

The AI Dividend and the Deadline — What They Mean for your Business

AI turbo-charges competition. The lead is won in the 2026–2028 window.

Abstract

“Practical AGI” (artificial general intelligence) has arrived and deployment is moving from pilot to production among the leaders. Now AI turbo-charges competition: every rival can cut cost and build new products fast. The advantage goes to whoever moves first. The lead is won inside the 2026–2028 window. For businesses, this is broken down in 6 chapters:

- **B1:** Rebuild business functions AI-first. This resets the cost base every rival competes on: operating cost reductions of ~21% are possible now, rising to ~48% as robotics and agents mature. Start with quick-win copilots that pay back in under a year.
- **B2:** Lead your sector by reinventing the product. The deepest AI gains create products that reshape whole industries: Robotaxi rides triple a year and the first AI-designed medicines reach the clinic. First movers out-innovate faster than rivals can copy. Build deep, narrow AI on your own data, not a generic copilot.
- **B3:** Close the AI implementation gap. The limit is execution, not the AI offering. 88% of firms run AI, yet ~95% of pilots show no profit in six months. Readiness decides the outcome, and mid-market firms can build it. One company AI platform, wired to your data and grading its own output, scales proven wins to return fast.
- **B4:** Compete in the agent-mediated economy: the platform that owns the agent owns the customer, because it holds the customer’s data and context. AI is expected to mediate \$20.9B of US retail in 2026. Ensure your best presence in the new agent platforms. Build a direct channel you own, so customer memory stays yours.
- **B5:** Quantify and drive the bottom-line impact. The prize is huge but mostly not the firm’s to keep long term. Competition eventually converts most of the cost reduction into lower prices, and by ~20 years the average firm earns near-zero additional economic profit. Speed wins the transition and justifies early investments. The moat wins the end: brand, proprietary data, customer memory, or patents.
- **B6:** Turn an understanding of the other players into business advantage: the AI firms you buy from, the people you employ and sell to, the governments that set your rules, and the world system. Competition within each player drives AI progress. This compresses a decade of innovation into the next 10 quarters, the remaining 2026–2028 window, where the lead is won.

Table of Contents

- [B1. Rebuild Business Functions AI-First](#)
 - [B1-1: Quick-win copilots that pay back in under a year](#)
 - [B1-5: AI's eventual cost reduction per function, 2023 vs now](#)
 - [B1-6: Cost reduction potential is ~21% today and ~48% eventually](#)
- [B2. Lead Your Sector or Lose It](#)
 - [B2-1: AI disruption mapped across the whole economy](#)
- [B3. Close the Implementation Gap](#)
 - [B3-1: Six avoidable traps between a pilot project and booked P&L](#)
 - [B3-2: Fifteen platform requirements a company AI platform must have](#)
- [B4. Compete in the Agent-Mediated Economy](#)
 - [B4-2: The methods that keep direct contact, sales, and data](#)
- [B5. Quantify and Drive the Bottom-Line Impact](#)
 - [B5-1: 20 years of competition erode the gain to near-zero economic profit](#)
 - [B5-2: The eight moats an agent can't arbitrage](#)
- [B6. Turn an Understanding of the Other Players Into Business Advantage](#)

Preface

This is a private synthesis of what matters in AI in mid 2026.

- **Initiative:** this is a personal, private project and freely shared.
- **Intent:** this distills what matters in AI in mid 2026 from a 360-degree perspective. Emphasis is to be quantitative even when uncertain. The aim is on how to reap the AI dividend while containing the associated risks.
- **Source:** initially built on systematic evaluation of approximately 19,000 AI-oriented news articles published in the past four months, supplemented by focused additional research where the news flow needed depth, checking, or quantitative grounding. AI itself was used, of course, to help assemble, structure, and refine the synthesis.
- **Style:** clear statements in a clear structure rather than continuous prose.
- **Caveat:** compiled in a few weekends. Please accept this as an imperfect, focused big-picture synthesis now, rather than a detail-honed article six months from now. The 2026–2028 window does not wait.
- **Disclaimer:** shared for information only. No advice. Without warranty.
- **Feedback:** is highly welcome, e.g. on [my Substack channel](#)

[↑ Back to Table of Contents](#)

Reader Overview

This document reads the 2026–2028 window through the business lens (B), the deployment layer where AI capability drives the bottom line. Five stakeholder lenses map the whole window, AI companies (A), businesses (B), individuals (I), governments (G), each racing their own rivals, and the world system (W). This view carries essential parts of lens B. The other four stakeholder lenses are already reinterpreted as seen from the business perspective in the final chapter (B6). For a business the consequences are steep: AI turbo-charges competition. The lead is won in the 2026–2028 window.

[↑ Back to Table of Contents](#)

B1. Rebuild Business Functions AI-First

Summary: Rebuilding business functions AI-first resets the cost base every rival competes on. Today's tools are worth ~21% of operating cost reduction, near-uniform across industries. By current forecasts, cost reductions eventually reach ~48% as robotics and agents mature. This forecast has climbed ~18% since 2023, as AI capabilities have developed further. About half of that cost reduction potential will materialize until 2035. One catch: AI-written code ships faster but tests have shown ~17% lower employee comprehension, so additional checks are necessary. Assign AI-agent governance to HR now, ahead of 2027–28 frameworks. Start with quick-win copilots that pay back in under a year.

Evidence:

- Software development is furthest along: one leading lab reports AI is involved in >80% of its merged code, and ~35% of merged pull requests were agent-written at Cursor by February 2026 ([InfoQ](#)).
- The counter-evidence matters: a 2025 trial found experienced developers 19% slower while still feeling ~20% faster, though a Feb-2026 update expects newer tools to reverse the slowdown ([METR](#)). A 52-junior-engineer study still found 17% lower comprehension ([Anthropic](#)): the gap between getting work done and actually learning.
- Adoption is broadening fast: Gartner projects 40% of enterprise apps will carry task-specific agents by end-2026 ([Gartner](#)) and 60% of HR tasks AI-augmented by 2030 ([Gartner](#)).
- The cost base rebuilds bottom-up. The economy sorts into four industry archetypes, physical/production, transactional, expert/knowledge, and human-service, each dominated by one kind of work and so exposed to one primary AI lever. Fourteen business functions, each cut to a single work type, sit under those archetypes, and each carries a different share of cost by archetype. Weighting per-function cost reduction (breakout B1-5) through those cost shares gives the firm-wide floor (breakout B1-6).

Breakout B1-1: Start with quick-win copilots on data you already own — they typically pay back in less than a year.

For a midsize firm (~200–2,000 employees), the question is not whether to adopt AI but which project to run first. This shortlist ranks 14 horizontal use cases by attractiveness: a large cut of the *addressed* process cost, a short build time, and a fast break-even. The top rows are quick-win copilots and document automations. They run on text a firm already owns and off-the-shelf tools, so they pay back in months. The lower rows add document and machine-learning automation. They return more but need cleaner data, so they break even nearer a year. Cost cut is of that one process's own cost, not the whole firm (breakout B1-6 sums the firm-wide effect). Ranges are indicative.

