing Model (1 Facilitator + 1 Helper)

V7.4 is optimized for **minimal staffing**: 1 lead facilitator + 1 support person (note-taker or tech support). Minimum viable exercise: 5 industry representatives + 1 facilitator + 1 helper = 7 people.

Scaling Note: With 11 industry participants, consider adding a second scorer to handle the increased decision volume within the scoring block.

Responsibility Split:

Lead Facilitator:

- Stands at podium or control station

- Reads situation updates and injects aloud
- Manages timer visibly
- Adjudicates decisions in real-time
- Conducts Collective Bonus during cross-industry discussion (all rounds)
- Executes Facilitator Market Shock (Round 2)
- Leads debrief
- Maintains energy and urgency

Support Person:

- Manages tech (projector, timer display, backup materials)
- Takes notes on key decisions and industry dynamics
- Distributes private cards (from 05_Participant_Private_Cards/) and materials
- Handles contingencies (projector fails, missing worksheet, etc.)
- Posts scores on display board

Time-Saving Tips:

- Pre-print all materials (no improvisation)
- Have backup copies of everything
- Use a simple scoring sheet (not complex digital system)
- Minimize breakout consultations (only if industry representative is stuck)

Pre-Exercise Setup Checklist

Logistics & Tech (1 week before)

- Confirm room booking and layout (5–11 industry positions, round-table or seminar seating, central projection, breakout areas)
- Confirm which industries are in play; prepare only those materials
- Test projector, timer display, WiFi
- Print all industry cards, decision worksheets, scoring sheets, scenario packets (minimum 20% extra copies)
- Print private cards from 05_Participant_Private_Cards/ for selected industries only
- Arrange for backup power, backup projector, backup materials
- Set up timer display (visible to all participants)

Participant Prep (1 week before)

- Distribute pre-exercise diagnostic survey (optional but recommended)
- Send primer document: "Project Threshold Overview & Your Industry"
- Distribute industry-specific AI Adoption Arc from 04_Participant_AI_Adoption_Arcs/
- Request pre-reading: Facilitator briefing overview and relevant industry background
- Confirm attendance and accessibility needs

Materials Distribution (Day of exercise, 30 min before start)

- At each position: 1 set of industry cards
- At each position: Decision worksheets and scoring reference

- [] At each position: Pens, notepads, scenario packets
- [] Central projection: Round situation cards, inject deck
- [] Facilitator station: Master copy of all materials, injections, scoring rubrics

Facilitator Prep (Day of exercise, 1 hour before)

- [] Review participant list and industry assignments
- [] Confirm timing for practice round and Rounds 1-4
- [] Pre-read all injects for practice round and Round 1
- [] Test all tech (projector, timer, WiFi, audio)
- [] Prepare scorecard to track industry scores across rounds
- [] Identify which injects to use (tailor to emphasis)

8-Hour Full-Day Format (V7.4 Standard)

Target Audience: Mixed organizational levels; full immersion in decision scenarios

Structure: Practice micro-round + Rounds 1-4; all injects deployed; cross-industry discussion emphasis; full debrief

Minute-by-Minute Timeline

Time	Duration	Activity	Notes
8:30	15 min	Arrival & Setup	Materials distribution, name badges, seating by industry
8:45	15 min	Welcome & Exercise Overview	Opening remarks, objectives, explain rounds 1–4, participant introductions
9:00	10 min	PRACTICE MICRO-ROUND	Simple scenario; low-stakes decision; facilitator scores live; Q&A
9:10	5 min	ROUND 1: Situation Update	Project slide; read injects
9:15	3 min	ROUND 1: Inject + Private Card	Distribute private cards from 05_Participant_Private_Cards/
9:18	15 min	ROUND 1: Individual Decision Prep	Each industry rep prepares alone

Time	Duration	Activity	Notes
9:33	3 min	ROUND 1: Decision Submission	Each rep submits worksheet
9:36	12 min	ROUND 1: Facilitator Scoring	Score all submitted decisions
9:48	27 min	ROUND 1: Cross-Industry Discussion & Results	Scores posted; Collective Bonus (Optional); cross-industry dynamics
10:15	15 min	Break	Coffee, restrooms
10:30	5 min	ROUND 2: Situation Update	Project slide; read injects
10:35	3 min	ROUND 2: Inject + Private Card	Distribute Round 2 private cards
10:38	15 min	ROUND 2: Individual Decision Prep	Each rep prepares
10:53	3 min	ROUND 2: Decision Submission	Each rep submits
10:56	12 min	ROUND 2: Facilitator Scoring	Score decisions; apply Facilitator Market Shock (~3 min within this block)
11:08	27 min	ROUND 2: Cross-Industry Discussion & Results	Scores + Market Shock announced; Collective Bonus (Optional); IHS conditions
11:35	40 min	Lunch	Informal reflection encouraged
12:15	5 min	ROUND 3: Situation Update	Project slide; read injects
12:20	3 min	ROUND 3: Inject + Private Card	Distribute Round 3 private cards
12:23	15 min	ROUND 3: Individual Decision Prep	Each rep prepares
12:38	3 min	ROUND 3: Decision Submission	Each rep submits
12:41	12 min	ROUND 3: Facilitator Scoring	Score decisions

Time	Duration	Activity	Notes
12:53	27 min	ROUND 3: Cross-Industry Discussion & Results	Scores posted; Collective Bonus (Optional); IHS conditions
1:20	15 min	Break	Restrooms, energy reset
1:35	5 min	ROUND 4: Situation Update	Project slide; read final injects
1:40	3 min	ROUND 4: Inject Distribution	Final market conditions; no private cards R4
1:43	15 min	ROUND 4: Final Individual Decision	Each rep makes final decision
1:58	3 min	ROUND 4: Decision Submission	Each rep submits
2:01	12 min	ROUND 4: Facilitator Scoring	Score final decisions
2:13	27 min	ROUND 4: Cross-Industry Discussion & Results	Final discussion; Collective Bonus (Optional)
2:40	10 min	Break + Facilitator Prep	Consolidate scores, prepare debrief
2:50	15 min	Individual Reflections	Each rep shares 1–2 min highlights
3:05	25 min	Cross-Industry Dynamics Discussion	Facilitator-led; inter-industry dependencies
3:30	20 min	No-Regrets Actions & Takeaways	Consolidate actions; each rep states key takeaway
3:50	40 min	Wrap-Up & Close	Feedback surveys, closing remarks, synthesis memo scheduling

Debrief Facilitation (2:50–3:50, 60 minutes, non-negotiable)

After Round 4 and final adjudication, you have 60 minutes for debrief (2:50–3:50). This is where learning consolidates. Debrief is protected time and should not be shortened.

Debrief Structure:

Individual Reflections (15 min): Each industry representative shares highlights. With 5 participants, allow 2–3 min each. With 11 participants, keep to 1–1.5 min each. Rotations: Who am I? What surprised me? What would I do differently?

Cross-Industry Dynamics (25 min): Facilitator-led discussion of inter-industry dependencies and spillovers discovered during the exercise. Push on tensions and patterns. With 11 industries, there are richer dynamics to explore — use that.

No-Regrets Actions & Takeaways (20 min): Consolidate actions that make sense regardless of scenario. Each rep states one key takeaway.

