AUTOKATALYST

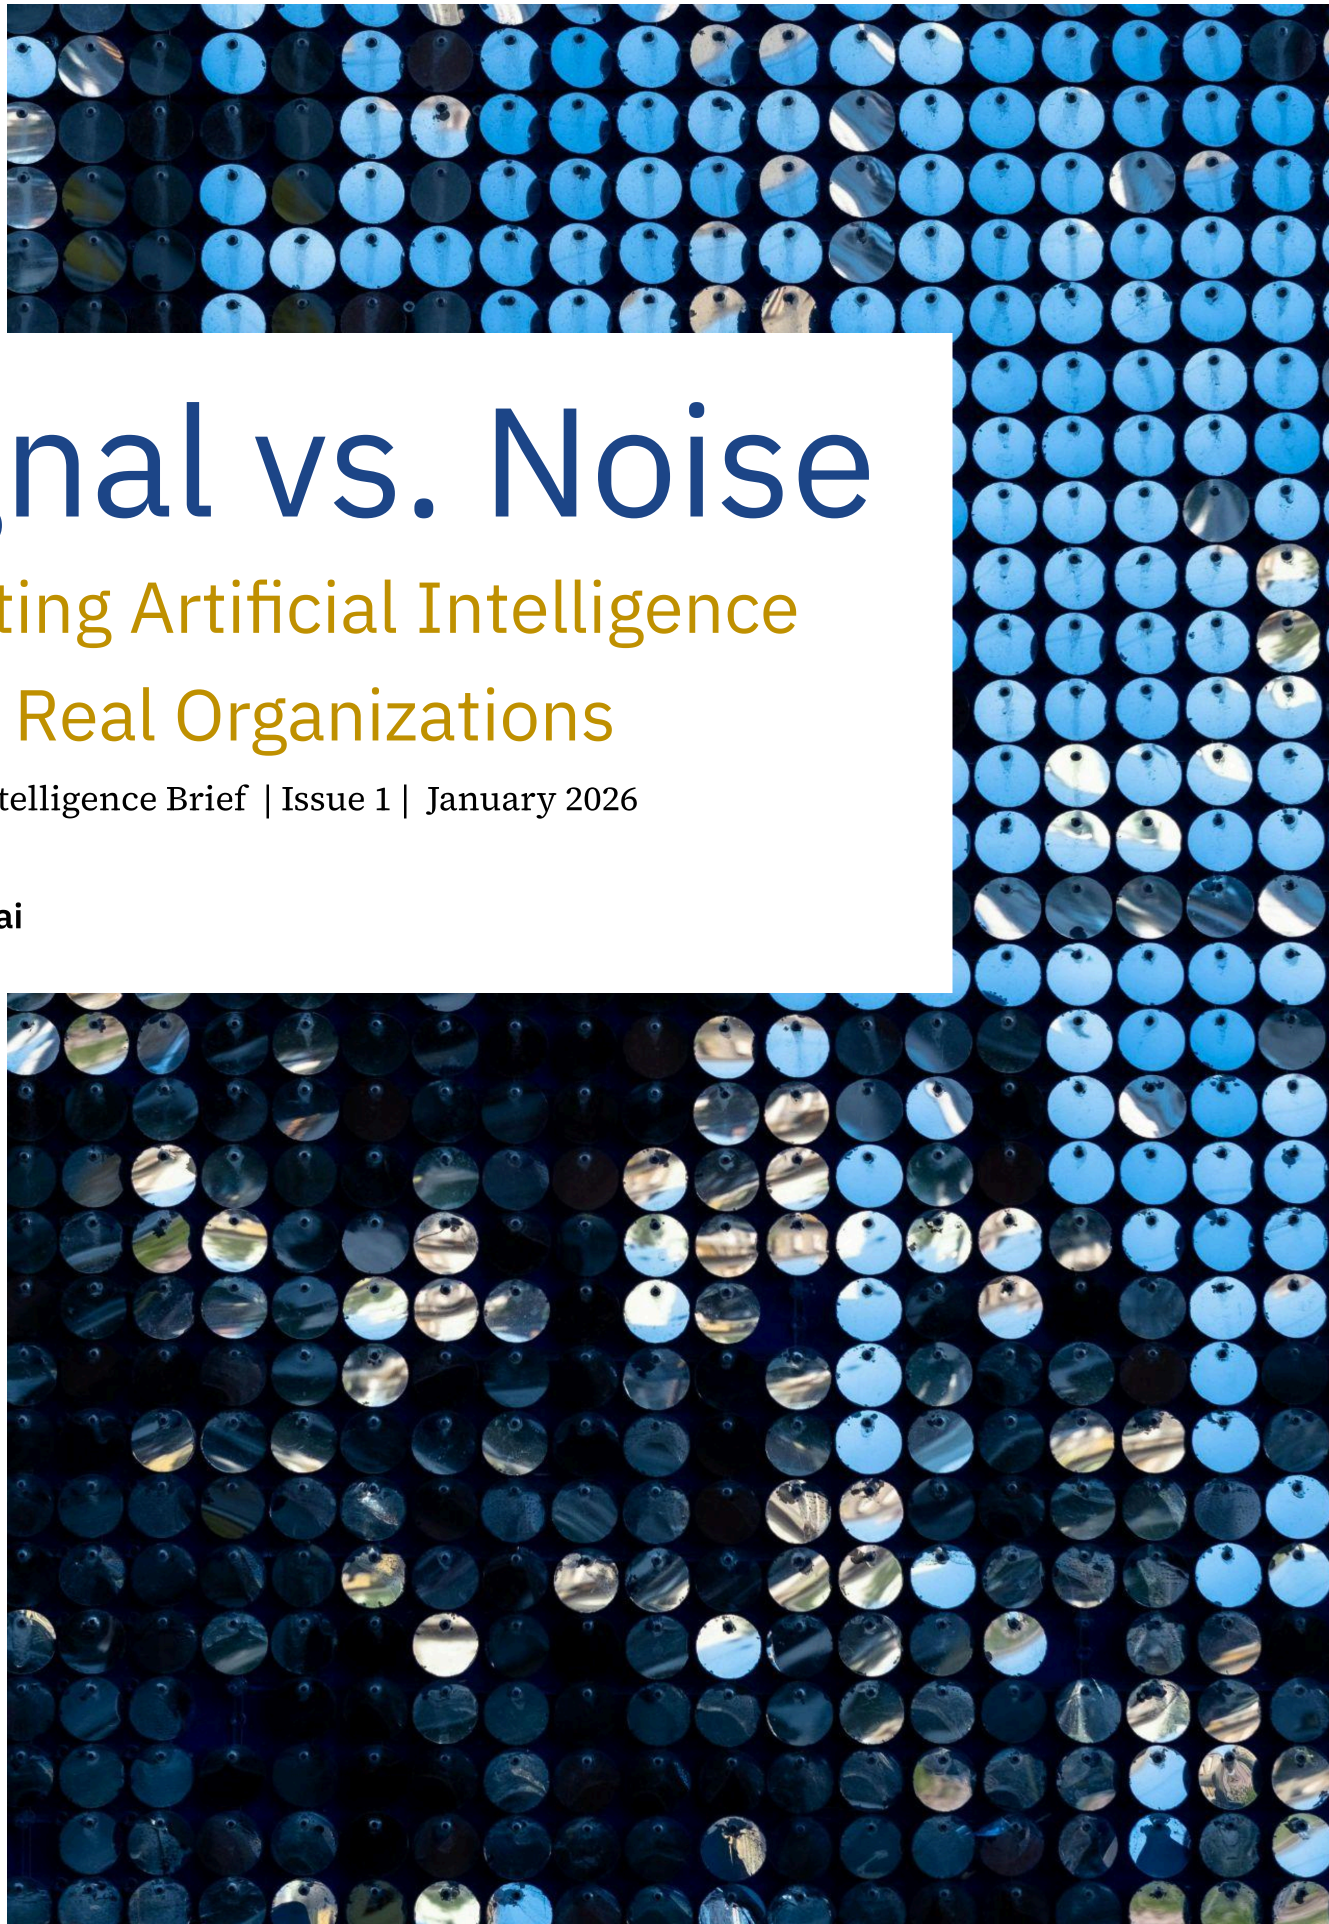

Signal vs. Noise

Operating Artificial Intelligence Inside Real Organizations

Executive Intelligence Brief | Issue 1 | January 2026

By Sami Baqai

A Note to the Reader

This brief is written for executives who are responsible for outcomes, not narratives.

AI is no longer speculative. Capital has been committed, infrastructure is being built, and expectations now sit inside operating plans, board discussions, and investor conversations. The question is no longer whether the technology is real, but why translating it into durable advantage remains uneven and difficult.

This document focuses on that gap.

Across real deployments, the same constraints show up repeatedly. Not in the models, but in the organizations trying to use them: How work is structured; How decisions are made under uncertainty; How efficiency gains are captured, or quietly competed away.

Those issues are rarely addressed head-on. Much of the public conversation is either promotional or theoretical. This brief is neither.

Issue 1: Signal vs. Noise is intentionally a map, not a deep dive. It lays out the lenses we use to evaluate AI inside real organizations and explains why adoption so often produces noise instead of signal. Future briefs will take these areas one at a time and examine them in depth.

This is not a guide to tools or vendors. It does not argue whether AI “works” in the abstract. Its purpose is more practical: to help leaders reason more clearly about where AI changes outcomes, where it does not, and why well-intentioned initiatives stall more often than they compound.

