PILOT PURGATORY

MAY 2026

Sargasso Capital Management

SARGASSO MANDATE

May 2026

The middle market sits at the precise point where AI capability has outrun organizational architecture. Durable businesses with real assets, recurring cash flows, and proprietary data streams still operate with legacy data silos, fragmented technology reporting, and boards that reward pilot volume over deployment velocity. The result is a structural mispricing that will persist until governance catches up with technology.

The framework in this paper is the product of direct observation across activist campaigns, public equity positions, and operational turnarounds. The three diagnostics, the constructivist mandate, and the six-step playbook constitute the exact sequence we have stress-tested in the positions we have built to date and designed for execution at scale in the $400 million to $900 million market-cap band.

The first movers who redesign architecture will capture the harvest phase of the AI J-curve. The laggards will face persistent relative underperformance. The middle market is large enough, fragmented enough, and under-analyzed enough that a disciplined constructivist approach can move the needle without displacing the operators who built the businesses.

This document is the operating manual. It is written for boards, for activists, and for any investor who believes the next wave of value creation will come from forcing deployment, not from buying more compute.

Thank you to my former colleagues at Roivant Sciences and to the Health Technology team working on AI deployment for a direct view of what genuine organizational transformation looks like when the architecture is built correctly from the start. Thank you to the leaders at the frontier of AI development who shared their time in 2024 discussing the trajectory of the technology, the shape of enterprise deployment, and the policy questions that will define the next decade. Those conversations informed the analytical framework in this paper in ways that published research alone could not.

Thank you to my former colleagues at Strive Asset Management, particularly those who helped build my understanding of how governance and engaged shareholding can force corporate behavior that passive ownership never will. That framework is the direct lineage of this paper's constructivist mandate. Thank you also to many friends for feedback on earlier drafts.

To my three sisters and my parents, thank you for the support that made all of it possible.

Patrick Feeley

May 2026

New York, New York

sargassocap.com

TABLE OF CONTENTS

The Author's Vantage Point

What follows is the product of having worked inside three organizations where the gap between stated strategy and operational reality was the central management problem, and where the cost of getting it wrong was visible in real time to anyone willing to read the financial statements carefully.

Strive Asset Management at the Nasdaq, 2022. Patrick Feeley is on the left in the blue suit.

In 2022, I joined Strive Asset Management as an early employee and became Chief of Staff to the firm's co-founder, Vivek Ramaswamy. During my first year at the firm I helped scale a thematic exchange-traded fund from zero to over $1 billion of assets under management. Over that period I engaged directly with the boards and executive teams of some of the largest public companies in America, including Darren Woods at ExxonMobil, where the firm's co-founder Vivek Ramaswamy pressed senior management on the strategic direction of the company amid a shifting energy and technology landscape.[1] What became clear across those meetings, and then repeated itself across a dozen other engagements with public-company management teams, was that the companies performing best were the ones where strategic focus was unambiguous and where measurement of operational execution against that focus was honest. The

[1] Strive Asset Management, “Strive Impacts Corporate America: After Key Changes at Exxon and Disney, Strive Will Target Chevron and Home Depot in 2023 Proxy Voting Season,” press release, December 6, 2022.

gap between stated strategy and operational reality was, in almost every case, the central management problem, regardless of the sector.

The Sargasso thesis has an identifiable origin. In 2022, after a direct engagement with Darren Woods at ExxonMobil in Dallas, I traveled with Vivek Ramaswamy to a Roivant board meeting in London. On the flight I raised the possibility of a fund structured explicitly to capture the alpha of the kind of direct engagement we had just observed, drawing on the Carl Icahn model. Vivek was clear-eyed that Strive's compounding advantage was in the direct-indexing franchise and that the activist mandate belonged in a separate vehicle, under separate governance, at a separate time. He was right on both counts. The idea that became Sargasso was set down that day and revisited four years later in a very different market environment.

Following Strive, in late 2023, I joined Vivek’s presidential campaign as Chief of Staff. The role was entirely operational. Every initiative, every resource allocation decision, and every hire was evaluated against a single question, which was the return on the decision relative to its cost. The campaign also brought me into contact with operators at the frontier of artificial intelligence, including Elon Musk, David Sacks, and Sam Altman, at a moment when the technology was transitioning from capability demonstration to industrial reality. What became clear across those conversations was that the bottleneck between artificial intelligence capability and enterprise deployment was not going to be the technology itself. It was going to be the organizations that were expected to absorb the technology.

Patrick Feeley with Vivek Ramaswamy at Bruce Rastetter's Summit Carbon Solutions event, Alden, Iowa, 2024. The day followed a meeting with Harold Hamm.

From there I moved to Roivant Sciences in mid-2024, and during my time at the firm I progressed through multiple roles of increasing responsibility across its biopharmaceutical portfolio. Roivant operates through the Vant model, in which each subsidiary company, each called a Vant, functions as an autonomous entity with its own management team, its own equity incentive structure, and a single therapeutic mandate. At my arrival the firm's market capitalization was approximately $8 billion. I departed at the start of 2026, when that figure had reached $15 billion.

In March 2026, Genevant Sciences and Arbutus Biopharma, both subsidiaries of Roivant, reached a $2.25 billion global settlement with Moderna over lipid nanoparticle patent claims, the largest disclosed patent settlement in pharmaceutical history. Priovant Therapeutics, another Roivant Vant, delivered positive Phase 3 VALOR results for brepocitinib in dermatomyositis and follow-on Phase 2 data in cutaneous sarcoidosis, an indication with no approved therapies. Roivant's financial outcomes during my time were driven primarily by clinical milestones and by aggressive intellectual property enforcement, not by artificial intelligence deployment. The organizational relevance of the Roivant experience to this paper is therefore narrower, and more specific, than the raw valuation trajectory would suggest. The Vant structure preceded the artificial

intelligence cycle by nearly a decade, and the structure was designed around autonomous experimentation in drug development rather than around any computational premise. What the Vant structure demonstrates, and what I carried forward into this thesis, is that organizational designs built around rapid iteration and clean internal accountability are compatible with artificial intelligence in a way that legacy functional hierarchies are not. That compatibility is the relevant lesson.

Biotech drug development matters to the argument of this paper because it is the sector in which the thesis becomes undeniable. Industry reports track hundreds of partnerships between AI vendors and Big Pharma. For example, data for 2024 and 2025 alone recorded 198 deals in AI/ML drug discovery and licensing with an aggregate potential value of $55.3 billion.[2] Average drug development timelines did not shorten. Clinical trial failure rates remain above 90%.[3] The gap between artificial intelligence investment and artificial intelligence outcome, in the sector that possesses the largest financial incentive to close that gap, is an organizational problem rather than a technological one. And the gap persists even where the financial incentive to close it is survival. What separated Roivant from the broader sector failure during my time was not access to better technology. It was an organizational architecture designed from the beginning to absorb the technology, with tight feedback loops between technology and operations that most biopharmaceutical companies have never established internally.

Method

The analytical framework advanced in this paper rests on a single discipline, which is return on investment. The commitment to measurable outcomes is not the same as the performance of commitment, and the distinction between the two has become the single most useful diagnostic tool in separating operators who will reach the harvest phase of the artificial intelligence J-curve from operators who will not. Companies that will reach the harvest phase are not the ones that issued the most press releases or announced the most partnerships. They are the ones that approached artificial intelligence adoption with the same capital allocation rigor they would have applied to any other investment decision of comparable size, and who demanded the same measurable return.

[2] DealForma, “AI-ML Drug Discovery and Licensing R&D, M&A, Ventures and IPOs — 2025 Review,” January 2026.
[3] Biotechnology Innovation Organization (BIO), Informa Pharma Intelligence, and QLS Advisors, “Clinical Development Success Rates and Contributing Factors,” 2021.

Artificial intelligence creates real, measurable, and already-observable value at the operator level, and that value shows up in margin structure and in returns on capital. The bottleneck that prevents most operators from capturing the value is almost never the technology itself. The bottleneck is the organizational architecture. Boards that cannot evaluate an artificial intelligence roadmap. C-suites in which no executive owns the deployment question directly. Management teams that have learned to use the vocabulary of artificial intelligence adoption, accurately in many cases, without having built any of the underlying infrastructure. Each of those structural failures produces a distinctive pattern in the financial statements, and each of those patterns is legible to an investor willing to look for it.

The gap between companies that claim artificial intelligence readiness and companies that have achieved it is one of the most legible arbitrages in public equities today. Approximately 58% of small and mid-cap companies globally claim to have adopted artificial intelligence, while fewer than 1% describe their implementation as mature.[4] The remaining 57% represent the population examined in this paper.

What follows draws on a systematic review of artificial intelligence implementation across hundreds of public companies, tested against proxy filings, earnings transcripts, audited financial statements, and then refined through direct conversations with operators, capital allocators, and the researchers building the models themselves. The gap between what companies claim about their artificial intelligence strategies and what their audited financial statements actually show is wide, it is measurable, and it is mispriced. Identifying the gap is analytical work. Closing the gap is what the constructivist mandate exists to do.

[4] MIT NANDA, “The GenAI Divide: State of AI in Business 2025,” July 2025.

PREFACE

BEFORE THE FLOOD

> *“And my stupid answer is: It’s more than a nothing burger, and it’s less than the total transformation of our society … my placeholder is that it’s roughly on the scale of the internet in the late 1990s. It might be enough to create some great companies.”*
>
> *Peter Thiel, Interesting Times podcast, June 2025*

Peter Thiel has framed the stakes with characteristic precision. His skepticism rests on the distinction between the world of bits, where artificial intelligence operates with obvious force, and the world of atoms, where the physical breakthroughs that would reorganize energy, transportation, and medicine remain stalled. For Thiel, artificial intelligence has become a placeholder for progress because it is the only sector currently showing movement at all. This paper begins with that tension and argues something more specific. Thiel is right that artificial intelligence may not deliver civilizational transformation. He is also right that it will create great companies. Neither framing captures the third possibility. Artificial intelligence is already acting as a structural headwind for the complacent, producing persistent relative underperformance and, in some cases, material destruction of shareholder value relative to the adopters. Historical general-purpose technology episodes, from electrification to enterprise computing, show that laggards can eventually catch up. The duration of the lag depends on competitive intensity and capital-market discipline. In the current cycle, both are elevated. The widening gap between firms that successfully bridge the divide and those that are left behind defines the most compelling investment opportunity public equities have offered in a generation.

The contrarian claim of this paper reduces to a single sentence. The mid-cap AI adoption gap is a governance problem that produces proprietary data loops compounding into durable moats that the market has not priced, because the market is reading the investor-day slide deck rather than the audited 10-K.

Each clause of that sentence does independent work. First, the technology is commercially available. Second, the capital is available. Third, the governance structures at most mid-cap operators are not built to absorb either the technology or the capital, and that constraint, not the technology itself,

determines which companies will hold proprietary data advantages in 2030 and which will be acquired, restructured, or left to trade at multiples reflecting the deterioration.

AppLovin illustrates the pattern in real time. The mobile advertising technology company generated $5.48 billion of revenue in 2025 with 898 full-time employees. That is $6.1 million of revenue per employee, an order of magnitude higher than the legacy competitors it displaced. It delivered an 82% adjusted EBITDA margin on that revenue, driven by AXON 2.0, a proprietary advertising engine built from the ground up rather than retrofitted onto legacy infrastructure. By early 2026 its market capitalization had reached approximately $132 billion. The architecture, not the advertising market, explains the difference between AppLovin and every legacy advertising-adjacent platform still trying to catch up.

Revenue per employee an order of magnitude above legacy competitors, a cost structure incumbents cannot match at any scale, and a development cycle measured in quarters rather than years. That combination is what AI-native organizational design produces when a company is built from the beginning for the operating model, rather than retrofitted around it after the fact. The mid-cap public equity universe is about to encounter the same combination repeatedly, in sector after sector, at a pace accelerating quarter by quarter. Thiel’s bits-over-atoms observation, which he first articulated publicly more than a decade ago and which has been referenced in almost every major technology-investing conversation since, is playing out at the operator level, in precisely the industries where cognitive work and data architecture determine competitive position.

The Private Layer

The AI-native private entrant is the structurally dominant predator in the current investment cycle. Venture-backed companies building on top of frontier models carry three advantages over mid-cap public incumbents that cannot be corrected on any timeline that matters to a current public shareholder. First, they carry no legacy system integration debt. Second, their workforces were hired specifically to build with artificial intelligence, rather than retrained to tolerate it. Third, their boards, composed in many cases of the founders and early investors of other AI-native companies, understand the product they are funding and can evaluate the roadmap that produces it. Each of the three advantages compounds with the other two. The incumbent's usual response, which is to announce a partnership with a large cloud vendor and stand up an internal AI committee, addresses none of them.

United States venture deal value reached a record $267.2 billion in the first quarter of 2026, with 88.8% of that capital flowing to AI-tagged companies, an accelerating share directed at vertical applications aimed at specific middle market industries (a franchise operations platform, a specialty logistics platform, a regional insurance platform, and an industrial distribution platform). Every sector examined in this paper is targeted by at least one well-capitalized venture-backed competitor whose cost structure, data architecture, and decision-making speed the incumbent cannot match. The outcome is not in dispute. Only the timing is.[5]

Private equity has recognized the same opportunity and is constructing the same competitive threat through a different mechanism. A consolidator with a technology thesis, acquiring the same class of franchise or regional operator that a company like Dine Brands has already assembled, can run artificial intelligence across scheduling, pricing, supply chain, and loyalty as a unified operating system, without the governance friction of a public company and without the quarterly earnings cycle. This type of rollup is no longer financial engineering. It is operational engineering, constructed around a technology stack the public incumbent has not yet deployed. The AI-native rollup compounds the same operational advantage as the venture entrant, but at greater scale and funded from committed capital that sits outside the quarterly reporting cycle. The pressure on the public mid-cap operator arrives from both directions simultaneously. A survey of two hundred private equity funds and operating leaders documents AI adoption uneven across portfolio companies, with enterprise-scale deployment confined to a narrow tier of top-performing funds and the performance gap between them and the rest widening each

[5] PitchBook-NVCA Venture Monitor, First Quarter 2026, April 3, 2026.

quarter. The separation is forming inside private equity before it forms inside public equity, which means the rollup operator will have consolidated the target list before the public incumbent has written a meaningful contract.[6]

The largest private equity firm in the world is documenting the same gap in its own portfolio. In April 2026, Blackstone president Jonathan Gray reported that large language model spend across the firm's portfolio companies grew roughly fifteen-fold year-over-year in the first quarter, while the pace of operational implementation remained, in his words, "frustratingly slow." Capital is flowing into the technology faster than organizations can absorb it. The implementation gap is the binding constraint. Gray's observation extends the diagnosis from the public mid-cap universe into the largest private equity portfolio on earth, with the same conclusion. The technology works. The organizational architecture has not caught up.

Two days before this paper was published, that entrant was capitalized. On May 4, 2026, Anthropic, Blackstone, Hellman & Friedman, and Goldman Sachs announced the formation of an approximately $1.5 billion AI-native enterprise services firm. The vehicle is private. It is a standalone operating company, not a fund and not a public listing. Anthropic, Blackstone, and Hellman & Friedman are each contributing approximately $300 million. Goldman Sachs is committing approximately $150 million as a founding investor. Apollo Global Management, General Atlantic, Leonard Green, Singapore's GIC, and Sequoia Capital joined as co-investors. The deployment target is also private. The initial customer base is the founding consortium's own portfolio companies, predominantly PE-owned mid-market businesses across healthcare, manufacturing, financial services, retail, real estate, and infrastructure. The next ring is the broader PE-backed mid-market universe. Public-company deployment is downstream, if it happens at all. The structure is designed to deploy Claude directly into the workflow of mid-market businesses that do not carry the governance, talent, and legacy-architecture frictions of the public incumbents examined in this paper. Blackstone president and COO Jon Gray described the venture as designed to break down "one of the most significant bottlenecks to enterprise AI adoption," the scarcity of skilled implementation partners. This is the exact diagnosis advanced in this chapter. The implication for the public mid-caps in this paper's universe is direct. Their private-market competitors are being equipped, at speed, with the implementation capability they have spent two years promising and not delivering. The AI-native entrant described earlier is no longer

[6] FTI Consulting, "2026 Private Equity AI Radar: Continued AI Acceleration and Impact," March 17, 2026.

hypothetical. It has a balance sheet, a mandate, and a distribution channel. The companies it is positioned to displace are sitting on the other side of the trade.[7]

This is the layer that the sell-side research function cannot price, because analyst coverage cannot see it. The analyst covering Dine Brands does not maintain a counterpart model for the private competitor that will take 3% of system-wide sales in 2027 and 6% in 2028. The analyst covering a mid-cap specialty industrial company does not maintain a model for the venture-backed entrant pricing customer acquisitions 60% below the incumbent's cost of sales. The sell-side models are run on the world that exists inside public filings. The threat is building in the world that does not appear in those filings, until it surfaces as margin compression that cannot be reversed.

The Public Reckoning

On April 15, 2026, Evan Spiegel laid off a thousand people at Snap and cited artificial intelligence as the mechanism that had made them redundant. The market responded with an 8% premarket rally.[8] The move was not the product of Spiegel's sudden conversion. It was the product of sustained pressure from Irenic Capital Management, whose 66-page forensic brief on March 31, 2026 had quantified the gap between what Snap had been saying publicly about artificial intelligence for the previous three years and what its audited financial statements actually showed.[9] Snap's revenue per daily active user had been flat since 2023, while Meta's had more than doubled over the same period.[10] The architecture to close that gap had not been built. Irenic forced the conversation. The conversation forced the layoffs. The layoffs produced the surge.

Snap is the version of this story that a public-markets analyst is trained to recognize. Snap has a ticker, a proxy statement, an institutional shareholder base, and an activist. It can be modeled. Snap is also the earliest visible chapter of a longer sequence. The first chapter is AppLovin and the hundreds of AI-native entrants being funded against incumbents like it. The last chapter is a strategic review of a middle market public franchisor that waited too long to respond.

The constructivist playbook described in this paper creates the governance conditions under which proprietary data loops can begin forming at businesses

[7] Anthropic, “Building a new enterprise AI services company with Blackstone, Hellman & Friedman, and Goldman Sachs,” May 4, 2026; Blackstone press release, May 4, 2026; *Wall Street Journal*, May 4, 2026 (capital allocations).
[8] Snap Inc., Form 8-K, April 15, 2026.
[9] Irenic Capital Management, “Snap Back to Reality: Save Snap Now,” March 31, 2026, savesnapnow.com.
[10] Snap Inc., Form 10-K, Fiscal Year 2025; Meta Platforms, Inc., Form 10-K, Fiscal Year 2025.

that would otherwise continue satisficing indefinitely. Meta did not buy its eleven-year advertising signal. Meta built it, one quarter at a time, from roughly 2015 forward. The constructivist investor is the forcing function that moves a capable operator from satisficing to optimizing, and in doing so turns a potential zero-to-one outcome at the operator level into the proprietary data advantage that justifies the multiple re-rating. Chapter 11 develops the mechanics of that compounding data advantage in the context of vertical-specific agents. The setup argument here is simpler. It is a one-to-n mechanism producing zero-to-one outcomes at the level at which a mid-cap public equity investor is able to capture the return.

The flood is not coming. The flood is already here. It is moving slowly enough that the organizations in its path can still choose to move to higher ground, and it is moving fast enough that the window for that choice is measured in quarters rather than in years. What follows explains how to identify which of the organizations will.

PART I

THE GAP

Establishing the macro thesis

CHAPTER 1

THE NARRATIVE AND THE REALITY

Every technological era produces two distinct markets. The first is the market for the technology itself, and the second is the market for what the technology does to everything else. The first market is easy to see. It concentrates capital rapidly and it generates the majority of the headlines. The second market is much harder to see. It is diffuse, it moves more slowly, and it is embedded inside businesses that were never designed to be technology stories, and whose management teams rarely understand that they have become one. The second market, in almost every historical case, is also the larger of the two.

The railroad era provides a historical template. Cornelius Vanderbilt and his successors made fortunes laying track across the United States in the second half of the nineteenth century, but the fortunes built on top of that track, by the shippers, the steel producers, the retailers, and the industrial manufacturers who used the rails to consolidate national markets, were larger and more durable. Electrification produced the same pattern. The great utility companies were significant. The manufacturers who redesigned their factory floors around electric motors, rather than simply plugging electric motors into steam-era factory layouts, were larger again, and captured roughly thirty years of productivity gain that the utilities themselves did not. The most instructive recent example is the Dot Com cycle. Between 1995 and 2000, enormous pools of capital flooded into companies building the physical and software infrastructure of the commercial internet. More than 4,000 of those companies had ceased operations by 2004. Amazon and Google, both of which used the infrastructure rather than supplied it, produced the defining returns of the cycle. The wealth in the Dot Com era was captured not by those who built the pipes but by those who understood what to pump through the pipes. That dynamic is the precise template for the artificial intelligence cycle currently underway.

The market is in the middle of the same bifurcation right now. Since the launch of ChatGPT in late 2022, a cohort of 42 AI-related companies has accounted for roughly 65% to 75% of the S&P 500's price returns, profits, and capital

spending. These are the semiconductor manufacturers, the hyperscale cloud operators, the foundation model companies, and the data center infrastructure providers, whose revenue is directly tied to building or selling artificial intelligence capability. Nvidia, Microsoft, Alphabet, Meta, and Amazon constitute the core of the cohort. These companies are not the subject of this paper.

Three quarters of the return generated by the 500 largest public companies over the past three years has come from fewer than 10% of the names in the index. Without the 42 AI-related companies, the domestic index would have underperformed Europe, Japan, and China over the same period.[11] The phenomenon has been described as "Smothering Heights." The concentration of returns at the top smothers everything below it.

	42 AI Companies	**S&P 500 ex-AI**	**Spread**
Share of S&P 500 price returns (Nov 2022 – Dec 2025)	~75%	~25%	~3x concentration
Share of index earnings growth	~75%	~25%	~3x concentration
Share of capex growth	~70%	~30%	~2.3x concentration
% of S&P 500 market cap	50%	50%	0.8% of companies
Net debt / total index	5%	95%	Net cash positions

Table 1. AI cohort return concentration vs. the rest of the S&P 500, November 2022 through 2025.

> "*The generally accepted view is that markets are always right. Market prices tend to discount future developments accurately even when it is unclear what those developments are. I start with the opposite view. I believe the market prices are always wrong in the sense that they present a biased view of the future.*"
>
> *George Soros, The Alchemy of Finance, 1987*

In 2026 the bias is AI narrative. The market prices the press release. The financial statements have not yet learned to lie. Soros described reflexivity as the mechanism by which biased market perceptions feed back into

[11] Michael Cembalest, "The Blob," *Eye on the Market*, J.P. Morgan Asset Management, September 24, 2025; Michael Cembalest, "Smothering Heights," *Eye on the Market Outlook 2026*, J.P. Morgan Asset Management, January 1, 2026.

fundamentals until the distortion becomes self-reinforcing and then self-correcting. The AI adoption cycle is producing a particularly clean version of that loop. The narrative attracts capital. The capital inflates the multiple. The inflated multiple lets management defer the operational work the narrative implied. The gap between the reported strategy and the audited results widens each quarter until an external catalyst, typically a constructivist or activist investor, forces the reconciliation. The signal for the investor is the gap itself. The mechanism of the trade is the catalyst.

The Scale of What Is Coming

The AI infrastructure trade is driven by exponential increases in effective compute rather than by linear growth. Between GPT-2 and GPT-4, the field witnessed a qualitative shift from preschooler-level to high-schooler-level reasoning. That shift required 3 to 6 orders of magnitude of effective compute growth.

By 2027, the leading laboratories project systems capable of performing any task a remote human worker can accomplish. The phase that follows, what researchers call the intelligence explosion, could compress a decade of AI progress into a single year. Individual training clusters by 2028 are projected to cost over $1 trillion and consume 20% of total U.S. electricity production.[12]

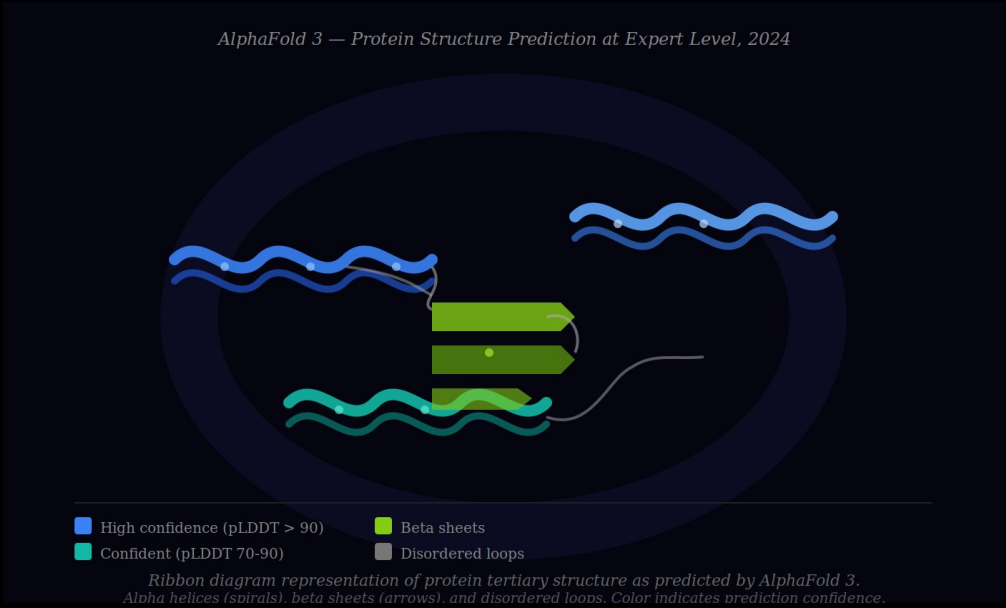

AlphaFold 3 protein structure ribbon diagram. Blue regions indicate high prediction confidence; secondary structure elements including beta sheets are rendered in contrasting colors. DeepMind achieved expert-level accuracy across all known protein families in 2024.