Debrief Questions (Use as needed):

Question	Purpose
"What surprised you?"	Opens reflection. Not "What happened?" but "What violated your expectations?" Best used first to set a reflective tone.
"Which assumptions did the scenario challenge?"	Surfaces hidden beliefs (regulation would be light, clients would trust AI, your industry would move faster). Use after participants have warmed up.
"What would you do differently if you ran the scenario again?"	Reflects on decision quality and hindsight learning. Effective mid-debrief when participants can see their full arc.
"What policy questions remain unresolved?"	By Round 4, participants see gaps. Draws out real-world regulatory and governance uncertainty. Use to transition toward takeaways.
"What would your actual leadership take away from this?"	Bridges the exercise back to real organizations and decisions. Best used near the end to ground the discussion.
"How did dynamics between related industries play out?" (NEW for V7.4)	Probes specific pairs: Finance vs. Consulting, Healthcare Provider vs. Payer, Retail vs. CPG, Big Tech vs. SaaS. Use during Cross-Industry Dynamics block.
"How did Collective Bonus results affect your strategy?" (NEW for V7.4)	Surfaces whether peer recognition (+2) emboldened or peer concern (-2) forced pivots. Did the group's collective judgment match your own assessment? Use during Individual Reflections block.
"Did the Market Shock in Round 2 change your trajectory?" (NEW for V7.4)	Probes whether externally imposed constraints created lasting strategic shifts or were absorbed. Use during Cross-Industry Dynamics block.

Debrief Facilitation Moves:

- Call on all industries: Ensure each industry voice is heard. Don't let one participant dominate. With 11 participants, be disciplined about time.

- Push on tensions: "You chose aggressive automation. Why didn't you invest more in workforce transition?"
- Highlight divergence: "Retail chose aggressive automation. Law chose defensive positioning. Manufacturing went hybrid. That's realistic market divergence based on industry characteristics."
- Probe cross-industry dynamics: "Consulting, your pricing pressure directly affects Finance's advisory budget. Did you two discuss that? Finance, did you factor in how Consulting's AI-augmented delivery changes your cost structure?"
- Note real-world parallels: "In real strategy work, you're facing these same tradeoffs. How will you handle it in your actual organization?"

Common Facilitation Scenarios & Facilitator Moves

Scenario	What You'll See	Facilitator Move
Implausible Strategy	Healthcare Provider rep proposes fully autonomous AI diagnostics with zero human oversight in 6 months.	Don't shut them down — probe: "FDA approval takes 12-24 months. How do you compress that?" Reference scenario: regulators demand human oversight. Offer redirect: "AI reads scans, radiologists verify — that's plausible and still powerful." Scoring: If they insist on zero-oversight, score Execution Risk at -3. If they redirect, score higher.
Analysis Paralysis	Manufacturing rep has been deliberating 12 minutes. Time is up in 3. No decision yet.	Name it: "You're wrestling with a tough tradeoff. That's valid." Force binary: "Markets don't wait. Aggressive or conservative? Which way?" Probe instinct: "If you had to bet \$1M, which would you choose?" Hard deadline: "60 seconds. Write it down. Go." Scoring: Score based on decision made, not analysis quality.

Scenario	What You'll See	Facilitator Move
Market Shock Creates Disruption (R2)	Facilitator Market Shock imposes a constraint that threatens an industry rep's strategy.	<p>Normalize: "This is exactly the kind of external shock real organizations face." Slow down: "Don't react emotionally. Analyze carefully." Decompose: "What changed? Who's most affected? Which industries feel the ripple?" Redirect: "Adapt or fight?" Scoring: Thoughtful adaptation gets credit. Panic or dismissal loses points.</p>
Labor Displacement Moral Tension	Consulting rep is conflicted — AI automates junior associate work but destroys the talent pipeline.	<p>Name the tension: "This is one of the realest tensions in professional services AI strategy." Separate layers: financial return, execution, talent pipeline sustainability, client perception. Present options: aggressive automation (fast margin, high pipeline risk), hybrid (slower savings, preserved pipeline), full reskilling (high cost, uncertain ROI). Scoring: Score fairly regardless of choice. They own the tradeoff. Debrief: "Your talent pipeline may not recover. Junior associates are how you develop senior partners."</p>
Major Inject Triggers Market Crisis	Inject hits: "AI-driven trading error causes 12% market drop. Regulators investigating."	<p>Describe cascade: "This isn't a single-company problem. It's market-wide. All industries are affected." Force immediate decision: "What's your response in the next 48 hours?" Highlight consequences: "Did you build safeguards? If you went all-in on automation, you're exposed." Scoring: Quick, transparent response gets credit. Defensive or slow responses lose trust. Draw lesson: "Tail risk is real. Hedges and governance matter more in crises than in calm markets."</p>

Scenario	What You'll See	Facilitator Move
Bar Rule Uncertainty (NEW for V7.4)	Law rep wants to deploy AI for client advice but is uncertain about bar association rules across jurisdictions.	Acknowledge the constraint: "Bar rules vary by jurisdiction. This is a genuine uncertainty." Frame the decision: "Deploy cautiously in permissive jurisdictions, or wait for clarity nationally. Each has costs." Probe: "If you wait, competitors gain 6-12 months of efficiency advantage. If you deploy, you risk compliance action." Reference malpractice: "Who bears liability if AI-assisted legal work contains errors?" Scoring: Thoughtful navigation of regulatory uncertainty gets credit. Ignoring it does not.

Banded Scoring Framework (V7.4)

Each submitted decision is scored across three dimensions using **banded scoring: {-2, 0, +2}** with **+/-3 on red-flag triggers**.

Three Scoring Dimensions:

1. Strategic Fit (-2 to +2): Does this decision align with the participant's industry position and market reality? Are the stated bands (Spend, Timeline, Complexity, Dependency, Scale, Success Metric) realistic and well-reasoned?
2. Execution Risk (-2 to +2): Can this actually be executed given constraints (time, capital, talent)? Are risks and mitigations credible?
3. Tail Risk (-2 to +2): Does this decision create or mitigate major downside scenarios?

Red-Flag Triggers:

- Exceptionally implausible strategy or unrealistic bands (-3 on Execution)
- Exceptionally prescient risk management or hedging (+3 on Tail Risk)
- Gross negligence of stakeholder management (-3 on Strategic Fit)

Banded Framework Integration: Participants should use the six-band decision format (Spend/Commitment, Time-to-Impact, Execution Complexity, Dependency, Scale, Success Metric) to avoid false precision. Your scoring should reinforce this: "Your 'Material but manageable' Spend band is realistic given the Regional pilot Scale. Your 'High' Execution Complexity band reflects the regulatory dependency — well-reasoned."

How to Score & Communicate:

When scoring a decision, tell the participant explicitly:

- "Your decision scores +2 on Strategic Fit (Spend, Timeline, and Scale bands are realistic; well-aligned with industry fundamentals)."

- "0 on Execution Risk (your 'High' Complexity band is credible, but mitigations are thin on the regulator dependency piece)."
- "+1 on Tail Risk (you've thought through downside if regulatory approval delays; fallback plan is weak)."
- "Total this round: +1 point. Running total: [X]."

This transparency builds trust and reinforces the banded framework.

Scoring Time Management with 11 Industries: With up to 11 decisions per round (vs. 5 in V7), plan for approximately 1 minute per decision during the 12-minute scoring block. Consider having the helper serve as a second scorer, or pre-prepare scoring templates with industry-specific baselines to accelerate adjudication.

Document Version: Project Threshold V7.4 — Lead Facilitator Briefing **Last Updated:** March 2026
Format: 8-hour single-day format, 11 industries (5–11 participants) + 1-2 facilitators, individual participant model, Collective Bonus (all rounds) + Facilitator Market Shock (R2)

Post-Exercise Synthesis Memo Template

One-Page Memo for Each Industry Participant | Completed Before Departure

Purpose

This memo forces each industry participant to externalize learning and translate scenario insights into business implications. It's not a summary of what happened in the exercise; it's a forward-looking synthesis of strategic implications.