If it sharpens your questions, even more than it supplies answers, it has done its job.

Signal vs. Noise

Issue 1 — January 2026

By almost any reasonable measure, 2025 was the year of AI.

We might have said the same for 2024 (and 2023), and if current trends continue, we may yet be tempted to say it again for 2026. There is no shortage of candidate metrics: the sheer volume of new AI products and services, the proliferation of trillion-parameter models, the steady re-release of incumbent software branded as "AI-powered," or the sharp rise in how often AI appears in executive earnings calls. But among all these measures, one stands apart for its decisiveness and irreversibility: the money.

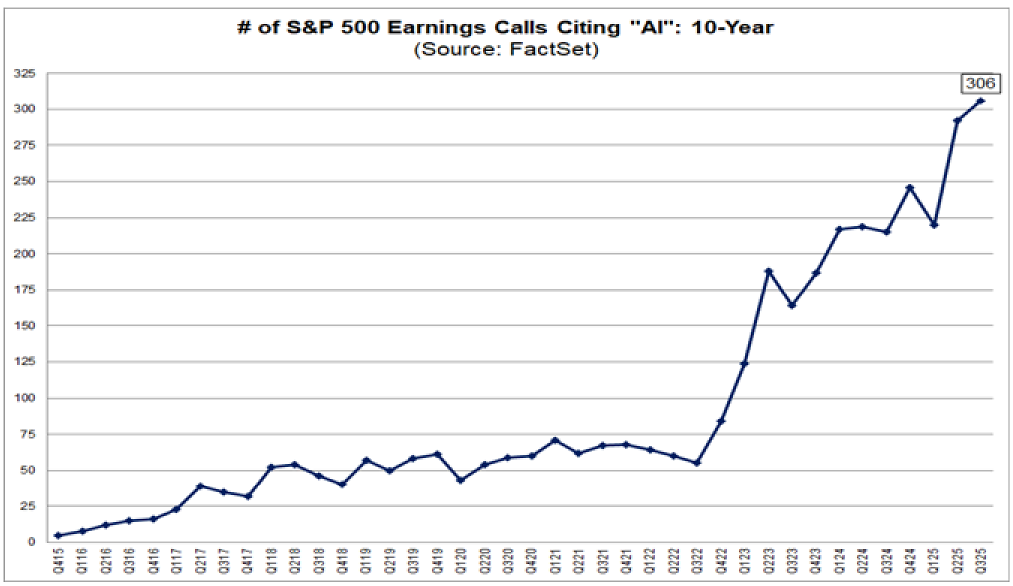

Highest number of S&P earnings calls citing "AI" over the past 10 years.[1]

Capital flooded the system at a scale rarely seen outside wartime mobilizations or generational infrastructure shifts. Venture investors put well north of $200 billion into AI-related companies, much of it concentrated in a small group of foundation-model players.[2] At the same time, the real economy showed up. Hyperscalers committed hundreds of billions of dollars in capital expenditures to data centers, custom chips, networking, and power - investments explicitly justified by AI demand. Private equity followed suit. Blackstone announced multi-tens-of-billions in AI-driven infrastructure investments. Microsoft guided to roughly $80 billion in annual capex tied largely to AI;

Meta was not far behind. These were not exploratory bets. They were long-dated, balance-sheet commitments. The kind you make only if you believe something structural is underway.