#	Business process	Research findings — business lever · AI technology · sources	Cost cut (of addressed process)	Impl. (mo)	Break-even (mo)
1	Marketing and sales content	Lever: faster drafts, localization, variants, less agency spend. Tech: AI copywriting + brand-tuned prompts, image generation. Src: BCG	30–60%	1–3	2–6
2	Customer support (Tier-1)	Lever: deflect and auto-resolve high-volume repetitive contacts, 24/7. Tech: AI chatbot + retrieval, agent assist. Src: McKinsey , Klarna	30–50%	2–4	3–9
3	Software development	Lever: faster coding and boilerplate, fewer context switches. Tech: AI pair-programmer in the code editor. Src: GitHub 55%, independent trial ~20–26%	10–30%	1–3	3–9
4	Sales development (AI rep)	Lever: auto research, enrich, qualify, and outreach so reps focus on closing. Tech: AI agent + data enrichment + CRM. Src: Salesforce	30–50%	1–3	3–9
5	IT / employee service desk	Lever: auto-resolve access and how-to tickets, deflect level-1. Tech: AI virtual agent + retrieval + service-desk integration. Src: Fin AI	20–40%	2–4	4–9
6	HR employee self-service	Lever: deflect benefits, leave, and policy queries. Tech: AI assistant + retrieval on policies + HR-system link. Src: Moveworks	30–50%	2–4	4–10
7	Software QA / testing	Lever: auto-generate and maintain tests, faster triage, shift-left. Tech: generative test generation + self-healing automation. Src: Katalon , workplace trial	30–50%	2–4	4–10
8	Knowledge-worker copilot	Lever: recover time on email, meetings, and documents. Tech: Company AI platform + speech-to-text summarization. Src: Microsoft , MS Research trial	5–15%	1–3	6–12
9	Accounts payable	Lever: touchless capture, 3-way match, straight-through processing. Tech: document processing + machine-learning matching, agentic exceptions. Src: Parseur , Rossum	60–80%	3–6	6–12
10	Accounts receivable / collections	Lever: auto cash-application and prioritized dunning, cutting days-sales-outstanding by 15–25 days. Tech: machine-learning matching + generative outreach + predictive risk. Src: Billtrust	25–40%	3–6	6–12
11	Contract review and management	Lever: extract clauses and obligations, flag risk, speed review (errors ~80%). Tech: AI model + document intelligence. Src: JPMorgan COIN	30–60%	3–6	6–12
12	Recruiting / screening	Lever: auto-screen, match, and schedule, cutting time-to-hire 50–90%. Tech: machine-learning matching + AI screening + async video. Src: Unilever	40–70%	3–6	6–12
13	Procurement / spend analysis	Lever: classify spend, surface savings (5–15% of addressed spend), auto-source. Tech: machine-learning classification + generative analytics + agents. Src: McKinsey	5–15% (of spend)	3–6	6–12
14	Customer churn / retention	Lever: flag at-risk accounts and target retention offers (revenue protected, not a cost cut). Tech: machine-learning classifier + CRM triggers. Src: peer-reviewed	churn –10–25%	3–6	6–12

- **Prerequisites: pick for volume, a standard process, clean data, and a measured baseline.**

The winners share four traits: 1. high transaction volume, because return is per-transaction times count. 2. a stable, documented process, since AI automates a defined path. 3. accessible clean data, above all a knowledge base for retrieval, labeled images for vision, clean history for machine learning. 4. A known cost-per-unit today to prove savings against.

- **During the build: the failure mode is adoption and accuracy, not the AI model.** Ground every generated answer in the firm's own data, and keep a person in the loop for anything customer- or money-facing. Bought licenses are not adoption. Gains concentrate in less-

experienced staff and appear only with training and workflow redesign ([NBER](#)). Model the fully loaded cost, guard against vendor lock-in, and budget to the low end of each range.

- **In operation: run it as a product and bank the freed time.** Track accuracy, deflection, and satisfaction against the baseline every week, with a fixed test set to catch model drift, and feed human corrections back so quality climbs. Savings bank only if freed hours are redeployed or headcount plans change. Prove one process, then reuse the platform and guardrails to make each next use case cheaper.

Breakout B1-5: AI's eventual cost reduction per business function increased significantly (~18% overall) between 2023 and now.

For each business function, AI cost-reduction potential is researched three ways: achievable today (2025–26 reports), the ultimate floor with every foreseeable lever including humanoid robots and near-AGI agents (2025–26 reports), and that same floor as estimated in 2023 (2023 reports only). Sources are whole-function, so a sub-task's headline does not stand in for the function, and both eventual columns are floors, so 2025/26 sits at or above 2023 everywhere. Each triplet gives a consensus range between different sources and its midpoint value. The full evidence, ~180 sourced findings, is in the [detailed research report](#).

Function	Today range	Today value	Eventual ('23 docs) range	Eventual ('23 docs) value	Eventual ('25/26 docs) range	Eventual ('25/26 docs) value	Increase '23 → '25/26
Production / Manufacturing	18–28%	23%	30–44%	37%	50–60%	55%	+18%
Logistics, warehousing & distribution	5–20%	13%	8–16%	12%	25–42%	34%	+22%
Field service, installation & maintenance	12–30%	21%	26–34%	30%	32–48%	40%	+10%
Accounting operations	12–20%	16%	25–40%	33%	35–50%	43%	+10%
Purchasing operations	19–40%	30%	25–40%	33%	35–40%	38%	+5%
HR & workforce administration	19–30%	25%	30–42%	36%	45–55%	50%	+14%
Order processing & back-office operations	15–20%	18%	25–35%	30%	52–70%	61%	+31%
R&D, engineering & product design	15–40%	28%	20–37%	29%	18–48%	33%	+4%
IT & software development	10–22%	16%	15–35%	25%	45–50%	48%	+23%
Marketing & brand	5–20%	13%	8–22%	15%	20–30%	25%	+10%
Legal, compliance & corporate strategy	12–22%	17%	20–40%	30%	52–66%	59%	+29%
Sales & business development	15–20%	18%	22–36%	29%	48–62%	55%	+26%
Customer service & success	20–35%	28%	30–45%	38%	60–65%	63%	+25%
Executive leadership & general management	15–30%	23%	25–40%	33%	30–50%	40%	+7%

- **The value is the midpoint of the researched range, and the range spans the sources.** “Today” counts tools deployable now, at a typical firm, not a top-decile leader. “Eventual” assumes every foreseeable lever with full agentic redesign, and no time limit.
- **The biggest jumps since 2023 land where agentic redesign unlocked whole-function work:** order/back-office (+31%), legal (+29%), sales (+26%), customer service (+25%), IT & software (+23%). In 2023 the tools assisted. By 2025–26 agents run the end-to-end workflow, so the whole-function floor climbs, not just a sub-task's.
- **The smallest jumps are where 2023 already understood the floor:** R&D (+4%), purchasing (+5%). The estimate barely moved, because the levers were already foreseen.
- **These are cost cuts per unit of functionality, not guaranteed budget cuts.** Each figure holds the work fixed, i.e. the cost to deliver the same output. Cheaper output invites more of it,

so a function's total spend can fall less than its per-unit cost, and even rise as the work expands ([NPR](#)).

Breakout B1-6: Cost reduction potential is ~21% today and ~48% eventually.

Multiply each function's cost weight by its consensus cost reduction (breakout B1-5), and sum over the 14 functions. The result is the weighted AI cost-reduction potential for each archetype, and for the world, under each framing.

Framing	I. Physical/ production	II. Transactional	III. Expert/ knowledge	IV. Human service	World avg
Achievable today (2025/26)	21	19	20	23	21
Eventual floor, 2023 view	31	30	29	32	30
Eventual floor, 2025/26 view	47	51	46	50	48
Increase '23 → '25/26	+16	+21	+17	+18	+18

- **The industry members for each archetype:** physical/production covers agriculture, manufacturing, construction, utilities, and transport & logistics. Transactional covers banking, insurance, and capital markets. Expert/knowledge covers software & IT and professional & business services. Human-service covers retail & wholesale, healthcare, education, hospitality & food, and public administration.
- **Today's cost reduction potential is ~21%, and near-uniform across archetypes.** Every archetype already has ~21% of cost-reduction potential, because each mix contains some highly-exposed functions. The potential barely changes by industry archetype.
- **The ultimate cost reduction floor is ~48% worldwide.** Transactional industries sit highest, where routine processing automates most deeply. Expert/knowledge sits lowest, held down by its R&D pool and a durable senior-judgment core.
- **The eventual estimate rose from ~30% to ~48% between 2023 and 2025/26 reports,** up 18 points worldwide. Most of the potential is still ahead: today's ~21% is under half of the ~48% floor, reached over the long run as agent reliability and humanoid dexterity mature.

What This Means:

- **Quick win projects excel on two different metrics.** Back-office automation delivers the deepest cost cut. Text-based quick wins in marketing, HR, and sales pay back fastest. Rank the portfolio on both return and speed, not on visibility.
- **HR is becoming the default owner of AI-agent governance,** before law or practice recognizes the role. Autonomous agents need identity, permissions, audit, compliance and shut-off, the expertise HR already runs for employees.
- **Most of the cost reduction is still ahead, and it arrives at different timelines.** AI is being adopted faster than the PC or the internet were ([Bick, Blandin & Deming](#)), yet some realized savings lag adoption because workflow redesign takes years. The bulk of the ~48% floor is unrealized, and about a third of it is robotics-gated work, e.g. production, logistics, and field service, which diffuses at the slow pace electrification took. Half the floor will be realized around 2035 and three-quarters around 2041 (modeled).