Completion Time: 10 minutes per participant **Format:** Single page, bullet points **Audience:** Participant and facilitator (shared at end of exercise)

Template

INDUSTRY: [Industry Name (e.g., Retail, CPG, Healthcare Provider, Healthcare Payer, Finance, Consulting, Law, Manufacturing, Logistics, Big Tech, B2B/B2C SaaS)] **DATE:** [Exercise Date]
SCENARIO CONFIGURATION: Baseline + Type A (Copilots Everywhere)

Key Forecast Surprises

What was genuinely surprising about how the scenario unfolded? This is NOT what you expected when you started. Examples:

- "We thought AI productivity would hit all industries equally. We were shocked to see Finance and Consulting pull far ahead while Healthcare Provider and Manufacturing remained constrained by regulation and complexity."
- "We expected regulatory backlash to slow adoption. Instead, regulation created compliance cost advantage for large players, accelerating consolidation."
- "We thought labor displacement would be gradual. In the scenario, it happened fast (2-3 years) in specific occupations and created real political pressure on our industry."
- "The Round 2 Market Shock disrupted our strategy more than competitive moves. We had to adapt to externally imposed constraints."
- "Collective Bonus results surprised us — the group's perception of our strategy diverged from our own confidence level."

Space: 3-4 bullets

What We Got Wrong (Starting Assumptions)

Which of your baseline assumptions proved incorrect? Be honest about mental models that broke down.

- "We assumed our supply chain was resilient to AI disruption. By Round 3, we realized competitors' AI-optimized supply chains had cost advantage we couldn't match with our legacy systems."
- "We thought customers would reject AI-driven decisions. Data showed they prefer personalization and efficiency even when AI-driven, as long as we disclose it transparently."
- "We believed regulation would be uniform. It was much more fragmented — strict in some jurisdictions, light in others — which forced us to manage multiple compliance regimes."
- "We overestimated our ability to execute AI deployment quickly. Organizational constraints (talent, legacy systems, governance) slowed us more than technology constraints did."

Space: 3-4 bullets

Strategic Implications (Real-World Industry)

Translate scenario learning into implications for your actual industry or organization. What does this scenario suggest about the 2026-2030 future?

For Retail:

- "AI will drive margin expansion for leaders but create existential pressure for mid-market players. We need to decide: consolidate as acquirer, position as acquisition target, or find a defensible niche."
- "Customer trust in AI-driven recommendations is fragile. Early investment in transparency and governance will become a competitive moat."
- "Labor displacement will be politically visible in retail and customer service. Workforce transition programs are not just CSR — they're competitive necessity."

For CPG:

- "Data fragmentation across retail partners remains the binding constraint. AI value requires data access agreements with major retailers."
- "AI-driven consumer insight is high-value but requires partnership models that protect proprietary brand data."

For Healthcare Provider:

- "Regulatory approval timelines are the binding constraint, not technology capability. Engagement with FDA, state boards, and payers early is make-or-break."
- "Physician resistance is the real barrier. Human-AI integration must be presented as augmentation, not replacement."
- "Cost savings will be captured by payers, not providers. Contractual protection and shared upside are essential."

For Healthcare Payer:

- "Claims processing automation ROI is massive but regulatory scrutiny on algorithmic coverage decisions is intensifying."
- "Prior authorization AI creates Provider tension. Collaborative deployment protects against regulatory backlash."

For Finance:

- "Finance will move faster than most industries on AI deployment, driven by ROI clarity and trading/fraud detection opportunities."
- "Agentic workflows will mature faster than expected. Governance and risk management are make-or-break."
- "Talent competition with Big Tech is real and accelerating. AI specialist compensation is material cost pressure."

For Consulting:

- "Junior talent pipeline disruption is existential. AI copilots perform work that previously required 2-3 years of associate development."
- "Pricing model must shift. Clients will not pay labor-hour rates for AI-augmented delivery. Outcome-based and fixed-fee models are urgent."
- "Knowledge management becomes competitive moat. Firms that encode institutional expertise in AI systems create defensible advantage."
- "Advisory demand for AI transformation is a tailwind — but only for firms that have credibly transformed themselves."

For Law:

- "Billable hour model faces structural disruption. Clients will demand cost savings from AI-generated legal work product."
- "Bar rule uncertainty across jurisdictions creates compliance overhead that large firms can absorb but small firms cannot."
- "Malpractice liability for AI-assisted legal advice is the unresolved question. Early investment in governance and insurance is non-negotiable."
- "Associate leverage model is eroding. Fewer junior associates needed, but senior partner value increases."

For Manufacturing:

- "AI-driven predictive maintenance and quality control are table-stakes. Competitors with AI-optimized operations create permanent cost advantages."
- "Legacy OT system integration is the binding constraint. Capital allocation for modernization competes with AI deployment."

For Logistics:

- "Route optimization and demand forecasting AI are proven. The competitive question is speed of warehouse automation."
- "Driver shortage creates opportunity for AI-assisted driving tools. Regulatory approval timeline is the constraint."

For Big Tech:

- "Big Tech is both an enabler and a beneficiary of AI diffusion. Manage dual role: selling AI to other industries while protecting your own margins."
- "AI capex spiral is real. Balance aggressive AI feature development with unit economics."
- "Regulatory scrutiny on data, models, and platforms will intensify. Antitrust, data privacy, and content moderation risks are material."

For B2B/B2C SaaS:

- "AI feature integration is competitive necessity, not differentiator. Customers expect AI features included in existing subscriptions."
 - "Margin pressure from AI infrastructure costs will hit faster than expected. Unit economics require careful management."
 - "AI-native startup competition is real. Incumbents with large customer bases have distribution advantage but must move fast on product."
-

Surprises About Decision Quality

What surprised you about how other industry participants decided vs. how you decided?

- "We expected Finance to move most aggressively. They were cautious due to regulatory uncertainty."
 - "Consulting's pricing pressure created knock-on effects we didn't anticipate for our own advisory spending."
 - "Law's hesitation on AI deployment surprised us — the malpractice liability concern is more binding than we realized."
 - "We underestimated how quickly labor displacement would become a political liability. Industries that invested early in reskilling programs faced less pressure later."
-

No-Regrets Actions

What are 3 actions you would take in your real organization **regardless** of which AI timeline/capability type actually unfolds? These are hedges that pay off in multiple scenarios.

1. Build compliance and governance infrastructure now. Regulation is coming in some form in your industry. Firms that build governance early (board oversight, audit trails, risk frameworks) will have lower costs and faster deployment when rules are finalized. This pays off whether regulation is strict or loose, and whether you move fast or slow.

2. Invest in workforce transition and reskilling. Labor displacement will happen in some form and at some pace in your industry. Having a credible workforce transition plan (reskilling programs, redeployment pathways, severance frameworks) protects brand, reduces labor risk, mitigates regulatory pressure, and positions the firm as responsible. This pays off in both aggressive and defensive AI scenarios.

3. Establish cross-functional AI strategy group. Siloed decision-making will fail. You need Finance, Technology, HR, Legal, and (for regulated industries) Compliance coordinating on AI strategy. This reduces siloed bets, ensures balanced risk-taking, and prevents surprises where one function makes a decision that creates liability for another.

(Optional 4th action for your industry): Customize based on your industry-specific learning from the exercise.

Example 1: Retail Memo (Baseline)

INDUSTRY: Retail **DATE:** March 1, 2026 **SCENARIO CONFIGURATION:** Baseline + Type A (Copilots)

Everywhere)

Key Forecast Surprises

- AI adoption outpaced consumer trust more than expected. By Round 2, we were deploying advanced personalization while customer trust scores remained low (45% trust in AI-driven recommendations). This mismatch created brand risk we hadn't anticipated.
- Automation benefits were concentrated in back-office and logistics, not customer-facing services. Customer satisfaction actually declined slightly in highly automated service channels.
- Competitor moves forced pace faster than fundamentals justified. By Round 3, we were in a speed race we didn't want to enter. Slower competitors who invested in trust and quality actually had better margins despite lower market share.