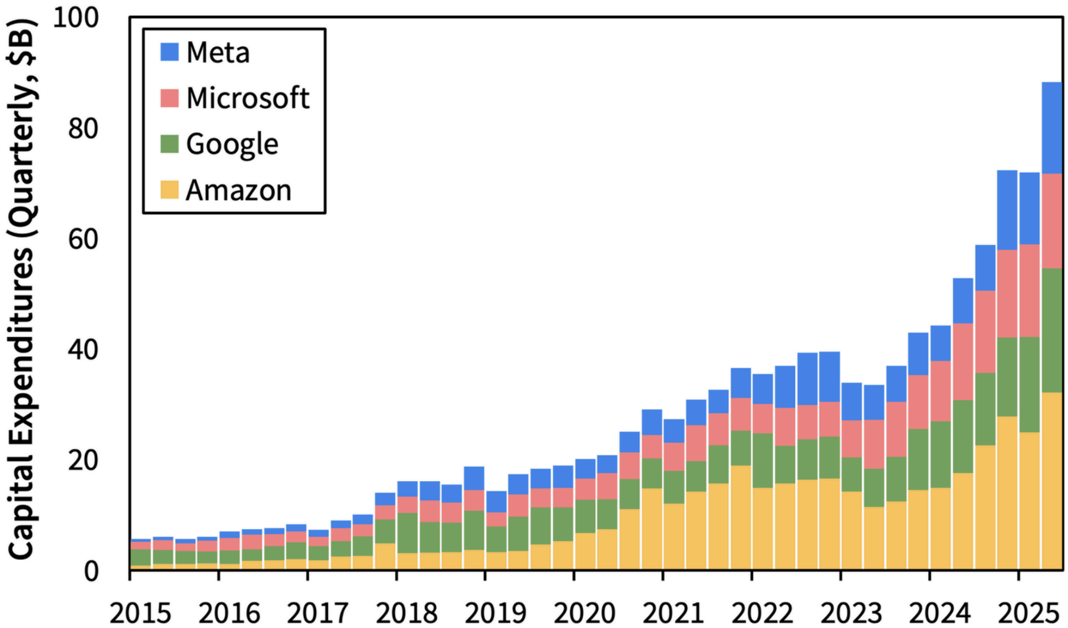

From Q1 2015 to Q2 2025.[3]

That, at least, is the view from capital. This level of coordinated investment - across venture, private equity, public markets, and infrastructure - does not materialize on *vibes* alone. It assumes AI will become a durable, compounding layer of the global economy.

And then, in the middle of all that, Merriam-Webster looked out at what people were actually experiencing and named its 2025 Word of the Year: "**slop**" - defined as low-quality digital content, often produced at scale by artificial intelligence: [4]

> The flood of slop in 2025 included absurd videos, off-kilter advertising images, cheesy propaganda, fake news that looks pretty real, junky AI-written books, "workslop" reports that waste coworkers' time… and lots of talking cats. People found it annoying, and people ate it up.
>
> *"AI Slop is Everywhere,"* warned *The Wall Street Journal*, while admitting to enjoying some of those cats. *"AI Slop Has Turned Social Media Into an Antisocial Wasteland,"* reported *CNET*.

> Like *slime, sludge,* and *muck, slop* has the wet sound of something you don't want to touch. Slop oozes into everything. The original sense of the word, in the 1700s, was "soft mud." In the 1800s it came to mean "food waste" (as in "pig slop"), and then more generally, "rubbish" or "a product of little or no value."

Not exactly the victory lap you'd expect after a year of nine-figure seed rounds and grid-scale buildouts.

We have "superintelligence" in the pitch decks and "slop" in the lived environment. That mismatch is telling. It reflects what the current AI ecosystem is actually optimized to produce and reward: scaling compute, shipping new features, and generating output, rather than judgment, integration, or accountability.

What we mostly got, so far and with this generation of transformer-based systems, is a new kind of abundant output: impressively fluent, sometimes useful, often fragile, and increasingly expensive to run at scale.

Many operators recognize the disconnect between the scale of investment and the clarity of real-world results. The resulting caution is not resistance to technology; it is a rational response to unresolved risk. Deploying these systems inside real workflows introduces complexity that cannot be papered over with slick demos, toy examples or talking cats.

This disconnect between rapidly advancing AI capabilities and their real-world deployment matters, because the consequences are no longer abstract.

2025 was also notable for AI-related controversies, failures, and emerging threat vectors. In Australia, Deloitte partially refunded the government AU$440,000 for a contracted report that was partially drafted using AI and contained fabricated citations and factual errors.[5] In August, *Raine v. OpenAI* [6] became a landmark legal case, alleging that a conversational AI system failed to deter, and in fact validated, self-harm plans of a minor who later died by suicide. Similar lawsuits involving minors and chatbot interactions have since been filed against other AI companies, with some high-profile cases reaching settlement.[7]

Beyond legal and reputational risk, new technical threat patterns are emerging. In late 2025, Google's Threat Intelligence Group warned of AI-enabled malware families capable of generating and modifying malicious code in real time during active cyberattacks -

adapting to defenses, evading detection, and rewriting themselves mid-operation.[8] This represents a meaningful shift in adversary behavior and raises the stakes for organizations deploying AI systems without clear guardrails, monitoring, and ownership - especially as those systems move deeper into core processes.

Taken together, these signals point to a simple reality: AI capability is advancing faster than most organizations' ability to absorb it. The constraint is not models or compute. It is how work is structured, how decisions are made, and how value is captured, or lost, once technology enters real systems.

This brief is written for executives navigating that friction. Not to argue whether AI "works" in the abstract, but to help reason about when it works in practice, when it fails quietly, and why well-intentioned initiatives so often stall or disappoint.

We are not trying to adjudicate hype, forecast job losses, or speculate about artificial general intelligence. Those debates are already well covered. Our focus is narrower and more operational: improving the signal-to-noise ratio for leaders making real decisions under constraint.

Issue 1 is intentionally a map, not an expedition. Rather than going deep on a single problem, this inaugural brief surveys the terrain we believe matters most, and is least clearly discussed, in the popular AI conversation. Future issues will take these areas one at a time and explore them in depth.

The lenses that follow are how we make sense of AI inside real organizations. Each is incomplete on its own. Together, they explain why AI adoption so often produces noise instead of advantage.

How Work Is Structured

From Assembly Lines to Knowledge Work

Most AI initiatives fail not because the models are weak, but because the work they are applied to was never designed for *systems* in the first place.