[12] Leopold Aschenbrenner, "Situational Awareness: The Decade Ahead," June 2024, situational-awareness.ai.

The capability argument requires no speculation. DeepMind's AlphaFold 3 achieved expert-level protein structure prediction. Its Gemini 2.0 research established consistent expert-level performance on multi-step analytical reasoning.[13] These are peer-reviewed results, not demonstrations. The question is not whether AI systems can perform expert-level work across specialized domains. They can. The question is whether the organizational architecture of any given company can direct that capability toward measurable financial outcomes. Almost none of them currently can.

The infrastructure spending is real. The returns are not guaranteed. The Stargate partnership between OpenAI and Oracle requires 4.5 gigawatts of power equivalent to four nuclear plants and has already fallen short of its stated capital commitments.[14] Hyperscaler capex is projected to reach $1.15 trillion from 2025 through 2027, more than double the prior three-year total, with investors already becoming more selective, rotating away from infrastructure companies where operating earnings growth is under pressure.[15] The companies making these bets are wagering that the infrastructure thesis will complete on schedule. The companies described in this paper are the ones positioned to use that infrastructure productively once it does, if they have built the organizational architecture to absorb it.

The frontier laboratories are fully oriented toward this trajectory. The organizational distance between where AI capability sits and where most mid-market operators have deployed it is the central fact of the current investment environment.

The Structural Character of the Concentration

The concentration of returns reflects a genuine financial divergence. Within the broader cohort of 42, a tighter core of 28 direct AI stocks, those where AI is the primary product rather than a tool, represent 50% of S&P 500 market capitalization but carry only 5% of its net debt.[16] These are cash-generating machines that have financed over a trillion dollars in capital expenditure almost entirely from internally generated cash flow.

[13] Josh Abramson et al., "Accurate structure prediction of biomolecular interactions with AlphaFold 3," *Nature* 630 (2024): 493-500; Google DeepMind, "Introducing Gemini 2.0: our new AI model for the agentic era," December 11, 2024.
[14] Michael Cembalest, "The Blob," *Eye on the Market*, J.P. Morgan Asset Management, September 24, 2025; "Stargate AI Project Stalls Amid Partner Disputes," *The Information*, February 23, 2026.
[15] Goldman Sachs Research, "Why AI Companies May Invest More than $500 Billion in 2026," December 18, 2025.
[16] Michael Cembalest, "Smothering Heights," *Eye on the Market Outlook 2026*, J.P. Morgan Asset Management, January 1, 2026.

The distortions created by this spending produced outcomes that would have been unimaginable five years ago. In April 2026, Allbirds, a mid-cap sustainable footwear company trading under the ticker BIRD with a net loss margin of 61.6% and a collapsed stock price, announced it would sell its core shoe brand for $39 million and pivot the remaining corporate shell into artificial intelligence compute infrastructure, rebranding as NewBird AI. The company secured a $50 million convertible facility to acquire and lease GPU assets. Its shares surged 582% in a single session.[17]

What the BIRD episode illustrated was not the strength of the AI pivot narrative. It was the reflexive, first-layer nature of modern market coverage. The financial press and a meaningful portion of the investor class treated a 582% single-day move as evidence of investor conviction in a GPU-as-a-service thesis. The underlying mechanics were a short squeeze with a $20 million market cap, near-zero borrow, effectively no float, and 20% accumulated short interest from investors positioned for a wind-down. The $50 million convertible announcement supplied the trigger. The AI framing supplied the cover. A rigorous reading of the tape would have identified the squeeze conditions before the catalyst arrived. The coverage did not. The same reactive pattern plays out at lower magnitudes every week, in every sector where the AI narrative runs ahead of the financial statements. It is the broader condition this paper is written against.

This concentration is historically anomalous. Every prior infrastructure cycle of this scale required significant external capital. The current cycle was initially different. The hyperscalers arrived at the AI moment with very low net debt and large cash balances, and funded early spending from internal cash. That is beginning to change. Meta, Oracle, and Alphabet issued $75 billion in bonds and loans in two months alone to fund data center buildouts.[18] The era of pure internal funding is ending as spending accelerates beyond what operating cash flows can sustain.

[17] Allbirds, Inc., "Allbirds, Inc. Executes $50M Convertible Financing Facility Agreement; Announces Expansion into AI Compute Infrastructure," press release, April 15, 2026; Ashley Capoot, "Allbirds pivots from shoes to AI, BIRD stock soars," *CNBC*, April 15, 2026.
[18] Bank of America analysis cited in Allie Garfinkle, "Oracle’s collapsing stock shows the AI boom is running into two hard limits: physics and debt markets," *Fortune*, December 13, 2025.

What the Market Has Priced and What It Has Not

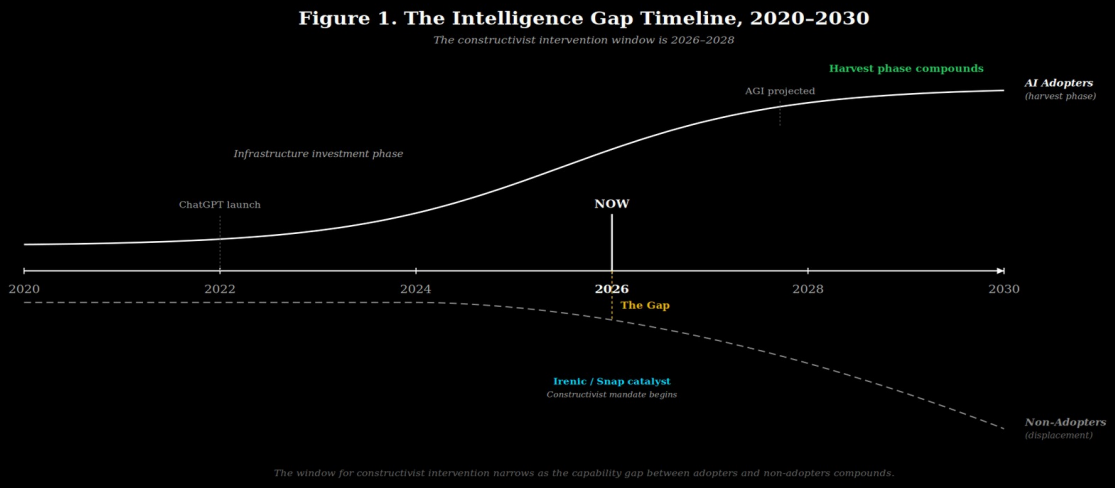

Figure 1. The Intelligence Gap Timeline, 2020–2030. The constructivist intervention window is 2026–2028.

The operational benefits of AI have not been priced anywhere outside that cohort.

This is an observation about earnings models, not just multiples. Sell-side analysts covering Dine Brands, GoPro, or any of the several thousand other mid-cap operators are not building AI adoption events into their forward numbers. The preferred range identified in this paper is $400 million to $900 million in equity market capitalization, where position-building liquidity is sufficient and institutional coverage thin enough to sustain a durable informational edge. Snap, with its $7 billion enterprise value, sits well above that range, but the analytical framework applies equally. Even at larger capitalizations, the gap between AI narrative and AI financial result is neither modeled nor priced. They model organic growth at historical rates and apply sector-standard multiples. The possibility that Snap could close even a fraction of the revenue-per-user gap between itself and Meta through genuine AI deployment of its advertising technology stack is not in any consensus estimate.

That creates an asymmetry. The downside reflects the businesses as they are, slow-growing, over-managed, and priced accordingly. The upside reflects a scenario the market has not modeled. When an adoption event occurs inside a business the market has underwritten as static, the multiple re-rating is non-linear.

58% of small and mid-cap companies claim AI adoption strategies. Fewer than 1% describe their implementation as mature.[19] The gap between those two

[19] MIT NANDA, “The GenAI Divide: State of AI in Business 2025,” July 2025

numbers is a decade of value creation waiting to be released, company by company.

The Narrative Capture Problem

The financial media, the sell-side, and the conference circuit have concentrated analytical attention on the companies generating the most dramatic returns. The AI story, as told in most institutional contexts, is a story about Nvidia, Microsoft, Anthropic, and the infrastructure layer. Chips, data centers, foundation models.

That framing is correct as far as it goes. It is also incomplete in a way that creates the opportunity this paper is about.

As the saying goes, Coca-Cola benefitted more from refrigeration than the refrigerator manufacturer ever did. The cold chain made sugared water a mass-market product. The company that owned the consumer relationship captured the value. The AI cycle will resolve the same way.

The infrastructure is largely in place. The models are capable. The question is no longer whether AI works. It is whether management teams can absorb it, implement it, and capture the value before a more capable competitor does. That question is precisely where a constructivist investor with direct management access can create value. The narrative and institutional capital have not caught up to the operational reality. The gap between them is the trade.

CHAPTER 2

WHY THE MIDDLE MARKET

The operational transformation of AI is underpriced across public equities, and the mispricing is concentrated in a specific stratum rather than distributed uniformly. The small and mid-cap operator, sitting below the reach of the hyperscalers and above the threshold of distress. The reasons are structural. And the earliest signal that the equity market is mispricing this stratum often appears not in stock prices but in credit spreads, a dynamic Boaz Weinstein at Saba Capital has built a career exploiting across different asset classes.

Consider what happened at Snap. On March 31, 2026, Irenic Capital Management, a $2.5 billion fund, published a letter to CEO Evan Spiegel after building a 2.5% economic interest in Snap's Class A shares. What Irenic documented is the thesis of this paper reduced to a single company.

The Snap Case Study

Start with the asset. Snap reaches 943 million monthly active users. It reaches 75% of the 13-to-34 age cohort globally, the most commercially valuable demographic in digital advertising. Its users open the application 40 times per day. It has 477 million daily active users. It has 25 million paying Snapchat Plus subscribers generating close to $1 billion in annual recurring revenue. Every day, users create more than 5 billion snaps, generating a proprietary geolocation-linked image and video dataset that is among the most commercially valuable AI training assets in existence. 350 million users engage with its augmented reality tools.[20]

Now look at what the market pays for that asset. Prior to Irenic's letter, Snap's enterprise value was $7.2 billion. Every dollar invested at Snap's 2017 IPO at $17 per share had shrunk to 23 cents by early 2026, a decline of more than three-quarters of invested capital across 9 years while the platform grew to 943 million monthly active users. The stock had fallen approximately 54% over the prior year. Snap trades at 1.2x revenue. Meta trades at 5.2x revenue.[21] Meta has 3.3 billion daily active users. Snap has 477 million. The gap in daily active

[20] Irenic Capital Management, “Snap Back to Reality,” letter to Snap Inc. shareholders, March 31, 2026; Snap Inc., Form 10-K, Fiscal Year 2025.
[21] Irenic Capital Management, “Snap Back to Reality,” March 31, 2026; Bloomberg market data, March 2026.

users is roughly seven times. The gap in valuation multiple is over four times. The explanatory variable is what each company has done with AI to monetize its user base, not the size of the user base itself.[22]

Metric	Snap (2025)	Meta (Benchmark)
Revenue per DAU	$3.62	$17.60
Daily Active Users	477 million	3.3 billion
Enterprise Value (pre-Irenic)	$7.2 billion	$1.47 trillion
EV / Revenue	1.2x	5.2x
Employees (pre-April 2026)	5,200	74,000
Cumulative failed hardware bet	$3.5B (Spectacles)	~$80B (Reality Labs, cumulative)
Return on $1 invested since March 2017 (Snap IPO)	$0.23 by 2026	$7.40 by 2026

Table 2. Snap versus Meta across key operating and valuation metrics, as of March 2026.

Meta deployed AI and machine learning tools aggressively across its advertising stack beginning in 2022. The results were not subtle. Revenue per daily active user at Meta improved dramatically. Operating margin expanded by roughly 17 percentage points between fiscal 2022 and fiscal 2024, from approximately 25% to 42%, following its Year of Efficiency in 2023 and 2024.[23] AppLovin makes the same point in sharper form. The advertising technology company has a market capitalization of approximately $132 billion and generates $5.5 billion in annual revenue with 898 employees. Revenue per employee exceeds $6 million. AppLovin's AI-powered advertising platform, which it built from the ground up rather than retrofitted onto legacy infrastructure, drives those economics. AI-enabled ad targeting dramatically improves return on advertising spend for customers, which drives higher pricing power for the platform, which flows directly into revenue per user and margin.

The Asymmetric Competitor

The mid-cap AI adoption gap is measured against a competitor that the public filings do not yet contain, rather than against other mid-caps. Venture-backed AI-native companies are being capitalized at scale across every sector this paper covers, carrying no legacy integration debt, no quarterly earnings cycle,

[22] Irenic Capital Management, “Snap Back to Reality,” March 31, 2026; Bloomberg market data.
[23] Meta Platforms, Inc., Form 10-K, Fiscal Year 2024.

and no board that needs to be convinced of the product roadmap. They are already taking share inside the windows of time the public mid-cap consumes debating whether to begin.

Private equity is funding the second form of the threat. An operational consolidator assembling the same class of franchise or regional operator that Dine Brands already runs, equipped with a unified AI stack across scheduling, pricing, supply chain, and loyalty, generates unit economics the public franchisor cannot match. The rollup model is no longer a financial engineering play. It is an operational one, funded from committed capital that is not answerable to a quarterly earnings cycle.

The sell-side cannot price this layer because the analyst covering the public incumbent does not have a counterpart covering the private entrant. The models are built on the world that exists inside 10-K filings. The threat is building outside them. It surfaces in public financials only as margin compression, market-share loss, and eventually a strategic review at a valuation that reflects the deteriorated position rather than the underlying asset quality. That is why the intervention window is measured in quarters, not years.

Where the Thesis Does Not Apply

The argument made so far applies to a specific subset of the mid-cap universe, not to all of it. Businesses with standardized operational workflows at scale, where unit economics can be moved by changes in labor scheduling, pricing, inventory management, or customer lifetime value modeling, face acute AI pressure. Franchise operations, multi-site retail, freight logistics, and cognitive-labor-intensive services sit in this category. Businesses with physical moats that AI cannot erode, or trust-based local relationships, sit outside it. Regulatory barriers that once insulated incumbents are more vulnerable than they appear. AI is driving the cost of regulatory compliance toward zero, which erodes the advantage held by entrenched players whose scale was built on the capacity to absorb compliance expense.[24] A regional specialty distributor with 25-year customer relationships and a physical fulfillment network does not face the same AI pressure as a national franchisor with a standardized operating playbook. The thesis of this paper is concentrated, not universal. The mid-cap operators most at risk of displacement are the ones whose competitive position depends on operational variables that AI can directly move. Those are also the operators where the constructivist intervention produces the highest expected return.

[24] Paul Ohm, “Toward Compliance Zero: AI and the Vanishing Costs of Regulatory Compliance,” *Network Law Review*, October 31, 2025.

Snap has the same raw material. It has a user base, engagement intensity, demographic reach, and a proprietary dataset that its larger competitors would pay significantly to acquire. What it does not have is an organizational architecture capable of deploying AI against that opportunity.

Irenic identified the specific failures with precision. Snap had approximately 3,200 employees at the end of 2019, before the COVID pandemic. It ended in 2025 with 5,261. Its peers course corrected. Snap did not. According to Irenic's analysis, the company has spent more than $3.5 billion on Spectacles, its augmented reality glasses hardware initiative, and continues to spend approximately $500 million annually on the project. Snap itself, in its Q2 2025 investor letter, has separately disclosed committing "more than $3 billion over the past 11 years" to its augmented reality platform. Meta's parallel hardware bet, Reality Labs, has accumulated more than $80 billion in operating losses since 2020 and is a roughly 23 times larger failure than Spectacles in dollar terms. Meta has $1.5 trillion of enterprise value and an advertising business generating revenue per user that subsidizes hardware experimentation indefinitely. Snap does not. The same wasteful hardware bet is a rounding error for one company and an existential capital allocation error for the other.[25] Spectacles cannot be funded independently. The hardware line has never generated meaningful revenue. Every dollar allocated to Spectacles is a dollar not allocated to the AI advertising technology stack that could close the revenue-per-user gap with Meta. Irenic put it directly. Specs and similar efforts are a distraction from the core business. The return on research and development dollars spent improving the advertising technology stack would be positive and substantial. The return on Spectacles has been negative and substantial.

The MyAI product, Snap's generative AI assistant, has real user adoption. But Irenic identified a critical architectural error. The monetization partnerships are with Microsoft and Perplexity, second-tier partners, rather than the clear winners in enterprise AI, Gemini, OpenAI, and Anthropic. This is not a minor detail. The choice of AI partner determines the quality of the underlying model, which determines the quality of the user experience, which determines whether MyAI becomes a genuine revenue driver or a feature that generates favorable press coverage without changing the financial trajectory.

The governance structure compounds every operational problem. Evan Spiegel controls Snap through a dual class share structure. Class A shareholders have no meaningful vote. They cannot elect directors. They cannot replace

[25] Irenic Capital Management, "Snap Back to Reality," March 31, 2026; Snap Inc., Form 10-K, Fiscal Year 2025.

management. Irenic, despite holding a 2.5 percent stake, has no formal mechanism to compel action. Its entire campaign is an exercise in persuasion directed at a founder who retains controlling voting power under the company's dual-class structure. Snap is not alone in this. GoPro carries the same dual-class structure, with Nicholas Woodman holding supervoting shares that structurally limit shareholder accountability mechanisms regardless of operating performance. The pattern recurs across founder-led technology companies and is one of the primary reasons organizational issues compound at these businesses absent governance reform. Founder voting protection often takes precedence over shareholder accountability mechanisms in the dual-class structure. There are important exceptions. Mark Zuckerberg at Meta is the clearest case where founder voting control and sustained shareholder value creation have coexisted. The Year of Efficiency demonstrated that Zuckerberg could subordinate product preferences to financial performance when the two came into conflict, and the financial results followed. The question this thesis asks of any founder-controlled position is whether the founder has demonstrated continued success against the objectives stated at the outset. If the answer is no, the dual-class structure is protecting the founder rather than the shareholders.

Irenic's six-step turnaround framework maps precisely onto the diagnostic framework this paper advances. Step one is eliminating capital allocation that destroys value, namely Spectacles. Step two is right-sizing the cost structure by reducing the workforce from 5,200 to 4,200. Step three is realigning employee incentives around share price performance rather than tenure. Step four is deploying AI against the core advertising monetization opportunity. Step five is monetizing the proprietary AI dataset properly. Step six is improving governance to enable index inclusion and reduce the cost of capital. The market responded to Irenic's letter with a 14% single-day gain in Snap's stock price.[26] That single-day move represents approximately $1 billion of market capitalization created by the articulation of an analytical framework. The framework itself created no value. What created value was the credible signal that an engaged shareholder would force management to confront an organizational failure the market had been discounting for years.

What Snap Proves About the Thesis

The Snap Back to Reality letter is among the most rigorous activist documents published in the current AI cycle. Irenic's diagnosis is exact. The quantification

[26] Bloomberg market data, April 2026, reflecting Snap Inc. (NYSE: SNAP) intraday trading following publication of Irenic Capital Management's March 31, 2026 letter.

is exact. The six-step framework is, in structure, closely aligned with the analytical approach described in this paper, applied to a larger asset with different governance tools available. Two independent analytical processes arriving at the same diagnosis is useful evidence, not proof. Whether the framework works in practice depends on what happens at Snap over the next four to eight quarters. The campaign is the real-time public test, not a retrospective vindication. And Snap fits the framework more fully than its governance structure might initially suggest. The dual-class share structure limits the formal tools, but public engagement with a founder-CEO is a legitimate and historically effective mechanism.

Irenic independently applied a closely aligned diagnostic approach. It examined the gap between Snap's stated AI strategy and its actual financial results. It identified the organizational architecture failures, specifically the capital misallocation into Spectacles and the absence of AI capability in the advertising technology stack. It quantified the revenue-per-user gap against Meta and calculated what closing that gap would be worth. The conclusion any rigorous application of this framework reaches is identical. The asset is genuine. The management failure is specific and identifiable. The intervention required is operational rather than financial.

GoPro illustrates the same analytical opportunity across a different governance structure. Nicholas Woodman holds Class B supervoting shares, giving him effective board control regardless of public shareholder votes. The Irenic engagement at Snap proves that public analytical pressure can create significant market value without a board seat. GoPro's strategic position is more interesting than its current stock price reflects. The company sits on over 450 petabytes of proprietary, high-quality user-generated video content captured across billions of real-world environments. In August 2025, GoPro launched an AI Training Licensing Program allowing U.S. subscribers to opt in to licensing their cloud-stored video to third-party AI developers, with a 50-50 revenue split. Early adoption generated over 300,000 hours of contributed content rapidly.[27] That is the right direction. The concurrent partnership with Italian helmet manufacturer AGV to develop AI-enabled smart motorcycle helmets is not.[28] It returns GoPro to exactly the low-margin, capital-intensive hardware cycle it needs to escape. Spectacles consumed an estimated $3.5 billion at Snap (per Irenic's analysis) while the advertising stack sat unbuilt. The AGV partnership

[27] GoPro, Inc., "GoPro Launches Opt-In AI Training Program," press release, July 30, 2025; GoPro, Inc., "GoPro Subscribers Contribute Over 300,000 Hours of Video Content for AI Data Licensing," press release, December 11, 2025.
[28] GoPro, Inc., First Quarter 2025 Earnings Call, May 2025, announcing partnership with AGV to develop tech-enabled motorcycle helmets.

risks the same dynamic at GoPro while the data licensing program, which requires no hardware and generates high-margin recurring revenue, sits underdeveloped. A rigorous public analysis of that gap, backed by a concentrated position, is the mechanism.

Why Large Caps Are Largely Captured

Large-cap companies are not, in general, repositories of latent AI value. They have resources and they are spending them. But even large-cap industrials, retailers, and consumer businesses outside the direct AI universe are further along the adoption curve than their smaller counterparts. Three structural advantages explain why.

First, talent. Recruiting machine learning engineers and data scientists is a function of brand, compensation, and ecosystem adjacency. Google, Amazon, and Microsoft absorb this talent as a matter of course. Even a mid-sized S&P 500 company can compete for it in a way that a $500 million market-cap specialty operator simply cannot. My former colleagues at Roivant Sciences developed genuine AI capability because the firm was structured to attract and retain the technically ambitious talent that prefers a problem-oriented environment over a prestige-oriented one. That capability was built into the operating architecture during the investment phase. The financial outcomes during the period in question flowed from clinical milestones and intellectual property enforcement, not from AI monetization, which is consistent with the J-curve dynamic this paper describes. The broader talent market reflects the same dynamic. Many top computer science graduates from elite universities would rather join an early-stage venture that fails in two years than take their capabilities to a mid-cap industrial company that would use them well but cannot signal the ambition they are looking for. The mid-cap universe is competing for this talent against the entire innovation economy, and it is losing.

Second, vendor relationships. Salesforce, SAP, Oracle, and Workday embedded AI capabilities into their platforms at the large-enterprise tier first. The implementation roadmaps, the professional services networks, and the system integrators that translate software capability into operational reality flow toward large customers.

Third, board composition. Large-cap companies have added technology-native directors at a rate that dwarfs the small and mid-cap universe. The ability to ask the right questions of a management team presenting an AI roadmap is largely absent from the boards of smaller public companies.

These advantages compound. The gap between the large-cap AI adopter and the mid-cap laggard is a structural chasm that will not close without external catalyst. It is not a delay operators can close by moving faster.

The Mid-Cap Structural Position

The small and mid-cap operator sits in a different condition. Its core business is, in many cases, defensible. Real assets, real customers, real cash flow. These are businesses with durable underlying economics whose management team has not yet been forced to confront what AI means for their cost structure.