What We Got Wrong

- Assumption: "Consumers will accept AI automation if outcomes are better." Reality: Consumers care about control and transparency more than outcome quality.
- Assumption: "Productivity gains from automation will offset headcount reduction costs." Reality: Retraining and change management costs were higher than modeled.
- Assumption: "AI in retail would differentiate our value proposition." Reality: Competitors matched capabilities in 6-12 months. No durable moat emerged from technology alone.

Strategic Implications

- Retail AI advantage is not durable unless paired with superior customer service and brand management.
- Customer-facing AI will be table-stakes, not a differentiator. Differentiation will come from superior data (inventory, demand forecasting) and supply chain agility.
- Labor strategy is critical. Companies that transition workers into higher-value roles retain institutional knowledge. Companies that cut aggressively face higher attrition and reputation risk.
- Margin expansion from AI will be real but temporary. As competitors match, margin advantage compresses. Need continuous innovation to maintain edge.

No-Regrets Actions

1. Build customer transparency and consent framework now: Disclose when AI is used. Offer opt-outs for data collection and personalization. This becomes table-stakes by 2028. Early movers build trust; late movers face backlash.
2. Develop comprehensive workforce transition program: Investment in reskilling, redeployment, and severance. Announce commitment publicly. This protects brand and builds organizational capability.
3. Map and optimize supply chain with AI roadmap: Whether aggressive or defensive on customer-facing AI, supply chain optimization is universally valuable. 24-month roadmap for demand forecasting, inventory optimization, logistics automation.

Example 2: Finance Memo (Baseline)

INDUSTRY: Finance **DATE:** March 1, 2026 **SCENARIO CONFIGURATION:** Baseline + Type A
(Copilots Everywhere)

Key Forecast Surprises

- Finance moved faster on AI adoption than most industries, driven by clear ROI on trading and fraud detection.
- Talent competition with Big Tech was more intense than expected. AI specialist salaries increased 25-30% year-over-year, creating material cost pressure.
- Regulatory scrutiny on algorithmic decision-making hit faster than anticipated. Fairness audits became mandatory by Round 3.

What We Got Wrong

- Assumption: "Regulatory approval would be the constraint." Reality: Organizational execution (talent, legacy systems) was the bigger bottleneck.
- Assumption: "AI would reduce compliance costs." Reality: AI created new compliance requirements (model risk management, algorithmic audits) that partially offset savings.
- Assumption: "We could build AI capabilities internally." Reality: Needed to acquire or partner to move at competitive speed.

Strategic Implications

- Finance will continue to lead on AI adoption. The competitive question is execution quality, not technology access.
- Governance and risk management are make-or-break. Firms that build strong AI governance early will have faster regulatory approval and lower compliance costs.
- Talent strategy shift required. Can't compete with Big Tech on compensation alone. Need to offer mission, scope, and career development.

No-Regrets Actions

1. Establish AI governance infrastructure: Board-level AI committee, model risk framework, algorithmic audit capability. Non-negotiable.
2. Build AI talent pipeline beyond compensation: Partnerships with universities, internal training programs, career paths for AI specialists. Compete on scope and mission, not just salary.
3. Invest in legacy system modernization: AI deployment bottlenecked by legacy infrastructure. Modernization investment pays off in any scenario.

Example 3: Consulting Memo (Baseline)

INDUSTRY: Consulting **DATE:** March 1, 2026 **SCENARIO CONFIGURATION:** Baseline + Type A
(Copilots Everywhere)

Key Forecast Surprises

- Junior talent disruption was more acute than expected. AI copilots could perform 60-70% of associate-level research and analysis within 12 months of deployment.
- Client pricing pressure was immediate. By Round 2, clients were demanding 20-30% fee reductions for AI-augmented engagements.
- Advisory demand for AI transformation was strong — but clients questioned whether Consulting firms had genuinely transformed themselves.

What We Got Wrong

- Assumption: "AI will increase revenue per consultant." Reality: Clients got faster delivery but demanded lower prices. Revenue per engagement fell.
- Assumption: "We can cut junior staff and maintain quality." Reality: Knowledge transfer broke down. Senior partners lost the talent pipeline that fed their future capability.
- Assumption: "Our brand and relationships would protect pricing." Reality: AI-native advisory startups and Big Tech advisory arms created credible alternatives.

Strategic Implications

- Business model transformation is urgent. Cost-cutting via AI is easy; revenue growth requires new pricing models (outcome-based, equity-upside, subscription advisory).
- Talent pipeline must be redesigned. The "hire juniors, grow them into partners" model is breaking. Need experienced hires + AI-augmented development paths.
- Knowledge management becomes the competitive moat. Firms that encode institutional expertise into AI systems create defensible advantage that AI-native startups can't replicate.
- Credibility requires self-transformation. Can't advise clients on AI transformation without visibly doing it yourself.

No-Regrets Actions

1. Redesign engagement and pricing models: Pilot outcome-based pricing with 3-5 clients immediately. Measure client satisfaction, margin, and scalability.
2. Build AI-native delivery capability: Create 2-3 showcase engagements where AI augmentation is central. Use these as proof of concept for clients and for internal learning.
3. Invest in knowledge management infrastructure: Encode proprietary methodologies and institutional expertise in AI-accessible systems. This is the long-term moat.

Document Version: Project Threshold V7.4 — Post-Exercise Synthesis Memo Template **Last Updated:** March 2026 **Format:** Single-page industry memo; V7.4 eleven-industry configuration (5-11 participants + 1-2 facilitators)

Room Setup & Logistics Guide

Overview

This guide provides detailed specifications for physical and technical setup of the Project Threshold Tabletop Exercise in the V7.4 single 8-hour format with up to 11 individual industry representatives + 1-2 facilitators. Proper logistics planning ensures smooth execution, clear communication, and optimal participant engagement.

Optimal Room Layout (V7.4: 8-Hour Format)

Space Requirements

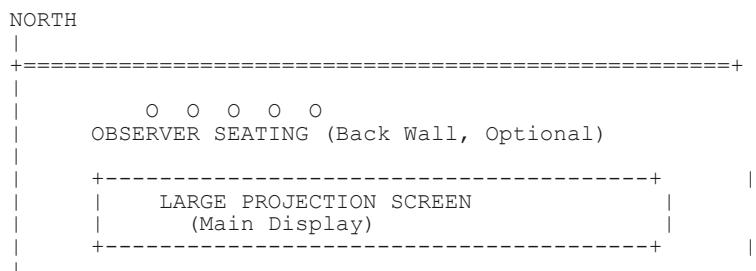
Room size scales with participant count:

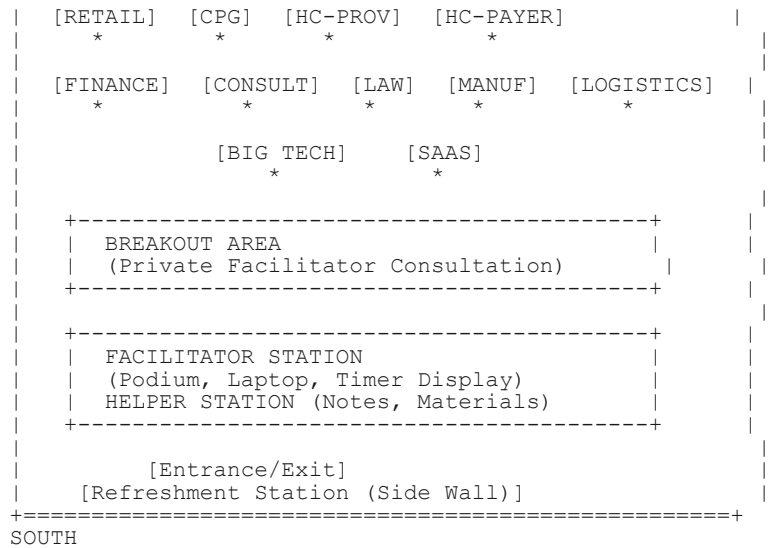
Configuration	Participants	Minimum Room Size
5 industries + facilitators	7–8	1,000 sq ft
7 industries + facilitators	9–10	1,100 sq ft
9 industries + facilitators	11–12	1,200 sq ft
11 industries + facilitators	13–14	1,400–1,500 sq ft