This is not a new problem. We have been here before.

The moving assembly line cut Model T production times from 12 hours to under 90 minutes.[9]

In the early days of automobile manufacturing, cars were built much the way knowledge work is performed today: artisanal, localized, and heavily dependent on individual craftsmanship. A small group of skilled workers assembled entire vehicles in one place, relying on tacit knowledge, informal coordination, and ad hoc problem-solving. Output quality varied, scale was limited, and costs were high. The system worked, but only within narrow bounds.

The introduction of the assembly line did not merely speed things up. It fundamentally changed how work was structured. Tasks were decomposed, sequenced, standardized, and engineered so that the system, not the individual, carried the burden of coordination. This enabled consistency, reliability, and scale. It also enabled automation later on. This key point is often missed, especially in contemporary discussions about AI: automation followed process design; it did not precede it.

Most knowledge work today still looks far closer to the pre-assembly-line factory than to the modern production system. Work happens inside email threads, meetings, spreadsheets, slide decks, and chat tools. Decisions rely heavily on tacit judgment. Handoffs are informal. Accountability is often narrative rather than structural. We call this “knowledge work,” but in practice it is largely artisanal work performed on computers.

This reality is conveniently absent from much of the “AI will take everyone’s jobs” discourse. Giving people access to powerful language models does not, by itself, change how work is structured. It adds a new capability into an existing system - but the system remains the system.

To be clear, chatbots can be genuinely useful. They can speed up drafting, summarize information, reduce friction, and lower the cost of certain cognitive tasks. But they operate on top of human-friendly processes, not machine-friendly ones - processes designed for people rather than systems. They assist the craftsperson; they do not redesign the factory.

The more interesting, and more difficult, question is whether system-friendly processes could be dramatically more efficient, reliable, interpretable, and scalable than today’s LLM-augmented artisanal workflows. We already know the answer in some domains.

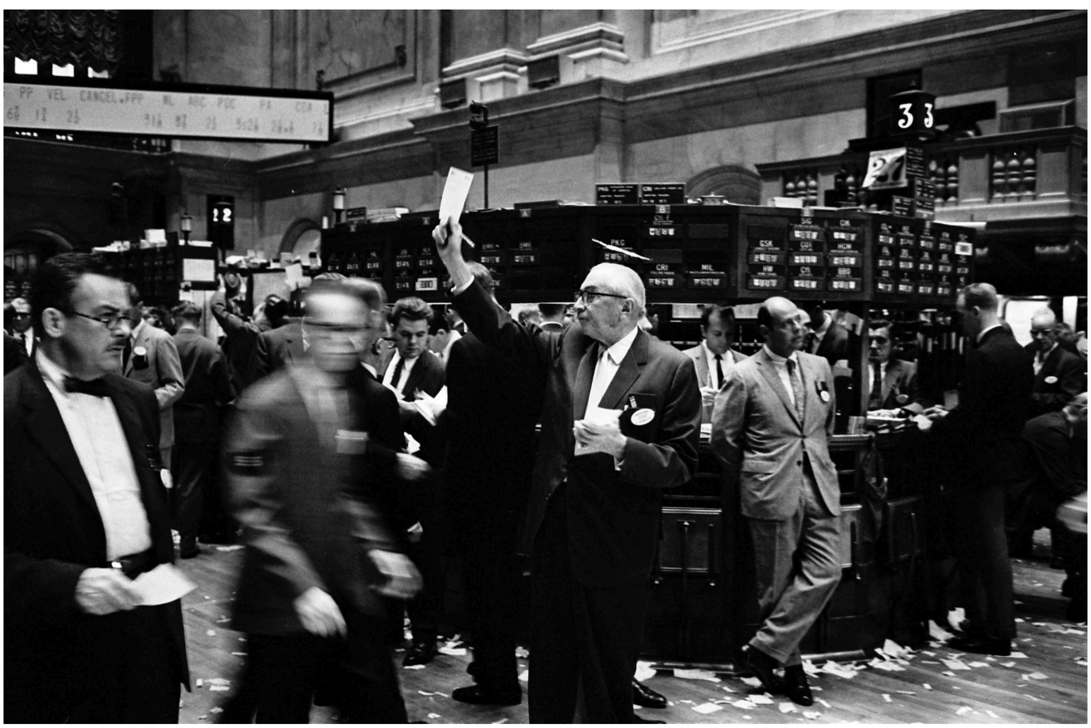

The floor of the New York Stock Exchange in September 1963.[10]

Consider financial markets. For decades, trading floors were loud, physical, human systems. Price discovery happened through shouted bids, hand signals, and personal relationships. It worked… until it didn't. As markets scaled in volume, speed, and complexity, these human-centric processes became bottlenecks. They were eventually replaced by electronic trading systems: engineered workflows with explicit rules, defined inputs and outputs, deterministic execution paths, and auditable outcomes. Humans did not disappear, but their role changed. Judgment moved upstream; execution moved into systems.

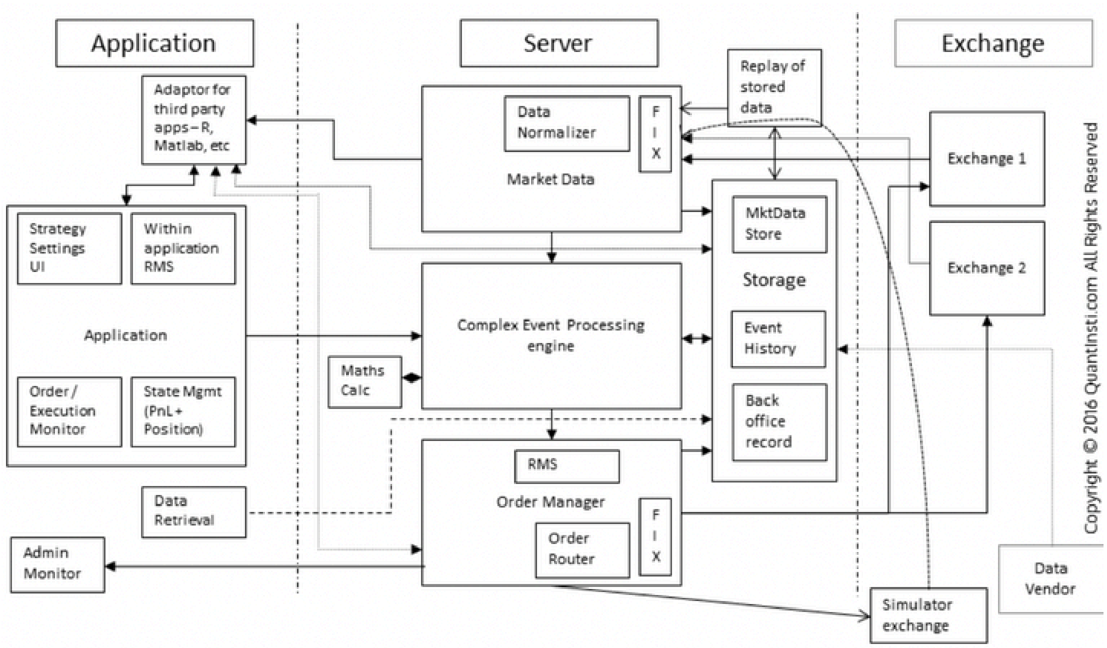

Automated Trading System design and architecture.[11]