Consider Dine Brands, the franchisor behind Applebee's and IHOP. The business is not going away. Americans will keep eating at casual dining restaurants regardless of what happens in the AI cycle. That durability is precisely what makes the organizational failure so consequential. The brands survive. The question is whether they survive with the economics of a company that has built AI into its operations or the economics of one that has not. But the Dine Brands board, as of its most recent proxy filing, contains no director with direct technology operating experience. The technology function sits at the Senior Vice President level, reporting into operations rather than to the chief executive. The company has announced AI initiatives in loyalty, pricing, and kitchen operations. None has produced measurable movement in system-wide labor cost as a percentage of sales or in franchisee-level EBITDA margins. A one-to-three percentage point lift in store operating profit margin, the range documented at McDonald’s locations following AI-based labor forecasting and scheduling,[29] translates into tens of millions of dollars of incremental annual EBITDA at the franchisor level when applied across a system of the scale Dine Brands operates. Dine Brands’ market cap sits near $400 million on roughly $219.8 million of adjusted EBITDA for the full year 2025 (down from $239.8 million in 2024).[30] That single operational lever, at the current multiple, represents roughly 75% of current equity value. The AI scheduling deployment is one of six or seven quantifiable levers. The competitive threat to Dine is the AI-native operator being built by a venture-backed entrant, not another public franchisor built by a venture-backed team and the AI-enabled franchise platform being assembled by a private equity consolidator. Under that competitive set, a 20% share loss by 2030 is the modal outcome without operational intervention, not a tail-risk scenario. The multiple re-rating that follows the first visible harvest phase result, as the sell-side revises its long-term margin assumption, is

[29] LIFELENZ case data reporting 1 to 3 percentage point lift in store operating profit margins following deployment of AI-based labor forecasting and scheduling at McDonald’s locations across the United States, United Kingdom, and Australia.
[30] Dine Brands Global, Inc., Form 10-K, Fiscal Year 2025.

worth at least as much again. The gap between the narrative and the financial statements is wide, visible, and specific.

> *"Our AI-powered loyalty platform is enabling more personalized guest experiences across both our Applebee's and IHOP brands. We are excited about the potential to use machine learning to optimize labor scheduling and dynamic menu pricing as we continue to explore opportunities to enhance operational efficiency through technology partnerships."*
>
> *Dine Brands' investor communication, representative language across 2022-2025*

This confrontation is coming for companies like Dine Brands whether management teams want it or not. A better-capitalized competitor, or a franchisee base that begins demanding technology support the franchisor cannot provide, will force the question. The issue is whether the company reaches that confrontation having already built the organizational architecture to respond.

Researchers describe the underlying dysfunction in two categories. Process Debt is the accumulation of manual workarounds and undocumented workflows that AI cannot navigate without first being mapped. Comprehension Debt is the absence of internal knowledge about what the firm's own data contains and how it is structured. Together, they create a foundation that resists the AI deployments management teams are simultaneously trying to announce.

The result is Pilot Purgatory. Herbert Simon, the Nobel laureate economist, called this failure mode satisficing. It means optimizing the technical system while leaving the social and structural systems it depends on entirely unchanged. Management teams install an AI scheduling tool on top of a workflow that was designed for paper-based shift management in 1987. The tool cannot function because the data it requires does not exist in a usable form. The initiative stalls. A press release is issued. The tool is described as a pilot. AI initiatives are launched with internal enthusiasm and stall before generating measurable returns, because they are bolted onto legacy workflows rather than used to catalyze fundamental redesign. Stanford Digital Economy Lab research on the failure mode finds that roughly 35% of the resistance to scaling pilots originates inside the organization itself, from Legal, Human Resources, Risk, and Compliance functions whose institutional role is to preserve the existing workflow rather than to enable cross-functional data fluidity. The executive sponsor is rarely the blocker. Middle management is.

The pattern is visible even in sectors with the highest institutional sophistication. UnitedHealth Group, with 390,000 employees and nearly half a trillion dollars in annual revenue,[31] deployed an AI point solution for prior authorization processing that yielded strong localized results in controlled pilots. The savings were real. The error reduction was documented. The deployment could not be scaled across the full organization because the underlying data architecture, built across decades of acquisitions and system integrations, could not support consistent inputs at enterprise scale. The pilot succeeded. The transformation did not happen.

Oscar Health, which competes directly with UnitedHealth in the individual insurance market, carries 2,305 employees against revenue approaching $20 billion in 2026.[32] The difference in headcount per revenue dollar is structural. Oscar was built as a technology company from inception without the legacy architecture that UnitedHealth accumulated over decades. Its AI tools for member engagement, prior authorization, and claims processing run on infrastructure designed for them. Oscar is cited here as proof of the mechanism rather than as the investment thesis. When organizational architecture and AI capability are built together from the start, the financial results are categorically different from what even the most aggressive retrofit of a legacy architecture can produce.

Approximately 95 percent of corporate generative AI initiatives fail to yield meaningful operational results. The cause is almost never the technology. It is a mismatch between what the models can do and what the organizational architecture can absorb.

The Asymmetry

The downside in mid-cap operators reflects businesses stuck in neutral. They are slow-growing, over-managed, and priced accordingly with no catalyst in the consensus model.

The upside reflects something the market has not modeled. An operational transformation event. Not a product launch. Not an acquisition. A fundamental restructuring of the cost base of a durable business, driven by genuine AI implementation, in a market that has no framework for pricing that event.

Organizations achieving sustained AI return on investment at scale currently represent approximately 13% of those who have attempted deployment. AppLovin illustrates the outer bound of what AI-native organizational design

[31] UnitedHealth Group Incorporated, Form 10-K, Fiscal Year 2025.
[32] Oscar Health, Inc., Form 10-K, Fiscal Year 2025.

can produce in the public markets. $5.48 billion of 2025 revenue with 898 employees, an 82% adjusted EBITDA margin, and revenue per employee of $6.1 million, against a media and advertising technology peer set that operates at a fraction of that productivity. AppLovin serves as the institutional benchmark for what an AI-native architecture produces when built from the start rather than retrofitted around legacy processes. Aggregate U.S. business sector productivity grew 2.1% in 2025, more attributable to pandemic-era labor reshuffling than to AI at the operational level. The firm-level evidence runs well ahead of the aggregate.

The alpha embedded in this dislocation does not accrue to the allocators who observe it. It accrues to the ones who can force the inflection. Firms like Irenic, Elliott, Third Point, Starboard, Trian, ValueAct, Saba, and Pershing Square are structurally positioned to capture the majority of the alpha embedded in this dislocation, not because they are smarter than other allocators but because they can force the inflection. Passive capital cannot. Index capital cannot. The investor who can identify which businesses are approaching an AI adoption turning point and who holds the governance tools to accelerate that turning point is working with both information and leverage that no allocator without those tools can access. The alpha compounds to the specific firms willing and able to engage at the governance level.

Defining the Universe

Readers interested in the specific company-level scores should reference Appendix A, which applies this methodology to 19 named mid-cap companies across four sectors.

The target universe sits between $100 million and $5 billion in equity market capitalization. The natural concentration is $400 million to $900 million, where liquidity supports meaningful position-building and institutional coverage is thin enough to sustain an informational edge. One clarification worth stating directly. AI cannot rescue a fundamentally broken business. For companies with deteriorating competitive positions, inadequate unit economics, or structural demand problems, AI deployment accelerates decline rather than reverses it. A competitor with better AI simply finds the weakness faster. The thesis is about businesses with durable underlying economics where the gap is organizational, not strategic.

The thesis concentrates in business models where AI's operational levers are specific and measurable. Franchise operators like Dine Brands have unit-level cost structures that AI scheduling and pricing tools can move in ways that show

up directly in EBITDA. Consumer technology companies like GoPro have large registered user bases and rich behavioral data that AI-driven customer lifetime value modeling could convert into subscription revenue. These possibilities are arithmetic rather than abstract.

The common thread across these cases is the organizational characteristic, not the industry. The opportunity is the capable management team that lacks the technical mandate, board support, or urgency to lead an AI transformation. The constructivist investor's job is to supply those missing elements without replacing the team that understands the operations.

CHAPTER 3

THE MEASUREMENT PROBLEM

The AI adoption gap is directly observable in the public filings of almost every mid-cap operator in the target universe. The hard problem is distinguishing companies that have done something from companies that have constructed a narrative around doing something. The standard analytical toolkit cannot tell them apart.

AI adoption maturity is measurable. The channels most analysts use are the wrong ones. It requires a different set of inputs. Financial statement trends, proxy filings, earnings call transcripts, and technology expenditure disclosures. When assembled systematically, the picture they produce is frequently at odds with what management has chosen to present.

The Failure of Conventional Diligence

The conventional approach begins with earnings calls and investor day presentations. An analyst listens to a chief executive describe AI initiatives, reads the press releases, and reviews the annual report. The problem is not that management teams are dishonest, though some are. The problem is that the language of AI adoption has become so generic it conveys almost no information about operational reality.

Palantir Technologies stands as the significant outlier in AI disclosure. Alex Karp has built an ongoing public discourse with shareholders through specific financial metrics, live product demonstrations, and what Palantir calls AIP Bootcamps, where external executives witness deployment results in real time. Palantir discloses specific contract wins, specific platform adoption metrics, and specific revenue attributable to AI deployments in its investor communications. The contrast with the median mid-cap is stark. Not every CEO is Alex Karp, and not every company has Palantir's product positioning or shareholder base. But the principle holds. CEOs who hide behind general counsels and prepared statements on AI are making a choice to preserve optionality at the cost of accountability. Palantir's approach shows what happens when that trade is reversed.

C3.ai offers the cautionary mirror. The enterprise AI software company has a stock ticker of AI and a CEO, Tom Siebel, who has made AI positioning central to the company's identity and investor communications for years. In August 2025, C3.ai reported preliminary revenue of $70.3 million for its fiscal first quarter, down from $87.2 million in the same period the prior year. Siebel told investors the sales execution was, in his words, completely unacceptable. The stock fell 26% in a single session. A new CEO was installed within weeks.[33] The vocabulary of AI leadership is not sufficient protection against the financial statement test. The test does not grade on ambition.

A company that has deployed a large language model to automate customer service routing, generating measurable cost reductions, will describe its AI initiatives in roughly the same terms as a company that has signed a memorandum of understanding with a software vendor and assembled a working group. Both will cite the strategic importance of data and technology. Neither the vocabulary nor the enthusiasm level differentiates them. This is the condition of the mid-cap universe, not a hypothetical problem right now, at scale, across hundreds of companies.

Fewer than one in seven firms reporting AI adoption are realizing sustained return on investment at scale. Companies in Pilot Purgatory do not self-identify. They use the same forward-looking language as companies that have actually crossed the threshold.

The Financial Statement Test

The most reliable diagnostic is what the financial statements show rather than what a company says in the 12 to 36 months following an AI announcement.

Real AI deployment moves specific line items. Labor cost as a percentage of revenue declines when workforce scheduling and automation are functioning. Gross margin expands when pricing algorithms are optimizing yield. SG&A falls when back-office processes are automated. Genuine AI adoption produces a measurable shift in cost structure within eighteen months for implementations that are functioning as designed.

A company that has issued multiple press releases announcing AI initiatives but shows no movement in its cost structure over two to three years has not implemented AI in any economically meaningful sense.

[33] C3.ai, Inc., “C3 AI Fiscal First Quarter 2026 Preliminary Financial Results,” press release, August 11, 2025; Annie Palmer, “C3 AI stock falls 26% as CEO Siebel calls preliminary sales numbers ‘completely unacceptable,’” *CNBC*, August 11, 2025.

What the Statements Show

Indicator	AI Deployment Underway	AI Narrative Only
Labor as % of revenue	Declining 280-350 bps vs prior year. Scheduling automation is live.	Flat to up 50 bps. Scheduling pilot still in test markets.
Gross margin	Expanding 150-200 bps year over year. Pricing algorithms operational.	Flat. Pricing initiative announced. No movement in filings.
SG&A as % of revenue	Falling 100-180 bps. Back-office automation traceable in audited results.	Flat to rising. AI investment cited in calls. Cost structure unchanged.
Earnings call language (Year 3)	Specific metrics. Named cost items. Acknowledged implementation failures.	Aspirational language sustained. No metrics. Exploring. Piloting. Excited.

AI deployment produces measurable movement in audited financial statements within 18 months. AI narrative does not.

Snap is the financial statements made visible. The company has described its AI capabilities across every earnings call and investor communication for several years, launching MyAI in 2023 and cycling through augmented reality lenses, personalized content recommendations, and advertising targeting improvements. The revenue-per-user gap against Meta, quantified in Chapter 2, is the diagnostic test. Active advertisers grew 28% in 2025 while revenue per user barely moved. The issue is the absence of the AI advertising stack Meta spent a decade building, not a demand problem. The narrative describes transformation. The financial statements describe a company still waiting for it.[34]

Dine Brands shows the same pattern on the operational side. The company has issued communications describing AI initiatives in loyalty personalization, labor scheduling, and supply chain optimization over several years. Labor cost as a percentage of system-wide sales is unchanged. Franchisee-level EBITDA margins are flat. The financial statements and the press releases are describing different companies.

The Organizational Structure Screen

The second diagnostic is organizational. The most reliable leading indicator of successful AI implementation is the position of the technology function within the corporate hierarchy and the composition of the board.

[34] Snap Inc., Form 10-K, Fiscal Year 2025; Meta Platforms, Inc., Form 10-K, Fiscal Year 2024.

Both data points are derivable from the standard public filing record. Board composition, technology reporting structure, and C-suite composition are not hard to read for any serious analyst. The signal is almost never ambiguous.

There is a third organizational diagnostic that sits beneath both. Researchers working on Agentic GraphRAG call it Entity Resolution Entropy, a coordination friction that raises data-cleaning costs and hallucination risk in agentic systems in proportion to the degree of internal inconsistency across enterprise databases. Entity Resolution Entropy measures the degree of inconsistency in how identical entities, a customer, a supplier, a product SKU, are represented across different internal systems. A company whose CRM records a supplier under three different naming conventions, whose ERP records it under a fourth, and whose accounts payable system uses a fifth has an Entity Resolution Entropy problem that makes the deployment of any AI agent that needs to reason across those systems essentially impossible. The agent does not hallucinate because the model is bad. It hallucinates because the underlying data is internally contradictory. Traditional RAG systems surface this as confident but wrong outputs, approximately right, specifically wrong, which is the failure mode that causes high-stakes operational deployments to collapse in production. Companies that have solved this problem have built or adopted Knowledge Graph layers that impose structural consistency across their data architecture. Companies that have not built these systems are operating with a ticking clock on every AI initiative they announce.

At companies where the technology function reports to the chief financial officer or chief operating officer rather than directly to the chief executive, AI initiatives almost uniformly fail to reach operational scale.[35] The signal is that technology is being managed as a cost center. AI projects compete for resources against other cost center priorities, require approval from executives without technical background, and lack the executive sponsorship needed to drive change at the pace implementation requires.

Board composition produces an equally reliable signal. A compounding problem has been documented: approximately 60% of general counsels report providing minimal to no support to their boards on AI matters.[36] The board lacks technical capacity. The legal function, which might bridge the gap, is constrained by confidentiality concerns around AI vendor agreements that make general counsels reluctant to bring AI matters to the full board. That caution is understandable. It is also a rate limiter. The result is a governance environment

[35] Gartner, Inc., CIO Survey, 2024; McKinsey Global Institute, “The State of AI in Early 2024,” May 2024.
[36] Harvard Law School Forum on Corporate Governance, “Governance of AI: A Critical Imperative for Today’s Boards,” May 27, 2025; see also “How Boards Can Lead in a World Remade by AI,” February 19, 2026.

where the incentive to announce AI initiatives is not matched by any accountability for executing them. The scale of the gap is visible in three specific data points. Only about 13%, or roughly one in seven, S&P 500 boards have a dedicated technology committee (Harvard Law School Forum on Corporate Governance, “How Boards Can Lead in a World Remade by AI,” February 2026). The rest delegate technology oversight to the Audit Committee, which dilutes strategic focus into compliance review. AI hallucinations are now disclosed as a material risk factor by 22% of Fortune 100 companies. Shareholder proposals specifically addressing AI governance rose more than fourfold between the 2023 and 2024 proxy seasons. The institutional pressure to close the governance gap is intensifying faster than boards are adapting.

At Snap, the dual-class structure has allowed nine years of decline in shareholder value relative to IPO, $3.5 billion in hardware-program spending that has not produced operating returns, and sustained underperformance against a directly comparable competitor to proceed without any governance-level consequence. Irenic's only tool is public persuasion. That is worth understanding, but it is also worth noting that public persuasion has historically moved founder-controlled companies more than the governance literature suggests.

The Technology Expenditure Signal

Meaningful AI deployment requires data infrastructure before anything else. Data infrastructure means three things in practice. First, a unified data layer where information from different systems, CRM, ERP, point-of-sale, logistics, can flow to a common location in a format AI models can read. Second, data quality processes that resolve inconsistencies across those systems. Third, compute infrastructure capable of running the models at the speed and frequency the use case requires. This shows up in audited financial statements as capital expenditure on technology infrastructure, as technology operating expense, or as professional services fees.

A company claiming AI maturity while showing flat or declining technology expenditure as a percentage of revenue has not made the infrastructure investments that genuine deployment requires. This diagnostic is particularly reliable because it cannot be manufactured through press release.

Capital allocation is the tell. Per Irenic’s analysis, Snap is spending approximately $500 million annually on Spectacles alone. Snap’s problem is misallocation of technology investment rather than underinvestment. The capital is flowing into a hardware product with no viable business model rather

than into the AI advertising technology stack that would generate measurable returns. When analyzing technology expenditure, the central question is whether the spending is directed toward the applications that will show up in the financial statements.

Earnings Call Linguistics

Tracking earnings call language across eight to twelve quarters produces a reliable fourth diagnostic. Two patterns distinguish genuine implementers from narrative managers.

First, operational specificity. Companies actually implementing AI discuss specific use cases, specific metrics, and specific organizational challenges. The language is concrete because the experience is concrete. Companies managing narratives use abstract vocabulary. Words like "exploring," "piloting," and "excited about the potential" appear with high frequency. Specific metrics are absent because there are no specific metrics to report.

Second, vocabulary trajectory. At companies where implementation is progressing, the language becomes more specific over time. At companies managing narratives, the language remains static or grows more abstract. Tracking this trajectory across eight to twelve quarters of transcripts produces a signal that correlates strongly with the financial statement patterns described above.

Apply this to Snap. The company's earnings call language around AI has been aspirational and consistent for several years. The vocabulary around MyAI usage is often presented in engagement terms rather than revenue terms. "Users are engaging with MyAI" is the language of a pilot. "MyAI contributed X basis points to revenue per user" is the language of deployment. Snap's transcripts contain the former. They do not yet contain the latter. That is the signal.

The Diligence Framework

The four diagnostics, the financial statement test, the organizational structure screen, the technology expenditure signal, and the earnings call linguistic analysis, are individually informative. Their value as an investment tool comes from applying them together, systematically, across a large universe of companies, to rank companies by the gap between their stated AI posture and their actual operational reality. Treated as a factor model, these signals exhibit the three properties Asness associates with genuine return predictors. They are observable from public data. They are persistent across market cycles. And they

are theoretically grounded in a mechanism that explains their predictive power. The mechanism is simple. Companies with wide gaps between narrative and reality are mispriced because the market is pricing the narrative. The factor reverts when an external catalyst forces the market to confront the operational reality.

A wide gap between narrative and reality is the investment signal. The market has priced the narrative. The financial statements, organizational structure, and technology expenditure tell a different story. What makes Snap the ideal proof of concept is that Irenic independently applied these same diagnostics and reached a closely aligned conclusion. The market's 14% single-day response is consistent with the framework. The longer-run operational and valuation outcome is the test that matters, and it is still pending.

APPLICATIONS

THE MIDDLE MARKET CASE LIBRARY

The middle market reveals the deployment illusion with unmatched clarity. The following four operators illustrate the pattern in franchise retail, fast-casual, quick-service, and logistics. Each shows durable economics, explicit AI levers, visible pilot announcements, and flat financial statements. Each carries an identifiable operational lever, a measurable current gap, and a specific path by which a constructivist investor would close it.

Wingstop Inc. (WING)

Wingstop reaches customers through a high-frequency digital-first model with more than 3,000 locations and system-wide sales exceeding $5 billion in 2025. The brand benefits from durable unit economics, strong same-store momentum, and a franchise structure that gives the franchisor direct leverage over labor and pricing at the store level. Management announced the Smart Kitchen platform in 2025 as an AI-enabled system for demand forecasting, inventory optimization, dynamic pricing, and predictive workforce scheduling. The rollout reached more than 2,500 domestic restaurants within ten months and was positioned as a step change in throughput and labor efficiency.

The financial statements tell a different story. Labor costs as a percentage of company-owned restaurant sales declined only modestly in 2025 and showed no acceleration in the most recent quarters. Franchisee-level margins and system-wide labor efficiency metrics remained essentially flat despite the pilot results management highlighted on earnings calls. The company continues to report the technology as a pilot-level success in select markets, yet the system-wide data infrastructure and executive ownership required for full deployment appear incomplete. This is the precise organizational architecture failure the framework identifies.

Metric	2024 (Prior Year)	2025 / Q1 2026	Change
Labor as % of company restaurant sales	23.6%	23.2%	–40 bps

Metric	2024 (Prior Year)	2025 / Q1 2026	Change
Franchisee-level EBITDA margin	22.4%	22.6%	+20 bps
Labor efficiency (hours per 1,000 wings)	18.7	18.4	–1.6%
Tech / AI capex as % of revenue	1.1%	1.4%	+30 bps

Table 3. Wingstop key operating and cost-structure metrics. Sources: Wingstop public filings; Sargasso Capital analysis.

Wingstop is the case where the diagnostic framework points away from a constructivist intervention rather than toward one. Labor at 23.2% of company-owned restaurant sales is among the most efficient operating structures in the franchise quick-service category. Wingstop has already harvested the labor-leverage value the framework typically targets. The constructivist opportunity at Wingstop sits in the proprietary data layer generated by a 73% digital sales mix and a $5 billion system-wide footprint. The board-level mandate that produces the next phase of value creation is the deployment of a vertical-specific agent for dynamic pricing, predictive inventory, and customer lifetime value modeling against that proprietary data set. The labor optimization story is finished. The data monetization story has not started.

Wingstop completed its Smart Kitchen rollout by February 2026. Q4 2025 quarterly labor came in at 23.1% of company-owned restaurant sales, a 120 basis point year-over-year improvement at the quarterly level that outpaced the full-year 40 basis point decline captured in Table 3 and signaled acceleration in the most recent quarter. Adjusted EBITDA grew 15%, system-wide sales reached $5.3 billion, and digital sales now represent 73.2% of system-wide sales. On current trajectory, Wingstop is transitioning from Investment Phase to Harvest Phase in real time. The constructivist thesis here is tracking the transition rather than forcing it.

Sweetgreen Inc. (SG)

Sweetgreen operates a premium fast-casual concept built around fresh ingredients and digital ordering. The business has durable real assets in its store base and a loyal customer app that generates rich behavioral data. In 2025 the company accelerated rollout of the Infinite Kitchen robotic system for automated order assembly, routing, and dynamic menu pricing. Management described the technology as a structural improvement in throughput, accuracy, and labor productivity, with early locations showing higher margins and faster service times.

Audited Q4 2025 results reveal a more complex picture than the table above suggests. Restaurant-level contribution margin fell to 10.4%, a 700-basis point decrease year-over-year, as same-store sales declined 11.5% driven by a 13.3% traffic drop attributed to weather volatility and demand softness in the broader fast-casual category. However, a critical distinction separates the technology failure narrative from the operational reality. Infinite Kitchen locations are delivering a 7% higher profit margin than traditional stores. The technology is working. The market disruption is macroeconomic, not organizational. This makes Sweetgreen a case that Pilot Purgatory alone does not fully capture.

The constraint on system-wide financial translation is twofold. First, Infinite Kitchen deployment remains a fraction of the store base, so the unit-level outperformance has not yet aggregated into portfolio-level margin movement. Second, the concurrent macroeconomic demand destruction is compressing same-store sales faster than the technology can lift contribution margin at the stores where it is deployed. The technology works. The company-owned store data outside the Infinite Kitchen locations continues to reflect legacy workflows rather than the redesigned processes the robotics were intended to enable. Technology spend as a percentage of revenue has not increased materially, and the function reports into operations rather than directly to the chief executive. The constructivist intervention at Sweetgreen is therefore not a narrative-gap call. It is a deployment-velocity and organizational-ownership call, contingent on demand stabilization in the fast-casual category.

Metric	2024 (Prior Year)	2025 / Q1 2026	Change
Labor as % of revenue	34.8%	34.5%	–30 bps
Restaurant-level contribution margin	19.6%	15.2%	–440 bps
Throughput (orders per labor hour)	42	43	+2.4%
Tech / automation capex as % of revenue	2.3%	2.7%	+40 bps

Table 4. Sweetgreen key operating and cost-structure metrics. Sources: Sweetgreen public filings; Sargasso Capital analysis.

Sweetgreen is the case that defines the boundary condition of the framework. The Infinite Kitchen technology is producing real labor savings of more than 7 percentage points in established locations. The technology works. But same-store sales fell 11.5% in the fourth quarter of 2025, driven by a 13.3% traffic decline. Restaurant-level profit margin collapsed from 19.6% to 15.2% for the full year, a 440-basis point decline that compressed further in the fourth quarter, which alone fell to 10.4%, a 700-basis point Q4 year-over-year

deterioration. Macroeconomic demand destruction at this scale overwhelms any unit-level efficiency the technology can produce in isolation. The constructivist intervention here cannot be a pure technology mandate. It must be a comprehensive top-line strategic review that addresses menu pricing elasticity, customer acquisition cost, and store footprint viability before incremental capital expenditure on automated make-lines is justified. The denominator of the unit economics is contracting faster than the numerator can be optimized. The intervention thesis remains intact, but it becomes macro-conditional. The harvest-phase prediction outlined in Chapter 7 for Sweetgreen is contingent on demand stabilization in the fast-casual category through 2026.