Recommendations:

- **Sequence the rebuild by business case, not visibility: back-office first for ROI, customer-facing next.** Back-office returns land in 6–12 months on clean data and metrics. Front-office matters for brand but pays back slower.
- **Require explicit comprehension checks before AI-written code is used.** Controlled studies show faster surface delivery but 17% lower comprehension and up to 19% longer end-to-end. Have engineers explain AI-generated code before merge, or the learning gap accumulates.
- **Assign AI-agent governance to HR explicitly in 2026.** Agents that run autonomously need an owner for identity, permissions, audit, compliance, and termination: the functions HR already performs for employees. Build it now, ahead of the 2027–28 frameworks.
- **Rebuild the entry rung as an AI-supervisory apprenticeship, rather than only cutting it.** Automating the tasks juniors once learned on removes how the next seniors are trained. Redesign the junior role around directing and checking agents, so juniors still build judgment while doing higher-value work, and you keep the training ground rivals are dismantling.

[↑ Back to Table of Contents](#)

B2. Lead Your Sector or Lose It

Summary: Lead your sector by reinventing the product, not just cutting its cost. The deepest AI gains create products and services that did not exist without AI and reshape the industry around them. Robotaxi ride numbers are roughly tripling a year (Waymo, Baidu Apollo Go), and the first drug with an AI-discovered target and AI-designed molecule has reached the clinic. AI now closes the design-make-test-learn loop, so it compounds innovation itself, and a first mover out-innovates rivals faster than they can copy. These are the largest growth openings in each industry, and among AI's biggest public gains. Two moves capture them: sequence admin → operations → R&D, and build deep, narrow AI for your industry, not generic copilots. The choice is to lead, or be led.

Evidence:

- The canonical anchor is healthcare: in an NVIDIA survey, 70% of organizations now actively use AI (up from 63% in 2024), with 85% of executives self-reporting revenue gains and 80% cost reductions ([NVIDIA](#)). Breakout B2-1 traces its adoption sequence across every regulated sector.
- New AI-native products scale fast. Waymo's paid robotaxi trips more than tripled to over 14 million in 2025 ([Waymo](#)) and Baidu's Apollo Go rides rose over 200% year on year in late 2025 ([Baidu](#)), while software, IT, and professional services face an analyst-estimated 25–40% margin squeeze over 36 months as pricing shifts from time-and-materials to outcomes.
- The consequence, once a sector reinvents: headcount-priced work reprices. India's IT-services index fell ~21% in February 2026, erasing \$70B+, amid AI-disruption concerns, as AI at \$20–100/month began substituting for work costing \$10–25K per head ([Moneycontrol](#)). Salesforce

says it replaced ~4,000 support roles with no drop in customer-satisfaction scores ([IT Pro](#)); the tools absorbing the work are 10–100× cheaper than the labor they replace, but require skilled use.

Breakout B2-1: AI disruption mapped across the whole economy — each industry’s two biggest AI markets by 2030, and some products already shipping in each.

Twelve sectors cover the world economy, listed alphabetically with each one’s approximate share of global GDP. For each, the two AI products or services with the largest estimated 2030 market lead the row, then a named product already shipping and the AI techniques underneath. Sizes read as current global market, then compound annual growth, then the 2030 estimate. All are analyst-modeled and vary by scope. Two 2030 sizes are extrapolated from longer-dated forecasts: self-driving software and agentic commerce (a 2033 endpoint).

Industry (~% of world GDP)	Top 2 AI products/ services (current · CAGR → 2030)	Example products/ services	Enabling AI
Agriculture & food (~4%): farming, livestock, food processing, forestry	Autonomous farm robots — \$17.7B (2025) · 26% → \$56.3B (MarketsandMarkets). AI food inspection & smart factory — \$10.8B (2025) · 29.6% → \$50.6B (Research and Markets).	John Deere See & Spray · Monarch Tractor MK-V · DJI Agras	Computer vision · Edge AI · Multimodal predictive machine learning
Automotive (~3%): car making, EVs, self-driving	Robotaxis (autonomous ride-hailing) — \$0.4B (2023) · ~92% → \$45.7B (MarketsandMarkets). Self-driving & driver-assist software — \$10B (2024) · 21.2% → ~\$32B (Global Market Insights).	Waymo One · Mobileye · Baidu Apollo Go	Vision + sensor fusion · End-to-end driving models · Reinforcement learning
Construction & real estate (~13%): building, infrastructure, architecture, property	Generative design & AI project management — \$2.9B (2023) · 26.9% → \$17.0B (Grand View). Autonomous construction equipment — \$4.4B (2024) · 14.2% → \$9.8B (MarketsandMarkets).	Autodesk Forma · Bentley iTwin · Built Robotics	Generative design · Computer vision · Digital twins
Energy, utilities & mining (~7%): power, oil & gas, grids, renewables, mining	AI wind/solar forecasting & optimization — \$20.6B (2025) · 25.7% → ~\$64B (Precedence). AI smart-grid management — \$8.9B (2024) · 36.9% → \$58.7B (MarketsandMarkets).	GE Vernova GridOS · Schneider EcoStruxure · Cat MineStar Command	Time-series forecasting · Reinforcement learning · Vision + autonomous stacks
Financial services & insurance (~8%): banking, payments, markets, insurance	Robo-advisors (automated wealth management) — \$14.3B (2025) · 29.6% → \$67.8B (2031) (Mordor). AI fraud detection — \$6.8B (2025) · 24.6% → \$49.5B (2034) (TrendX).	FICO Falcon · Betterment · Upstart	AI-model agents · Gradient-boosted trees + anomaly detection · Graph neural networks
Government, education & public sector (~12%): public admin, defense, schools	Citizen-service chatbots & case automation — \$8.3B (2024) · 41.4% → \$65.8B (Grand View). AI tutoring — ~\$7–8B (2025) · 31–43% → \$32–41B (Grand View).	Khan Academy Khanmigo · Palantir AIP · Anduril Lattice	AI tutors & agents · Computer vision · Autonomous systems + predictive analytics
Healthcare & life sciences (~9%): hospitals, pharma, diagnostics, devices	AI-designed medicines (drug discovery & design) — ~\$2–3B (2025) · ~30% → ~\$9.2B (Grand View). AI diagnostic imaging — \$0.76B (2025) · 24.5% → \$2.3B (MarketsandMarkets).	Insilico — rentosertib · Isomorphic Labs · Aidoc	Protein-structure prediction + molecule design · Image segmentation · Multimodal foundation models
Information technology, media & telecom (~7%): software, cloud, security, content, networks	Generative-AI apps & coding assistants — \$37.1B (2024) · ~35% → ~\$220B (ABI). AI cybersecurity — \$25.4B (2024) · 24.4% → \$93.8B (Grand View).	GitHub Copilot · Claude Code · Microsoft Security Copilot	Large language models · Agentic orchestration · Diffusion image/video models
Manufacturing & industrials (~14%): factories, machinery, electronics, chemicals	Factory digital twins — \$21.1B (2025) · 47.9% → \$149.8B (MarketsandMarkets). AI vision inspection & industrial robots — \$33.9B (2024) · 9.9% → \$60.6B (Grand View).	NVIDIA Omniverse · Siemens Xcelerator · Cognex In-Sight 2800	Vision defect detection · Robot foundation models · Physics-based digital twins
Professional & business services (~7%): legal, accounting, consulting, marketing	Generative marketing & content tools — \$21.5B (2025) · 29% → \$77.2B (Research and Markets). AI legal, tax & customer-service agents — \$12.1B (2024) · 25.8% → \$47.8B (Grand View).	Harvey · Thomson Reuters CoCounsel · Intuit Assist	Large language models · Retrieval-augmented generation · Agentic orchestration
Retail & consumer (~11%): e-commerce, stores, consumer goods, advertising	Recommendation & personalization engines — \$9.2B (2025) · 33.1% → \$38.2B (Mordor). AI shopping assistants (agentic commerce) — \$5.7B (2025) · 35.7% → \$65.5B (2033) (Grand View).	Amazon Alexa for Shopping · Salesforce Agentforce · Shopify Sidekick	Recommender systems · AI shopping agents · Computer vision
Transport & logistics (~5%): freight, shipping, warehousing, delivery	AI supply-chain & route optimization — \$7.3B (2024) · 42.7% → \$63.8B (Strategic MR).	Symbotic · project44 · Amazon Robotics	Route & network optimization · Vision + sensor fusion · Demand forecasting

Industry (~% of world GDP)	Top 2 AI products/ services (current · CAGR → 2030)	Example products/ services	Enabling AI
	Warehouse robots & automation — \$19.2B (2023) · 18.7% → \$59.5B (Grand View).		