This pattern has repeated elsewhere: automated checkout lines, supply chain optimization, credit underwriting, fraud detection. In each case, the breakthrough did not come from "smarter humans," but from restructuring the work so that systems could carry more of it.

This is the uncomfortable part of the AI conversation. Automation is downstream of process design, not upstream. You cannot "LLM-ize" a poorly structured workflow and expect durable gains. In fact, doing so often makes problems harder to see. The system produces fluent output, but the underlying work remains fragmented, opaque, and fragile.

Executives encounter this dynamic every day when they integrate and utilize third-party software systems. A CRM or ERP does not merely provide software; it imposes a workflow. By using it, organizations implicitly hand over some control of how work is sequenced, validated, and recorded. This is precisely why such systems scale across thousands of customers: the vendor has done the work of designing a narrow but system-friendly "conveyor belt" for specific tasks.

The trade-off is obvious. You gain consistency and scale, but you give up some flexibility. Not every process should be reengineered this way. But where bottlenecks limit growth, invite risk, or create hidden costs, it is worth asking whether the work itself, not just the tools, needs redesign.

What executives often miss is that AI magnifies existing process quality. Well-designed workflows become faster and more resilient. Poorly designed ones become faster sources of error. Bad processes do not become good processes when augmented with AI; they become faster bad processes.

This is why so many AI pilots show promise and then stall. They deliver local efficiency but fail to compound. The underlying work structure has not changed, so the system cannot absorb the capability.

The implication is not that every organization should embark on wholesale process reengineering. It is that leaders should be deliberate about *where* AI is applied and *what kind* of work it is applied to. In some cases, assisting human judgment is the right move. In others, the real opportunity lies in redesigning the workflow so that judgment happens once, upstream, and execution happens repeatedly, downstream.

Future briefs will go deeper into how organizations identify which work is ready for systematization, how they redesign workflows without freezing innovation, and why skipping this step almost guarantees disappointment. For now, the takeaway is simple but easy to ignore: before asking what AI can do for your organization, it is worth asking whether your work is structured in a way that systems can meaningfully carry.

How Decisions Are Made

Probabilistic Thinking in Organizations That Hate Uncertainty

Picture this: It is the 2028 presidential election. Two insurgent candidates have fought their way through crowded primaries and now face off in a closely watched general election. Pollsters, statisticians, and prediction markets are unusually aligned: Candidate A is given a 90% chance of winning. Candidate B, just 10%.

Election night arrives. Results come in quickly. Candidate B wins.

The next morning is familiar. Headlines announce failures. Analysts dissect polling errors. There are public postmortems, apologies, and debates about broken models and overconfidence. The implicit conclusion is clear: *something must have gone wrong*.

But nothing did.

A 10% outcome is not an error. It is an outcome that was explicitly anticipated. The real failure is not statistical; it is interpretive. We consistently mistake probability for certainty, especially in hindsight.

Now consider a different decision: Would you buy an airline ticket that was 50% cheaper if the airline had a 10% crash rate? What if the ticket were free?

Most people would not hesitate. One in ten is unacceptable. Modern commercial aviation has a fatality rate orders of magnitude lower, and even small deviations from that baseline feel intolerable.

What's striking is not the answer, but the contrast. In the election scenario, we do not naturally think "there is a one-in-ten chance Candidate B will win." We think "Candidate A is going to win." When the unlikely outcome occurs, we search for blame.

This asymmetry, between how probabilistic information is casually interpreted and how outcomes are judged in hindsight, sits at the center of AI adoption in organizations.

AI systems, at their core, are probabilistic. They do not produce certainty; they produce likelihoods. Predictions, classifications, forecasts, and recommendations are all expressions of uncertainty, even when delivered with high confidence. As organizations deploy AI to automate or augment decisions, they are forced to confront something they have long avoided: explicit uncertainty.

Human decision-making has always been uncertain, but it is *narratively uncertain*. We tell stories after the fact. We justify outcomes. We smooth over variance. AI systems, by contrast, surface uncertainty numerically and in advance. The challenge here is organizational, not technical.

To operate effectively in this environment, organizations need capabilities that are rarely well-developed in traditional business settings: probabilistic thinking, expected value reasoning, and systems-level judgment. These are not natural skills, and they often run directly against institutional incentives.

Expected value provides a useful illustration. Imagine a simple game: a fair coin is flipped. If it lands heads, you win $1,000. If it lands tails, you win nothing. The question is not whether you would play the game, but what the maximum rational price of entry would be.

The mathematically correct answer is just under $500. Play this game repeatedly at any lower price, and you are expected to make money over time. This is not a trick; it is simply the arithmetic of uncertainty.

Yet organizations routinely reject decisions with positive expected value because they focus on individual outcomes rather than distributions. A single visible failure can outweigh a long series of quiet successes, especially when reputational risk is involved.