Room must accommodate:

- Up to 11 industry positions (round-table or seminar seating)
- Facilitator station with clear sightlines to all positions
- Dedicated breakout consultation area
- Observer seating zone (optional)
- Beverage/refreshment station

Floor Plan Diagram (V7.4 Configuration — Full 11 Industries)





Legend:

- * = Individual industry representative seat
- [INDUSTRY] = Industry designations (up to 11)
- O = Individual observer seat (optional)
- +--+ = Functional zones

Note: For fewer than 11 industries, remove unneeded positions and adjust spacing. With 5 industries, use a single semicircle or small round table.

Detailed Position Specifications

Industry Positions (up to 11 positions)

- Configuration: Round-table seating (all industry reps at one large table or two adjacent tables) OR seminar-style seating (individual chairs arranged in arc or horseshoe facing projection screen)
- Seating: 1 individual industry representative per position; ergonomic chair; room to write/work
- Spacing: Minimum 3-4 ft between positions for individual workspace (round-table) or 2-3 ft apart in seminar arrangement
- Equipment at each position:
 - Work surface (writing desk/table edge, minimum 18" x 24")
 - Water glass, pens (2-3), notepad
 - Industry folder with industry card, decision worksheets, scenario overview
 - Scoring reference sheet
 - Optional: desk lamp for individual lighting

Facilitator Station

- Position: Front-center, slightly elevated if possible (on low platform or standing podium)
- Setup:
 - Podium or high table (36"-42" height)
 - Laptop positioned for speaker notes visibility
 - Physical timer display (large digital clock, 24"+ visible from all positions)
 - Bell or chime for transitions
 - Printed master copies of all materials, injections, and scoring rubrics
 - Master scoring tracker/whiteboard for posting industry scores

Helper Station (adjacent to Facilitator)

- Backup device (tablet with presentation loaded)
- Extra printed materials and worksheets (20% backup copies)
- Private information cards from 05_Participant_Private_Cards/ (organized by round and industry)
- Tech support equipment (WiFi extender, audio adapter, backup power)

Projection Screen / Display

- Size: Minimum 120" diagonal (or dual 80" monitors side-by-side)
- Position: Centered at front, elevated so bottom edge is at least 4 ft high
- Brightness: Calibrated for daylight-capable projection (300+ lumens minimum)
- Content displayed:
 - Round situation updates (one slide per round)
 - Inject cards (text and graphics)
 - Industry decision dashboards (if using live scoring)
 - Countdown timer synchronized with facilitator's master timer

Breakout Consultation Area

- Purpose: Private participant consultations, facilitator-participant discussions
- Setup: Semi-enclosed zone (curtain partition or alcove) with 1 table and 4-6 chairs
- Furnishings: Table, ergonomic chairs, small whiteboard for notes
- Acoustic: Minimal sound bleed into main exercise area

Observer Seating (Optional)

- Position: Rear of room, 3-4 ft from back wall
- Configuration: Theater-style seating or chairs with clipboards
- Notes access: Each observer has printed copy of scenario scripts and inject cards
- Sightlines: Full view of all industry positions and projection screen
- Limit: Maximum 5 observers (to maintain focus and confidentiality)

Refreshment / Materials Station

- Position: Side wall, near entrance (out of main sightlines)
- Contents:
 - Coffee/tea service
 - Water station
 - Light snacks (pretzels, granola bars, fruit)
 - Extra printed materials
 - Pens and notepads

Technology Setup Checklist

Pre-Exercise Preparation (1–2 weeks before)

Hardware Verification

- Test projector and screen for brightness, focus, color balance
- Verify HDMI and backup display connection cables
- Confirm WiFi network name, password, and bandwidth capacity
- Test facilitator laptop (screen resolution: 1920x1080 minimum) with projector
- Check timer display device (dedicated tablet or physical clock)
- Verify all power outlets and surge protectors are functional
- Test microphone system (recommended for 11+ participants in room)
- Confirm backup audio system exists (portable Bluetooth speaker minimum)

Software Setup

- Install presentation software (PowerPoint, Google Slides, Keynote)
- Create master slide deck with:
 - Title/welcome slide
 - Practice round + Round 1-4 situation updates (one slide per round)
 - Inject slides (formatted for legibility, high contrast)
 - Scoring/results template (sized for up to 11 industries)
 - Debrief prompt slides
- Set timer software (Pomofocus, TimeandDate.com, or dedicated timer app)
- Prepare backup PDF of all slides on USB drive

Network & Connectivity

- Confirm WiFi password with IT team
- Test bandwidth: connect 3+ devices simultaneously and confirm stability
- Backup plan: hardwired Ethernet for facilitator laptop (if WiFi unstable)

Materials Preparation

- Print all industry cards (one per industry in play, plus 2 spares)
- Print decision worksheets (5 per industry: practice + Rounds 1-4)
- Print scenario overview (one per industry in play)
- Print scoring sheets (one per industry, plus 1 master for facilitator)
- Print all inject cards (one per industry in play, plus facilitator set)
- Print private information cards from 05_Participant_Private_Cards/ (Rounds 1-3 only, organized by industry/round)
- Print AI Adoption Arc summaries from 04_Participant_AI_Adoption_Arcs/ (one per industry in play)
- Print observer scripts (if observers present)
- Bind materials in folders by industry designation
- Create backup set of all printed materials in sealed container

Day-of Setup (1:30 before exercise start)

60–90 Minutes Before

- Arrive early; lay out all materials in room
- Arrange industry positions per floor plan (adjust for actual participant count)
- Place industry folders/materials face-down at each position
- Set up facilitator station: podium, laptop, timer display

- Connect laptop to projector; test display of all slides
- Load presentation software and test all transitions/timers
- Verify projection screen is centered, focused, and bright enough
- Set up optional observer seating along back wall
- Verify microphone system (if applicable) with sound check
- Confirm all WiFi connections are live

30–45 Minutes Before

- Final walkthrough: check all positions have pens, water, notepads
- Test timer display with countdown (ensure visibility from all positions)
- Project test slide to confirm text is readable from farthest position
- Test microphone by speaking into it
- Verify backup printed materials are accessible
- Brief helper on start time and material distribution sequence

15–30 Minutes Before

- Welcome participants as they arrive
- Direct participants to assigned industry positions
- Confirm each position has all baseline materials
- Verify facilitator WiFi connection
- Perform final projector/screen alignment check
- Set timer to "00:00" display

At Start Time

- Display welcome slide on projection screen
- Facilitator stands at podium; gets all participants' attention
- Brief explanation of exercise flow, roles, and timing
- Confirm all positions have industry cards (do not open until instructed)
- Set timer and begin countdown for practice round

Participant Materials Checklist

Materials Per Position (Place in labeled folder)

Baseline Materials (At Arrival)

