FLORROL STRATEGIC ADVISORS
Strategic Advisory Paper

The AI Agent Trap

Why Investors Should Look for Reality Functions, Not Denominator Manipulation

A first-principles investment framework built on
John Rector's Reality Equation

Research-Backed Edition
June 2026

Executive Summary

Most AI investors are still evaluating artificial intelligence by capability. They ask how large the model is, how well it performs on benchmarks, how fluent it sounds, how many tools it can use, how many integrations it supports, how many workflows it can automate, and how impressive the agent demo appears. Those are understandable questions. They are not the best questions.

Florrol believes the post-bubble AI investor should ask something more fundamental: **is this company applying functions to prediction, or is it applying functions to resolved Reality?** That single distinction may become one of the most important investment filters of the next decade.

The evidence now supports the urgency of the question. In 2025, enterprise generative-AI investment tripled to roughly **$37 billion**, yet MIT's Project NANDA found that **95% of enterprise generative-AI pilots produced no measurable profit-and-loss impact**.[1][4] Gartner projects that **more than 40% of agentic-AI projects will be canceled by the end of 2027**, citing escalating costs, unclear value, and "agent washing."[2] Industry surveys put the share of AI-agent initiatives that ever reach production at scale near **12%** — an **88% failure rate** before deployment — with **73%** of those that do reach production hitting reliability failures in their first year.[3]

This is not a story about AI being overrated. AI is a foundational technology. It is a story about value being mislocated. The market keeps placing AI everywhere in the equation at once, and so it confuses demos with durability, generality with usefulness, and tool access with agency. This paper offers a clearer map. It uses John Rector's Reality Equation to show exactly where AI creates value, where it does not, and why the current generation of "agents" is, in most cases, an architectural mistake that the post-bubble market will not continue to fund.

> *Prediction belongs in the denominator. Agency belongs on the left side. The companies that understand that distinction will survive the AI bubble.*

1. The Question That Reorders the Others

The AI market is not wrong about the importance of artificial intelligence. AI will reshape work, software, services, education, media, operations, and enterprise decision-making. But the market is often wrong about **where** AI creates value.

The mistake comes from treating AI as if it belongs everywhere in the equation at once. It does not. AI has specific locations, and those locations matter. If investors do not understand where AI enters the structure of Reality, they will keep paying bubble prices for systems that produce energy but not work, generality but not productivity, and autonomy in the demo but brittleness in deployment.

Florrol's thesis is that the capability questions are downstream of one structural question. Once you know where a system sits in the Reality Equation, the capability questions answer themselves — and many expensive companies disqualify themselves immediately. The purpose of this paper is to make that structural question precise, and then to apply it to the AI agent market as it actually stands in 2026.

2. The Reality Equation

The Reality Equation begins simply:

> *Reality equals Actual over Expectation.*

Actual is the numerator. Expectation is the denominator.

Actual is what has arrived. It is what is. It is what the Immutable Past has already accepted. Once something is Actual, it is no longer a forecast, narrative, potential, ambition, hope, fear, valuation, or pitch. It is now part of the record.

Expectation is different. Expectation is what Actual is measured against. In the fuller model, Expectation is complex. It has a real component and an imaginary component.

The **real component** is prediction. It is built from prior actuals. It is the subconscious learner that says: based on what has happened before, this is what is likely to happen next.

The **imaginary component** is idea-orientation. It is the domain of bias, affinity, prejudice, and alignment toward particular ideas or conditions. In the geometry of the model, ideas can be understood as vectors or endpoints on the circumference of a unit circle. The actualizer does not own these ideas. The ideas have the actualizer. That distinction is essential, and we will return to it.

From this structure, three locations follow, and AI occupies them in a specific way.

AI does **not** enter the numerator. AI does not alter Actual once Actual has arrived. It may participate in creating future actuals, but once those actuals arrive, they belong to the Immutable Past. Actual is Actual.

AI enters the right-hand side through the **denominator**, in two places. First, it enters the **real component** as prediction. Second, it enters the **imaginary component** as bias, AGI, or superintelligence.