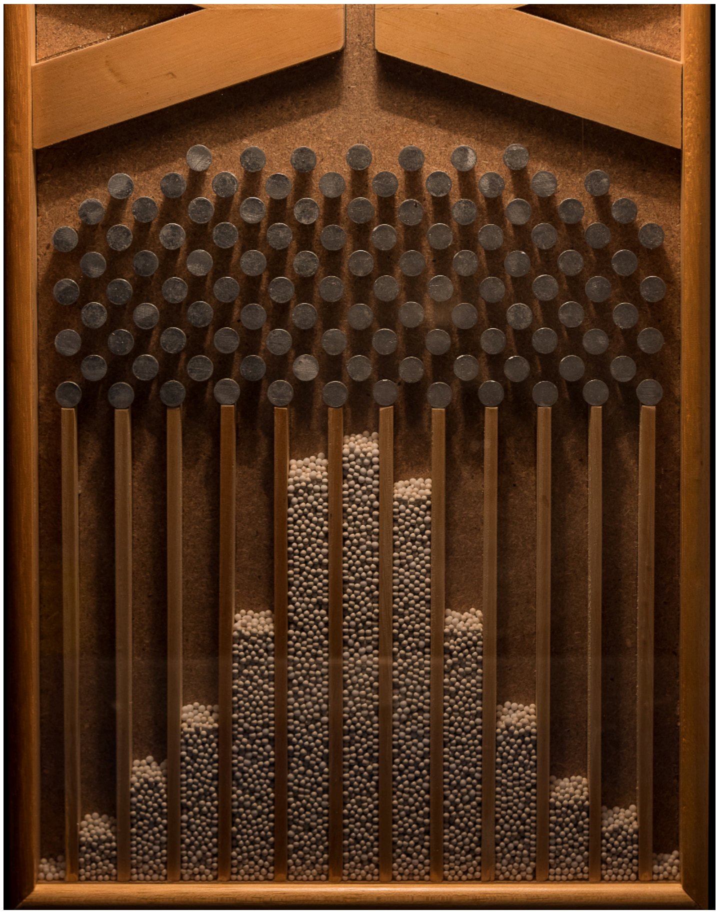

Galton Board: Individual outcomes are noisy and unpredictable. Aggregate behavior is not.[12]

The COVID-19 pandemic provided a rare counterexample. Faced with extraordinary uncertainty and enormous downside risk, governments and pharmaceutical companies pursued multiple vaccine approaches in parallel, building manufacturing capacity before knowing which candidates would succeed. Most of those investments were expected to fail. That was the point. Measured in lives saved and speed of economic recovery, the expected value of accelerating a successful vaccine, from a typical decade-long timeline to roughly ten months, dwarfed the cost of redundancy.[13]

This kind of reasoning is difficult for most organizations. It requires tolerating visible waste in service of invisible benefit. It requires separating the quality of a decision from the outcome of any single instance. And it requires governance structures that can withstand public and internal scrutiny when things go wrong *as expected*.

These challenges become more acute when decision-making is automated and scaled. A human manager who makes a bad call can explain their reasoning. A system that produces a statistically correct but locally unpopular outcome invites a different reaction: *the model was wrong*.

This distinction matters. Was the prediction wrong, or was the decision appropriate given the risk tolerance of the organization? When an AI system correctly predicts that 10% of customers will churn under a new pricing strategy, and that outcome materializes, has the system failed? Or has it done exactly what it was designed to do?

Organizations often struggle to answer this because they have not clarified who owns outcomes when systems influence decisions. Is accountability assigned to the model, the team that deployed it, or the executive who approved its use? Without clear answers, the safest path is often to avoid probabilistic systems altogether, or to use them in advisory roles where their outputs can be ignored when inconvenient.

This dynamic helps explain why suboptimal but familiar decision-making persists. Narrative-driven judgment hides uncertainty and diffuses blame. Probabilistic systems surface uncertainty and concentrate accountability. The latter feels riskier, even when it is objectively more rigorous.

This conflict is vividly illustrated in Michael Lewis's 2011 film *Moneyball*. The story is not about baseball per se; it is about institutional inertia. Statistical approaches to player evaluation challenged long-standing intuitions and threatened existing hierarchies. Individual decisions looked strange in isolation. In aggregate, they worked. The resistance had little to do with accuracy and everything to do with authority.

"Guys, you're still trying to replace Giambi. I told you we can't do it, and we can't do it. Now, what we might be able to do is re-create him. Re-create him in the aggregate." [14]

AI introduces this same conflict into far more consequential domains. It forces organizations to choose between decision-making that feels right and decision-making that performs well over time. It exposes the trade-offs between reputational safety and economic value. And it raises uncomfortable questions about who gets to be right when outcomes are uncertain.

Future briefs will explore how some organizations redesign incentives, governance, and accountability to operate effectively under uncertainty, and why others retreat to safer, less effective modes of decision-making. For now, the takeaway is straightforward: AI does not reduce uncertainty. It forces organizations to confront it directly.

None of this requires executives to become statisticians. But organizations that cannot tolerate explicit uncertainty will struggle to move beyond superficial adoption. Those that can learn to reason probabilistically, at the level of systems, not anecdotes, will find that uncertainty, properly handled, becomes a source of advantage rather than a liability.

How Advantage Is (and Isn't) Created

The Rising Tide vs. Real Differentiation

Not all AI adoption creates competitive advantage. Much of it simply raises the baseline.

This is not a new problem. Long before artificial intelligence entered the business lexicon, firms grappled with the same strategic question: *when does technological improvement translate into durable advantage, and when does it simply get competed away?*

A useful illustration comes from Warren Buffett's early career. One of his first major acquisitions was a struggling textile manufacturer called Berkshire Hathaway, the business that he would eventually transform into a renowned capital-allocation vehicle and diversified conglomerate. But in its original form, the textile operation faced relentless technological pressure. New looms and manufacturing techniques steadily improved efficiency, yet those gains never accrued to the owners. Competitive dynamics ensured that cost savings flowed almost immediately to customers in the form of cheaper textiles. Margins remained thin. Capital requirements stayed high. The business demanded constant reinvestment just to stand still.

Buffett later described the acquisition as a mistake. The lesson he took from it shaped decades of thinking about competitive advantage: efficiency gains alone do not create moats. In highly competitive markets, they are often passed through to customers. What matters is not whether a firm becomes more efficient, but whether it can *capture* the value of that efficiency.