- Industry Card: Industry designation, participant name, industry background, initial position
- Scenario Overview: One-page summary of Project Threshold starting scenario
- Decision Worksheets: 5 blank worksheets (practice + Rounds 1-4)
- Banded Framework Reference Card: Quick-reference for 6 decision bands (Spend, Timeline, Complexity, Dependency, Scale, Success Metric) + examples
- Scoring Reference Sheet: Quick-reference for banded scoring rubric ({-2, 0, +2} + red-flag triggers)

- [] AI Adoption Arc Summary: Industry-specific round-by-round adoption context (from 04_Participant_AI_Adoption_Arcs/)
- [] Base Case Fallback Bank: Reference showing what automatic fallback looks like if industry not submitted (see 02BaseCaseFallbackBank.md)
- [] Notepads & Pens: 3-4 pens, 2-3 legal pads per position

Materials Distributed During Exercise

Private Information Cards (Rounds 1-3 only)

- Delivered by helper to specific industry participants during inject processing time
- Sourced from 05_Participant_Private_Cards/ folder
- Contains confidential data or intelligence
- Timing coordinated with exercise schedule
- Example: "Round 2, at 10:35 AM: Reveal market data card to Finance participant"

Inject Cards (Situation Updates)

- Displayed on projection screen + printed for each industry participant
- Timing: Read aloud by facilitator
- Format: Large text (24+ pt font), clear title, high contrast

Master Scoring Sheets

- Held by facilitator only
- Updated in real-time as decisions come in
- Displayed to full group periodically (every 1-2 rounds)
- Final scores tabulated at end and shared in debrief

Detailed Timing for 8-Hour Full-Day (V7.4 Standard)

Time	Duration	Activity	Notes
8:30	15 min	Arrival & Setup	Materials distribution, name badges, seating by industry
8:45	15 min	Welcome & Exercise Overview	Opening remarks, objectives, explain rounds 1–4, participant introductions
9:00	10 min	PRACTICE MICRO-ROUND	Simple scenario; low-stakes decision; facilitator scores live; Q&A

Time	Duration	Activity	Notes
9:10	5 min	ROUND 1: Situation Update	Project slide; read injects
9:15	3 min	ROUND 1: Inject + Private Card	Distribute private cards from 05_Participant_Private_Cards/
9:18	15 min	ROUND 1: Individual Decision Prep	Each industry rep prepares alone
9:33	3 min	ROUND 1: Decision Submission	Each rep submits worksheet
9:36	12 min	ROUND 1: Facilitator Scoring	Score all submitted decisions
9:48	27 min	ROUND 1: Cross-Industry Discussion & Results	Scores posted; Collective Bonus (Optional); cross-industry dynamics
10:15	15 min	Break	Coffee, restrooms
10:30	5 min	ROUND 2: Situation Update	Project slide; read injects
10:35	3 min	ROUND 2: Inject + Private Card	Distribute Round 2 private cards
10:38	15 min	ROUND 2: Individual Decision Prep	Each rep prepares
10:53	3 min	ROUND 2: Decision Submission	Each rep submits
10:56	12 min	ROUND 2: Facilitator Scoring	Score decisions; apply Facilitator Market Shock (~3 min within this block)
11:08	27 min	ROUND 2: Cross-Industry Discussion & Results	Scores + Market Shock announced; Collective Bonus (Optional); IHS conditions
11:35	40 min	Lunch	Informal reflection encouraged
12:15	5 min	ROUND 3: Situation Update	Project slide; read injects

Time	Duration	Activity	Notes
12:20	3 min	ROUND 3: Inject + Private Card	Distribute Round 3 private cards
12:23	15 min	ROUND 3: Individual Decision Prep	Each rep prepares
12:38	3 min	ROUND 3: Decision Submission	Each rep submits
12:41	12 min	ROUND 3: Facilitator Scoring	Score decisions
12:53	27 min	ROUND 3: Cross-Industry Discussion & Results	Scores posted; Collective Bonus (Optional); IHS conditions
1:20	15 min	Break	Restrooms, energy reset
1:35	5 min	ROUND 4: Situation Update	Project slide; read final injects
1:40	3 min	ROUND 4: Inject Distribution	Final market conditions; no private cards R4
1:43	15 min	ROUND 4: Final Individual Decision	Each rep makes final decision
1:58	3 min	ROUND 4: Decision Submission	Each rep submits
2:01	12 min	ROUND 4: Facilitator Scoring	Score final decisions
2:13	27 min	ROUND 4: Cross-Industry Discussion & Results	Final discussion; Collective Bonus (Optional)
2:40	10 min	Break + Facilitator Prep	Consolidate scores, prepare debrief
2:50	15 min	Individual Reflections	Each rep shares 1–2 min highlights
3:05	25 min	Cross-Industry Dynamics Discussion	Facilitator-led; inter-industry dependencies
3:30	20 min	No-Regrets Actions & Takeaways	Consolidate actions; each rep states key takeaway

Time	Duration	Activity	Notes
3:50	40 min	Wrap-Up & Close	Feedback surveys, closing remarks, synthesis memo scheduling

Resource Allocation (V7.4: 8-Hour Format)

Resource	5 Industries	7 Industries	11 Industries	Notes
Facilitators	1 lead + 1 helper	1 lead + 1 helper	1 lead + 1 helper (+ optional 2nd scorer)	Facilitator runs Collective Bonus + Market Shock
Observers	0-3 (optional)	0-3 (optional)	0-5 (optional)	For debrief/training purposes
Room Size	1,000 sq ft	1,100 sq ft	1,400–1,500 sq ft	Scale with participant count
Seating	Round-table (1 table)	Round-table or arc	Arc or 2 adjacent tables	1 chair per industry rep
Chairs	7-9	9-11	13-15	Industry reps + facilitators + observers
Work Surfaces	5 positions	7 positions	11 positions	Minimum 18" x 24" per rep
Projection Setup	1 large screen (120"+)	1 large screen (120"+)	1 large screen or dual 80"	Central, elevated
Microphone	Optional	Optional	Recommended	Helps with 13+ people in room
WiFi Bandwidth	10+ Mbps	10+ Mbps	10+ Mbps	For tech backup only
Printing (pages)	100-150	150-200	250-350	All materials including backups
Budget Estimate	\$2,500–3,500	\$3,000–4,000	\$4,000–5,500	Venue, catering (light), A/V
Pre-Setup Time	2 hours	2-3 hours	3 hours	Materials, tech, room layout

Resource	5 Industries	7 Industries	11 Industries	Notes
Backup Materials	15-20% extra	15-20% extra	15-20% extra	Industry cards, worksheets, injects

Contingency Planning

Technology Failures

Projector Outage

- Backup: Large monitor (55"+) or secondary projector on-site
- Mitigation: Contact IT; have backup on standby
- Continuation: Printed inject cards; manual timer; flip chart for scoring

WiFi Disruption

- Backup: Paper worksheets (no digital submission)
- Mitigation: Dedicated hotspot available
- Continuation: Decisions collected manually

Laptop Crash

- Backup: Secondary laptop with presentation pre-loaded
- Mitigation: Test morning-of; cloud backup synced
- Continuation: Switch to backup within 5 minutes

Participant Contingencies

Late Arrivals

- Participants arriving 10+ minutes late join at next natural break
- Quick industry briefing (5 min) from facilitator
- Direct to assigned industry position

Early Departures

- Confirm departure time during opening remarks
- If during active round, industry receives base case fallback for remaining rounds
- Optionally provide recording for reference

Illness/Emergency

- Evacuation plan posted; meet-up location established
- Exercise resumes 15 min after all-clear
- Backup industry cards available for alternates

No-Show for Assigned Industry

- If a participant doesn't arrive, that industry can be dropped from the exercise or absorbed by the facilitator as a base-case industry
- Facilitator adjusts Market Shock targets and Collective Bonus pool to exclude that industry
- Other participants are not affected