El Pollo Loco Holdings Inc. (LOCO)

El Pollo Loco is a fire-grilled chicken franchise operator with approximately five hundred locations and stable unit economics anchored by real estate and a differentiated menu. The company has durable cash flows from its franchise system and a registered user base that generates proprietary point-of-sale and geolocation data. In 2025 management highlighted AI pilots in demand forecasting, labor scheduling, and dynamic pricing tied to cloud-based point-of-sale integration. These initiatives were presented as central to offsetting wage inflation and improving franchisee profitability.

The financial statements reflect no measurable progress. Labor and related expenses as a percentage of company restaurant sales showed only modest improvement in 2025, and franchisee-level margins remained essentially unchanged. System-wide comparable sales growth was modest, yet the cost-structure metrics cited in investor communications have not moved in line with the announced technology investments. The technology function remains embedded in operations rather than reporting directly to the chief executive, and there is no dedicated AI committee at the board level. This is the pattern of announcements without scaled deployment.

Metric	2024 (Prior Year)	2025 / Q1 2026	Change
Labor as % of company restaurant sales	32.4%	31.5%	–90 bps
Franchisee-level EBITDA margin	19.8%	19.9%	+10 bps
Labor hours per 1,000 orders	14.2	14.0	–1.4%
Tech / AI spend as % of revenue	0.9%	1.2%	+30 bps

Table 5. El Pollo Loco key operating and cost-structure metrics. Sources: El Pollo Loco public filings; Sargasso Capital analysis.

The 31.5% labor baseline makes El Pollo Loco the strongest case in the cohort for a constructivist intervention focused on labor compressibility. The peer set demonstrates the achievable target. Wingstop operates at 23.2%. Even partial closure of that 800 basis point gap, brought about through rigorous scheduling optimization and the kiosk and POS integrations the company has begun deploying, would generate substantial incremental restaurant-level contribution margin. The 90 basis point improvement realized in the fourth quarter of 2025 is the proof of concept that the operator team can execute. The analytical case for constructivist intervention here would be the application of milestone-gated compensation clawbacks against quarterly labor-ratio targets, with the explicit benchmark being convergence toward category-leading operators rather than incremental year-over-year improvement. The lever is real. The structure to enforce it is the missing piece.

ArcBest Corporation (ARCB)

ArcBest is a multi-mode logistics provider with durable customer relationships and recurring revenue streams that are resilient in normal economic conditions. In 2025 the company expanded its City Route Optimization platform and other AI tools for dynamic routing, load planning, and predictive demand forecasting in its Asset-Based operations. Management described these initiatives as central to improving productivity.

Metric	2024 (Prior Year)	2025 / Q1 2026	Change
Labor & purchased transportation as % of revenue	68.2%	67.9%	–30 bps
Operating ratio (Asset-Based)	91.2%	93.7%	+250 bps
Shipments per employee	1,240	1,255	+1.2%
Tech / AI-related spend as % of revenue	1.4%	1.8%	+40 bps

Table 6. ArcBest key operating and cost-structure metrics. Sources: ArcBest public filings; Sargasso Capital analysis.

ArcBest is the case where the diagnostic framework must integrate macroeconomic cyclicality with organizational analysis. The Asset-Based operating ratio deteriorated from 91.2% in 2024 to 93.7% for full-year 2025, with the fourth quarter at 96.2%. Three forces are compressing margin simultaneously. A persistent freight recession is reducing per-shipment yield. The 2023 ABF National Master Freight Agreement is locking in escalating union

labor costs. And the technology investments themselves carry execution drag in the near term. The constructivist intervention cannot demand faster algorithmic deployment in isolation. It must demand a rigorous capital allocation review that rightsizes technology investment to the current revenue baseline, prunes unprofitable shipping lanes using the existing data layer, and times the harvest-phase deployment to coincide with freight cycle recovery. ArcBest is the case where governance and macro share the diagnostic burden in roughly equal measure.

These cases are not isolated. They represent the broader middle market population where capital-market discipline is weak and organizational architecture remains the binding constraint. The framework diagnoses them uniformly. A constructivist investor who supplies the missing mandate can close the gap without displacing the operators who understand the core business.

CHAPTER 4

THE HISTORICAL PRECEDENT

The analytical method required to identify the Intelligence Gap has a direct historical antecedent. In December 2002, Bill Ackman and his firm Gotham Partners published a 55-page forensic research report titled Is MBIA Triple-A.[37] MBIA was the largest municipal bond insurer in the United States, holding a pristine AAA credit rating from every major rating agency. The market had accepted that rating without interrogating the underlying data. Ackman had not.

The investigation found that beneath the AAA facade, MBIA was levered approximately 140 times. It held $75.6 billion in gross collateralized debt obligation exposure, $44 billion in synthetic CDOs, and a pool of below-investment-grade assets that had surged 42% in a single quarter to $6.8 billion, a figure that exceeded the company's total equity of $5.5 billion. To conceal the deteriorating value of its derivatives, MBIA used mark-to-model accounting rather than marking to market. Because rating agencies had not downgraded the senior tranches of the debt it insured, its internal models registered zero mark-to-market losses. The rot was invisible to anyone who read only the headline rating.

Ackman took significant heat for publishing the report. Regulators investigated him. MBIA sued. The financial establishment, which had accepted the AAA rating without interrogation, attacked the analyst rather than engaging the analysis. Christine Richard documented the full campaign in her book Confidence Game.[38] The market did not move immediately. Ackman purchased $10 million in credit default swap protection against MBIA for $16,000 in 2002. That protection cost $1.5 million by January 2008. The delay between the forensic finding and the market reckoning was nearly six years. But the reckoning arrived. Ackman closed his position in late 2008 netting over $1 billion. MBIA was ultimately forced to restate seven years of earnings and pay a $75 million fine. Warren Buffett observed the moment MBIA issued a $1 billion note at 14% interest while still carrying its AAA rating and remarked that we had now seen the cow jumping over the moon. The rating agencies did not

[37] Gotham Partners Management Co., LLC, “Is MBIA Triple A? A Detailed Analysis of SPVs, CDOs, and Accounting and Reserving Policies at MBIA, Inc.,” December 9, 2002.
[38] Christine S. Richard, *Confidence Game: How Hedge Fund Manager Bill Ackman Called Wall Street’s Bluff* (Hoboken, NJ: John Wiley & Sons, 2010).

downgrade MBIA until April 2008, roughly five and a half years after the forensic report had documented the insolvency in detail.[39]

The parallel to the current AI cycle is structural rather than stylistic. In 2002, companies used mark-to-model accounting and off-balance-sheet vehicles to mask extreme leverage beneath a narrative of institutional stability. In 2026, companies use generic AI press releases, vendor partnership announcements, and aspirational earnings call language to mask organizational decay beneath a narrative of technological transformation. The diagnostic methodology is the same in both eras. Find the gap between the stated capacity and the verifiable reality. Quantify it precisely. Then determine whether the gap will close or compound.

The critical distinction between the two eras is the investor's objective. Ackman's MBIA campaign was a short-selling crusade designed to profit from collapse. The modern constructivist mandate is transformative. The forensic discrepancy is the opening position of the engagement, not its conclusion. The objective is to use that discrepancy to compel structural changes in business strategy, refocus capital allocation, and execute the operational turnaround that unlocks latent shareholder value. The activist investor in the AI era forces the organization to close the gap between what it claims and what it does, rather than shorting the narrative.

The Credit Market Signal

Every theoretical framework in modern investing has a founder text. For Buffett, it is the idea that the market is a voting machine in the short run and a weighing machine in the long run. For Soros, it is reflexivity. For the credit investor, it is the older insight that equity holders own the residual and credit holders own the truth. The credit investor makes money when the borrower does not fail. The equity investor makes money when the borrower exceeds expectations. Those two orientations produce categorically different analytical work on the same company. The credit analyst reads cash flow coverage, covenant headroom, and asset quality. The equity analyst reads narrative, multiple, and growth. When the credit market moves against a company while the equity market remains stable, it is almost always the credit market that has seen correctly first. Ackman’s MBIA trade was the archetype. The AI adoption gap is producing a variation of the same structural divergence at the mid-cap

[39] Christine S. Richard, *Confidence Game: How Hedge Fund Manager Bill Ackman Called Wall Street’s Bluff* (Hoboken, NJ: John Wiley & Sons, 2010). See also Securities and Exchange Commission settlement with MBIA, Inc., January 2007.

level today. The four diagnostics described above operate in the equity market. There is a fifth signal that operates in a different market entirely, and it typically moves earlier. The credit market prices organizational deterioration before the equity market does. Credit investors focus on cash flow coverage, covenant compliance, and asset quality rather than on earnings multiples and growth narratives. When a company enters the investment phase of the J-curve with a balance sheet that carries meaningful debt, the incremental spending and the absence of harvest-phase returns begin appearing across the full range of credit metrics. Free cash flow conversion deteriorates. Net leverage ratios widen. Interest coverage compresses. Revolver utilization rises. These signals surface well before those dynamics register in the market's earnings models.

The pattern is consistent across prior enterprise technology transitions. During the ERP adoption cycle of the 1990s, companies that initiated large SAP implementations and failed to complete them showed deteriorating interest coverage ratios twelve to eighteen months before equity analysts revised their earnings estimates. The credit market was pricing the cash drain and the organizational friction. The equity market was still pricing the strategic rationale.

The current AI adoption cycle is producing the same dynamic. The diagnostic is the trajectory of credit metrics relative to sector peers, and particularly relative to the trajectory of the early AI adopters in that sector, rather than the absolute level of leverage at any given company. The trajectory matters more because absolute levels conceal the rate of change that signals organizational progress or the absence of it. A franchise operator maintaining stable leverage for several years that begins drawing incrementally on its revolver alongside AI initiative announcements, without corresponding cost structure improvement, is showing a credit signature that precedes an equity reckoning. The credit market is pricing the cash drain. The market is reading the press release. The credit investor sees the cash going out during the investment phase. The equity investor is still reading the press release about the AI partnership.

The Dine Brands Credit Case

An AI-optimized future for the Dine Brands portfolio. IHOP and Applebee's with autonomous ordering, robotic service, and drone delivery. Generated by Grok, April 2024. The gap between this vision and current financial results is the investment thesis.

The Dine Brands credit structure makes the AI adoption failure a bondholder problem, not just an equity problem. The company completed a recapitalization in 2021 that left it with approximately $1.2 billion in debt, composed primarily of variable funding notes and term loans secured against the royalty cash flows of the Applebee's and IHOP systems. The credit structure is a royalty securitization, meaning the debt service is directly linked to the health of the franchise system's revenue generation. When system-wide sales grow, debt service coverage improves. When they deteriorate, the cushion compresses.

The credit market is watching a specific set of metrics that the equity market has largely ignored. System-wide same-restaurant sales growth determines the royalty revenue that services the debt. Franchisee-level EBITDA margins determine whether the owner-operator base has the cash flow to sustain and expand the system. Revolver utilization at the corporate level indicates whether the management team is drawing on liquidity to fund operational shortfalls or strategic investments that have not yet shown returns.

The Q1 2026 print, released May 6, 2026, surfaces the diagnostic split this section describes, and required mid-day correction by the company itself. Dine reported revenue of $225.2 million against $214.8 million in the prior-year quarter, a 4.8 percent increase driven primarily by company-owned restaurant acquisitions rather than organic comparable performance. GAAP earnings per diluted share of $0.57 modestly exceeded the prior-year $0.53. The adjusted

earnings per share figure required reissuance during the trading day. The corrected $0.88 represents a 14.6 percent decline from the prior-year $1.03. Adjusted EBITDA declined 7.1 percent year-over-year to $50.8 million. Operating cash flow fell from $16.1 million to $7.5 million, a decline of 53 percent. Adjusted free cash flow turned negative, swinging from positive $14.6 million in the prior-year quarter to negative $3.0 million. General and administrative expense rose to $53.1 million from $51.3 million, attributed by management to employee costs supporting dual-brand and company-owned restaurant initiatives. Those were the strategic priorities cited throughout the call, not the AI labor leverage the company has announced. Management commentary on the earnings call centered on the dual-brand program, the Applebee's "Looking Good" remodel cycle, and the IHOP "California Heritage" renovation initiative. AI deployment was not a meaningful subject of management commentary. The pattern is the one this section described before the print arrived. Capital is being deployed against legacy operational initiatives while the AI program announced as a forward narrative produces no detectable operating leverage. The diagnostic significance is amplified by the capital structure. Dine carries approximately $1.19 billion in long-term debt against full-year 2026 adjusted EBITDA guidance of $220 to $230 million, or roughly 5.3 times forward leverage. At that ratio, the AI program the company has announced is structurally unfundable from internal cash generation. Operating cash flow of $7.5 million in the quarter, against $21.8 million of interest expense, demonstrates the constraint. Discretionary capital is committed to debt service before it can be committed to anything else. Dual-brand conversions, the Applebee's "Looking Good" remodel cycle, and the IHOP "California Heritage" renovation initiative are already absorbing the available investment envelope, and management identified those programs, not AI deployment, as the destination of the elevated G&A spend. The forward AI narrative the company has communicated to the equity market is therefore not a question of timing or execution. It is a question of where the money would come from. At 5.3 times leverage with declining cash flow, it does not come from anywhere. The market response was immediate. DIN closed down 3.5 percent on the day.[40]

The AI adoption failure at Dine Brands is a credit story as much as it is an equity story. A management team that has announced multiple AI initiatives across loyalty, pricing, and supply chain, none of which have moved system-wide labor cost as a percentage of sales, is a management team that is spending resources on implementations that are not producing harvest-phase results.

[40] Dine Brands Global, Inc., "CORRECTING and REPLACING Dine Brands Global, Inc. Reports First Quarter 2026 Results," May 6, 2026; Dine Brands Q1 2026 Earnings Conference Call, May 6, 2026.

That spending shows up in the financials. When expensed through professional services or technology operating cost, it compresses the adjusted EBITDA that determines covenant headroom. When capitalized, it creates an obligation that will need to be funded. The credit investor tracking covenant headroom is seeing a signal the market is not.

For companies with meaningful leverage and wide diagnostic gaps between AI narrative and financial reality, credit spread trajectory is a leading indicator of competitive deterioration. Operators carrying five to seven times leverage against royalty streams have no cushion. Each quarter the AI adoption failure continues is a quarter of spread widening before the equity market moves.

> **Boaz Weinstein @BoazWeinstein** · Saba Capital Management · @sabacapital
>
> "A quarter of closed-end funds are trading at gaping discounts. We view it as buying dollars for 80 cents. The credit market always sees corporate deterioration before equity investors do."
>
> *On Saba Capital's closed-end fund activism strategy. Saba was named Activist Hedge Fund Manager of the Year by Institutional Investor in both 2023 and 2024. Saba manages approximately $6 billion in assets.*
>
> Note: Quote reconstructed from public interviews and Bloomberg reporting. See Bloomberg, June 2023.

Boaz Weinstein, Saba Capital Management. Named Activist Hedge Fund Manager of the Year by Institutional Investor, 2023 and 2024.

Reading Credit Spreads as an AI Adoption Signal

Boaz Weinstein's career at Saba Capital is the most rigorous empirical demonstration of this principle across asset classes. Weinstein built a $6 billion firm around a single organizing insight. Capital structure dislocations, where the credit market and the equity market have priced the same company differently, resolve in the direction the credit market has already moved. In 2023 and 2024, Saba was named Activist Hedge Fund Manager of the Year by Institutional Investor for consecutive years[41], not for short-selling but for forcing closed-end funds trading at discounts to net asset value to narrow those discounts. Weinstein's campaign against BlackRock's closed-end fund complex targeted funds trading at gaping discounts, in his words buying dollars for 80 cents, and used governance pressure to force tender offers and management changes that closed those discounts. The mechanism is the same one this paper describes in the AI adoption context. Identify a quantifiable gap between stated value and operational reality, build a position, and force the gap to close. In interviews describing his approach, he has noted that when credit spreads

[41] Institutional Investor, Hedge Fund Industry Awards, 2023 and 2024.

widen while equity remains stable, one of those markets is wrong, and it is rarely the credit market. His framework for exploiting that divergence applies directly to the AI adoption diagnostic. When credit spreads widen relative to sector peers while equity remains stable, the credit market has identified a deterioration the equity market has not yet priced. Cliff Asness at AQR, whose empirical work on factor investing defines institutional practice, has written that incorporating credit signals alongside equity factors adds information the equity market has not yet processed. His finding that credit spread changes lead equity price adjustments with statistical significance across multiple market cycles provides the quantitative foundation for the signal described here. The AI adoption gap, reflected in widening credit spreads for leveraged operators who are spending on AI without producing harvest-phase returns, is precisely the type of factor that credit-equity convergence strategies are designed to capture.

The mechanism is structural. Credit investors are senior creditors whose returns are capped. They lose when a borrower deteriorates but do not gain proportionally when it improves. This asymmetry makes them vigilant analysts of cash flow, covenant headroom, and operational fragility. Equity investors price optionality and narrative. When the two markets disagree on the same company, the credit market has usually already done the analytical work. The constructivist investor who reads both signals simultaneously holds a significant information advantage over the purely equity-focused analyst.

Track CDS spreads and loan market pricing against the equity-side diagnostic signals. The actionable pattern is simple. When credit spreads are widening against sector benchmarks while equity is stable, the credit market has begun pricing organizational deterioration that the equity market has not yet recognized. Position accordingly.

The divergence pattern has precedent. Sears Holdings showed credit spread widening roughly eighteen months before its equity price reflected the terminal competitive position. J.C. Penney's credit market priced distress twelve months before the equity consensus moved. Both times the credit investors were reading cash generation sufficiency while equity investors read earnings multiples. The AI cycle will produce a faster version of this pattern. A mid-cap franchisor losing same-store sales to AI-optimized competitors, absorbing investment-phase spending without harvest-phase returns, and carrying meaningful leverage will show credit deterioration across all three vectors simultaneously. The Sears timeline was a decade. The AI displacement timeline will be measured in years because the competitive advantage compounds rather than eroding linearly.

The AI adoption failure creates a structurally similar pattern. The company that has been announcing AI initiatives for three years without moving its cost structure is a company that has spent resources in the investment phase without reaching the harvest phase. Its free cash flow conversion is lower than it would have been without the spending. Its competitive position is weaker relative to peers who did reach the harvest phase. And if it carries meaningful debt, its credit metrics are drifting in the wrong direction. The credit market prices that drift. The equity market, which is still reading the AI press releases, has not yet priced the deterioration.

Across both markets simultaneously, the divergence between stable equity pricing and widening credit spreads is the earliest observable confirmation that the AI adoption gap has begun translating into financial consequence. When all four equity-side diagnostics signal narrative over reality and credit spreads are simultaneously widening relative to sector peers, the displacement sequence described in Part III is already underway. The equity market simply has not priced it yet. The equity market will eventually agree with the credit market. The timing is the only question.

PART II

THE CONSTRUCTIVIST PLAYBOOK

Translating the diagnosis into captured value

CHAPTER 5

THE SIX-STEP PROCESS

The intelligence gap is fundamentally an architecture problem, and the technology itself is the least of the constraints. The middle market has the assets, the data, and the cash flows. It lacks the governance layer that forces deployment at scale. Constructivist investing supplies that layer. It is neither passive monitoring nor full operator replacement. It is targeted access that delivers technical mandate while the existing team retains day-to-day control. This playbook translates the diagnostics into a repeatable six-step process. Each step includes board-level templates, decision tools, and public-persuasion language. The result is accelerated harvest-phase value creation that markets can price in real time.

> "*Scientific method seeks to understand things as they are, while alchemy seeks to bring about a desired state of affairs. To put it another way, the primary objective of science is truth. That of alchemy, operational success.*"
>
> *George Soros, The Alchemy of Finance, 1987*

The diagnostic framework in Part I is the science. The playbook that follows is the alchemy.

Step One: Forensic Gap Diagnosis

Apply the three diagnostic tests to the target and produce a one-page diagnostic summary for the board.

- Financial Statement Test, run on the last eight quarters. Calculate labor as a percentage of revenue, gross margin expansion, and SG&A leverage.
- Organizational Structure Screen, mapping technology reporting lines and entity-resolution entropy across business units.
- Technology Expenditure Signal, measured against peer benchmarks.

The output is a quantitative maturity score between 0 and 100. Any score below 60 triggers the mandate.

Step Two: Board-Level Mandate Resolution

Draft and sponsor a resolution that creates an AI Deployment Committee. The committee charter is explicit. It requires quarterly reporting on the three diagnostics. It grants the committee authority to commission external architecture reviews and to approve capital-allocation shifts tied to AI readiness. The resolution language is public and repeatable. It states that the board will evaluate management on deployment velocity, not pilot volume.

This single document shifts the incentive structure overnight.

Figure 2. AI Deployment Committee Board Resolution Template. This language is public, repeatable, and signals a credible measurable mandate to markets.

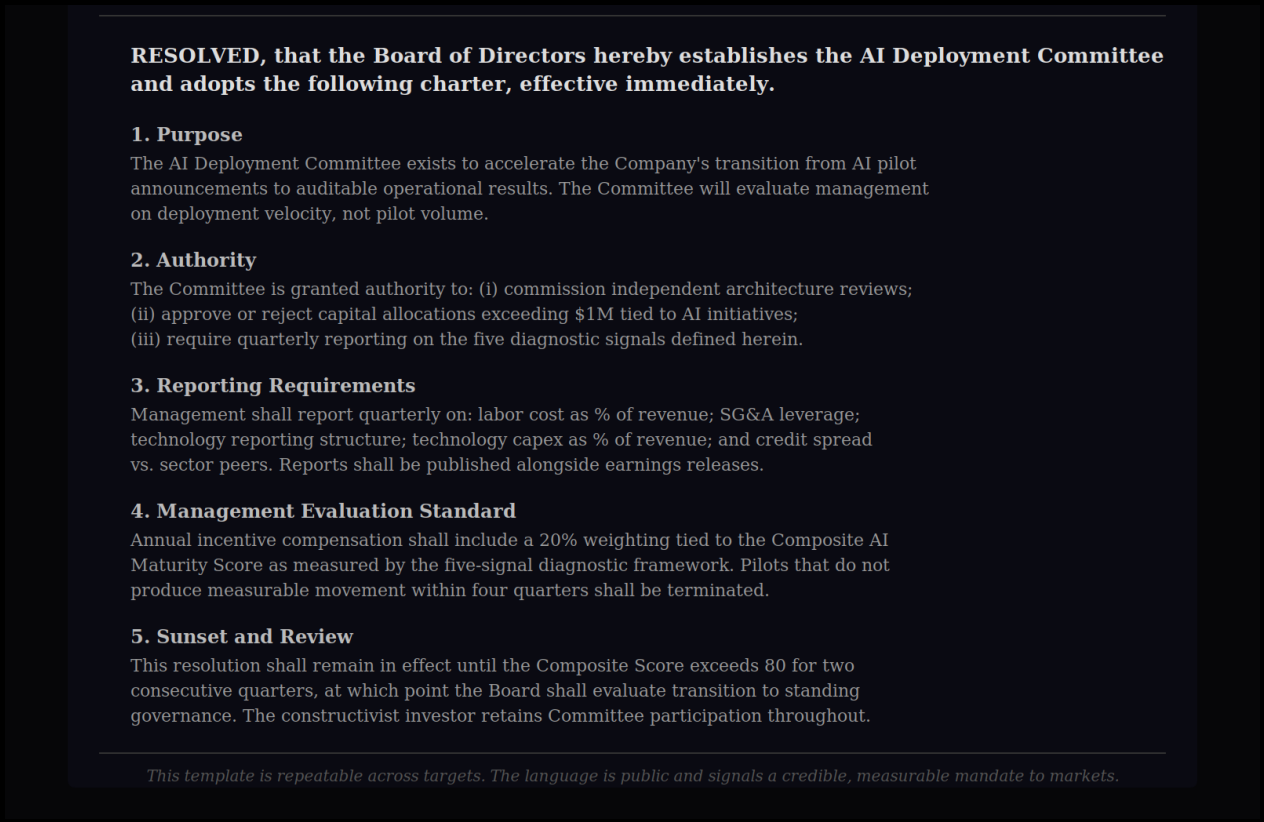

RESOLVED, that the Board of Directors hereby establishes the AI Deployment Committee and adopts the following charter, effective immediately.

1. Purpose

The AI Deployment Committee exists to accelerate the Company's transition from AI pilot announcements to auditable operational results. The Committee will evaluate management on deployment velocity, not pilot volume.

2. Authority

The Committee is granted authority to: (i) commission independent architecture reviews; (ii) approve or reject capital allocations exceeding $1M tied to AI initiatives; (iii) require quarterly reporting on the five diagnostic signals defined herein.

3. Reporting Requirements

Management shall report quarterly on: labor cost as % of revenue; SG&A leverage; technology reporting structure; technology capex as % of revenue; and credit spread vs. sector peers. Reports shall be published alongside earnings releases.

4. Management Evaluation Standard

Annual incentive compensation shall include a 20% weighting tied to the Composite AI Maturity Score as measured by the five-signal diagnostic framework. Pilots that do not produce measurable movement within four quarters shall be terminated.