- **One shared engine for most levers.** Across most rows the same few methods recur: language models and agents, computer vision, forecasting and optimization. The wrapper differs (a tractor, a contract, a scan). The engine does not. AI is one horizontal technology sold as many vertical products.
- **Some levers are industry-specific and do not transfer.** An AI-designed medicine rests on protein-structure prediction and generative chemistry over proprietary biology. A robotaxi rests on models of real-world traffic, trained on driving data a general model never sees. Neither is bought off a shelf, which is where deep, narrow players win.
- **The 2030 prize scales with how digital the product is.** The largest, fastest markets are pure software: generative-AI apps (~\$220B), factory digital twins (\$150B), marketing content (\$77B), robo-advisors (\$68B), all compounding growth at 30–48% a year. Physical, safety-gated products are smaller and grow 10–26% a year.
- **Within a sector, adoption sequences admin, then operations, then R&D.** Healthcare set the template: back-office paperwork first (payback in 6–12 months), clinical operations next, drug R&D last. Regulated sectors from finance to legal bind on that same sequence.

What This Means:

- **Reinventing the product is the deepest competitive move, and a CEO question.** The biggest gains rebuild the product and the discovery process around AI, closing the design-make-test-learn loop. A first mover out-innovates and underprices at once, which compounds faster than rivals can copy. That makes it a product strategy question, not an efficiency project for operations.
- **When a cheap tool absorbs headcount-priced work, the value moves to the AI supplier.** A \$20–100/month tool replacing work that cost \$10–25K per head reprices the whole sector (India IT). This is a value-chain shift, not just a cost cut: revenue migrates to whoever supplies the AI, the white-collar version of humanoids displacing physical labor.
- **These products are lead-or-lose and a public good at once.** For an industry sector, this decides who leads product innovation, a strategic stance. For society they are among AI's largest gains, e.g. earlier cancer detection, AI-designed drugs, safer road traffic. Competitive necessity and public benefit point the same way.

Recommendations:

- **Sequence AI adoption as admin first, then operations, then R&D.** Map each major function to the phase it resembles and plan investment and reskilling on that timeline.

- **Professional-services firms: shift client contracts from time-and-materials to outcome-based pricing at the next renewal**, before agent autonomy lands in probably late 2027 and you negotiate from weakness against 25–40% margin compression.
- **Build deep, narrow AI capability for your industry, not generic horizontal tools.** The billion-dollar vertical-AI firms win by going deep on one sector's workflows, data, and regulation. Hire or partner for that depth.
- **If your sector touches science, materials, or climate, invest in AI-for-science now on a multi-year horizon.** Breakthroughs (Chai-2, Iambic's 24-month path to a clinical trial) pay off in the 2030s. Treat it as a lead-or-lose frontier even when returns lag.

[↑ Back to Table of Contents](#)

B3. Close the Implementation Gap

Summary: Most AI value is lost between pilot and P&L, not in the AI technology. 88% of enterprises run AI somewhere, yet only 1% call their strategy mature and 95% of pilots show no profit within six months. The failures are organizational: the wrong project, no baseline, unready data, a tool that never improves, licenses mistaken for adoption, and over-building what could be bought. Each trap maps to a missing capability. The return is won in measurement and evaluation, which most deployments skip. So gate every pilot on the six traps before funding it, and concentrate on two or three high-return functions first. Rebuild each workflow end-to-end on one shared company AI platform. Govern employee AI rather than ban it: 90%+ use personal AI tools whatever the policy. Readiness is a durable advantage. Mid-market firms can build it faster than sprawling incumbents.

Evidence:

- Adoption is wide but shallow. 88% of enterprises now run AI in at least one function, up from 78% a year earlier, yet only 1% rate their strategy mature ([McKinsey](#)). Independent government data is starker: only ~1 in 5 US firms use AI at all ([Census BTOS](#)).
- The gap is widening, not closing. ~88% of proofs-of-concept never reach production, and the share of firms abandoning *most* of their AI initiatives jumped from ~17% to ~42% in 2025 ([S&P Global](#)).
- Firms split near evenly on build-versus-buy — ~47% build, ~53% buy ([Menlo Ventures](#)), even as success rates diverge sharply between the two (breakout B3-1).
- Shadow AI is pervasive and invisible. 90%+ of workers use personal AI tools regardless of policy, and ~74% of workplace AI-chatbot use runs through non-corporate accounts the firm cannot see ([Cyberhaven](#)).

Breakout B3-1: Pilots fail on organization, not AI technology — six avoidable traps between a pilot project and booked P&L.

The quick-win copilots in breakout B1-1 pay back in months on paper, yet most stall in practice. Six failure modes recur between a green-lit pilot and booked profit. Each sits at one stage of the project: pick, prove, ready the data, keep improving, adopt, and source. The top example in each row is the most-cited version of that failure.

Common pitfall	Description — what commonly fails (top failure mode)	Safeguard — what to do / ensure
Visibility over ROI	Budgets chase demo-friendly front-office pilots, not the back office where the return sits. Top mode: over half of generative-AI spend goes to sales and marketing, yet the largest, fastest returns are in back-office automation (MIT/Fortune).	Rank the portfolio by payback, not by visibility. Fund the back-office automations first (breakout B1-1).
No baseline, no P&L	Success is never defined, so the pilot cannot prove value and dies at review. Top mode: 95% of pilots show no measurable P&L within 6 months (MIT/Fortune).	Set one KPI and a measured cost-per-unit baseline before building. Track weekly against it.
Data not ready	Scattered, dirty, or ungoverned data starves the model. Top mode: poor data quality is a leading reason Gartner projected ~30% of generative-AI projects abandoned after proof-of-concept by end-2025 (Gartner).	Pick processes with accessible, clean data and permissions already in place. Stand up retrieval on text you own first.
Static tool (learning gap)	The tool never retains feedback or adapts, so quality plateaus and users drift away. Top mode: this “learning gap” is the main barrier separating the ~5% that scale from the rest (MIT/Fortune).	Run it as a product. Keep a feedback loop, a fixed test set to catch drift, and corrections fed back so quality climbs.
Licenses, not adoption	The tool is bought and switched on, but the workflow is never redesigned, so usage and gains never appear. Top mode: the value gap traces to workflow redesign and readiness, not the tool (BCG).	Redesign the workflow end-to-end and train the staff. Gains concentrate in less-experienced staff, and only with training (NBER).
Build-it-yourself bias	Teams over-build custom systems that stall before production instead of buying proven tools. Top mode: purchased tools reach success ~67% of the time, versus ~33% for internal builds (MIT/Fortune).	Partner for speed on commodity capability. Build in-house only what differentiates: governance, integration, edge workflows.

- **The common pitfall is organizational, not technical.** Five of the six traps sit in choosing, measuring, and adopting, not in model quality. The firms that scale and the firms that stall often run the same models. Readiness is what separates them.
- **Every safeguard is set before the build, not bolted on after.** A ranked portfolio, a measured baseline, clean data, and a redesigned workflow are all decided up front. A pilot launched without them cannot be rescued at review.
- **The traps compound, so treat each of them as a gate.** A wrong pick, on bad data, with no baseline fails three times over. Clear all six before green-lighting, then reuse the cleared guardrails to make each next use case cheaper.

Breakout B3-2: The quick wins and the six traps reduce to one capability spec — fifteen requirements a company AI platform must have

Breakout B1-1’s quick wins only pay if the platform behind them can do the work, and four of breakout B3-1’s six traps are each a missing platform capability. They resolve to the same fifteen requirements, grouped by who holds control: the platform integrates the firm’s data, operates the work, then hands the human the skill to use it and the sign-off on risk. The From column cites the source, a B1-1 quick win by row number, its build or operate guidance, or a B3-1 trap by number.