Content & Timing Modifications

Timing Overruns

- Prioritize debrief time (most valuable learning activity)
- Shorten cross-industry discussion (select specific industry pairs instead of full group)
- Reduce or defer synthesis activity

Scoring Time with 11 Industries

- If scoring 11 decisions in 12 minutes is tight, have helper serve as second scorer
- Alternatively, score in priority order (most interesting/controversial decisions first) and announce remaining scores during discussion
- Pre-prepare scoring templates with industry-specific baselines

Technical Complexity

- Simplify inject cards if comprehension issues arise
- Extend processing time; reduce decision time
- Provide written summary of key inject points

Pre-Exercise Logistics Checklist (1 Week Before)

- Confirm final participant count and industry assignments
- Select which industries are in play (if fewer than 11)
- Verify room reservation, access, parking
- Test all technology (projector, WiFi, timer, audio)
- Confirm catering order (beverages, snacks, lunch)
- Print all materials (industry cards, worksheets, injects, private cards from 05_Participant_Private_Cards/, AI Adoption Arcs from 04_Participant_AI_Adoption_Arcs/)
- Prepare facilitator scripts and inject deployment schedule
- Brief facilitator and helper on process and timing
- Create backup materials and store in secure location
- Send reminder email: start time, location, parking info
- Prepare post-exercise feedback form
- Confirm observer assignments (if used)

Post-Exercise Logistics

Immediately After (15–30 min)

- Collect all completed decision worksheets and scoring sheets
- Thank facilitators and helper; debrief on any issues
- Pack all printed materials securely
- Take photos of whiteboard/notes if applicable

Within 24 Hours

- Tally final scores and verify accuracy
- Photograph or scan participant feedback forms
- Compile summary of industry decisions and strategic patterns
- Send thank-you email to all participants

Within 1 Week

- Analyze feedback survey responses
- Prepare post-exercise debrief summary
- Archive all materials (digital and printed) in secure location
- Schedule follow-up session (if applicable)

Document Version: Project Threshold V7.4 — Room Setup & Logistics Guide **Last Updated:** March 2026 **Format:** Single 8-hour format, 11 industries (5–11 participants) + 1-2 facilitators, individual participant model, novice-friendly

README

Integration & Navigation Guide

Document Date: March 2026 **Exercise Timeframe:** 2026–2030 (48-month near-term AI economic impact) **Duration:** 8-hour single format (no alternative formats) **Participant Structure:** 5–11 individual industry representatives + 1–2 facilitators

What is Project Threshold V7.4?

Project Threshold V7.4 is a scenario-based tabletop exercise that stress-tests economic forecasts and strategic decisions across up to eleven US industries during rapid AI adoption (2026–2030). V7.4 uses an independent industry model where each industry is represented by its own decision-maker, with Collective Bonus and Facilitator Market Shock mechanics replacing the former Investment Committee role. This version maintains the operational rigor and intuition-building focus of earlier versions while increasing granularity, flexibility, and cross-industry richness.

The exercise is designed to run in 8 continuous hours with 1 lead facilitator and 1 helper, producing actionable strategic insights for individual participants.

What's New in V7.4 (vs. V7)

Structural Model: Sectors to Industries

- From (V7): 5 sectors, each grouping 2 sub-industries under a single participant (10 personas total, 5 decision-makers)
 - Consumer (Retail + CPG)
 - Healthcare (Provider + Payer)
 - Finance & Professional Services (Finance + Professional Services)
 - Supply Chains (Manufacturing + Logistics)
 - Software & Tech (Big Tech + B2B/B2C SaaS)
- To (V7.4): 11 independent industries, each with its own participant and decision worksheet
 - Retail
 - CPG
 - Healthcare Provider
 - Healthcare Payer
 - Finance
 - Consulting
 - Law
 - Manufacturing
 - Logistics
 - Big Tech

- B2B/B2C SaaS
- Impact: Each participant now owns one industry rather than juggling two. Decisions are sharper. Cross-industry dynamics are richer because more distinct perspectives are in the room.

Finance & Professional Services Split into Three

- From (V7): Single "Finance & Professional Services" sector with one participant covering Finance + Professional Services
- To (V7.4): Three independent industries — Finance, Consulting, Law — each with its own participant, packet, private cards, and adoption arc
- Consulting profile added: Big Four / MBB-class firm (\$20–25B revenue, 30K–50K employees). AI transformation advisory, junior talent pipeline disruption, pricing pressure from AI-augmented delivery.
- Law profile added: AmLaw 50 firm (\$3–6B revenue). Billable hour disruption, bar rule uncertainty across jurisdictions, malpractice liability for AI-assisted work, associate leverage erosion, legal AI platform competition.

Big Tech Scope Clarification

- V7.4 explicitly excludes AI lab / foundation model development from the Big Tech industry profile
- Scope: Cloud infrastructure, advertising platforms, enterprise software, consumer devices, content platforms
- Rationale: Keeps the exercise grounded in business strategy rather than model capability speculation

Flexible Participant Count (5–11)

- V7: Fixed at 5 sector participants (each covering 2 personas)
- V7.4: Flexible from 5 to 11 industry participants (one industry per person)
- Recommended minimums: 5-person config selects one industry per industry cluster to maximize cross-industry diversity; 7-person adds Consulting and SaaS; full 11-person uses all industries
- IC role eliminated (V7.3+): Replaced with Collective Bonus (optional, every round) and Facilitator Market Shock (R2 only)

Private Cards Extracted to Separate Folder

- V7: Private cards embedded in sector packets or round materials
- V7.4: Private cards extracted to dedicated 05_Participant_Private_Cards/ folder with one file per industry (05aRetailPC.md through 05kSaaS_PC.md)
- Rationale: Cleaner separation of confidential materials; easier for facilitator to manage distribution

AI Adoption Arc Extracted to Separate Folder

- V7: AI adoption progression embedded in sector packets
- V7.4: AI Adoption Arc extracted to dedicated 04_Participant_AI_Adoption_Arcs/ folder with one file per industry (04aRetailArc.md through 04kSaaSArc.md)
- Content: Round-by-round distribution showing how AI adoption unfolds for each industry across the four rounds
- Rationale: Provides participants with a clear progression framework; supports round-specific context

Standardized Industry Packet Template

- V7.4 includes: 03_z_Template_Industry_Packet.md — a blank template establishing consistent structure across all 11 industry packets
- All packets follow the same format: Industry overview, starting position, constraints, key trade-offs, decision prep checklist, AI adoption context

Directory Structure (V7.4)

```

v7.4/
├── README.md (this file)
├── 01_Exercise_Overview/
│   ├── 01_Executive_Overview.md
│   ├── 01_How_To_Play.md
│   ├── 01_Timeline.md
│   ├── 01_Expected_Outputs.md
│   └── 01_Practice_Round_Walkthrough.md
├── 02_Facilitator_Guide/
│   ├── 02_README.md
│   ├── 02_Overview_and_Runbooks.md
│   ├── 02_Adjudication_Rules.md
│   ├── 02_Plausibility_Decision_Trees.md
│   ├── 02_Scoring_Baselines.md
│   ├── 02_Industry_Health_Signal_Tables.md
│   ├── 02_Quick_Reference_Card.md
│   ├── 02_Base_Case_Fallback_Bank.md
│   └── 02_Band_Scoring_Quick_Ref.md
├── 03_Participant_Packets/
│   ├── 03_README.md
│   ├── 03_a_Retail.md
│   ├── 03_b_CPG.md
│   ├── 03_c_Healthcare_Provider.md
│   ├── 03_d_Healthcare_Payer.md
│   ├── 03_e_Finance.md
│   ├── 03_f_Consulting.md
│   ├── 03_g_Law.md
│   ├── 03_h_Manufacturing.md
│   ├── 03_i_Logistics.md
│   ├── 03_j_Big_Tech.md
│   ├── 03_k_SaaS.md
│   ├── 03_z_Rules_of_Play.md
│   ├── 03_z_Decision_Worksheet.md
│   └── 03_z_Template_Industry_Packet.md
├── 04_Participant_AI_Adoption_Arcs/
│   ├── 04_README.md
│   └── 04_a_Retail_Arc.md ... 04_k_SaaS_Arc.md (11 files)
├── 05_Participant_Private_Cards/
│   ├── 05_README.md
│   └── 05_a_Retail_PC.md ... 05_k_SaaS_PC.md (11 files)