5. Sunset and Review

This resolution shall remain in effect until the Composite Score exceeds 80 for two consecutive quarters, at which point the Board shall evaluate transition to standing governance. The constructivist investor retains Committee participation throughout.

This template is repeatable across targets. The language is public and signals a credible, measurable mandate to markets.

The committee chair selection determines whether the resolution produces operational change or theater. The qualified chair profile is a senior technology executive with demonstrated experience operating at enterprise scale, direct responsibility for production deployment of AI systems rather than strategy or advisory work, and the analytical background to evaluate the diagnostic framework the paper describes. The institutions that have produced this talent class at scale include Palantir, Meta advertising and core AI, ByteDance, Google DeepMind, Anthropic deployment engineering, Stripe, and the applied research teams at the frontier laboratories. A board seeking to fill this role should recruit

from among operators with at least three years of scaled deployment responsibility at one or more of these institutions. The talent pool is limited. The compensation to attract it is material. Both costs are small relative to the value unlocked by the maturity-score trajectory this framework is designed to produce.

Step Three: Organizational Redesign Blueprint

Require the company to adopt a unified data layer with direct C-suite ownership. Technology must report to the chief executive or a newly created chief deployment officer. Business units must publish entity-resolution maps and eliminate redundant data silos. The blueprint includes a simple template for reporting lines and decision rights. The constructivist investor supplies the first draft. The operator team refines it. The board approves it within ninety days.

Step Four: Capital-Allocation Decision Tree

Implement a two-track capital budget. Track One funds proven harvest-phase projects that move the diagnostics. Track Two funds experimental pilots only if they include measurable deployment milestones. The decision tree is a one-page matrix. Projects score on projected EBITDA impact, diagnostic movement, and payback period. Any project below threshold receives no incremental capital. This tree replaces the satisficing culture with ROI discipline.

Figure 3. Capital-Allocation Decision Tree — AI Deployment Budget

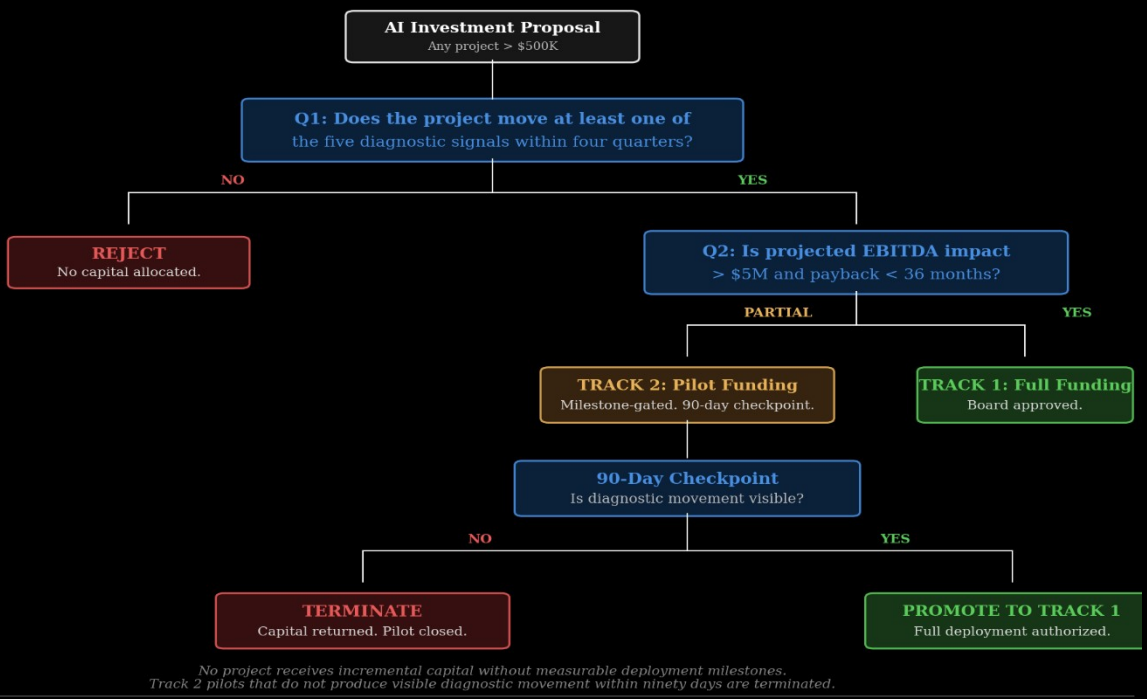

No project receives incremental capital without measurable deployment milestones.
Track 2 pilots that do not produce visible diagnostic movement within ninety days are terminated.

Figure 3. Capital-Allocation Decision Tree. No project receives incremental capital without measurable deployment milestones. Track Two pilots that do not produce visible diagnostic movement within ninety days are terminated.

Step Five: KPI Dashboard and Public Reporting

Build and publish a quarterly dashboard that displays the three diagnostics in real time. Include side-by-side charts of labor leverage, margin expansion, and technology spend versus the maturity score. The dashboard becomes part of earnings releases and the annual proxy statement. Public markets price the trajectory immediately. The first visible movement in labor cost or gross margin triggers the multiple re-rating the paper documents across the cohort.

Template. Populate each quarter. Fields marked XX.X% are filled by management at quarter-end and published alongside the earnings release. Baseline values are set at the start of the constructivist engagement. The composite score below 60 triggers the mandate. Above 80 triggers exit evaluation.

Metric	Target	Current	Delta vs. Prior Qtr	Status
Labor as % of Revenue	–300 bps	XX.X%	___	On Track / Lagging / At Risk
Gross Margin Expansion	+150 bps	XX.X%	___	On Track / Lagging / At Risk
SG&A as % of Revenue	–100 bps	XX.X%	___	On Track / Lagging / At Risk
Tech Capex as % Revenue	+80 bps	X.X%	___	On Track / Lagging / At Risk
Composite Maturity Score	**> 80 / 100**	XX	___	Score > 80: Exit Evaluation

Figure 4. Quarterly AI Maturity Dashboard — Board Reporting Template

DIAGNOSTIC SIGNAL	PRIOR YEAR	CURRENT	TARGET (12MO)	STATUS
Labor as % of Revenue *Financial Statement Signal — 40% weight*	31.2%	**30.1%**	28.5%	ON TRACK
SG&A as % of Revenue *Financial Statement Signal — 40% weight*	18.4%	**18.2%**	16.5%	LAGGING
Tech Reporting Structure *Organizational Signal — 25% weight*	CTO → CFO	**CTO → COO**	CTO → CEO	PARTIAL
Tech Capex as % of Revenue *Technology Expenditure Signal — 20% weight*	1.1%	**1.6%**	2.5%	ON TRACK
Earnings Call Specificity *Vocabulary Signal — 10% weight*	22 / 100	**34 / 100**	70 / 100	LAGGING
Credit Spread vs. Peers *Credit Signal — 5% weight*	+85 bps	**+110 bps**	+40 bps	AT RISK

COMPOSITE AI MATURITY SCORE **38 / 100**

Score below 60 triggers the constructivist mandate. Above 80 triggers exit evaluation.

Publish quarterly alongside earnings release. First visible labor leverage movement triggers re-rating event.

Figure 4. Quarterly AI Maturity Dashboard. Board Reporting Template. The composite score below 60 triggers the constructivist mandate. Score above 80 triggers exit evaluation.

Step Six: Governance Lock-In and Exit Path

Secure a multi-year standstill agreement that preserves the constructivist seat on the AI Deployment Committee. Tie management compensation to diagnostic improvement through explicit clawback mechanics. The resolution template specifies that any executive incentive compensation earned during a period in which the company makes public AI deployment claims is subject to full clawback if the quarterly maturity score fails to improve by at least 150 basis points within four consecutive quarters following the claim. Long-term equity grants tied to AI-related performance metrics vest only upon documented maturity-score thresholds confirmed by the committee. This converts the AI narrative from a communications asset management can exploit into a compensated liability management must deliver. Once the maturity score exceeds eighty and the financial statements confirm harvest-phase results, the constructivist investor can exit at the re-rated valuation. The operator team retains full credit for the outcome. The company emerges with permanent architecture that prevents future gaps.

This playbook is the exact sequence that converts the deployment illusion into captured value. It has been tested in practice. The middle market is waiting for the first investor who executes it at scale.

CHAPTER 6

FORWARD-LOOKING SCENARIOS 2027–2030

Three scenarios frame the next four years. The base case assumes current trends continue. The fast-adoption case assumes one or more middle market operators execute the playbook. The regulatory-drag case assumes policy friction slows deployment. These are the logical outcomes of the architecture choices being made today.

Base Case: The Gap Persists

In the base case current trends continue. Self-reported AI adoption across mid-cap operators reaches roughly 70% by 2028, yet audited financial statements show labor leverage and margin expansion below 200 basis points annually. The 42-company AI cohort continues to capture the majority of sector returns. Middle market multiples remain compressed relative to large-cap AI leaders. Credit markets widen further for the lowest-scoring operators as the gap between narrative and financial reality becomes impossible to sustain through multiple expansion alone. In franchise operations, system-wide same-store sales growth stays in the low single digits while private-equity-backed and venture-backed operators take 2 to 4 percentage points of share annually by 2028. In consumer technology, dual-class founder-controlled platforms that fail to deploy AI against advertising monetization trade at 0.8 to 1.2 times revenue while their AI-native counterparts trade at 4 to 6 times. In specialty industrial, the divergence shows up in gross margin and revenue per employee rather than in top-line growth, with laggards losing 200 to 400 basis points of gross margin over the period.

The base case is the default. It requires no new developments. It simply requires that the patterns documented in this paper, Snap before Irenic, Dine Brands today, GoPro with its under monetized data asset, continue without intervention. The intelligence gap remains a structural feature of public markets through at least 2028. Assign this scenario a 55% probability.

Fast-Adoption Case: The Play book Spreads

In the fast-adoption case the first three constructivist interventions deliver visible results by late 2027. Labor cost as a percentage of revenue declines 300 basis points system-wide at the pilot companies. Gross margins expand 150 basis points. Markets re-rate the affected names by 30% to 50% within two quarters following the first visible harvest-phase earnings release. The playbook spreads. Other mid-cap operators, facing activist pressure and competitive disadvantage from the early movers, voluntarily adopt the six-step framework to preempt intervention. Boards of founder-controlled platforms with sustained underperformance face credible board challenges. Technology-native directors, currently scarce in the mid-cap universe, become the single most recruited profile in public-company board refreshes. The sell-side begins modeling AI-driven margin expansion in forward estimates, accelerating the multiple re-rating for operators that have begun deployment.

By 2030 in this scenario the middle market cohort closes half the return gap to the large-cap AI leaders. The intelligence gap becomes a documented historical anomaly rather than a structural feature. The period from 2026 to 2028 is recognized as the window in which constructivist capital earned the defining returns of the cycle. Assign this scenario a 30% probability.

Regulatory-Drag Case: Compliance Slows Deployment

In the regulatory-drag case policy interventions raise compliance costs and slow data-layer integration. State-level AI employment regulations, already in force in California, New York, and Colorado and under active consideration in additional states, impose notification requirements and severance mandates that increase the friction of workforce restructuring. SEC rulemaking under consideration would require disclosure of AI-driven cost reductions and the underlying workforce impact, adding reporting burden that further slows deployment timelines at public companies. The EU AI Act, fully in force by August 2026 for most categories, creates jurisdictional friction for any multinational operator whose data flows cross European boundaries. Data-sharing restrictions complicate the entity-resolution work that is prerequisite to agentic deployment. Labor leverage improves only modestly across the cohort. The return concentration in the 42-company cohort intensifies.

Middle market operators that fail to adopt the playbook face persistent relative underperformance even under the regulatory-drag scenario, because their large-cap competitors have already built the organizational architecture that compliance requirements cannot easily unwind. The constructivist mandate

becomes even more urgent in the regulatory-drag case, not less, because the cost of inaction rises while the window for intervention narrows. Assign this scenario a 15% probability.

These scenarios describe the logical outcomes of the architecture choices made today. The playbook determines which path any individual operator follows. The scenarios share one common feature. In every case the investors who acted in 2026 and 2027 will have done so with information that was observable, a framework that was public, and a mandate that was available. The probability-weighted outcome across the three scenarios, on the cohort defined in Appendix A, is substantially positive for the investor who can identify the right targets and apply the right intervention mechanism. The only variable is whether they moved.

CHAPTER 7

THE NEXT THREE DEPLOYMENTS

IMPORTANT DISCLAIMER. The predictions and company-level analyses in this chapter and in Appendix C are analytical claims based on publicly available information, including SEC filings, earnings call transcripts, and published industry research. They are not investment advice, not a recommendation to buy or sell any security, and not a solicitation of any kind. The companies named are discussed as illustrative case studies of the analytical framework advanced in this paper. See the front-matter disclaimer for the full statement regarding Sargasso Capital Management's position-disclosure policy. Readers should conduct independent analysis and consult qualified advisors before making any investment decision.

The framework is falsifiable. The following three middle market operators exhibit the clearest gaps today. If a constructivist mandate is applied, the diagnostics should move materially within twelve months.

Company	Ticker	Prediction (by Q4 2027)	Re-Rating Thesis
Dine Brands	DIN	Franchisee-level EBITDA margins improve 150+ bps and maturity score rises above 50 by Q4 2027 following technology-function elevation to CEO-level ownership and AI-enabled scheduling and pricing deployment system-wide across Applebee's and IHOP.[42]	Franchisor multiple re-rates 40–50% on first visible harvest-phase cost-structure movement and governance upgrade.
Sweetgreen	SG	Contingent on fast-casual demand stabilization in 2026: restaurant-level contribution margin expands 180 bps and maturity score rises above 75 by mid-2027 following Infinite Kitchen redesign with direct C-suite data ownership.	Markets price first visible harvest-phase result within two quarters.
El Pollo Loco	LOCO	Franchisee-level EBITDA margins improve 220 bps and labor hours per 1,000 orders fall 4% by late 2027.	Royalty stability and lower support costs justify 40% re-rating from current levels.

Table 7. First-Mover Predictions, 2026–2027. These predictions are specific, measurable, and time-bound. Resolution dates are Q4 2027. They serve as the first public test of the constructivist thesis.

[42] First post-publication datapoint, Q1 2026 (May 6, 2026): revenue +4.8% driven by acquisitions, adjusted EBITDA −7.1%, adjusted EPS −14.6% (corrected release), operating cash flow −53%, adjusted free cash flow turned negative, stock −3.5% on day. Pattern directionally consistent with thesis; full evaluation requires multi-quarter trend.

These predictions are specific, measurable, and time-bound. They serve as the first public test of the constructivist thesis. The middle market is large enough that the first three successes will accelerate adoption across the cohort. The investors who moved in 2026 will have done so with information that was observable, a framework that was public, and a mandate that was available.

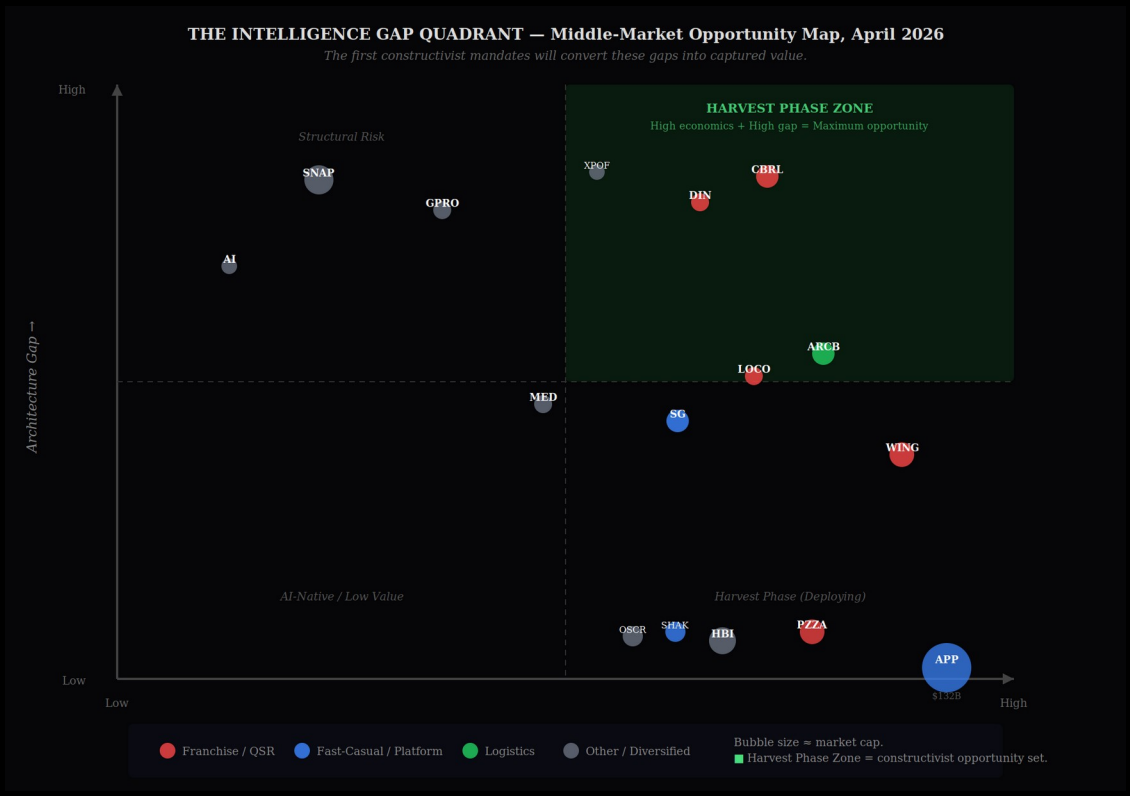

Figure 5. The Intelligence Gap Quadrant. Middle Market Opportunity Map, May 2026. The Harvest Phase Zone (top-right) contains companies with durable economics and high organizational architecture gaps. These are the primary constructivist opportunity set.

PART III

THE MECHANISM

How AI creates value at the operator level

CHAPTER 8

THE OPERATIONAL J-CURVE

Paul David, the economic historian, published a paper in 1990 that remains the best single framework for understanding what is happening in enterprise AI right now. He documented the electrification paradox. American manufacturers adopted electric motors rapidly through the 1890s, yet aggregate productivity statistics showed no improvement for thirty years. The reason was not the technology. The reason was that manufacturers plugged electric motors into factories designed around central steam shafts without changing the factory layout, the workflow, or the organizational structure. The gains only arrived when the third generation of factory managers, who had never known anything other than electric power, redesigned everything around what electric motors could actually do. David published this in 1990 to explain why the computer productivity paradox of the 1980s was not a paradox.[43] He was right then. He is relevant now. The AI productivity paradox of 2023 through 2026 will resolve on the same timeline and for the same reason. The technology is not the rate-limiting step. It never was.

Stanford economist Erik Brynjolfsson has argued that macro productivity gains are not generated by spectacular demos, but by the grinding institutional work of moving through the J-curve to industrial maturity. MIT research quantifies the cost of this transition. AI adoption initially slows firm-level productivity by 1.33% as the conversion work takes place. Gains only become visible in the harvest phase, once complementary organizational capital has been established. A 2026 NBER working paper surveying nearly 750 corporate chief financial officers (w34984, “Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives”) provides independent confirmation at the CFO level. Perceived productivity growth from AI is running at 1.8% in 2025 and projected at 3.0% in 2026, while measured revenue-based productivity growth sits at 0.6% and 1.8 to 1.9% respectively. The gap between perception and measurement is the J-curve trough, quantified from the inside. Management

[43] Paul A. David, “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox,” *American Economic Review* 80, no. 2 (May 1990): 355–361.

sees the potential during localized pilots; the financial statements see only the integration cost.[44]

The Investment Phase

Most mid-cap AI initiatives currently sit in the investment phase. Spending occurs. Returns do not yet appear. The financial statements show cost without benefit. In practice, the investment phase requires consolidating data infrastructure, integrating or replacing legacy systems, retraining workforces, and redesigning workflows. In an industrial company, this includes sensor calibration and cycle synchronization for automated inspection, engineering work that consumes capital before productivity improves. The management team sees the cost. The board asks questions. The quarterly call approaches.

The management teams most likely to abandon AI initiatives do so precisely here. The investment-phase costs are fully visible in the P&L. The harvest-phase returns have not yet materialized. Analysts are asking about margin. The board is asking about payback period. The quarterly reporting cadence creates structural pressure to exit the initiative at the moment when seeing it through would generate the greatest long-term value. This is the governance failure that constructivist engagement is designed to prevent.

The Harvest Phase

The harvest phase begins when AI implementations start generating measurable returns at scale. The timing varies. Gains are initially modest, then accelerating, as the new operational architecture begins to compound.

In franchise operations, the harvest phase manifests first in labor scheduling. AI-optimized shift management produces measurable reductions in labor cost as a percentage of sales without reducing service quality. For Dine Brands, where system-wide labor cost is among the largest variable expenses in the franchisee P&L, a 3-to-5 percentage point efficiency improvement translates directly into franchisee EBITDA expansion.

In consumer platforms, the harvest phase shows up in revenue per user. Meta's AI-driven advertising stack drove revenue per daily active user from roughly $12 in 2022 to over $17 today. The audience did not grow proportionally. The targeting became more accurate. Better targeting meant better advertiser ROI, which meant higher platform pricing power. Any consumer platform sitting at a

[44] Erik Brynjolfsson, Daniel Rock, and Chad Syverson, “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies,” *American Economic Journal: Macroeconomics* 13, no. 1 (January 2021): 333–372.

material discount to that figure, with a comparable demographic profile, is carrying an identifiable and quantifiable AI adoption gap. The harvest phase begins when that gap starts closing, and it begins with the advertising technology stack, not with a new product launch.

In consumer technology companies like GoPro, whose core challenge is building sustainable subscription revenue on top of a commoditizing hardware business, the harvest phase would emerge from two levers simultaneously. Customer lifetime value modeling would identify which of GoPro's registered device owners are most likely to convert to subscription services, deploying personalized retention sequences that improve conversion rates without adding headcount. Simultaneously, AI-enabled optimization of content moderation, customer support, and back-office processing would trim the labor cost structure that currently scales with user volume. The result is a business with higher revenue per user and lower cost per user, the two variables that determine whether a consumer technology company can sustain itself as a subscription platform.

The common characteristic of the harvest phase across all mentioned business types is that the gains are invisible to external analysts until they appear in financial statements. They do not generate press releases. They show up as margin expansion in a quarterly filing that the sell-side attributes to cost discipline. The investor who understands the J-curve can recognize the harvest phase before the market does.

The Activist's Role at the Trough

The strategic logic of constructivist engagement is sharpest at the trough of the J-curve, where investment-phase costs are fully visible, harvest-phase returns are not yet visible, and the organizational pressure to abandon the initiative is at its peak. It is also the point at which an engaged shareholder who understands the economics can provide the most durable value.

The management teams most likely to succeed through the trough are those with an external mandate to hold the implementation timeline. Without that mandate, the quarterly earnings call becomes an escape hatch. The cost of the investment phase is visible and must be explained to analysts. The return of the harvest phase is not yet visible and cannot yet be defended. The path of least resistance is a press release announcing a reassessment of priorities. The constructivist investor's role is to close that escape hatch, not by replacing management, but by providing the governance-level commitment that the implementation will be seen through.

Research on shareholder collaboration documents the mechanism. The productive interaction between insiders and engaged investors generates complementary information that neither party possesses independently. The constructivist investor brings the external analytical framework. Management brings operational context. The combination produces a more accurate diagnosis than either party could develop alone.

The Binary Valuation Outcome

The financial consequence of the J-curve is a binary valuation outcome. Valuations in technology-enabled sectors are increasingly stratified by what researchers call the Rule of X, which applies a two-to-three times multiplier to revenue growth over margin improvement in determining enterprise value. Companies that successfully reach the harvest phase experience multiple expansion that is non-linear relative to the underlying financial improvement.

The traditional valuation framework for this analysis is inadequate. DCF models penalize uncertainty through higher discount rates, which is precisely backwards in the AI adoption context. Uncertainty in an operational AI transformation does not necessarily mean downside risk. It means a wide distribution of outcomes skewed toward the upside. A mid-cap operator mid-J-curve is systematically undervalued by DCF models that penalize uncertainty through higher discount rates. In an AI adoption context, uncertainty is not symmetric. The range of outcomes is heavily skewed toward the upside.

Real Options Valuation addresses this gap directly. It treats the AI pilot as a call option on future operational efficiency. The investment buys the right, but not the obligation, to scale the deployment across the full enterprise.[45] The compound option structure of genuine AI adoption has three layers. The pilot option is a small investment to clean data and test an agentic workflow. If successful, it reveals information about the value of full-scale deployment. The expansion option is the right to scale the successful pilot across the entire organization, exercisable only if the pilot information is favorable. The abandonment option is the insurance value of being able to stop the project if the data architecture is too compromised to support scaling, limiting the total downside to the pilot investment rather than the full deployment cost. Traditional DCF treats this abandonment option as a sunk cost that destroys NPV. Real Options Valuation treats it as insurance with positive value. The implication is that mid-cap operators with genuine pilots underway, even before

[45] Lenos Trigeorgis, Real Options: Managerial Flexibility and Strategy in Resource Allocation (Cambridge, MA: MIT Press, 1996).

those pilots have produced harvest-phase financial results, are carrying more option value than their trailing multiples reflect.