Band	Id	Requirement	From
Integrate	I1	Adopt: offer an interface intuitive enough that staff use it over shadow tools	B3-1.5
	I2	Connect: link file storage, email, ERP, CRM, HR and service-desk systems	B1-1.5,.6; B3-1.3
	I3	Govern: enforce roles, rights and permissions on every data source and action	B3-1.3
	I4	Reference: ground answers in the firm's own specific data; retrieval meets the quick wins, a knowledge graph goes beyond	B1-1.2,.5; B3-1.3
	I5	Choose: switch freely between model providers, avoiding lock-in	B1-1 build
	I6	Reuse: serve every use case from one shared platform and guardrails	B1-1 op
Operate	O1	Author: draft, revise and summarize documents	B1-1.1,.11
	O2	Perceive: work in images and speech, generating visuals and transcribing audio	B1-1.1,.8
	O3	Extract: read scanned inputs with document processing	B1-1.9,.11
	O4	Automate: run agents that act across the connected systems	B1-1.4,.9
	O5	Predict: match, classify and score risk with machine learning	B1-1.10,.12,.14
	O6	Measure: track a KPI, a cost-per-unit baseline and weekly accuracy against target	B1-1 op; B3-1.2
	O7	Evaluate: hold quality with a feedback loop, fixed test set and drift monitoring	B1-1 op; B3-1.4
Collaborate	C1	Empower: train users to experiment safely and learn, building AI fluency	B3-1.5; B1-1 build
	C2	Verify: gate customer- and money-facing output with a human check	B1-1 build

- **These are requirements, not a wish list.** Each row traces to a source: a B1-1 quick win that needs it, or a B3-1 trap that punishes its absence. The set is the floor to make the quick wins pay and to clear the four traps a platform can close.
- **The bands hand control down one level at a time.** Integrate gives the platform the data, Operate gives it the work, Collaborate keeps the person on the skill to wield it and the sign-off on risk.
- **The value is won in Operate.** Integrate is the enabler. Within Operate, Measure and Evaluate are the no-baseline and learning-gap traps, the two that separate the ~5% pilots that scale and deliver profit. Project ranking and buy-versus-build stay human-owned without necessary support of the AI platform.

What This Means:

- **Shadow AI is the tell-tale signal.** Employees get value from personal AI tools while the official pilot stalls, so the gap sits in the workflow around the model, not the model itself.
- **Readiness is a capability gap, and the return is won in the “Operate” band most skip.** Each of the six traps maps to a missing capability (breakout B3-2). Integration is the enabler. Measurement and evaluation, where a pilot becomes P&L, separate the ~5% pilots that are successful.
- **Readiness is a durable, compounding advantage, and mid-market firms can build it.** Discipline, not capital or scale, decides the outcome, so a focused mid-market firm can gate and redesign faster than a sprawling incumbent. The laggards do not automatically catch up, which cuts against “AI only rewards incumbents.”

Recommendations:

- **Gate every pilot on the six traps before you fund it.** Treat breakout B3-1 as a go/no-go checklist: a payback-ranked pick, a measured baseline, ready data, a feedback loop, a redesigned workflow, and buy-versus-build judged on what differentiates. Cut a pilot that fails the gate before spend, not at the six-month review.

- **To start, concentrate on two or three high-return functions and rebuild them end-to-end.** Target where AI transforms unit economics — finance back-office, support, software development (breakout B1-1).
- **Stand up one company AI platform covering the B3-2 requirements, so proven wins scale and reach ROI quickly.** Cover the operate band, not just integrate. Measurement and evaluation are where a pilot becomes P&L. Reuse the platform across use cases, and build only what differentiates.
- **Company-license and govern employee AI rather than ban it.** Most workplace AI use is invisible to the firm, so a ban only drives it further out of sight. Allowed tools with clear data rules and audit logging are the only measure that actually cuts shadow-AI risk.

[↑ Back to Table of Contents](#)

B4. Compete in the Agent-Mediated Economy

Summary: The platform that owns the agent owns the customer, because it holds the customer's data and context. This holds in consumer and business-to-business selling alike: buyers deploy agents to research and negotiate, sellers deploy agents to qualify and close. AI platforms are expected to mediate \$20.9B of US retail in 2026, nearly 4× the prior year, as traditional search falls ~25%. Inside each platform, often only its own agent transacts end-to-end, and the fight is already fierce: in March 2026 Amazon won an injunction blocking Perplexity's shopping agent from accessing Amazon. Compete in the agent layer or risk losing the customer entirely with these three moves: optimize your content so agents understand your products better than rivals', keep product feeds compatible with several agent protocols, and build a direct channel you own so customer memory stays yours. It is better to own the relationship than rent it.

Evidence:

- AI platforms are expected to mediate \$20.9B in US retail spending in 2026, nearly 4× 2025 ([eMarketer](#)).
- AI-referred traffic rose +805% year on year on Black Friday 2025, and AI-sourced shoppers were 38% more likely to buy ([Adobe Analytics](#)). Traditional search volume is projected to fall 25% by 2026 as share moves to AI chatbots and answer engines ([Gartner](#)).
- The economics are consolidating. OpenAI charges merchants a 4% transaction fee on Instant Checkout and retired single-merchant checkout in March 2026 ([CNBC](#)), and Amazon won a March 2026 injunction blocking Perplexity's shopping agent while walling off ~47 bots and pushing its own Rufus ([CNBC](#)): managed-marketplace business is arriving.
- The shift spans business-to-business, not only consumer retail: buying agents research, shortlist, and negotiate, while selling agents qualify and close. AI sales-development agents already automate research, enrichment, and outreach ([Salesforce](#), B1), and agentic procurement cuts ~5–8% of addressable spend ([Bain](#)). Agent memory is likely to become the

dominant defensibility through 2026–27, and antitrust scrutiny of the layer arrives on the same timeline.

Breakout B4-2: Keep a direct line to the customer when agents mediate the sale — the methods that hold contact, sales, and data, tagged by system.

Agent-mediated selling can cut a seller off from its own customer. The buyer's agent negotiates the price down, and the platform's agent keeps the sign-up, the payment, and the memory. Two systems split the agent economy, and the difference decides your defense as a seller. In the **US** model, an agent buys from each seller directly, calling your own checkout, so the market stays open and contestable, and you keep the checkout, the feed, and your answer-engine presence. In **China**, only a super-app's own agent transacts end-to-end inside its closed platform (login, offer data, and payment all closed), faster to assemble but captive, so the move is to win the app agent's recommendation and then convert that reach into a private channel. The *Agent type* column tags where each method works: *both* systems, *US-only*, or *CN-only*. This breakout is about reach; breakout B5-2 ranks which of these positions also holds a margin.