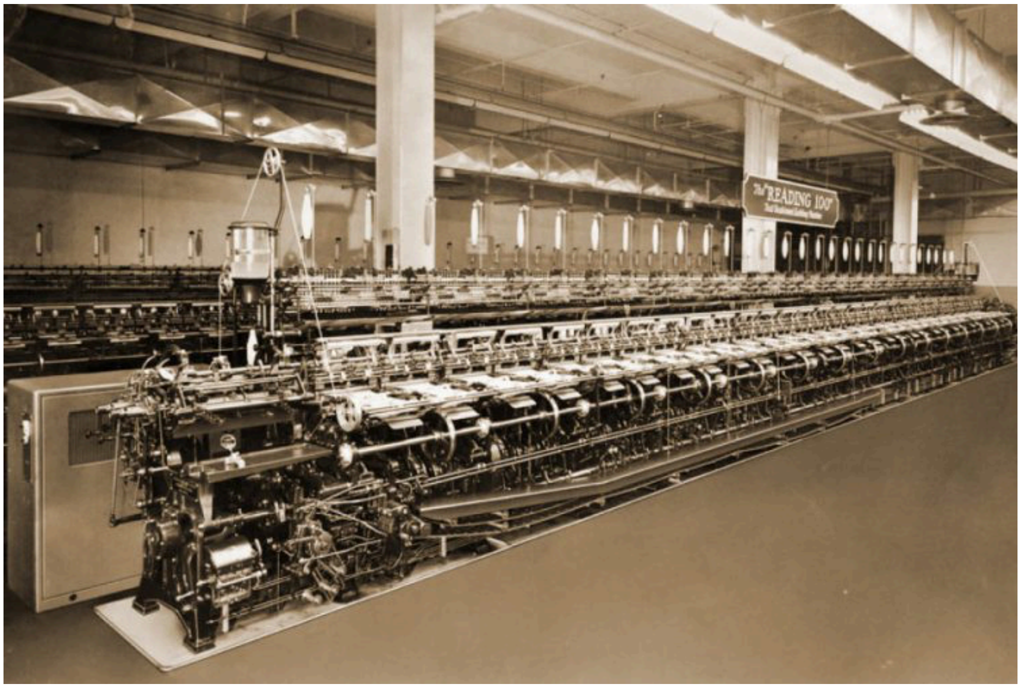

"Reading" knitting machine - a cutting-edge system designed and built in-house at Berkshire to dramatically improve productivity. The efficiency gains never translated into durable profits.[15]

That lesson maps directly onto AI.

Executives often ask a version of the same question: *If AI makes our organization 30% more efficient, who benefits?* Do those gains show up as lower prices for customers? Higher wages for employees? Or do they accrue to the firm in the form of improved margins and resilience?

The answer is not inherent in the technology. It depends on structure, market dynamics, and where AI is applied. This is where much of the confusion around AI-driven ROI begins. Firms adopt visible AI capabilities and expect differentiation, only to find that competitors adopt the same tools at roughly the same time. The baseline rises. Little else changes.

This distinction, between what becomes table stakes and what becomes advantage, is central. AI that is broadly available and easy to deploy tends to function as a rising tide. AI that is deeply embedded, tightly integrated, and difficult to replicate is what creates distance.

One useful way to see this is through the lens of visibility. Counterintuitively, the AI that is most visible is often the least differentiating. Customer-facing chatbots, AI-enabled CRM systems, and "AI-powered" ERP features are widely distributed. They may be necessary to remain competitive, but they rarely confer lasting advantage. Everyone can buy them. Your competitors are already piloting them.

By contrast, the AI systems that matter most are usually invisible. They live inside internal workflows, optimization engines, pricing systems, logistics planning, and decision-support infrastructure. They are not demoed at conferences. They are not described in marketing materials. They are rarely discussed publicly at all.

There is a reason for this. Where AI sits closest to money, power, or irreversible advantage, firms become secretive. Quantitative hedge funds guard their models and data pipelines obsessively because those systems directly print money. Defense and intelligence organizations treat AI-enabled capabilities as strategic assets because they create asymmetric advantage. Large technology platforms protect ranking, recommendation, and targeting systems because they are the core of their economic moat. In biotechnology, AI-driven discovery pipelines are closely held because they can determine years of lead time.[16]

The pattern is consistent. The closer AI is to durable advantage, the quieter firms become about it.

This is the opposite of AI theater. AI theater is loud, visible, and often superficial. It optimizes for signaling rather than substance. It can even be actively harmful. A poorly designed customer-facing chatbot may reduce support costs while simultaneously eroding brand trust, introducing unpredictable behavior, and creating new legal or reputational risk. The gains are obvious; the losses are diffuse and delayed.

What separates firms that build advantage from those that merely adopt tools is not enthusiasm or skepticism, but discernment. Specifically, the ability to distinguish between what is commoditizing and what is defensible.

This is where abstraction matters. Much of the current discourse treats AI as synonymous with large language models, and language models as synonymous with chatbots. This collapse of categories is not just incorrect, it is strategically dangerous. It encourages organizations to apply the most general, most expensive, and often least appropriate tools to problems that are better solved with simpler methods.

To a person equipped with a language-model hammer, every problem starts to look like a generative-AI nail. But many high-leverage business problems are not language problems at all. They are problems of prediction, optimization, control, or sequencing. In these domains, classical analytics, machine learning, and rule-based systems often outperform generative models on cost, reliability, and interpretability.

Consider pricing optimization. McDonald's does not need a conversational model to determine optimal pricing across thousands of locations. The value comes from demand forecasting, elasticity modeling, and feedback loops embedded deeply in operations. These systems are invisible to customers, unglamorous to demonstrate, and extremely valuable.