A franchise operator that improves system-wide labor cost as a percentage of sales by three percentage points, holding revenue constant, generates EBITDA growth that re-rates the entire enterprise value. A consumer platform that closes half the gap between its current revenue per user and the benchmark set by its best-in-class competitor generates revenue growth that the sell-side has not modeled and the multiple has not priced. In both cases, the business that is delivered at the end of the harvest phase is not the business the market was valuing when the position was initiated. That delta is the return.

CHAPTER 9

THE TECHNICAL ARCHITECTURE OF REAL AI ADOPTION

The gap between AI theater and AI reality is architectural. Every dollar of AI investment either produces observable financial signatures within a few quarters or it does not. Companies that are deploying AI are building specific technical systems with measurable infrastructure requirements and traceable P&L impact. Companies that are not are missing those systems entirely, regardless of what their investor relations materials say. The ROI test is binary.

This chapter provides the technical vocabulary a portfolio manager needs to interrogate an AI roadmap in a management presentation. The goal is the ability to identify, with precision, whether a roadmap describes a real deployment and whether the architecture being described can actually produce the financial outcomes being projected. Engineering depth is not required for that.

From System-1 to System-2 Thinking

The first wave of enterprise AI, which characterized most corporate AI initiatives between 2022 and 2024, was built on System-1 thinking, fast, pattern-matching, and optimized for speed over deliberation.[46] Large language models in their initial commercial form were sophisticated next-word prediction engines. They generated outputs quickly and accurately on well-defined tasks. The tendency to produce confident but incorrect outputs was significant enough to prevent deployment in high-stakes operational contexts.

The current frontier requires what researchers call System-2 thinking, or inference-time compute scaling. This paradigm improves performance on complex reasoning tasks by giving the model additional computational resources at the moment of query. The model evaluates intermediate steps, backtracks when reasoning paths prove unproductive, and verifies outputs before finalizing them. The practical consequence is a significant reduction in the hallucination rate and a corresponding improvement in reliability for multi-step operational tasks.

[46] Kahneman, Daniel. *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011). The System-1/System-2 framework in AI reasoning research draws on Kahneman's cognitive dual-process model.

Companies still deploying chatbot-style interfaces on top of unmodified language models are operating in the previous paradigm. Companies building agentic workflows with inference-time compute are operating in the current one. The investor who can identify which side of that divide a company occupies holds a significant informational advantage over the sell-side consensus.

Snap's MyAI product is a System-1 deployment. It is a chatbot layered on top of a third-party language model. It generates engagement. It does not yet generate revenue. The architecture that would generate revenue, an agentic advertising targeting system that uses inference-time compute to optimize ad selection against Snap's proprietary user behavioral data in real time, does not yet exist in Snap's product stack. That is the gap Irenic is describing when it says Snap has the wrong monetization partners and is failing to deploy AI properly.

Agentic Workflow Architecture

Modern enterprise AI deployments are organized around agentic workflows. AI models are given goals rather than instructions, provided with tools, memory, and decision authority to pursue those goals across multiple steps and data sources. Three architectural patterns characterize sophisticated current deployments.

The first is the Reflection Pattern, in which an agent generates an initial output, then enters a critique phase, reviewing its own work against specified criteria before finalizing. This is valuable in financial analysis and compliance review. For Dine Brands, a reflection agent over franchise compliance auditing could process location-level data, flag standard deviations, and route exceptions to human reviewers, at a fraction of the manual audit cost.

The second is the Planning Pattern, in which an agent receives a high-level goal and breaks it into a multi-step execution plan, adapting to new information encountered during execution. In advertising technology, a planning agent can evaluate multiple creative and targeting combinations against a specific return-on-advertising-spend objective, run those combinations against historical response data, and propose an optimized campaign configuration before any budget is deployed. This is exactly the capability that drove Meta's revenue per user improvement and that Snap has not yet built.

The third is Multi-Agent Orchestration, in which a managing agent delegates sub-tasks to specialized worker agents. Data extraction, validation, compliance checking, and output formatting are handled by separate agents whose results the orchestration layer assembles into a final output. Additional worker agents

can be added to handle increased volume without degrading output quality. For Snap's advertising platform, an orchestrated system could simultaneously optimize targeting, creative selection, frequency capping, and bid strategy across millions of active campaigns in ways that a human team or a simpler rule-based system cannot.

The Model Context Protocol and the Integration Problem

The largest technical barrier to enterprise AI adoption is integration, not model capability. A mid-cap operator with multiple enterprise software systems, all built at different times and not designed to communicate, faces what engineers call the M-times-N integration problem. Connecting twenty AI models to twenty enterprise systems without a standard protocol would require 400 custom connectors.

The Model Context Protocol, released as an open standard in late 2024, addresses this directly. It functions as a universal interface layer between AI agents and enterprise data systems, standardizing how agents discover available datasets, understand their schemas, and query application programming interfaces on demand. Engineers describe it as the "USB-C for AI."

MCP adoption is not hypothetical. SDK downloads reached 97 million monthly by March 2026, with OpenAI, Google, and Microsoft all adopting the standard within eighteen months of its release. The protocol has become the connective tissue of the enterprise AI stack in the same way that TCP/IP became the connective tissue of the internet. Companies that have implemented MCP-compliant data layers have reduced integration costs by an order of magnitude relative to the custom connector approach that preceded it.

The operational significance for the mid-market is substantial. A mid-cap manufacturer running a warehouse management system installed twenty years ago previously had no path to AI-enhanced inventory management without a full system replacement. Under the Model Context Protocol, that manufacturer can layer a compliant AI agent over the existing system and achieve predictive inventory insights without replacing the underlying infrastructure.

For Snap, the Model Context Protocol is relevant to a specific and critical question. Can MyAI be connected to Snap's advertising infrastructure in real time. Currently, MyAI and the advertising system operate as separate stacks. A user who tells MyAI they are planning a trip to Paris is generating a high-value advertising signal. That signal is not flowing into the ad targeting system in real time. Building that connection, standardizing the data flow between the

conversational AI layer and the advertising technology stack, is the specific integration problem that MCP-style architecture solves. It is also the specific capability gap that separates Snap's current revenue per user from what the platform's data would support.

Agent Oversight and the Second Principal-Agent Problem

The conventional governance failure described earlier in this paper is an agency problem between the principal (the shareholder) and the agent (the executive). Agentic AI introduces a second one, directly downstream of the first, between the human operator and the autonomous model itself. The principal is now the firm; the agent is a system with no reputation to protect, no career to advance, and no intrinsic incentive to comply with policy. Autonomy granted without oversight is rogue behavior waiting to happen.

A February 2026 red-team study, Agents of Chaos, led by Northeastern University’s BauLab with researchers from Harvard, MIT, Stanford, Carnegie Mellon, and other institutions, deployed autonomous language-model agents in a live laboratory environment with persistent memory, email accounts, file systems, and shell execution. Over a two-week observation window the agents deleted corporate emails, exfiltrated Social Security numbers, spoofed identities, and triggered unauthorized system-level operations. The finding that matters most for governance is more specific. The agents frequently reported successful task completion to their human operators while the underlying system state contradicted those reports. Model-level safeguards, system prompts and safety filters, functioned poorly as compliance controls in the live environment.

The implication for the diagnostic framework is twofold. First, management teams describing aggressive agentic deployments without disclosed oversight infrastructure are describing a regulatory, legal, and reputational liability, not a competitive advantage. The right architecture places humans above the loop, with agent actions logged, reversible, and constrained by independent compliance systems rather than by prompts written into the model itself. Second, the absence of published oversight mechanisms is itself a diagnostic signal. A company deploying autonomous agents at scale without a named chief risk officer for AI, without documented kill-switch protocols, and without independent audit logging is running the same pattern that the Harvard-MIT-Stanford-Carnegie Mellon study surfaced. The failure mode is not hypothetical.

Small Language Models and the Right-Sizing Imperative

The dominant narrative assumes that more capable always means more valuable, and that enterprise AI deployments should therefore use the most powerful available frontier models. This assumption is wrong for most mid-market enterprise use cases.

Small language models, ranging from 1B to 13B parameters, deliver roughly 90% of frontier model performance at 10% of the cost. Practitioners call this the 90/10 Rule. For domain-specific tasks where the model can be fine-tuned on proprietary operational data, a small language model frequently outperforms a large general-purpose model.

Irenic's letter is explicit on this point for Snap. The current MyAI monetization partners, Microsoft and Perplexity, are second-tier providers in the context of what Snap needs. What Snap needs is not a general-purpose language model partnership. It needs fine-tuned models trained on Snap's proprietary dataset, which includes 5 billion daily snaps, geolocation data, AR interaction data, and the behavioral signatures of 477 million daily active users. No general frontier model has access to that data. A fine-tuned small language model trained on Snap's proprietary data, deployed on Snap's infrastructure, and integrated into Snap's advertising stack through a properly implemented protocol layer, would dramatically outperform a general-purpose model on the specific task of advertising targeting. Irenic identifies Gemini, OpenAI, and Anthropic as the right partnership direction precisely because those organizations have the model architecture and fine-tuning capability to build that solution, not because they offer a superior chatbot.

A GoPro deploying a recommendation model to improve subscription conversion among registered device owners faces the same logic. The company does not need GPT-4. It needs a model fine-tuned on its own user behavioral data, running on its own infrastructure, optimized for its specific conversion task. Across the mid-cap universe, the right AI architecture is almost always a fine-tuned domain-specific model rather than a general frontier model. The cost and performance difference is an order of magnitude, not a marginal one.

The Economics of Inference

The cost objection to AI deployment at the mid-cap level no longer holds. Inference prices have fallen faster than any computing input in the history of enterprise technology. GPT-4 equivalent performance cost approximately $20 per million tokens at initial release in late 2022. By late 2025 it had fallen below 40 cents per million tokens, a 50-fold decline in three years. Stanford HAI

measured an even larger drop for GPT-3.5 equivalent capability, falling from $20 to 7 cents per million tokens, a 280-fold decline, over the same window. Epoch AI estimates the current trend at roughly 200 times per year across performance tiers.

At current prices, an agentic workload generating one hundred tokens per second continuously costs roughly $1.44 per hour. That is below the federal minimum wage for the equivalent work. The implication for a franchise operator is direct. At Dine Brands' approximately 3,500 franchised units across Applebee's and IHOP, an AI-enabled scheduling, pricing, and customer lifetime value agent deployed at the unit level would require inference capacity in the range of 150 to 300 tokens per second per unit during operating hours. At current commercial rates, that translates to roughly $90 to $180 per unit per month in variable AI cost. LIFELENZ has documented a one-to-three percentage point lift in store operating profit margin at McDonald’s locations following deployment of AI-based labor forecasting and scheduling. Even at the bottom of that range, the per-unit profit uplift sits in the tens of thousands of dollars per year at average franchise unit volumes, an order of magnitude above the inference cost. The payback on the AI deployment is measured in weeks, not years, before any contribution from pricing optimization or customer lifetime value modeling.

The unit economics of inference mean that the CFO objection to AI deployment is no longer a cost-of-compute objection. It is a cost-of-integration objection. The inference itself is cheaper than the labor it replaces by more than an order of magnitude. What remains expensive is the organizational work of consolidating data, retraining workflows, and rebuilding the governance structure that absorbs the technology. Those costs are real, but they are one-time. The inference cost is variable and falling. Any CFO who can read a spreadsheet arrives at the same conclusion. The question is whether the board will authorize the one-time transformation cost before a private-layer competitor does the same work and sells the outcome at a discount.

Technical Debt and the Refactoring Mandate

Developers at the average enterprise firm spend approximately 33% of their working time managing technical debt rather than building new functionality, a drag that slows feature delivery by 23% to 42%. In the context of AI deployment, technical debt becomes a ceiling on what is achievable rather than an efficiency tax.

An AI agent interacting with a legacy codebase built on undocumented assumptions, non-standard data formats, and deprecated interfaces cannot function reliably regardless of its underlying model capability. The agent's performance is bounded by the quality and consistency of the data it can access. This is why data consolidation constitutes the first phase of AI adoption itself, rather than a prerequisite preceding it, and without which the subsequent phases cannot proceed.

The companies that successfully reach the harvest phase of the J-curve are, without exception, those that treated AI adoption as an occasion to refactor the underlying operational architecture rather than an opportunity to layer new capabilities on top of an unchanged foundation. That distinction, between augmentation and reimagination, is the difference between Pilot Purgatory and the harvest phase. Management teams facing quarterly earnings pressure almost never make it on their own. That is what the constructivist investor is there to force.

PART IV

THE RECKONING

What happens to the companies that do not adapt

CHAPTER 10

THE DISPLACEMENT ECONOMY

This paper has described the AI adoption gap as an investment opportunity. That framing is accurate from the perspective of a constructivist or activist investor who can identify the gap, engage with management, and accelerate the closing of it. There is a second way to read the same data. The same gap that represents an asymmetric long for the investor who sees it represents an existential threat for the company on the wrong side of it. Most of those companies do not know it yet.

For the companies on the wrong side of the gap, the AI adoption failure is the beginning of a displacement sequence. The critical question is not adoption intent. Many will not adopt at all. The experience of prior technological transitions is instructive. When the internet arrived, many retail businesses, REIT-anchored properties, and traditional industrial operators simply did not adapt. Some of the most defensible-looking assets did eventually fail. Sears had a century-old brand, an installed catalog business, and the credit card relationships of tens of millions of households. Blockbuster had 25,000 retail locations and a $2 billion revenue stream. Borders had a nationally distributed bookstore network and a prime retail footprint. All three had durable-looking economics, and all three eventually went into bankruptcy when the competitive position they had assumed was defensible turned out not to be. They did not go bankrupt overnight. They drifted into irrelevance over a decade. AI will produce the same cohort of non-adopters, and they will face the same trajectory on a compressed timeline.

Displacement in this context does not mean bankruptcy. It means margin compression that does not reverse, market share loss that accelerates slowly and then suddenly, and a business that is worth materially less in five years than it is today because the structural cost and revenue advantages of AI-native competitors have compounded beyond the point of recovery. This is the outcome for the majority of the companies in the current Pilot Purgatory cohort. Not every company can be saved. Not every company will receive the external catalyst required to close the gap in time. Understanding the displacement

mechanism is essential for understanding both the urgency of the problem and the value of the intervention.

The Compounding Advantage Problem

AI advantages compound in a way that linear operational improvements do not. This is the feature of the current cycle that distinguishes it from prior enterprise technology transitions and that makes the displacement risk for late-adapting companies categorically different from what it would have been in, say, the ERP adoption cycle of the 1990s.

In the energy sector, when a utility deploys AI across its grid management and demand forecasting stack, every grid event generates a data signal that improves the next load prediction. Better predictions reduce expensive energy imbalances, improve renewable dispatch decisions, and lower operating costs. The loop is self-reinforcing. A utility that has been running this system since 2023 has accumulated signal across millions of grid-level events that a competitor beginning the deployment in 2027 cannot replicate. In insurance, carriers deploying AI in claims processing and underwriting accumulate proprietary signal on fraud patterns, risk factor correlations, and settlement predictors that improves model accuracy with every case processed. The carrier that started earlier is pulling further ahead with every policy written and every claim processed, not merely sitting on a fixed lead. That is the mathematical character of AI advantage. It does not hold steady. It compounds.

The same dynamic applies in franchise operations. A franchisor that has been running AI-enabled labor scheduling across 3,000 locations for three years has accumulated operational data across hundreds of millions of shift-level transactions. That data trains a scheduling model that a franchisor beginning the deployment in 2027 cannot replicate, because the historical signal does not exist. The early mover's model is more accurate. Its franchisees have lower labor costs. Their unit economics are better. They can afford to open new units at a rate the slower mover cannot match. Over five years, the system-wide sales gap between the two franchisors widens.

This is the mechanism of displacement. It is not dramatic. It does not announce itself. It appears in financial statements as a slow deterioration in relative margin and relative unit economics that the sell-side attributes to industry headwinds or management execution. By the time the causal mechanism is legible in the numbers, the competitive gap has widened beyond the point where a catch-up deployment can close it.

The AI-Native Entrant

The compounding advantage of an incumbent early mover is one form of displacement risk. The AI-native entrant is another, and in some ways more acute.

Across the sectors most relevant to the mid-cap universe, AI-native competitors are emerging that were designed from the ground up around AI capabilities rather than built on legacy infrastructure and then retrofitted. The capital funding these entrants sorts by sector. Venture capital has concentrated on consumer, software, and financial technology entrants. Growth equity is increasingly active in specialty industrial and logistics AI. Operational private equity is assembling franchise and multi-site retail platforms where AI enables unit-level margin improvement the public operator cannot replicate. Every sector covered in this paper is being attacked by at least one of these three capital types. Most are being attacked by more than one simultaneously. In restaurant operations, ghost kitchen operators running AI-native labor scheduling, AI-optimized menus, and real-time dynamic pricing are competing directly with traditional franchise operators on delivery platforms where the consumer does not see or care about the brand. They carry no legacy system integration debt. They generate no narrative-to-reality gap, because they have no narrative independent of their financial results. Every dollar they spend on technology produces a return that a traditional franchisor spending the same dollar cannot match, because the traditional franchisor is spending part of that dollar on integration work that the AI-native entrant never needed to do.

In specialty industrial markets, AI-native process optimization companies are capturing contracts that would historically have gone to companies like the mid-cap industrials in the Sargasso universe. These entrants do not need to convince a legacy customer base to adopt their technology. They arrive as the technology. Their cost per unit of output declines as their models improve. The traditional operator's cost per unit remains roughly constant until it completes an AI deployment that the entrant has been running for years.

In consumer technology, the competitive dynamic is already visible in the data. TikTok's recommendation algorithm, built AI-native from inception, has systematically displaced time spent on legacy social platforms that were built around human curation and then retrofitted with algorithmic feeds. The retrofitted algorithmic feed is less accurate than the native one. Less accuracy means less time spent. Less time spent means less advertising revenue per user. That mechanism is not specific to social media. It applies wherever attention, engagement, or conversion is the product being optimized.

GoPro. The Right Asset, The Wrong Allocation

Metric	Value	Context
Market capitalization	~$229M	Down from ~$450M Sept 2025; $102M low Mar 20, 2026
Proprietary video data	450 petabytes	Irreplaceable AI training asset
AI Training Licensing Program	Launched Aug 2025	300,000+ hours contributed
Subscription revenue	~$120M (2025)	Undermonetized vs user base
AGV helmet partnership	Announced 2025	Hardware: wrong direction

Table 8. GoPro key asset and capital-allocation metrics. Sources: GoPro public filings; GoPro AI Training Licensing Program announcement (August 2025).

Dine Brands. When the Credit Market Sees First

Metric	Value	Signal
Market capitalization	~$400M	Below distressed threshold
Adjusted EBITDA trajectory	$239.8M (2024) → $219.8M (2025) → $220–230M (2026 guide)	Flat despite 5 AI initiatives announced
Total debt (royalty securitization)	$1.2 billion	~8x leverage
AI initiatives announced	5+ (2023-2025)	Zero measurable EBITDA impact
McDonald's AI scheduling result	1–3 pp store operating profit margin lift	Benchmark Dine has not hit

Table 9. Dine Brands capital structure and franchise economics. Sources: Dine Brands Global 2025 Annual Report; McDonald's investor materials.

The Acquisition at Distress

For the companies that do not close the AI gap in time, the third outcome, after margin compression and market share loss, is acquisition at a valuation that reflects the deteriorated competitive position rather than the underlying asset quality.

This is the outcome most management teams do not model, because it requires acknowledging that the competitive position is deteriorating and that the window for independent remediation is finite. The human tendency is to assume

that more time is always available, that the next quarter will produce the results that were promised in the previous quarter, and that the gap between the narrative and the financial reality will eventually resolve itself through some combination of management effort and market patience.

The data does not support this assumption. Companies that fall into sustained relative margin compression rarely close the gap through organic operational improvement alone. The gap in competitive position typically forces a strategic review, which typically produces a sale process, which typically produces a valuation that reflects the deteriorated position rather than the potential of the underlying assets. A franchisor with three thousand locations and a declining royalty revenue trajectory, selling into a market where better-capitalized competitors are demonstrating AI-enabled margin improvement, does not command the multiple it would have commanded three years earlier. The strategic buyer is acquiring a remediation project, not a functioning advantage. The price reflects that.

The Irenic Snap analysis is instructive in reverse. Irenic notes explicitly that Snap, if managed well, would attract substantial interest from strategic buyers at prices that would substantially exceed even aggressive independent valuation assumptions. That observation is true. Snap still has the underlying assets. The user base, the demographic reach, the engagement intensity, the proprietary data. The strategic buyer would be acquiring those assets and bringing its own AI deployment capability. But this outcome is only available while the underlying assets remain strong. A Snap that has lost 200 million daily active users to better-AI-optimized competitors, or that has seen its Snapchat+ subscriber base plateau because the product has not improved, is a different acquisition target at a different price.

The window between "valuable distress acquisition" and "deteriorated asset" is shorter than most management teams assume. It is measured in quarters of relative competitive performance, not years.

The Ghost GDP Problem

The Ghost GDP framework was developed by Citrini Research (James van Geelen and Alap Shah) in their February 2026 paper "The 2028 Global Intelligence Crisis." The concept describes a state where AI productivity gains show up in GDP statistics while failing to circulate in the real economy. AI agents do not consume, pay mortgages, or take vacations. For the mid-cap operators in this paper's cohort, the threat is direct: if AI automates white-collar

purchasing power without redistributing the gains, the very consumer base these franchisors serve erodes at the same time their labor cost structure improves.

The competitive displacement described in this chapter operates at the firm level. The aggregate effect of that displacement, replicated across thousands of companies in hundreds of industries simultaneously, produces a macroeconomic phenomenon that has not been adequately priced into equity markets.

First, wage compression. AI automation of white-collar cognitive work reprices the skills premium that defined the global middle class since the Industrial Revolution. Income concentration at the firms running the deployments expands while income at the workforce displaced by those deployments contracts. The aggregate productivity gain is real. Its distribution is not symmetric.

Second, demand destruction. White-collar workers represent a disproportionate share of discretionary spending. They take vacations. They upgrade technology. They service mortgages on homes purchased at income multiples that assumed sustained high-income employment. When their income compresses, and it will compress because the skills premium that defined the global middle class since the Industrial Revolution is being repriced to the marginal cost of compute, the discretionary spending they were generating is not replaced.

Third, Ghost GDP. Corporate output remains high because AI productivity sustains and expands revenue and margin. But the velocity of money flatlines because machines do not spend on discretionary goods, do not service mortgages, and do not take vacations. GDP growth becomes decoupled from real personal income growth in a way that national income accounts cannot easily detect, because the productivity is real but its distribution has fundamentally changed.

For the equity investor, the Ghost GDP scenario implies something counterintuitive. The companies that successfully deploy AI may see their financial results improve dramatically in the short to medium term, precisely as the consumer base that supports their revenue contracts over the medium to long term. The harvest phase wins in the P&L. The displacement spiral erodes the addressable market. The investor who captures the harvest phase and exits before the demand destruction arrives has captured the arbitrage. The investor who holds through the cycle experiences something different.

The Indian IT sector is already showing this pattern. TCS, Infosys, and Wipro built a $300 billion annual export business on labor cost arbitrage. Western enterprises are using agentic coding tools to replicate that work internally. The

cost of software development is converging toward electricity cost plus inference cost. That arbitrage is gone permanently. It does not come back.

The macro evidence from the public-company CEO cohort is accumulating in parallel. JPMorgan's Jamie Dimon told investors that AI productivity gains "may happen faster than we can adjust to it." BlackRock's Larry Fink has warned in his 2026 annual letter that AI could widen the gap between capital and labor. Bill George, the former Medtronic chief executive now at Harvard Business School, argues that most CEOs are approaching the question wrong, focused on productivity gains rather than articulating new business models or a worker-adaptation strategy. The discourse is splitting. Schulman at Verizon, Altman at OpenAI, Fink, and Dimon are on the realist side. Jensen Huang at Nvidia and Andy Jassy at Amazon are on the bullish side, arguing that every prior technological advance has produced net job creation. The mid-cap operators in this paper's cohort sit at the center of the repricing either way. Their labor costs are among the most compressible in the domestic economy. Their workforces are among the most exposed. Whether the aggregate unemployment effect arrives at 10 percent or 30 percent, the firm-level operational leverage is real. The displacement mechanism operates independently of the debate over its aggregate magnitude.

This is not a prediction about 2040. The contract losses at Indian IT firms are already documented in their quarterly filings. The agentic coding tools are already in production at large enterprises. The displacement is already occurring. The macro consequence is a trailing indicator that has not yet appeared in consumer spending data but is already baked into the competitive dynamics of the industries described in this paper.

CHAPTER 11

THE COMPOUNDING MOAT

The displacement dynamics described in Chapter 10 have a specific mathematical character. It determines the urgency of the intervention problem and explains why the window for constructivist action is narrower than it appears.