Method	Agent type	Description - what it is and why it works	Prerequisite	Creation — build & maintain
Brand and trust (specify-by-name)	both	A name buyers ask the agent for <i>by name</i> , skipping the shortlist. As products commoditize and content goes synthetic, trust turns scarce and priceable.	A repeated purchase where quality or authenticity is hard to judge up front.	Consistent quality, a signature experience, content provenance, human relationships rivals can't fake.
Product exclusivity or differentiation	both	If the item is unique or exclusive, the agent can't substitute a cheaper rival — the test that decides whether a moat survives.	A genuinely differentiated or exclusive product, bundle, IP, or design.	Own the differentiation. Withhold it from marketplaces that would commoditize it.
First-party data flywheel	both	Unique operational or customer data makes your personalization beat a rival on the same base model — the durable software-era moat (B5-2 #2).	A data exhaust from your own operations others can't reproduce.	Instrument every touchpoint. Keep feeds exclusive, private, and refreshed. Feed it back into your agent and offers.
Loyalty, membership, or subscription	both	A recurring, logged-in relationship gives a standing reason to deal direct and a stream of first-party behavior.	A repeat-purchase category and a benefit worth returning for.	Tiered rewards, member pricing, subscription refills. Run it as a product and refresh the benefit.
Direct-to-customer storefront	both	Sell direct so buyer, data, and margin route through you. US: a store you host. China: a WeChat Mini-Program store (\$500B+ channel, 4.3× year on year), more owned than a marketplace listing.	Fulfillment and a reason to come direct: price, range, or exclusivity.	Stand up owned commerce, drive traffic to it, keep checkout and sign-up yours.
Physical, local, or experiential delivery	both	Bind the sale to a place, person, or experience an agent can't ship in from a cheaper rival (B5-2 #4) — the clearest shield.	Delivery tied to a location, a person, or an in-person moment.	Anchor in stores, showrooms, service, events. Use AI to cut back-office cost while keeping the physical front.
High-touch human relationship selling (B2B)	both	For complex, high-value deals, a human account relationship the buyer's procurement agent can't replace. It intermediates only the transactional layer.	Considered, high-ticket, or trust-dependent purchases.	Named account managers, co-design, service commitments. Let agents quote, keep humans on the close.
Answer-engine optimization	US	Structure content and data so agents surface and correctly understand you. Discovery has moved from search to agents (search -25% by 2026); content an agent can't parse is invisible.	An indexable web presence and clean, structured product attributes.	Publish machine-readable specs, schema, provenance, and Q&A. Track presence in AI answers as a core metric.
Multi-protocol product feeds	US	Keep feeds compatible with the emerging agent-commerce standards (merchant guide) so any agent can transact you without your picking a platform first.	A structured catalog and a checkout interface.	A feed layer mapping products to each standard. Adopt the leading two to three and re-sync as standards shift (no consolidation before 2027).
Own negotiating selling agent (agent-to-agent)	US	Your agent qualifies, answers, and negotiates back against the buyer's agent, competing on structured value, not list price.	First-party pricing logic and a place to host the agent.	Build it on your data. Arm it with verifiable value signals: bundles, terms, stock. Join the agent-to-agent economy.
Owned checkout and payment	US	Keep the transaction on a checkout you operate, so the payment record, receipt, and buyer identity stay yours. In the US the agent calls your checkout directly.	A direct checkout you run.	Own the payment relationship, resist managed-marketplace defaults, keep a direct fallback.
In-app flagship and catalog optimization	CN	The super-app's agent recommends from its own catalog, not the open web. Dominate inside one app: official flagship, in-app recommendation optimization, in-app membership.	A committed presence in one super-app (Tmall, JD, Douyin, WeChat).	Run an official store, optimize listings and reviews for the app's agent, treat the app as rented reach, not owned.
Private-domain community (私域)	CN	A direct channel you own outright inside WeChat/WeCom groups and Mini-Programs, a relationship no platform algorithm can remove. 5–10× lower cost-per-conversion than platform ads.	A base of customers you can move off the public feed, plus staff to nurture them.	Move buyers into WeChat/WeCom, one-to-one service, Mini-Program membership. Sustain with content and human contact.
Creator and livestream social commerce	CN	Trusted creators and merchant-led streams sell direct with in-stream checkout, holding demand inside a relationship, not a price-compare agent (~\$1.1T live commerce in 2026).	Access to relevant creators or in-house on-camera talent.	Douyin and WeChat Channels streams, a shift toward merchant-led streams you control, private-domain followers driven into each session.

- **One test runs down the whole list: can an agent import a cheaper substitute for what you sell?** Where it can't — trusted, exclusive, data-backed, physical, or relationship-bound transactions — the moat holds and often strengthens. The *both* rows are the endgame. Build

them first, because competition erodes everything else toward ~0 additional ai-based economic profit within ~20 years (B5).

- **In the US, keep the checkout, the product feed, and your answer-engine presence yours.** The window is closing as OpenAI's 4% Instant Checkout pulls economics toward a managed marketplace. Stay compatible with several agent protocols and keep your data portable.
- **In China, do not fight the closed login and payment.** Win the super-app agent's recommendation with an official flagship, then convert that reach into a private channel (🔒), the direct line no platform algorithm can take back.

What This Means:

- **The platform that owns the agent owns the customer.** Whether the platform shuts neutral agents out (China's super-apps) or pulls economics toward a managed marketplace (the West), the merchant competes in the agent layer or goes invisible. Breakout B4-2 catalogs the methods that keep direct contact, sales, and data.
- **The merchant-customer relationship is partly broken by a new middle-man.** The merchant stays seller of record but loses the email sign-up, the cross-selling, and control of the brand experience.
- **Platform power is consolidating even under "open" protocols.** OpenAI retiring single-merchant checkout is the tell tale signal: a merchant's access can narrow to a managed marketplace on the platform's terms.
- **The dynamic is not retail-only.** In business-to-business, whoever owns the buyer's procurement agent and its memory intermediates the account, so agent-legible content and an owned channel matter for business sellers too.

Recommendations:

- **Invest in answer-engine optimization in 2026.** Search volume is projected to fall 25% and agents are the new discovery layer. Treat it as the successor to search-engine optimization, and measure presence in AI answers as a core KPI.
- **Make product feeds compatible with several leading agent-commerce standards.** Consolidation is not expected before 2027, so you need to be present on several agents in parallel.
- **Build a direct customer channel you own (email, loyalty, branded app)** to keep customer memory outside the agent platforms that otherwise own the relationship layer driving repeat business.

[↑ Back to Table of Contents](#)

B5. Quantify and Drive the Bottom-Line Impact

Summary: The prize is huge. Cost reductions of ~21% are realistic now, rising to a ~48% floor once every lever is used and climbing further as AI advances (breakout B1-6). Most of it is not yet realized. Speed is where the money is. As the cost reduction matures, a 3-month earlier realization can yield up to a full year of additional EBIT over the following years. This mandates extra investment now for extra profit. Long-term, rivals and agents compress prices until most of the cost reduction converts into lower prices for customers, and by ~20 years the average firm earns near-zero additional economic profit. A durable margin sits only with a moat: brand and trust, proprietary data, customer memory, network effects, and patents. Spend the lead from AI-driven cost reduction on building it. Speed wins the transition. The moat wins the end.

Evidence: The financial case is one subtraction, in five steps:

- *Start from the cost reduction.* Today's tools, applied consistently across the cost base, are worth **~21% of operating cost** (breakout B1-6).
- *Add a margin edge from speed, for a while.* Faster cycles (time-to-market ~40%, breakout B5-1) let early movers launch more and sooner, a real edge until rivals catch up.
- *Subtract the cost of the AI.* Compute is small and falling, but integration, data, governance, and reskilling are large and front-loaded. That is the execution gap (B3), priced: despite near-universal adoption, only **~6% of firms** capture ≥5% of EBIT today, and those that do invest at scale and redesign workflows end-to-end (only ~21% have) ([McKinsey](#)).
- *Subtract what competition takes back.* As rivals deploy the same tools and buyers shop through agents, and outcome-based pricing displaces per-seat pricing (the agent-mediated economy, B4), most of the cost reduction reappears as lower prices, not retained margin.
- *Net it for the margin.* Savings pass both ways (suppliers' cost reductions reach you, yours reach your customers), so system-wide value is huge while the slice one firm keeps is thin without a moat.

Breakout B5-1: Over 20 years competition erodes the AI cost reduction to near-zero economic profit — only a moat survives (modeled).

Each row reads as an identity for a representative adopter (**net margin = cost reduction + innovation edge – added AI cost – price erosion**), every figure in *points of revenue* (per-unit margin) so the columns sum across. Cost reduction is what an adopter *realizes*. The achievable reduction runs higher (~21% of operating cost now → a ~48% ultimate floor, realized over time — breakout B1-6), and realization lags it early (the execution gap) before competition becomes what holds net down.