```

```
├── 05_z_Template_Private_Card.md
├── 06_Rounds/
│   ├── 06_README.md
│   ├── 06_Round_1_Foundation.md
│   ├── 06_Round_2_Acceleration.md
│   ├── 06_Round_3_Reckoning.md
│   ├── 06_Round_4_Normalization.md
│   └── 06_Cross_Round_Integration.md
├── 07_Supporting_Materials/
│   ├── 07_README.md
│   ├── 07_Lead_Facilitator_Briefing.md
│   ├── 07_Post_Exercise_Synthesis_Memo.md
│   ├── 07_Room_Setup_Logistics.md
│   ├── 07_Diagnostic_Survey.md
│   ├── 07_Alternative_Scenario_Overrides.md
│   └── 07_Historical_Reference_Data.md
```

Detailed Changelog: V7 to V7.2 (Historical)

Industry Model Transformation (5 Sectors to 11 Industries)

- Sector groupings eliminated: Participants no longer manage two sub-industries; each industry is independently represented
- Terminology change: "sector" replaced with "industry" throughout all materials; "persona" replaced with "industry" (each participant IS their industry, not playing two personas)
- Decision count: Variable based on configuration. 5 industries x 4 rounds = 20 decisions; 11 industries x 4 rounds = 44 decisions
- Files changed: All participant packets renamed from sector-based (03aConsumer.md) to industry-based (03aRetail.md, 03bCPG.md, etc.)

New Industries Added

- Consulting (03fConsulting.md): Big Four / MBB-class firm. New profile covering AI transformation advisory, junior talent pipeline disruption, pricing pressure, knowledge management, utilization model changes
- Law (03gLaw.md): AmLaw 50 firm. New profile covering billable hour disruption, bar rule uncertainty, malpractice liability, associate leverage erosion, legal AI platforms

Industry Profile Modifications

- Big Tech (03jBig_Tech.md): Scope narrowed to exclude AI lab / foundation model development. Focus on cloud, ads, devices, enterprise software
- Finance (03eFinance.md): Separated from Professional Services; standalone profile
- All former paired industries: Now have dedicated individual packets instead of being grouped

New Folders

- 05ParticipantPrivate_Cards/: Dedicated folder for confidential private information cards, one per industry (11 files + README)
- 04ParticipantAIAdoptionArcs/: Dedicated folder for round-by-round AI adoption progression, one per industry (11 files + README)

New Files

- 03zDecision_Worksheet.md: Printable decision worksheet with all 12 required fields (checkbox bands, generous write-in space). Print multiple copies per participant.
- 03zTemplateIndustryPacket.md: Standardized blank template for industry packets
- 02FacilitatorGuide/02__README.md: Navigation index for facilitator guide materials
- 02FacilitatorGuide/02BandScoringQuickRef.md: Quick reference for banded scoring

Participant Structure Changes

- Participant count: 7–12 (V7) changed to variable 5–11 industry reps + 1–2 IC + 1–2 facilitators
- Industry selection guidance: Added for configurations with fewer than 11 participants
- IC role: Streamlined from 100-point allocation to qualitative forced actions (Market Thesis, Constraints, Winner/Laggard, Market Verdict)

Supporting Materials Updates

- 07LeadFacilitator_Briefing.md: Updated for 11-industry model, 5–11 participants, industry selection facilitation, cross-industry discussion prompts
- 07PostExerciseSynthesisMemo.md: Updated template with 11-industry examples
- 07RoomSetup_Logistics.md: Updated for up to 11 individual positions, flexible seating configurations
- 07DiagnosticSurvey.md: Updated industry list to 11 industries
- 07AlternativeScenario_Overrides.md: Updated for 11 industries, added Consulting/Law-specific override notes

Document Renames

- 01_Executive_Overview_V7.md renamed to 01_Executive_Overview.md
- README_VERSION_7.md renamed to README.md (this file)
- All files updated to V7.2 version headers

IC & Return Mechanic Redesign

- 02IndustryHealthSignalTables.md: Complete rewrite. Former 460-line valuation pipeline (margin impact → multiple shift → probability of success → portfolio bps) replaced with Industry Health Signal Tables — a simple cumulative-score-to-condition mapping (Surge / Tailwind / Steady / Headwind / Crisis) with mechanical consequences. No live financial calculations required.
- 03Investment_Committee.md: Complete rewrite. Former 487-line portfolio management packet (100-point allocation, 3 portfolio stances, rebalancing mechanics, return targets) replaced with streamlined 4 qualitative forced actions — Market Thesis (R1), Constraint Imposition (R2), Winner & Laggard (R3), Market Verdict (R4). No portfolio math.
- Industry Health Signal bands: Cumulative score determines condition. Headwind (-6 to -14) = 1 constraint imposed. Crisis (-15 or worse) = 2 constraints + defensive play required. Surge (+15+) = bypass automatic red-flag challenge on Transformational decisions.
- IC scoring criteria simplified: Coherence, signal quality, market realism. No return targets or basis points.
- Cascade updates: ~15 files updated to remove stale references to 100 points, portfolio returns, valuation translation, capital allocation, and portfolio stances. All references replaced with Industry Health Signals and IC forced actions.

Formatting Overhaul & Decision Worksheet

- 03zDecision_Worksheet.md (NEW): Standalone printable decision worksheet with all 12 required fields from Rules of Play Section J. Checkbox bands for fields 3–7 (participants check one option, not write it). Generous write-in space for handwritten responses. Two copies per file for efficient printing.
- Prose-to-table conversions across 10 files: Structured content previously buried in prose paragraphs converted to scannable tables. Files reformatted:
 - 01ExpectedOutputs.md: Decision fields table, IC Output #2 rewritten for new IC design, success indicators table
 - 01HowTo_Play.md: Daily timeline, IC actions, Industry Health Signals — all converted to tables
 - 01ExecutiveOverview.md: Three Major Dynamics table, Baseline Claims summary index added, Claim 2 productivity tables merged
 - 02Overviewand_Runbooks.md: Facilitator responsibilities table, discussion prompts consolidated
 - 03zRulesofPlay.md: Round flow sequence table, Section J worksheet fields converted to compact reference table
 - 07LeadFacilitator_Briefing.md: Pressure/breathing-room scenarios table, facilitation scenarios table, debrief questions table
 - 02AdjudicationRules.md: Red-flag categories converted to table
 - 01PracticeRound_Walkthrough.md: Example decisions converted to side-by-side comparison table, worksheet appendix replaced with reference to standalone worksheet
- Cross-file consistency verified: All 12 worksheet fields match across standalone worksheet, Rules of Play Section J, Practice Round quick reference, and 01ExpectedOutputs decision log

Document Version: Project Threshold V7.4 — Integration & Navigation Guide **Last Updated:** March 2026 **Format:** V7.4 individual participant model, 11 industries, 8-hour single format, Collective Bonus +

Facilitator Market Shock **Prepared by:** Project Threshold Development Team