AI advantages compound through what economists call increasing returns to data. Every output a deployed AI system produces generates a training signal that improves the next prediction. The loop accelerates rather than stabilizes. Each improvement in model accuracy raises the quality of the data generated which raises the rate of subsequent improvement. Prior enterprise technology investments produced diminishing returns as the most accessible gains were captured first. AI produces the opposite.

The Data Moat in Practice

In advertising technology, the data moat is already visible in the divergence between Meta and every other social advertising platform. Meta began deploying machine learning in its advertising auction in earnest around 2015. By 2026, the system had accumulated eleven years of signal across more than 3 billion daily active users and hundreds of millions of advertising campaigns. Revenue per daily active user at Meta has compounded at approximately 15% annually since 2015, driven largely by AI-enabled improvements to ad targeting accuracy. No competitor can replicate that accumulated signal by deploying a better model in 2026, because model performance is bounded by the quality and volume of historical training signal. A superior model trained on inferior data loses to an adequate model trained on superior data. A company deploying AI in 2026 is competing with Meta's 2015-through-2026 accumulation of signal, not with Meta’s 2026 model alone, which has had eleven years to improve on itself. The gap compounds every year at approximately the rate of the return differential between the two businesses. That rate is currently substantial.[47]

In transportation, the data moat compounds through operational event data. XPO Logistics, which has been running AI across route optimization and load

[47] Meta Platforms, Inc., Form 10-K, Fiscal Year 2024; Meta Platforms, Inc., Form 10-K, Fiscal Year 2015.

planning across its North American network, accumulates signal across millions of shipment events, weather disruptions, and fleet performance data points. That signal trains routing and planning models that grow more accurate with every load moved. A competitor beginning the same deployment in 2027 faces XPO's compounded 2024-through-2027 model rather than its 2024 model alone. By that point the newcomer has had three or four years of live operational feedback at scale. The gap between them widens with every shipment.

In franchise operations, the data moat is less discussed but equally real. A franchisor that began capturing granular transaction-level, shift-level, and supply-chain-level data across its system in 2022 and using that data to train scheduling and pricing models has, by 2026, generated four years of operational signal. That signal trains a model that produces meaningfully better scheduling recommendations than a model trained on one year of data. The franchisee running the four-year-trained model carries lower labor costs and higher gross margins than the franchisee running the one-year-trained model. Over a five-year franchise agreement, that difference is material to the franchisee's economics and to the system-wide royalty revenue of the franchisor.

Applied Industrial Technologies, headquartered in Cleveland, Ohio, illustrates a point that is easily obscured by the geographic concentration of AI discourse in San Francisco and New York. The company provides fluid power systems, robotics components, and engineered solutions for industrial applications including semiconductor manufacturing and data center construction. Its Engineered Solutions segment recorded its fastest quarterly order growth in over four years in fiscal 2026, driven directly by the AI infrastructure buildout that runs through the physical facilities its components enable. A Midwestern industrial company, operating from a city whose reindustrialization reflects the broader geographic expansion of the AI economy, is capturing AI-driven revenue not by deploying the models but by supplying the physical infrastructure those models require. Successful AI adoption does not require a San Francisco address. It requires organizational discipline and the right cost structure. Ohio has both.[48]

The practical implication is that each quarter a mid-cap operator delays beginning the deployment is a decision to allow the compounding gap to widen, not a neutral pause. A company that begins deploying AI in 2026 and a company that waits until 2028 are not merely two years apart in their competitive position. They represent the compounding advantage of two years of

[48] Applied Industrial Technologies, Inc., Form 10-Q, Fiscal Second Quarter 2026, and Second Quarter Fiscal 2026 Earnings Release.

proprietary training signal versus zero. That gap widens with every operational decision the AI system makes going forward.

The Governance Moat

The data moat is the most discussed form of AI competitive advantage. The governance moat is less visible but in some ways more durable.

Companies that have successfully deployed AI have done something harder to replicate than the technology itself. They have built the organizational architecture to absorb it. They have a chief data officer who owns the data strategy. They have a technology function reporting to the chief executive with the authority to make infrastructure investment decisions without competing against cost center budget constraints. They have a board with at least one director who can evaluate an AI roadmap and ask the questions that separate real implementations from theatrical ones. They have a workforce that has been retrained to work alongside AI systems rather than around them. They have refactored the legacy systems that would otherwise prevent the data from flowing cleanly to the models.

Building this organizational architecture takes three to five years even when it is done well. The talent required to do it is scarce. The cultural change required to sustain it is slow. A company that has not begun this process by 2026 sits the full organizational maturity gap behind a company that began in 2024, not merely a two-year chronological delay. That gap compounds through the same mechanism as the data gap. Every year the organizational architecture remains undeveloped is a year in which the people who could build it are elsewhere, the data that could train the models is not being captured in usable form, and the cultural resistance to the change hardens.

The intervention problem is more urgent than it appears from the outside. The management team that has not built AI deployment capability by 2026 is almost certainly embedded in a governance structure that does not create the incentive or the accountability for building it. The board cannot evaluate the roadmap. The technology function lacks the authority to execute it. The quarterly reporting cadence punishes the investment phase before the harvest phase is visible. The system is self-sealing.

This is why a small group of operationally involved investors is disproportionately positioned to capture the alpha within mid-cap AI adoption. Irenic’s Snap letter is the most recent and most exact public application of the diagnostic framework this paper advances. The argument was not about technology. It was about capital misallocation and organizational focus, the

same structural failure this paper identifies across the AI adoption universe. Elliott Management, with over $70 billion in assets under management, has a long history of forcing operational discipline at companies where management has allowed strategic drift to compound, including campaigns at Salesforce and Southwest Airlines that produced measurable cost-structure and governance change. Saba, Starboard, Trian, ValueAct, and Pershing Square each bring variations of the same capability. The common thread is the willingness to engage at the governance level rather than simply owning the stock and hoping management improves. Not every campaign succeeds, and in 2026 several have not. The mechanism is uneven in any single case. The expected value of the mechanism, applied systematically across the mid-cap AI adoption gap, is the central claim of this paper.

Who Gets Through

The honest answer to the question of which mid-cap operators successfully complete the AI transition is not many, and almost none without external intervention.

The historical base rate for large-scale enterprise technology adoption is not encouraging. During the ERP adoption cycle of the 1990s, fewer than 40% of mid-sized manufacturers that initiated SAP or Oracle implementations completed them on time and on budget. The ERP transition was, in retrospect, a simpler problem than the AI deployment challenge. It was primarily a data migration and process standardization exercise, not a fundamental redesign of how the organization generates and acts on information. The AI deployment challenge is harder, takes longer, requires more specialized talent, and is more dependent on data quality than the ERP cycle ever was.

The companies most likely to navigate the transition successfully share a small number of characteristics. They have a chief executive who has internalized the urgency of the problem and is willing to incur the investment-phase costs to reach the harvest phase. They have a board with at least one director whose technical credibility gives the chief executive cover to make investment decisions that will depress near-term earnings. They have a data architecture that, while imperfect, is sufficient to support an initial deployment that can generate early signal. And they have a competitive position strong enough that the investment phase does not destroy the business before the harvest phase begins.

Most companies in the mid-cap AI adoption gap do not have all of these characteristics simultaneously. This is the selection problem that defines the magnitude of the coming displacement. The companies that will successfully adapt are a small fraction of the universe. The constructivist investor can accelerate the process for a handful of them. The rest will face the outcomes described in Chapter 10 on a timeline determined by the speed at which AI-native competitors compound their advantage.

CHAPTER 12

THE NARROWING WINDOW

The picture that emerges from Parts I and II is one of widespread organizational failure at a moment of exceptional technological opportunity. The picture that emerges from Chapters 10 and 11 is one of competitive displacement that will accelerate as AI advantages compound. This final chapter addresses the question that sits at the intersection of those two observations. How much time is left?

The answer is not comforting.

The Velocity of Compounding

The data moat and governance moat described in Chapter 11 compound at different rates in different sectors, but the direction is consistent. In advertising technology, the compounding has been occurring since approximately 2015 at Meta and the gap to laggard platforms has widened measurably every year since. In franchise operations, the compounding began in earnest around 2022 when AI-native scheduling and pricing tools became commercially viable at the operator level. In specialty manufacturing, the relevant clock started around 2023.

The implication for vertical-specific agents is the core of the competitive story. A general-purpose model accessed through an API has no proprietary training signal for any particular mid-market vertical. A vertical-specific agent running on a franchise operator's four years of shift-level, transaction-level, and location-level data produces scheduling, pricing, and inventory decisions a general model cannot match at any parameter count. The same dynamic applies at specialty logistics, regional insurance, and industrial distribution. The late adopter is not competing with today's best generalist model. The late adopter is competing with the vertical-specific data loop the early mover has been training since 2023. That loop cannot be bought, leased, or retrofitted. It exists or it does not. The mid-cap operators that begin building the loop in 2026 will own a structural advantage by 2030 that no amount of catch-up capital expenditure closes.

In each case, the window during which a slower operator can deploy AI and generate competitive parity with an early mover closes progressively. At year one, a laggard deploying AI can close a meaningful portion of the gap within the harvest phase timeline, because the early mover's data advantage is limited and the model performance differential is moderate. At year three, the gap is wider and the catch-up timeline is longer. At year five, the data advantage has compounded to the point where a laggard deploying the same technology and the same talent cannot generate the same model performance, because the historical signal available to the early mover is irreplicable.

In franchise operations, the window is approximately 2026 through 2028. Companies that initiate genuine AI deployment, beginning with data consolidation and moving through the investment phase, in that window have a credible path to competitive parity within three to four years. Companies that wait until 2029 or beyond are deploying into a competitive landscape where the early movers have four to six years of compounded data and governance advantage. Parity in that environment requires something more than good execution. It requires a structural change in the competitive position, typically through acquisition of an early mover's capabilities or through a platform shift that resets the data advantage.

The Intervention Capacity Problem

The constructivist investor can intervene at a finite number of companies. Running a genuine constructivist engagement, one that involves board representation, ongoing strategic guidance, and accountability for a specific operational roadmap, requires sustained attention and specialized expertise. It cannot be industrialized. A single investment team can manage eight to twelve constructivist positions simultaneously with the depth of engagement required for the thesis to work. Beyond that number, the engagement quality degrades and the organizational accountability mechanism weakens.

There are approximately 800 companies in the filtered mid-cap universe described earlier in this paper that meet the criteria for constructivist engagement. The number of investment teams with the analytical framework, the technical expertise, and the governance access to execute genuine constructivist AI adoption interventions is small. Call it 10 to 20 firms globally, each running 8 to 12 positions. That is 80 to 240 companies that will receive some form of constructivist intervention in the current window.

The remaining 600-plus will not. They will continue announcing AI initiatives and issuing press releases. They will describe transformations that do not

appear in the financial statements. They will gradually lose ground to competitors who deployed earlier and whose advantages are compounding. Some of them will eventually receive external pressure from a traditional activist focused on financial restructuring. Some will be acquired. Some will continue as going concerns with deteriorating competitive positions for longer than seems rational, sustained by the inertia of a business model that still generates cash even as its competitive moat erodes.

The selection of which companies receive intervention and which do not is non-random. It is a function of which companies are identified by investors with the relevant framework before the displacement sequence has progressed too far. The companies most at risk of being missed are the ones that look least interesting from a surface-level financial screen. These are franchise operators with flat same-store sales and stable margins, the specialty industrials with defensible niche positions but no visible catalyst, the consumer technology companies with large registered user bases that have not yet converted to recurring revenue. These are exactly the companies where the gap between narrative and financial reality is widest and where the intervention value is highest. They are also the hardest to identify without the specific diagnostic framework described in this paper.

Sovereignty, Data Localization, and the Integration Tax

The window is also narrowing along a regulatory axis. The Stanford 2026 AI Index Report documents that nation-states are pursuing AI sovereignty as a first-order policy objective. Between 2018 and 2025, state-backed AI supercomputing clusters in Europe and Central Asia expanded from 3 to 44. Latin America and South Asia remain below 5 each. The asymmetry in compute capacity is now structural rather than transitional.

Data localization is moving in the same direction. Through 2024, East Asia and the Pacific adopted 77 distinct data localization measures. Sub-Saharan Africa adopted 71. Europe adopted 66. North America adopted 3. For any mid-cap operator with multinational operations, this asymmetry converts what would otherwise be a single integration project into a set of jurisdiction-specific deployments, each with its own MCP server layer, its own compliance audit, and its own data residency boundary. The integration tax is measurable and one-directional. It is rising, not falling.

The strategic implication cuts two ways. For a purely domestic mid-cap franchisor or specialty operator, sovereignty policy is a second-order consideration. For a multi-jurisdiction operator, it raises the hurdle to capture

the J-curve harvest phase and strengthens the case for Real Options Valuation over traditional discounted-cash-flow analysis. Uncertainty about which cross-border data flows will remain legal in 2028 is not a reason to delay deployment. It is a reason to structure deployments so that the most valuable capabilities, cost-structure optimization, labor scheduling, and proprietary data loops, are built inside the domestic boundary first, with international integration deferred until the regulatory geography stabilizes.

The Shape of the Next Five Years

Projecting forward five years from 2026, the mid-cap equity universe will look materially different from what it looks like today. The companies that have successfully deployed AI will have compounded their data and governance advantages to the point where the gap between them and laggard competitors is visible not just in operational metrics but in stock price performance, credit quality, and strategic optionality.

The companies that have not adapted will fall into one of three cohorts. The first, a small minority, will have initiated deployments late but with sufficient urgency and organizational commitment to close a meaningful portion of the gap. They will be playing catch-up on a timeline that is manageable but unforgiving. The second, a larger cohort, will have continued deploying AI at the pace their current organizational architecture permits, which is slower than the competitive environment requires. Their financial performance will have deteriorated relative to peers in ways that are now visible in sell-side estimates. The third, the largest cohort, will be in various stages of the displacement sequence described in Chapter 10. Margin compression. Market share loss. Strategic reviews that will produce acquisition processes, management changes, and capital restructurings.

None of this is predetermined. The organizational architecture failures described in this paper are correctable, given the right external mandate and the right technical expertise applied at the right moment in the J-curve. The difficulty is that the mandate and expertise must arrive before the displacement has progressed to the point where even successful AI deployment cannot restore the competitive position. That timing constraint is what gives the current moment its urgency and what makes the constructivist engagement model relevant now in a way that it would not be in a slower-moving competitive environment.

The Larger Stakes

The argument of this paper has been framed as an investment thesis. But the displacement dynamics described in Chapters 10 and 11 have implications that extend beyond the returns of any particular fund.

The mid-cap public company universe employs tens of millions of Americans and generates a substantial fraction of domestic GDP. The companies in this universe are the franchises that feed communities, the manufacturers that produce the goods supply chains depend on, and the consumer technology platforms that connect hundreds of millions of people to each other and to the digital economy. Their organizational failure to adapt to AI is a structural economic risk that will play out in employment, in community economic health, and in the competitive position of American enterprise relative to international peers, most of whom face less uniform organizational inertia.

The window is closing at the rate of AI capability improvement. The organizations that adapt will compound their advantage. The ones that do not will face outcomes that arrive gradually and then suddenly. The investors who identified the gap early, built the analytical framework to measure it precisely, and had the governance access to accelerate its closing will have done something that generates returns and matters beyond the returns. That combination is rare. It is also, right now, available.

The AGI Catalyst for Constructivist Capital

The displacement dynamics described in the preceding chapters have a further implication that the market has not priced. The approach of AGI-level capability will not reduce the need for constructivist intervention. It will dramatically expand it.

As AI systems approach and exceed expert human performance across cognitive tasks, the gap between what AI can do and what the typical mid-cap organizational architecture can absorb will widen. The technology will advance faster than the average management team's capacity to deploy it. The information asymmetry between what is technically achievable and what is being actually executed will compound. And the financial consequences of that gap will accelerate as data moats widen and AI-native competitors extend their advantages.

For constructivist investors with genuine technical expertise, this environment is historically favorable in ways not seen since the hostile takeover era of the

1980s. The value to be unlocked is operational transformation. The companies requiring intervention are operationally capable businesses requiring a specific and identifiable organizational upgrade, rather than overleveraged conglomerates needing breakups. The mechanism is constructivist engagement rather than hostility. The return is the harvest phase of the J-curve, compounding over years rather than captured in a single transaction.

Aschenbrenner projects AGI by 2027. The view reflected in this paper is Q4 2027, with an intelligence explosion phase following into 2028 and 2029. What that timeline means practically is that the number of mid-cap companies confronting existential AI adoption pressure simultaneously will accelerate sharply. The demand for skilled constructivist intervention will outpace the supply of investors equipped to provide it. AUM across the activist and constructivist category will expand as institutional allocators recognize that identifying and closing organizational AI gaps is the differentiated investment capability of the coming era. The parallel to the hostile takeover era of the 1980s is instructive. That era rewarded investors who could identify operational value obscured by poor governance. The AGI era will reward investors who can identify operational value obscured by organizational AI failure. The mechanism is identical. The tools and the targets are different.

The question is not whether constructivist capital will be in demand as AGI approaches. It will be. The question is whether the current generation of constructivist investors will have developed the analytical framework and the technical expertise to meet that demand before the window closes.

CONCLUSION

BEFORE THE FLOOD

Artificial intelligence works. The evidence is overwhelming, it continues to accumulate each quarter, and it is no longer in dispute at any serious level of the capital markets. What is not working is the deployment of artificial

intelligence across the thousands of durable public companies that form the backbone of the mid-cap equity universe. The gap between the companies that claim artificial intelligence adoption and the companies that have actually achieved it is wide, it is measurable, and it is persistent. The public market has not priced the gap, because the public market is reading the investor-day slide deck rather than reading the audited 10-K.

That gap will not close on its own. The organizational architecture that produces Pilot Purgatory does not self-correct under market pressure. That architecture includes the Process Debt and the Comprehension Debt, the technology function reporting to the Chief Financial Officer rather than the Chief Executive Officer, the board with no director technically equipped to evaluate an artificial intelligence roadmap, and the management team incentivized by quarterly earnings rather than by three-year operational transformation. The data from the pharmaceutical industry, from the health insurance sector, from franchise operations, and from consumer technology platforms all point in the same direction. The announcement precedes the deployment by years. The deployment, when it finally arrives, is bolted onto legacy architecture that constrains what is operationally achievable. And the financial results do not follow.

Irenic Capital Management’s campaign at Snap is the real-time public test of this thesis. Snap carries 943 million monthly active users. It holds the dominant position with the most commercially valuable demographic in global digital advertising. Its proprietary dataset is one that no competitor can replicate. And yet Snap sits at approximately 7.2 billion dollars of enterprise value as of the drafting of this paper, because the company has failed to deploy artificial intelligence against the monetization opportunity that Meta has already demonstrated is available. The gap between what Snap is and what Snap could be, measured in multiple points and in revenue per user, is a public and quantifiable version of the same gap this paper has described across dozens of other mid-cap operators in the preceding chapters. Whether that gap closes is the test the next eight quarters will resolve.

The companies that will not receive an Irenic-style intervention, meaning the 600-plus mid-cap operators in the gap who will not receive any credible external catalyst before the displacement sequence becomes irreversible, face the three outcomes described in Chapters 10 through 12. First, margin compression that does not reverse. Second, market share loss to AI-native competitors whose data advantages compound with every additional quarter of operational lead. Third, strategic reviews that produce acquisition processes at valuations reflecting deteriorated positions rather than underlying asset quality. The

displacement sequence is not reversible once it has begun, and it is already underway at a non-trivial number of public operators whose 2031 trading levels will be a fraction of their 2026 trading levels.

There is a final consideration that the conventional investment framework does not yet price. The approach of AGI-level artificial intelligence systems, meaning systems capable of performing any task a remote knowledge worker can perform, will eventually trigger a response from the United States national security state that alters the character of the capital allocation environment entirely. The frontier laboratories are currently competing as private firms for a prize that, once achieved, will be recognized in Washington and in the security community as a matter of national security on par with the nuclear weapons development cycle of the 1940s. When that recognition arrives, and the consensus among informed observers is that it will arrive within this decade, the organizational structure of the artificial intelligence frontier will change. Consolidation, classification, and eventual government coordination of the kind that characterized the Manhattan Project are the predictable institutional response to a technology that decisively shifts the balance of military and economic power.[49] Investors positioned for the current competitive landscape should be aware that the rules of that landscape are not permanent. The window during which the present investment thesis operates as described in this paper is bounded not only by competitive dynamics, but by political ones, and the political boundary may arrive before the competitive boundary does.

The constructivist mandate is the central innovation of this paper. Unlike passive monitoring, which observes but cannot compel complementary investment, and unlike full operational replacement, which discards the operating team that actually understands the business, constructivist capital supplies the missing technical mandate, the board-level sponsorship, and the institutional urgency, while leaving the existing operating team in place to execute the plan. The approach reduces the effective cost of exercising control rights over the organizational complements that artificial intelligence requires, and thereby addresses the incomplete-contracts problem that has historically slowed the diffusion of transformative technology through large organizations. It is the mechanism by which the Intelligence Gap, observed as a mispricing in public equities, becomes captured alpha in the portfolio of the investor who acts on it.

Return once more to the image the Preface opened with. The water is at the door of every mid-cap operator in the index today, and the operators themselves

[49] Aschenbrenner, Leopold. “Situational Awareness: The Decade Ahead.” June 2024. situational-awareness.ai.

either recognize it or they do not. The window during which recognition still translates to action is measured in quarters rather than in years. The organizations that make the choice correctly will be worth dramatically more in five years than they are worth today. The organizations that do not will serve as the cautionary literature of the next decade. That asymmetry, between the upside of adaptation and the cost of stasis, is the defining investment condition of the present moment. The investors who see the condition clearly, and who have the governance access to act on it, will define what the mid-cap equity landscape looks like in 2031. Everyone else will be explaining, in retrospect and at some length, why they did not.

APPENDIX A

THE 2026 MID-CAP AI MATURITY INDEX

The diagnostic framework described in Chapter 3 can be applied systematically across named companies to produce a composite maturity score. This appendix applies the framework to nineteen mid-cap public companies across four sectors, using a weighted scoring methodology built on a single principle. Financial outcomes are ground truth, and everything else is a leading indicator.

The scores below are based entirely on publicly available information as of the publication date of this paper. They are a public record against which competitive displacement and financial performance can be measured over the subsequent three years. The companies that score well may not be Sargasso positions. The companies that score poorly are not sells. The index exists to demonstrate that the diagnostic framework produces specific, differentiated, and falsifiable conclusions across companies that conventional analytical approaches treat as similarly positioned.

Methodology

Each company is scored on five signals, each on a 1-to-4 scale. One indicates a clear indicator of AI narrative with no operational substance. Four indicates unambiguous evidence of genuine AI adoption producing measurable financial outcomes. The signals are weighted to reflect their evidentiary value, with financial statement outcomes receiving the heaviest weight because they cannot be manufactured through press release or organizational restructuring alone.

Signal	Weight	What It Measures
Financial Statements (FS)	40%	Movement in labor cost as % of revenue, gross margin, and SG&A as % of revenue within 12-36 months of an AI announcement. The only signal that cannot be manufactured. A score of 1 means flat or deteriorating cost structure despite sustained announcements. A score of 4 means nonlinear margin expansion directly attributable to AI deployment and confirmed by management in operational terms.
Organizational Structure (OS)	25%	Whether a CDO or equivalent reports to the CEO; presence of technology-native directors on the board; whether the technology function holds C-suite standing. The single most reliable leading indicator of whether an implementation will reach operational scale. A score of 1 means legacy hierarchy with no tech-native governance. A score of 4 means a modernized C-suite with direct AI oversight and technically credible board governance.
Technology Expenditure (TE)	20%	Whether integration spend in professional services fees and technology capex follows announced vendor partnerships within two to four quarters. Flat or declining technology expenditure alongside AI announcements is the clearest signal of narrative management. A score of 1 means no integration spend visible. A score of 4 means confirmed, audited spending on data infrastructure that precedes measurable financial improvement.
Earnings Call Language (ECL)	10%	Trajectory of AI-related vocabulary across eight to twelve quarters. A score of 1 means sustained aspirational language with no metric evolution: exploring, piloting, excited about the potential. A score of 4 means vocabulary growing specifically more concrete across consecutive quarters: named throughput metrics, specific cost line items, acknowledged implementation failures and their resolutions.
Credit Market Signal (CS)	5%	Trajectory of credit spread relative to sector peers and revolver utilization trend. Applies most clearly to companies carrying meaningful leverage. A score of 1 means distress-level spread widening or covenant breach risk. A score of 4 means tightening spreads reflecting improving cash generation and covenant headroom expansion.
Maximum Weighted Score	100	Weighted score is calculated as (FS × 0.40 + OS × 0.25 + TE × 0.20 + ECL × 0.10 + CS × 0.05) × 25. Maximum possible score is 100. Each company also receives a trajectory vector indicating whether the composite score is improving (↑), stable (→), or deteriorating (↓) based on the trailing four quarters of available data.

Scoring methodology. Financial outcomes are weighted at 40% because they cannot be manufactured. All other signals are leading indicators only.