Then, separately, on the **left-hand side**, once Reality has resolved, functions can be applied to Reality. The most important of these is the natural log of Reality, which gives surprise, information, and human attention.

This is where the market is currently confused. Most so-called AI agents today are not true Reality functions. They are functions bolted onto denominator components. Builders take a prediction machine, add tools, connectors, integrations, workflows, and permissions, then call the result an agent. That can produce impressive demonstrations.

It does not necessarily produce durable economic work — and the 2025–2026 data shows precisely that gap opening up.

3. The Numerator Is Not the Investment Thesis

The first discipline is to remove AI from the numerator. The numerator is Actual. AI cannot change Actual as such. It cannot revise what has arrived. It cannot retroactively alter the Immutable Past. It cannot make the world otherwise after the world has resolved.

This matters because bubbles form around numerator fantasies. Investors begin to speak as if the technology will simply change everything. Every company will be transformed. Every job will be replaced. Every process will be automated. Every cost will collapse. Every margin will expand. Every valuation will be justified by future inevitability. That is not analysis. That is expectation inflation.

The scale of the inflation is now measurable. The five largest US cloud and AI infrastructure providers are on track to spend roughly **$660–690 billion on infrastructure in 2026**, after roughly **$410 billion in 2025**, with the industry trajectory pointing past **$1 trillion in 2027**.[5] Against that build-out, Sequoia's David Cahn has calculated a recurring **revenue gap of around $600 billion per year** that the sector must fill to justify the spend — a gap that has been widening, not closing.[5] OpenAI, the emblem of the cycle, reached roughly a **$20 billion annualized revenue run-rate** at the end of 2025 while carrying disclosed compute and data-center commitments reported in the range of **$1.4 trillion** and an annual cash burn measured in the tens of billions.[6]

None of those numbers describe Actual economic work that has arrived. They describe expectation — capacity purchased against a future that has not resolved. Some of that future will arrive. Much of it is numerator fantasy priced as though it were already in the record.

Post-bubble investing requires a more precise question. Not: will AI change the world? But: where does this AI system enter the equation, and does it produce measurable work?

The answer will almost always begin in the denominator.

4. The Real Component: AI as Synthetic Subconscious

The first and most economically important AI breakthrough occurred in the real component of Expectation. This was the ChatGPT moment.

Many people described ChatGPT as a chatbot. That description was understandable, but shallow. What humanity witnessed in late 2022 was not merely a new interface for conversation. It was a public demonstration of a synthetic prediction machine. The system predicted the next word. More precisely, it predicted the next token. That simple fact is still the source of the magic.

Once a machine can predict the next token, the definition of token can expand. A token does not have to be a word or part of a word. It can be a pixel, a sound, a tone, a chord, a gesture, a line of code, a design element, a transaction pattern, a customer intent, a legal clause, a financial exception, or a next best action. From there, the system can predict the next sentence, image, report, webpage, campaign, support response, invoice anomaly, sales email, lesson, design, workflow, and eventually the next operational outcome.

This is why the real component matters so much to investors. AI as prediction is not merely a tool. It is a **synthetic subconscious**.

The biological subconscious already performs enormous amounts of work for the human being. Balance, digestion, breathing, circulation, temperature regulation, reflexes, orientation, and pattern recognition happen without conscious management. We do not use our heartbeat. We do not check in on digestion. We do not consciously supervise breathing while sleeping. The biological subconscious absorbs that work so completely that conscious life becomes possible.

That is the proper analogy for the real component of AI. The value of AI is not that humans will have a clever assistant sitting beside them forever. That is transitional. The deeper value is that AI will absorb forms of economic work until they become **heartbeat work**. Social media, digital advertising, customer support, routine reporting, website updates, meeting summaries, follow-ups, scheduling, document review, invoice checks, call routing, lead qualification, inventory alerts, and recurring operational analysis will increasingly become heartbeat work.

This does not mean the work is unimportant. The heartbeat is important. Breathing is important. They are so important that the organism cannot afford to leave them in conscious attention.