Horizon	Cost reduction (realized)	+ Innovation edge	– Added AI cost	– Price erosion (agent competition)	= Net margin
Now	+3–5 (achievable ≈21% of operating cost — B1)	+0–1 (personalization +40% on covered revenue — McKinsey)	--0.2 (AI ≈4% of tech budgets — Oxford Economics)	-2–3 (agentic retail \$20.9B, 4× year on year — B4)	+1–3 adopters; ~0 economy-wide (execution gap: ~6% at ≥5% EBIT — McKinsey)
~5 yr	+10–14 (toward ~25% of tasks automatable — Goldman)	+1–3 (time-to-market ~40% — McKinsey)	--0.5 (AI ~10% of tech budgets — Oxford Economics)	-8–11 (25–40% margin compression / 36 mo, partial — B2)	+3–6 leaders, ≤0 laggards
~10 yr	+22–26 (≈ full ~25% of tasks — Goldman)	+1–3 avg, moat-holders more (McKinsey)	--0.7 (AI ~15% of tech budgets — Oxford Economics)	-22–26 (25–40% margin compression realized — B2)	+1–3 average, wide dispersion
~20 yr	~+40 realized, approaching the ~48% floor (B1-6); physical-heavy firms lag, knowledge-heavy lead	+1–3 avg, moat-holders more	--1 (AI ~23% of tech budgets by 2035, industry-wide — Oxford Economics)	--40 (near-full pass-through of the cost reduction)	≈0 economic profit; moat-holders positive

Industry-wide, added AI cost converts from AI's share of enterprise tech budgets (~4% → ~23% by 2035, Oxford) at tech ≈4–5% of revenue, i.e. ~0.2% → ~1% of revenue. Figures are per-unit. Volume effects are in What This Means.

Breakout B5-2: Not every moat survives AI-driven competition — spend the lead on the eight an agent can't arbitrage.

B5 ends on the moat. Where breakout B4-2 asks how to reach the customer, this asks which of those positions holds a margin. Over 20 years only a moat holds net margin above zero, but not every moat survives AI. This ranks the eight that matter most for a non-AI business, a manufacturer, retailer, bank, or services firm that deploys AI, by how decisively each holds up in AI-driven competition. In the relevance column, arrows mark AI's effect: ↑ strengthened, ↓ eroded, ↗ split by sub-type.

#	Moat type	Description — what it is and why it earns margin	Prerequisite	Creation — build and defend	Relevance for AI-driven competition
1	Brand and trust	A name customers pay more for or default to, cutting search and perceived risk. Margin from pricing power and cheaper acquisition. Trust turns scarce as content and offers go synthetic.	A repeated purchase where quality or authenticity is hard to judge up front.	Consistent quality, a signature experience, content provenance on official media, human relationships rivals can't fake.	Very high † — trust and status brands strengthen as AI commoditizes products and deepfakes erode shared trust. Discovery and search-visibility brands collapse (Morningstar).
2	Proprietary data flywheel	Unique operational or customer data rivals can't get, making your AI better than a competitor running the same off-the-shelf model. Margin from a compounding capability gap.	A data exhaust from your own operations others can't reproduce, plus the pipes to use it.	Instrument products and processes to capture the data, lock exclusive first-party feeds, refresh it, keep it private and governed.	Very high † — proprietary data is a moat rivals can't buy, and the live signal from daily operations keeps it fresh (Bain). Frontier models and public data commoditize, so only branded or proprietary datasets endure (Morgan Stanley).
3	Customer relationship and agent memory	Owning the channel and the stored context the customer buys through, so the agent platform can't disintermediate you (B4). Margin from repeat business you control, not rent.	A direct channel (email, loyalty, app) and portable customer context.	Build the direct channel, own agent and customer memory, stay compatible with several agent protocols.	Very high † — the platform that owns the agent owns the customer (B4). Embedding AI on your own stack and customer data raises the bar for would-be disruptors (Morgan Stanley).
4	Physical, local, or inelastic position	Work delivered on-site or in person, or demand that won't chase price. An agent can't import a cheaper remote rival, so the cost reduction is not competed away. Margin from un-arbitrageable supply.	Delivery bound to a place or a person, or essential inelastic demand.	Anchor in local presence and hands-on service. Use AI to strip back-office cost while keeping the physical front.	Very high † — B5's key exception. The clearest shield for trades, in-person care, and hospitality.
5	Network effects	Each user makes the product more valuable to others, giving a lead rivals can't match, plus proprietary interaction data. Margin from a self-reinforcing position.	A multi-user or two-sided product.	Seed liquidity, raise density, capture and use interaction data, be the destination not the intermediary.	Very high † (where held) — the most AI-resilient moat, but destination platforms win while pure intermediaries lose to agents (Morningstar).
6	Regulatory license and compliance moat	A permit or accreditation only some clear, capping entry, plus first-mover advantage in AI governance inside regulated sectors. Margin from a legally limited supplier count.	A regulator that gatekeeps entry (finance, health, insurance, defense, utilities).	Win and hold the license, shape the standard early, build evaluation, provenance, and audit into the product.	High † — a non-digital gate AI can't arbitrage. AI-governance readiness becomes its own edge (B6, Lens G).
7	Physical, capital, or infrastructure cost advantage	Structurally lower unit cost from owned hard assets and scale, such as plants, logistics, energy, and distribution. Margin from undercutting or holding price at a fatter margin.	High fixed cost over large volume, or scarce assets and permits.	Own the low-cost physical base, integrate vertically, lock long-term supply, grow volume.	High † — capital and physical scale get heavier as AI raises entry barriers. Labor-arbitrage scale does the opposite, competed away (India IT –21%, B2).
8	Patents and applied-science IP	A legal right to exclude rivals from a product or process for a term. Margin from a state-granted temporary monopoly.	A novel, enforceable invention in a patent-effective domain.	File early and broadly, stack patents around a core, litigate. Keep trade secrets where filing would teach rivals.	High † — very high in applied science, where AI speeds discovery but patents lock rivals out (AI drugs, materials, hardware — B2). Near none for pure software.

• **The test is one question: can an agent import a cheaper substitute for what you sell?**

Where it can't, such as physical and local delivery, licenses, trust, network density, unique data, and agent memory, the moat holds and often strengthens. Where it can, such as commodity remote work, generic data, classic switching costs, and labor scale, it is competed to zero.

- **Two of the top three are new.** Proprietary-data flywheels and customer and agent memory were not moats a decade ago. For a non-AI firm they are the highest-leverage build, because they turn everyday operations into a defensible position without becoming an AI company.

What This Means:

- **The prize is huge, but mostly not the firm's to keep.** Generative AI could add **\$2.6–4.4T/yr** ([McKinsey](#)) and lift global GDP **~1–7% over a decade** ([Acemoglu](#) at the low end, [Goldman](#) near the high). That is *society's* gross gain, not the firm's. Most reaches people as cheaper, better goods, not company profit.
- **Speed is where the money is during the transition.** The eventual ~48% cost reduction lands on the thin operating margin, ~13% in the US and ~12% in the EU across listed firms ([Damodaran US](#), [Europe](#), Jan 2026). At the mature ~48% cost reduction, one year of accelerated savings is worth ~4 years of typical EBIT, so **the compound cost reduction of a 3-month earlier implementation yields about a full year of additional EBIT**.
- **The gains are temporary: without a moat, competition reprices them toward zero.** The ~48% cost advantage is reduced by rivals matching AI and agents compressing prices. By ~20 years the average firm earns ≈0 additional economic profit (breakout B5-1). Higher volume offsets only part of the squeeze, and it flows to the lowest-cost or moated player, not the average firm. So the cost reduction cannot be booked as steady-state margin.
- **Only a moat holds margin at the end.** Durable profit sits with the eight moats ranked in breakout B5-2, strongest in applied domains like AI-discovered drugs and materials. The one exception that needs no such moat is local, physical, and licensed work, where no agent can import a cheaper rival and no robot yet does the hands-on job.

Recommendations:

- **Invest ahead of the curve: extra spending now is justified by the short-term profit it buys.** The cost reduction is front-loaded, so deploying *earlier* pays back well. “Wait and see” forfeits profit and headroom. **Speed wins the transition.**
- **In the EU, request public support and plan as if it arrives.** Ask for capital, compute access, simpler rules, freed industrial data, and procurement contracts at €2T+/yr as anchor demand. The EU is increasingly acting on sovereignty pressure, mobilizing capital, compute, and demand to catch up (G5, W6). Anticipate increasing support rather than wait for it: public backing de-risks the early investment above, and public contracts seed the moat below.
- **Bank the ~21% cost reduction now, but don't book it as permanent margin.** Assume competition reprices most of it within a few years.
- **Budget the real cost (integration, data, governance, reskilling), not the token bill.** It is why most firms miss the EBIT, and the input you actually control.
- **Measure net, not gross.** Track AI-attributable EBIT alongside price erosion in your sector. The gap is your true competitive position.
- **Reinvest savings into what competition can't copy, or can't legally copy.** Brand, proprietary data, distribution, and customer memory hold net margin once everyone runs the same AI technology. Patent the products and processes where imitation would otherwise take months. **The moat wins the end.**

B6. Turn an Understanding of the Other Players Into Business Advantage

Summary: Winning the lead also includes reading the other players - AI companies, individuals, governments and the world system - and acting before foreseeable shifts actually happen: Towards AI companies, distribute your sourcing and maintain the ability to switch. Towards individuals, prepare for them to buy via agents and hire or retrain employees for future senior skills needed in AI-supervisory roles. Governments will increasingly issue AI regulations, providing opportunities for firms that treat AI compliance it as product, not paperwork, to ship faster when the rules bind. For the entire world system, competition within each player and AI progress compress a decade of innovation into the next 10 quarters, the remaining 2026–2028 window.