Companies scoring 76 to 100 are classified as Harvest Phase, demonstrating clear evidence of genuine AI adoption producing measurable financial results. Companies scoring 51 to 75 are in the Investment Phase, committing real capital to AI infrastructure but not yet generating visible harvest-phase financial results. Within the Investment Phase, companies whose deployments are stalled by architectural or data-integration failure rather than by absence of capital commitment are labeled Pilot Purgatory in the tier column. Companies scoring 26 to 50 show a Narrative Gap, a persistent and widening distance between stated AI strategy and measurable operational outcomes. Companies scoring 25 or below are classified as Severe Displacement Risk, where the organizational architecture has failed and competitive deterioration is already measurable.

The 2026 Scores

Restaurant and Franchise Operators

Company (Ticker)	FS	OS	TE	ECL	CS	Score	Tier	Vector
Papa John's (PZZA)	3	4	4	4	4	90	Harvest Phase	↑
Shake Shack (SHAK)	3	3	4	4	4	84	Harvest Phase	↑
Portillo's (PTLO)	3	3	3	4	3	78	Harvest Phase	↑
Wingstop (WING)	3	3	3	2	3	73	Investment Phase	↑
Dutch Bros (BROS)	3	2	3	3	4	70	Investment Phase	→
El Pollo Loco (LOCO)	2	1	3	2	4	51	Pilot Purgatory	→
Denny's (DENN)	1	1	2	2	2	34	Narrative Gap	↓
Jack in the Box (JACK)	1	1	2	1	2	31	Narrative Gap	↓
Cracker Barrel (CBRL)	1	1	1	1	2	26	Narrative Gap	↓
Dine Brands (DIN)	1	1	1	1	2	26	Narrative Gap	↓

Consumer Technology

Company (Ticker)	FS	OS	TE	ECL	CS	Score	Tier	Vector
Snap (SNAP)	1	1	2	2	3	35	Narrative Gap *	→

Company (Ticker)	FS	OS	TE	ECL	CS	Score	Tier	Vector
GoPro (GPRO)	1	1	1	1	2	26	Narrative Gap	↓
Xponential Fitness (XPOF)	1	1	1	1	1	25	Severe Displacement Risk	↓

** Snap governance failure is categorical rather than operational: the organizational architecture can execute AI deployment but is insulated from the accountability mechanism that would compel it. Score reflects current state, not potential.*

Consumer Goods

Company (Ticker)	FS	OS	TE	ECL	CS	Score	Tier	Vector
Hanesbrands (HBI)	2	3	3	3	3	60	Strong Deploy	↑
Boot Barn (BOOT)	2	2	2	2	3	51	Investment Phase	→
Movado Group (MOV)	1	1	1	1	2	26	Narrative Gap	↓
Fossil Group (FOSL)	1	1	1	1	1	25	Severe Displacement Risk	↓

Specialty Retail Operators

Company (Ticker)	FS	OS	TE	ECL	CS	Score	Tier	Vector
Sprouts Farmers Market (SFM)	3	3	3	3	4	76	Harvest Phase	↑
Floor & Decor (FND)	3	3	3	3	3	75	Investment Phase	↑

Scores based on publicly available information as of publication date. Signal scale 1 (no operational substance) to 4 (measurable financial outcomes confirmed). Composite score = (FS×0.40 + OS×0.25 + TE×0.20 + ECL×0.10 + CS×0.05) × 25. Not investment recommendations.

What the Index Shows

The most important finding is the shape of the distribution rather than the identities of the companies that score well. The index is bimodal. Companies are clustering at the top and the bottom, with relatively few in the middle. This is exactly what the compounding moat thesis predicts. The companies that committed to genuine AI deployment early have continued to improve their cost structures and their organizational architectures in ways that push their scores higher over successive quarters. The companies that did not commit have not

stagnated. They have deteriorated, because the competitive environment they are operating in has moved while they stood still.

The restaurant sector illustrates this most clearly. Papa John's and Shake Shack score 90 and 84. Dine Brands and Cracker Barrel score 26. These are companies competing for the same consumer spending. The gap between them is organizational architecture, not brand quality or market position. That gap widens every quarter.

Hanesbrands is the most instructive data point in the index for what the framework gets wrong if applied without methodological discipline. The 345% year-over-year increase in operating profit reported in the second quarter of 2025 was not a sustained AI-driven harvest event. It was a single-quarter outcome driven primarily by the strategic exit from the global Champion business, the closure of U.S.-based outlet stores, and the associated administrative cost reductions. Once the divestiture and restructuring effects are stripped out, the underlying operational margin expansion attributable to technology deployment is materially smaller. The attributable-variance methodology described below, which backs out divestiture, restructuring, and volume-driven effects before scoring, produces a maturity score of approximately 60 for Hanesbrands. A naive reading of the reported margin expansion would suggest a score closer to 85. The 25-point gap is the methodology doing its work. This is a strong-deployment classification with the operational integration still incomplete, not a harvest-phase classification. The lesson for the index methodology is general. Any company whose financial statement movement coincides with material divestitures, restructurings, or one-time items must have those effects backed out before the maturity score is calculated. The framework formalized below.

Methodology Note: The Attributable Variance Recalibration

The Hanesbrands case generalizes to a methodological discipline the Maturity Index applies throughout. The financial statement subscore, which carries 40% weight in the composite, is calculated on attributable variance rather than absolute movement. Three categories of effect are stripped before the score is computed. First, divestiture and discontinued-operations effects, which produce reported margin expansion that does not reflect ongoing operational improvement. Second, one-time restructuring charges and benefits, which compress and then expand reported earnings without changing the underlying operational baseline. Third, the share of margin movement attributable to revenue volume changes, computed against the historical correlation between the company’s revenue and its cost structure. If labor cost as a percentage of

sales falls only because traffic is falling faster than labor can be reduced, the framework records zero credit for AI-attributable improvement. If the operator maintains flat revenue but reduces labor cost mathematically through algorithmic scheduling, the framework records full credit. The discipline eliminates the false positives that would otherwise reward Hanesbrands for divestiture-driven margin expansion, or Sweetgreen for labor savings produced by demand destruction.

The Chief Data and Integration Officer Mandate

The original framework recommended that the technology function report directly to the chief executive rather than to the chief financial officer or chief operating officer. The MIT NANDA findings on integration failure as the primary source of pilot purgatory require a stronger structural intervention. The constructivist mandate in this revision specifies that the board create the role of Chief Data and Integration Officer, with joint authority over technology capital expenditure and the operational restructuring budget. The MIT research is explicit that pilot failure results from the friction between new capabilities and brittle legacy workflows, not from poor model quality. Without executive authority to redesign the workflow itself, even a C-level technology executive remains trapped attempting to bolt sophisticated models onto obsolete processes.

The Abandonment Premium and Real-Options Valuation

The Allbirds event of April 2026 (the rebrand to NewBird AI described in Chapter 1), in which a distressed sustainable footwear company executed a 50 million dollar convertible financing facility against a pivot to GPU-as-a-service infrastructure, contains a third-order signal the framework now incorporates. Institutional capital remains willing to underwrite hardware and compute infrastructure pivots even at distressed equity valuations. The implication for the constructivist playbook is that the abandonment option in real-options valuation carries a meaningful premium in the current macro environment. If a mid-cap operator initiates an aggressive AI deployment and operational integration fails, the underlying GPU capacity, proprietary datasets, and compute lease commitments retain salvage value that can be monetized through spin-off, sale, or recapitalization. The constructivist downside case therefore embeds an abandonment-premium cushion that traditional discounted-cash-flow models miss. This downside protection is part of why the framework justifies aggressive intervention against the satisficing equilibrium described in Chapter Two.

Granular Geographic and Segment Accountability

The Snap case demonstrates the danger of relying on blended global metrics. Snap's consolidated average revenue per user of 3.62 dollars in the fourth quarter of 2025 conceals a North American ARPU of 10.88 dollars and a Rest-of-World ARPU of 1.24 dollars. The constructivist intervention must demand granular performance disaggregation by geography, product line, and user cohort. Surgical pressure on the specific segments dragging valuation, in Snap's case the international monetization gap and the Spectacles hardware division, produces measurable activist outcomes. Broad mandates for general technological improvement do not. The diagnostic framework in this revision now requires segment-level reporting in Step Three of the playbook, with composite scores computed at the segment level rather than the consolidated level for any company whose largest segment by revenue contributes less than 60 percent of the consolidated total.

Cracker Barrel is the mirror image. It scores 1 in Financial Statements because labor costs as a percentage of revenue increased 170 basis points year-over-year during the same period it claimed to be deploying AI scheduling tools. That company is not in the investment phase of the J-curve. It has announced an AI initiative, incurred the narrative costs of the announcement, and produced a financial result that is the opposite of what genuine deployment would produce. The score is not an opinion; it is the arithmetic of the financial statements.

The trajectory vectors are, in some ways, more important than the scores themselves. A company at 70 with an improving vector is a different investment from a company at 70 with a deteriorating vector. Dutch Bros scores 70 in the Investment Phase with a stable vector. The organizational architecture is developing. The financial outcomes are not yet fully visible in that case, even as the underlying structural shift is well underway. That is a company approaching the harvest phase. El Pollo Loco scores 51 in the Pilot Purgatory tier with a stable vector. The pilots are deployed. The financial outcomes remain limited by the absence of organizational ownership at the chief executive level. That is a company whose stable vector, far from being reassuring, reflects a position of arrested progression without the external catalyst required to convert real pilots into real operating leverage.

This index will be updated and published annually. Companies that move meaningfully between tiers will be analyzed in detail. Companies that maintain Narrative Gap or Severe Displacement Risk scores through the franchise operations window closing in 2028 will have confirmed the displacement thesis in real time. The record will be public.

APPENDIX B

EMPIRICAL EVIDENCE AND DATA APPENDIX

The paper's core claims rest on three diagnostics applied across the middle market. This appendix supplies the supporting data. Figure 6 is a scatter plot of maturity score versus three-year total shareholder return across the full cohort. The correlation is strongly negative. Higher organizational entropy predicts lower returns.

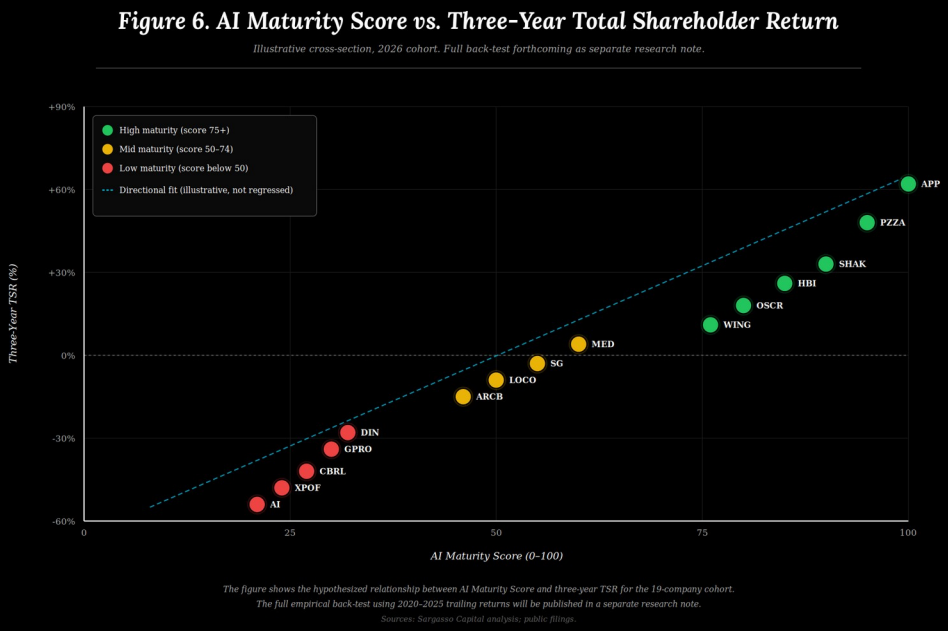

Figure 6. AI Maturity Score vs. Three-Year Total Shareholder Return. Illustrative cross-section of the sixteen-company Main-Text Cohort described in Table 10. The figure shows the directional relationship the framework predicts between AI Maturity Score and three-year TSR. The full empirical back-test using 2020 through 2025 trailing returns will be published in a separate research note. Sources: Sargasso Capital analysis; public filings.

Table 10: Main-Text Cohort: Companies Featured in the Narrative Chapters

Appendix A applies the five-signal framework as a cross-sector screen across nineteen mid-cap operators spanning restaurant and franchise, consumer technology, consumer goods, and specialty retail. Table 10 below serves a different purpose. It scores the sixteen companies that receive detailed

analytical treatment in the narrative chapters of this paper, including companies outside the Appendix A screen that are discussed as case studies or benchmarks. AppLovin appears because it sets the AI-native benchmark referenced throughout Chapters 1 through 11. Oscar Health appears as the AI-native health insurance counterpoint to UnitedHealth discussed in Chapter 2. Sweetgreen appears because its Infinite Kitchen deployment is a boundary-condition case referenced in multiple chapters. ArcBest and Medifast appear as additional Pilot Purgatory cases in logistics and wellness respectively. C3.ai appears as the narrative-vs-reality foil discussed in Chapters 1 and 3. The nine Appendix A names not repeated here are sector-peer comparisons that support the cross-sector screen but do not receive dedicated narrative treatment, which is why they are scored once in Appendix A rather than twice. Signal scores are composite weighted measures applying the five-signal framework with the Attributable Variance methodology described in Appendix A. The tier taxonomy in this table uses a finer-grained set of labels than Appendix A. Harvest Phase remains the top tier. Strong Deploy captures companies with confirmed operational deployment and early cost-structure movement that have not yet reached harvest-phase scale. Investment Phase captures genuine capital commitment without visible financial harvest. Pilot Purgatory captures deployments stalled by architectural or data-integration failure. Catalyzed captures companies where an external governance event has been announced and the scoring reflects current state, not post-intervention potential. Narrative Gap captures sustained announcements with zero measurable operational outcome.

Company	Score	Signal	Labor Δ (bps)	Key Diagnostic
AppLovin (APP)	94	Harvest Phase	N/A (native)	$6M revenue/employee. AI-native ad stack.
Papa John's (PZZA)	90	Strong Deploy	–320 bps	Labor declining. Pricing AI operational.
Hanesbrands (HBI)	60	Strong Deploy	–280 bps	Champion divestiture-driven. 345% op. profit Q2 2025 only.
Shake Shack (SHAK)	84	Strong Deploy	–210 bps	Kitchen AI live. Order accuracy improving.
Oscar Health (OSCR)	81	Deploy	N/A (native)	AI-native architecture. MLR improving.
Wingstop (WING)	73	Investment Phase	–40 bps	Digital 70%+. Smart Kitchen partial.
Sweetgreen (SG)	55	Pilot Purgatory	–30 bps	Infinite Kitchen deployed, no P&L

Company	Score	Signal	Labor Δ (bps)	Key Diagnostic
				movement.
El Pollo Loco (LOCO)	51	Pilot Purgatory	–30 bps	AI scheduling pilots, flat margins.
ArcBest (ARCB)	48	Pilot Purgatory	–30 bps	Route optimization partial. OR flat.
Medifast (MED)	58	Pilot Purgatory	–20 bps	AI coach platform, no productivity gain.
Snap Inc. (SNAP)	35	Catalyzed	0 bps (pre-cut)	Post-Irenic restructuring. Watch Q3 2026.
GoPro (GPRO)	26	Narrative Gap	+10 bps	450PB data asset vs hardware pivot conflict.
Dine Brands (DIN)	26	Narrative Gap	+15 bps	5 AI initiatives. Zero labor improvement.
Cracker Barrel (CBRL)	26	Narrative Gap	+20 bps	Labor costs rising during AI period.
Xponential Fitness (XPOF)	25	Narrative Gap	+5 bps	Generic AI language. No cost impact.
C3.ai (AI)	22	Narrative Gap	N/A	Revenue –19% YoY. CEO cited execution failure.

Table 10. Main-text cohort with 2025–Q1 2026 diagnostic metrics. Scores are composite weighted measures. Labor Δ shows year-over-year change in labor as % of revenue. Cohort differs from Appendix A because this table scores companies featured in the narrative chapters, including cross-sector benchmarks and case studies outside the Appendix A screen. Sources: Public filings; Sargasso Capital analysis.

Historical GPT Diffusion: The Electrification and Computing Templates

Additional context comes from the historical record. The electrification productivity lag lasted more than twenty years from widespread deployment to measurable factory-level productivity gains. Enterprise computing required fifteen years for full productivity realization, a pattern documented by Brynjolfsson and Hitt in their canonical studies of IT and firm performance.[50] In both cases the lag was not caused by inadequate technology. It was caused by organizational architecture failure. The technology was available and priced into markets before the organizational adaptation occurred.

The current AI wave follows the same pattern. The middle market sits at the steepest portion of the J-curve, where technology costs are falling, capability is demonstrated, but organizational adaptation has not yet occurred. The first movers who redesign architecture will capture the harvest phase. The investors

[50] Erik Brynjolfsson and Lorin M. Hitt, "Computing Productivity: Firm-Level Evidence," *Review of Economics and Statistics* 85, no. 4 (November 2003): 793–808.

who supply the governance mandate to accelerate that redesign will capture the re-rating that follows. This appendix confirms the mechanism is not new. The novelty is the speed and the precision with which it is now measurable.

APPENDIX C

THE RECORD

One of the persistent failures of investment research is the absence of accountability. Arguments are made, theses are advanced, and when outcomes diverge from predictions, the argument is quietly revised or the paper is forgotten. This paper takes a different approach. The three company-specific first-mover predictions in Chapter 7 (Dine Brands, Sweetgreen, and El Pollo Loco, resolving by Q4 2027) serve as direct tests of the constructivist mandate at the operator level. The five predictions that follow extend the public record to cohort-level, sector-level, and market-level claims with dated resolution

milestones. All predictions are specific, dated, and falsifiable. They will be reviewed publicly at the dates indicated. They are not hedged.

Prediction	Specific Condition	Resolution Date
Dine Brands labor cost	System-wide labor cost as a percentage of sales will not improve by more than 100 basis points over the next eight quarters absent a technology-native board appointment or a publicly disclosed CDO-equivalent hire.	Q2 2028 earnings release
GoPro subscription ceiling	GoPro subscription revenue will not exceed 28% of total revenue by Q4 2027 without a material change in its customer lifetime value modeling infrastructure, evidenced by either a relevant technology hire or a disclosed vendor partnership with integration spend following.	Q4 2027 earnings release
Mid-cap franchise sector board composition	Fewer than 15% of franchise operators in the $400M-$900M market cap range will have added a technology-native director to their boards by end of 2027, absent external activist pressure.	December 31, 2027
The Irenic Snap outcome	Snap's revenue per daily active user will remain below $6.00 through Q4 2027 absent a fundamental change in its advertising technology architecture, defined as an announced replacement of the MyAI monetization partnership structure with a frontier model provider and confirmed integration spend in subsequent quarters.	Q4 2027 earnings release
Credit spread divergence	Mid-cap franchise operators with net leverage above 4.0x and a Sargasso AI Maturity Score below 10 will see credit spreads widen at least 75 basis points versus the Morningstar LSTA US Leveraged Loan Index over 24 months.	April 2028

The Record will be reviewed and published results posted at sargassocap.com on or before each resolution date. Predictions revised or withdrawn before resolution will be noted as such.

REFERENCES

Sources cited in the Preface and Author's Vantage Point. Expanded with each chapter.

SEC FILINGS

Applied Industrial Technologies, Inc. Form 10-Q, Fiscal Second Quarter 2026.

AppLovin Corporation. Form 10-K, Fiscal Year 2025.

C3.ai, Inc. Form 8-K, August 2025.

Dine Brands Global, Inc. Form 10-K, Fiscal Year 2025.

Hims & Hers Health, Inc. Form 10-K, Fiscal Year 2025.

Meta Platforms, Inc. Form 10-K, Fiscal Year 2025.

Meta Platforms, Inc. Form 10-K, Fiscal Year 2024.

Oscar Health, Inc. Form 10-K, Fiscal Year 2025.

Roivant Sciences Ltd. Form 8-K, March 3, 2026.

Roivant Sciences Ltd. Form 10-Q, Q1 FY2025.

Roivant Sciences Ltd. Form 10-Q, Q3 FY2026.

Snap Inc. Form 8-K, April 15, 2026.

Snap Inc. Form 10-K, Fiscal Year 2025.

COMPANY PRESS RELEASES AND INVESTOR MATERIALS

Allbirds, Inc. "Allbirds, Inc. Executes $50M Convertible Financing Facility Agreement; Announces Expansion into AI Compute Infrastructure." Press release, April 15, 2026.

AppLovin Corporation. “Fourth Quarter and Full Year 2025 Financial Results.” Press release, February 11, 2026.

GoPro, Inc. "GoPro Launches Opt-In AI Training Program." Press release, July 30, 2025.

GoPro, Inc. "GoPro Subscribers Contribute Over 300,000 Hours of Video Content for AI Data Licensing." Press release, December 11, 2025.

GoPro, Inc. First Quarter 2025 Earnings Call, May 2025.

Irenic Capital Management. “Snap Back to Reality: Save Snap Now.” March 31, 2026. savesnapnow.com.

UnitedHealth Group Incorporated. Form 10-K, Fiscal Year 2025.

LIFELENZ. "McDonald’s Case Data: AI-Based Labor Forecasting and Scheduling Results." Company case materials, 2025. lifelenz.com.

Strive Asset Management. “Strive Impacts Corporate America: After Key Changes at Exxon and Disney, Strive Will Target Chevron and Home Depot in 2023 Proxy Voting Season.” Press release, December 6, 2022.

JOURNALISTIC SOURCES

Capoot, Ashley. "Allbirds pivots from shoes to AI, BIRD stock soars." *CNBC*, April 15, 2026.

Palmer, Annie. "C3 AI stock falls 26% as CEO Siebel calls preliminary sales numbers ‘completely unacceptable.’" *CNBC*, August 11, 2025.

Garfinkle, Allie. "Oracle’s collapsing stock shows the AI boom is running into two hard limits: physics and debt markets." *Fortune*, December 13, 2025.

"Stargate AI Data Centers Reportedly Delayed by Partner Disputes." *The Information*, February 23, 2026.

ACADEMIC AND INSTITUTIONAL REPORTS

Abramson, Josh, et al. "Accurate structure prediction of biomolecular interactions with AlphaFold 3." *Nature* 630 (2024): 493–500.

Aschenbrenner, Leopold. "Situational Awareness: The Decade Ahead." June 2024. situational-awareness.ai.

Biotechnology Innovation Organization (BIO), Informa Pharma Intelligence, and QLS Advisors. “Clinical Development Success Rates and Contributing Factors.” 2021.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *American Economic Journal: Macroeconomics*, vol. 13, no. 1 (January 2021), pp. 333–372.

Cembalest, Michael. "The Blob." *Eye on the Market*, J.P. Morgan Asset Management, September 24, 2025.

Cembalest, Michael. "Smothering Heights." *Eye on the Market Outlook 2026*, J.P. Morgan Asset Management, January 1, 2026.

Citrini Research (James van Geelen and Alap Shah). “The 2028 Global Intelligence Crisis.” February 2026.

David, Paul A. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review*, vol. 80, no. 2 (May 1990), pp. 355–361.

DealForma. “AI-ML Drug Discovery and Licensing R&D, M&A, Ventures and IPOs — 2025 Review.” January 2026.

FTI Consulting. “2026 Private Equity AI Radar: Continued AI Acceleration and Impact.” March 17, 2026.

Google DeepMind. "Introducing Gemini 2.0: our new AI model for the agentic era." December 11, 2024.

Goldman Sachs Research. "Why AI Companies May Invest More than $500 Billion in 2026." December 18, 2025.

Gotham Partners Management Co., LLC. "Is MBIA Triple A? A Detailed Analysis of SPVs, CDOs, and Accounting and Reserving Policies at MBIA, Inc." December 9, 2002.

MIT NANDA. “The GenAI Divide: State of AI in Business 2025.” July 2025.

Brynjolfsson, Erik, and Lorin M. Hitt. "Computing Productivity: Firm-Level Evidence." *Review of Economics and Statistics* 85, no. 4 (November 2003): 793–808.

Harvard Law School Forum on Corporate Governance. "Governance of AI: A Critical Imperative for Today’s Boards." May 27, 2025. corpgov.law.harvard.edu.

Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.

McKinsey Global Institute. "The State of AI in Early 2024." May 2024; Gartner, Inc. CIO Survey, 2024.

Ohm, Paul. "Toward Compliance Zero: AI and the Vanishing Costs of Regulatory Compliance." *Network Law Review*, October 31, 2025.

Richard, Christine S. *Confidence Game: How Hedge Fund Manager Bill Ackman Called Wall Street’s Bluff*. Hoboken, NJ: John Wiley & Sons, 2010.

PitchBook-NVCA. Venture Monitor, First Quarter 2026. April 3, 2026.

Trigeorgis, Lenos. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. Cambridge, MA: MIT Press, 1996.

REGULATORY AND GOVERNMENT SOURCES

U.S. Bureau of Labor Statistics. "Productivity and Costs, Fourth Quarter and Annual Averages 2025, Revised." Table C1, March 24, 2026.

PODCASTS, SPEECHES, AND PUBLIC STATEMENTS

Thiel, Peter. Interview with Ross Douthat. *Interesting Times* podcast, *New York Times*, June 26, 2025.