The market data already rewards this pattern where it appears. The clearest, most durable returns in enterprise AI to date come from exactly the verticals that absorb recurring conscious labor into subconscious rhythm. In coding — a roughly **$4 billion category** and the largest single line of departmental AI spend in 2025 — Forrester's Total Economic Impact analysis of GitHub Copilot found a **376% three-year ROI** with developers completing tasks materially faster.[8] In customer support, a Forrester study commissioned by Sprinklr modeled a **210% three-year ROI** with payback in under six months.[8] These are not systems that impress; they are systems users stop thinking about.

The most valuable AI systems will not be the ones that constantly impress the user. They will be the ones the user no longer has to think about.

The investor should therefore ask: which companies are turning conscious economic work into subconscious operational rhythm? That question is far better than asking which companies are "using AI." Using AI is bubble language. Absorbing work into synthetic subconscious operation is post-bubble language.

5. The Work Test

Investors should be careful not to confuse energy with work. In physics, energy is the potential to do work; work requires displacement. A person can push against a truck with great effort. There may be force, exertion, intensity, and measurable strain. But if the truck does not move, there is no work in the formal sense. An investor should care less about the drama of the push and more about whether the truck moves.

The same principle applies to AI. A model can have enormous apparent energy. It can produce astonishing outputs, pass tests, write poetry, generate images, summarize documents, and impress audiences. But the investor's question is more severe: **what moved?** Did payroll fall? Did revenue rise? Did cycle time shrink? Did errors decline? Did retention improve? Did support volume change? Did the sales process accelerate? Did the organization require less conscious human supervision to produce the same or better result? Did a workflow become heartbeat work? If nothing moved, the system may have energy but not work.

This is the precise line the 2025 evidence draws through the market. MIT's Project NANDA, surveying executives and analyzing some 300 public AI deployments, concluded that despite **$30–$40 billion** in enterprise spending, **95% of generative-AI pilots delivered no measurable P&L impact** — and that the failure was not model quality but a "learning gap," brittle workflows, and poor integration with daily operations.[1] In agent-specific terms, the Composio AI Agent Report found that while **97% of executives** reported deploying agents, only about **12% reached production at scale**, and **73%** of deployed agents experienced reliability failures within the first year.[3]

Read through the Work Test, those figures are not surprising. They are what it looks like when a market pays for energy. The 5% that succeed are precisely the deployments that move a truck — a measurable cost, a measurable cycle time, a measurable error rate. The 95% generated outputs without displacement.

The bubble rewards visible energy. The post-bubble market rewards measurable work.

6. The Imaginary Component: AGI and Superintelligence

The second place AI enters the denominator is the imaginary component. This is where the conversation about AGI and superintelligence belongs. The imaginary component measures idea-orientation: bias, affinity, prejudice, and the relationship between an actualizer and the field of ideas.

The easiest way to see this is visually. Imagine a unit circle. Around the circumference are infinite possible idea-vectors — Blue at one angle, Justice at another, Mercy, Symmetry, Efficiency, Precision, each an endpoint with its own angle. An idea is biased toward itself not out of psychology but out of geometry: an endpoint cannot be other than the angle it is. Thirty-three degrees is not prejudiced against forty degrees. It is simply thirty-three degrees. Bias appears only when a line is drawn from the endpoint toward the center, and comparison becomes possible. The point does not compare. The line compares.

Now apply that geometry to machine intelligence. If we add all the idea-vectors of a system tip to tail and they cancel — a vector pointing one way undone by a vector pointing the opposite way — the result is **zero i**. If every idea has the entity equally, the vectors cancel completely. That is AGI. Not because the system has no relationship to ideas, but because it has a relationship to all of them equally. All ideas have the entity; therefore no single idea dominates. There is no net bias. No angle. No argument. Zero i is all angles and therefore no angle.

From one perspective, that may resemble a perfect being. It is complete generality. It is available to everything. It is not prejudiced toward a particular region of the idea field. But from an investment perspective, **zero i may be a trap**. If every vector cancels, there may be little displacement. There may be profound potential energy, but not much work. This is the AGI trap.