Lens A — AI companies. AI firms race to build the intelligence stack and concentrate profit, compute, and the best models. AI capability has reached “practical AGI”, and prices fall as open-weight models follow the frontier within months at 5–30× lower cost. Reliability still trails, and compute stays scarce through 2027. For a buyer, that means real choice, a falling price floor, a narrowing supplier field, and a reliability gap to manage. **Derived recommendations:**

- **Distribute your AI sourcing.** Chinese open-weight models trail the US lead by few months at 5–30× lower price ([Artificial Analysis](#)), and the cost to reach a fixed capability falls ~10×/yr ([Stanford HAI](#)). Run commodity workloads on that cheap floor, reserve frontier capacity where it pays, and keep a warm second stack so no vendor can exploit your dependence.
- **Gate autonomous output on quality.** The best agent finishes only 16.1% of real paid freelance projects at client-acceptable quality ([CAIS](#)), and AI errors hide behind perceived competence. Put an evaluation and stake-sized human-approval gate on any agent facing customers or money, before a confident error ships.
- **Build the edge in the agent, not the base model.** Capability comes from the plan-act-observe loop, so owning the best model matters less than the best orchestration around it ([Anthropic](#)). Base-model quality commoditizes as prices fall. Invest in orchestration, tools, evaluation, and the proprietary data wrapped around models.

Lens I — Individuals. Individuals meet AI as workers and as consumers. Workers invest in AI fluency for a 62% wage premium while the ladder for entry-level jobs shortens. Consumers hand shopping to agents. For a firm, those individuals are its talent pool, its customers with a new buying interface, and a brand-trust risk at once. **Derived recommendations:**

- **Your buyer’s agent negotiates your margin.** In tests on real consumer-goods transactions, capable agents paid \$2.45 less and earned \$2.68 more than weak ones, unnoticed by the user ([Anthropic](#)). The agent filters before it recommends, so you now face a machine that optimizes

against your price. Meet it with your own selling agent and agent-legible value, as B4 sets out, not list price alone.

- **Hire and retrain from the churn.** AI fluency pays a 62% wage premium and keeps widening ([PwC](#)), while entry-level roles vanish first (workers aged 22–25 in exposed jobs –13% since 2022, [Stanford](#)). Recruit the AI-fluent early, and retrain displaced mid-career specialists into AI-supervisory roles before rivals bid them up.
- **Defend the brand against trust collapse through deepfake media.** Deepfakes rose ~16× toward ~8M in 2025 and deepfake-enabled fraud jumped ~11x ([Sumsu](#)). The default assumption on media content has flipped from “probably real” to “probably fake.” Adopt content provenance for official media, and add verification steps to payment approvals.

Lens G — Government. Governments try to govern at AI’s pace, from bounded self-interest: competition abroad, legitimacy at home. The likely path is sector regulators binding first, catastrophe-capable AI moved under licensing, and provenance and agent-identity rules arriving. For a firm, government sets the market-access rules and a large block of anchor demand. **Derived recommendations:**

- **Turn AI compliance into market access.** The EU AI Act’s labelling duties bind 2 August 2026 and high-risk rules 2 December 2027, with fines up to €35M or 7% of turnover ([AI Act](#)). Build evaluation, provenance, and audit into products now. Firms that treat it as product, not paperwork, ship faster when the rules bind.
- **Track your sector regulator, not horizontal law.** Sector regulators bind in 12–18 months while economy-wide AI law takes 5+ years. The FDA has already authorized 1,451 AI-enabled devices ([FDA](#)). Your binding AI rules come from your own industry regulator first. Engage it early to shape the standard, and plan investment on that 12–18 month timeline.
- **Migrate long-lived secrets to post-quantum encryption.** A “harvest now, decrypt later” threat puts today’s encrypted data at future risk, and standards are now finalized with deadlines through 2030–2035 ([NIST](#)). Inventory cryptographic dependencies now, and require post-quantum readiness.

Lens W — the World. This is the system view across lenses A, B, I & G. Competition in each of these actors drives runaway AI capability. Likely evolutions are a potential boom-bust in the capital cycle, an AI supplier field narrowing to 5–7 firms, and rising antitrust pressure. For a firm, that decides where profit survives, its financial exposure, and the terms it gets from suppliers. **Derived recommendations:**

- **Stand where the system lets durable value pool.** Most of the AI dividend passes to customers as lower prices. The profits settle with whoever holds a scarce layer, data, a network, or a physical or licensed position. That is a system-level behavior, not a sector quirk. Build toward one such layer now, before profit settles there and entry closes. The moat menu is breakout B5-2.
- **Stress-test for the boom-bust.** The likely near-term path is a capital-cycle bust, and the IMF and Bank for International Settlements now flag AI market concentration and circular vendor

financing as a stability risk ([IMF](#)). A crash would delay the build-out, not the technology, but it can still kill an over-extended supplier. Diversify AI-supplier and treasury exposure, and check that a key vendor could survive a funding freeze.

- **Plan for antitrust to reopen the supplier market.** The field is narrowing to 5–7 frontier firms as per-run training cost rises $\sim 2.4\times/\text{yr}$ ([Epoch AI](#)), which is drawing antitrust and interoperability pressure by 2027. Portability mandates or forced interoperability would cut your switching costs. Keep contracts and data portable now, so you can exploit any opening the moment it comes.
- **Competition within each player and AI progress compress a decade of innovation into the next 10 quarters, the remaining 2026–2028 window.** Several loops now each reach major innovations every quarter: 1. Increasing AI self-improvement (one lab reports AI involved in over 80% of its own code, [Anthropic](#)), 2. The cost decrease to reach a fixed capability ($\sim 10\times/\text{yr}$, [Stanford HAI](#)), and 3. Agent work-horizons (doubling every ~ 89 days toward a full work-week by late 2027, [METR](#)). Each compresses progress that once took more than a year into a quarter. They feed each other, so progress compounds rather than adds, and competition forces every rival to keep pushing. Overall innovation is driven by AI, squeezing a decade of innovation into the next 10 quarters.

[↑ Back to Table of Contents](#)

Coda

“Practical AGI” lets every rival cut cost and build new products quickly, so old advantages erode fast. One new advantage, moving first on AI, opens up. Competition will intensify not only between businesses, but also between AI companies, individuals and governments. All will use AI as the key lever to succeed, which will accelerate AI development. Overall innovation is driven by AI, squeezing a decade of innovation into the next 10 quarters.

AI turbo-charges competition. The lead is won in the 2026–2028 window.

[↑ Back to Table of Contents](#)

Stay tuned for more insights on AI from the perspective of the other lenses on [Robert Bruckmeier’s private paper channel](#).

Please cite this work as:

Bruckmeier, Robert. *The AI Dividend and the Deadline — What They Mean for your Business*. Robert Bruckmeier's private papers (July 2026). papers.robertbruckmeier.com/ai-dividend-and-deadline/business-view.html

Or use the BibTeX citation:

```
@article{bruckmeier2026aidividend,  
  title   = {The AI Dividend and the Deadline – What They Mean for your Business},  
  author  = {Bruckmeier, Robert},  
  journal = {Robert Bruckmeier's private papers, papers.robertbruckmeier.com},  
  year    = {2026},  
  month   = {July},  
  url     = {papers.robertbruckmeier.com/ai-dividend-and-deadline/business-  
view.html}  
}
```

© 2026 by Robert Bruckmeier

<https://papers.robertbruckmeier.com/ai-dividend-and-deadline/business-view.pdf>