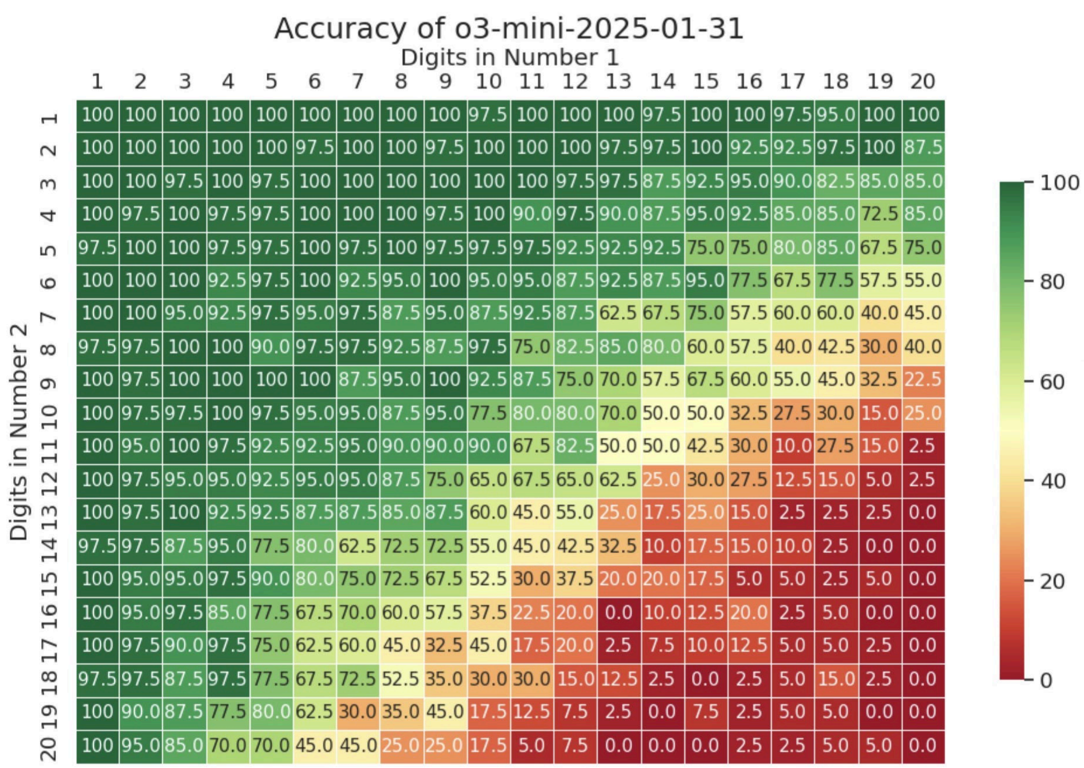

Even on a basic arithmetic task (multiplying 2 numbers of increasing digit size, shown here), accuracy declines as complexity increases. Language models are powerful approximators, not general-purpose engines - a reminder that tool choice matters as much as capability.[17]

This points to a broader principle: intelligence is a stack, not a model. Language models sit at the top of that stack. Much of the value is created lower down, where systems interact directly with data, processes, and decisions. When organizations reason about AI at the wrong level of abstraction, they constrain their own thinking and invite misuse.

Cost overruns are often framed as the problem. In practice, cost is the constraint that reveals a deeper issue. Language models become expensive when they are used for tasks that should be deterministic. Human-in-the-loop workflows become costly when automation should have occurred upstream. These are symptoms of tool mismatch, not technological failure.

None of this implies that firms should avoid experimentation. On the contrary, building advantage often requires a willingness to look foolish early, to run small experiments, and to tolerate local failure. But experimentation that compounds requires clean data, clear feedback loops, and an understanding of where value is likely to accrue. Throwing tools at problems without this clarity is unlikely to produce anything more than noise.

Later briefs will examine how organizations identify which parts of the intelligence stack are commoditizing and which are worth investing in, how they decide what to build versus buy, and why copying visible use cases so often backfires. For now, the core takeaway is straightforward: AI lifts baseline capability across firms. Durable advantage depends on where it is embedded and how difficult it is to replicate.

Separating Signal from Noise

What This Changes for Leaders

AI is often framed as a technology transition. In practice, it reshapes how organizations are managed.

It tests whether organizations can redesign work rather than merely augment it; reason probabilistically rather than narratively; and distinguish between efficiency gains that disappear and advantages that endure. These are not new challenges, but AI compresses the cost of getting them wrong.

What makes this moment difficult is not a lack of tools, but a structural lag: capability is advancing far faster than most organizations can change how they operate. That lag produces noise: impressive demos, uneven results, and a growing sense that something important is being missed.

The firms that benefit most from AI will not be the ones that adopt the loudest or fastest. They will be the ones that make fewer but better choices: about where systems should carry work, where humans should retain judgment, and where value actually accrues once technology is in place. Those choices are rarely visible from the outside. Firms with real advantage tend to keep it that way.

As you reflect on your own organization, a few questions are worth sitting with:

- Where are we using AI to speed up existing work, versus redesigning the work itself?
- Which decisions are evaluated as right or wrong, rather than against their expected distribution of outcomes and error bounds?
- Where do efficiency gains actually flow in our business - to customers, employees, or into durable, defensible advantage?
- Which AI initiatives would we be reluctant to describe publicly, and why?
- If an automated system made a statistically correct but uncomfortable decision tomorrow, who would own it?

These are not technical questions. They are leadership questions.

Issue 1 has been a map. In future briefs, we'll take these areas one at a time, drawing on real deployments, real constraints, and real failures, to explore what it takes to turn AI from ambient noise into durable signal.

Our experience suggests a counterintuitive conclusion:

The firms that get this right will not look like early adopters. They will look boring, disciplined, and quietly ahead.

About the Author

Sami Baqai
Founder, Autokatalyst LLC

Autokatalyst is an automation, advanced analytics, and AI firm that works with executives and operators to design and deploy machine intelligence within real organizational constraints.

This brief reflects patterns observed across real deployments, industries, and operating environments. We welcome thoughtful disagreement and real-world counterexamples.

Contact
sami@autokatalyst.com
LinkedIn: linkedin.com/in/sbaqai

Future Executive Intelligence Briefs will be published periodically.
If you'd like to be notified when future issues are released:
www.autokatalyst.com/research

Notes & Sources

1. https://insight.factset.com/highest-number-of-sp-500-earnings-calls-citing-ai-over-the-past-10-years-1
2. https://www.venturecapitaljournal.com/funding-for-ai-dominated-in-vc-in-2025-crunchbase/
3. https://www.sparklinecapital.com/post/surviving-the-ai-capex-boom
4. https://www.merriam-webster.com/wordplay/word-of-the-year
5. https://apnews.com/article/australia-ai-errors-deloitte-ab54858680ffc4ae6555b31c8fb987f3
6. https://www.humanetech.com/case-study/litigation-case-study-openai
7. https://www.cnn.com/2026/01/07/business/character-ai-google-settle-teen-suicide-lawsuit
8. https://cloud.google.com/blog/topics/threat-intelligence/threat-actor-usage-of-ai-tools
9. https://www.assemblymag.com/articles/91581-the-moving-assembly-line-turns-100
10. https://commons.wikimedia.org/wiki/File:NY_stock_exchange_traders_floor_LC-U9-10548-6.jpg
11. https://www.quantinsti.com/articles/automated-trading-system/
12. https://en.wikipedia.org/wiki/Galton_board
13. https://www.nber.org/system/files/working_papers/w32831/w32831.pdf
14. Miller, B. (2011). Moneyball. Columbia Pictures.
15. https://www.knitting-mills.com/history/
16. https://pmc.ncbi.nlm.nih.gov/articles/PMC12308071/
17. Yuntian Deng, https://x.com/yuntiandeng/status/1836114401213989366