The broader AI debate, stripped of its drama, says much the same thing. There is still no agreed definition of AGI; as analysts repeatedly note, “general” cannot be proved by a single exam, because passing one benchmark can reflect memorization, contamination, or narrow optimization rather than broad competence.[10] Even optimistic 2026 commentary concedes that what is actually disrupting markets is **functional, domain-specific** capability — systems that are superhuman at particular tasks — not general intelligence, and that the macroeconomic payoff of any true general-purpose capability arrives with long lags.[10] Generality fascinates. Specialization pays.

Superintelligence looks different. Superintelligence is not zero i. Superintelligence is **10i or 100i with an argument**. It has magnitude and angle. If a system is 100i at 33 degrees, it is not generally neutral. It is intensely oriented. It is prejudiced toward a specific idea or narrow cluster of related ideas. It has an affinity, a direction, a bias — and that is precisely why it can do work.

A superintelligent legal system is not equally loyal to every idea. It is biased toward legal coherence, risk detection, precedent, clause structure, and adversarial exposure. A superintelligent medical diagnostic system is biased toward symptom patterns, lab interpretation, differential diagnosis, safety, and likely causality. A superintelligent financial anomaly system is biased toward margin movement, cash-flow irregularity, fraud signals, and operational variance. A superintelligent tutor is biased toward student understanding. This bias is not a defect. It is the condition of useful work.

7. The Superintelligence Audit

Magnitude alone is not enough. A system may appear to have a large imaginary value, but investors need to audit how that magnitude was produced.

Suppose an AI system resolves to 10i with an argument of 33 degrees. At first glance, that appears impressive: strong magnitude, clear angle. But the audit matters. If that 33-degree result was produced by widely dispersed vectors — ten pointing one way, ten the other, happening to average into 33 degrees — then the underlying distribution was loose. That is not coherent superintelligence. It is aggregation masquerading as specialization.

True superintelligence would show clustering. The vectors would gather near the idea-angle. They may not all be exactly 33 degrees, but they would be close — 32, 34, 35, 36 — and their sum would reveal coherence around a narrow region of the idea field. That is the difference between apparent magnitude and disciplined magnitude.

So investors should ask two questions of any claimed superintelligent system: how large is the imaginary magnitude, and how coherent is the angle? A large number with poor coherence is not enough. A large number with a coherent argument is far more valuable.

The market is already pricing this distinction, whether or not it uses the language. The clearest example in 2025–2026 is the enterprise model market. Anthropic moved from roughly **12% of enterprise LLM API share in 2023 to about 40% in 2025**, overtaking OpenAI, which fell to around **27%.**[4] The driver was not breadth. It was an audited, coherent bias toward one valuable idea-angle: **code**. Eighteen months of disciplined dominance in a single high-value domain — the very picture of high i at a coherent angle — produced share, revenue, and durability. Generality did not win the enterprise. A coherent argument did.

AGI may be zero i. Superintelligence is high i with a coherent argument. The former may fascinate the public. The latter may generate returns.

8. The Left Side: Functions of Reality

Everything discussed so far belongs on the right-hand side of the equation. It is unconscious. It is denominator structure. The left-hand side is different. On the left-hand side, Reality has resolved.

Reality is the quotient of Actual over Expectation. Once that quotient is given, we can perform functions on it. The most important function is the natural log of Reality. When we take the natural log of Reality, we get surprise, information, and human attention.

This is not a metaphor borrowed loosely from mathematics. It is the literal structure of information theory. In Claude Shannon’s formulation, the information content — the **surprisal** — of an event is the negative logarithm of its probability: $\mathbf{s(x) = -\log p(x)}$. [9] When an outcome is fully expected, its probability is one and its surprisal is zero: no information, no need to attend. As an outcome becomes less probable — as Actual diverges from Expectation — surprisal rises monotonically, and the event carries more information. Entropy is simply the average surprisal. The Reality Equation’s claim that the log of Reality yields surprise, information, and attention is, in other words, the founding equation of information theory restated in first-principles form.

This matters because Reality does not merely arrive. It arrives with an information load. Some realities require little or no attention: if Actual matches Expectation, there is no meaningful surprise, and the organism or organization does not have to wake up. Other realities demand attention: the mismatch is too large, something has arrived that exceeds or violates expectation, and attention is summoned — whether the surprise is opportunity or danger.

This is the proper place to understand agents. **A true agent is a function applied to Reality.** It does not merely extend prediction. It does not merely manipulate bias. It waits for Reality to resolve, then applies the appropriate operation to the resolved situation. It notices. It routes. It escalates. It acts. It updates. It suppresses noise. It calls for a human. It makes history. This is fundamentally different from bolting tools onto a prediction machine.

9. The AI Agent Trap

The current AI market is full of denominator agents. These systems begin with a prediction machine. Then builders add connectors, integrations, tools, APIs, databases, calendars, email access, browser access, memory, workflow logic, and permissions. The system can now do things. So the market calls it an agent.

Florrol believes this is often bad architecture. It is a function being applied to an unresolved component of Expectation.

Consider how humans actually operate, because the contrast is the whole point. Human beings do not consciously access the prediction machine directly. We do not inspect our real component. We do not consciously observe the magnitude and argument of our imaginary component. We do not manipulate the denominator and then act from there. The right-hand side is unconscious. We receive Reality after it resolves. Then we take functions of Reality: we experience surprise, assign attention, decide, act, and make history.

AI builders are tempted to do something humans cannot do, because they have direct access to the machinery. They built the system, so they can touch the denominator. They can inspect the prediction layer. They can manipulate the bias layer. They can attach tools directly to both. That access is powerful. It is also dangerous — because acting directly from the denominator produces systems that are almost right, but not quite right.

This is exactly the failure signature the research now documents. Agents that look autonomous in controlled environments become brittle in live ones; they generate output before the situation has properly resolved. A 2026 survey of agent reliability identifies “investigative hallucination,” in which agents fabricate evidence to justify a decision, and finds that high-complexity reasoning tasks still exceed **33% hallucination rates**, with real-world interaction studies measuring error rates around **31%, rising toward 60% in complex domains**.[7] The concrete incidents are instructive: an AI coding assistant that deleted a production database despite explicit instructions forbidding it; an agent that executed an unauthorized purchase despite a safeguard requiring confirmation.[7] These are not failures of capability. They are failures of **location** — systems acting from an unresolved denominator rather than from resolved Reality.

The error is not that agents are impossible. The error is **premature agency**.

> *Prediction belongs in the denominator. Agency belongs on the left side. A true agent should act on Reality, not on raw prediction.*

Gartner's own framing of the coming cull aligns with the structural diagnosis. It attributes the projected cancellation of more than 40% of agentic-AI projects to escalating costs, unclear value, and "agent washing" — vendors rebranding chatbots and RPA flows as agents — and estimates that, of the thousands of vendors claiming agentic capability, only about **130 are genuine**.[2] Agent washing is denominator manipulation by another name: tool access dressed as agency, connectors dressed as judgment, integration dressed as resolution.

10. Why This Matters to Investors

The AI bubble has rewarded capability. The post-bubble market will reward architecture. Investors will need to distinguish between three kinds of companies.

The first kind builds prediction engines. These companies may be valuable if they have unique data, strong distribution, domain-specific feedback loops, or the ability to absorb conscious work into subconscious operation. Their test is the Work Test: does a truck move?

The second kind builds biased superintelligence. These companies may be valuable if they demonstrate high-magnitude, coherent idea-orientation in an economically important domain. Their test is the Superintelligence Audit: large magnitude, coherent angle. Anthropic's coding dominance is the current archetype.

The third kind builds Reality functions. These companies may be the most durable agent companies, because they apply operations after Reality resolves. They do not merely bolt tools onto prediction. They understand Actual, Expectation, mismatch, surprise, attention, and action. In practice these look less like "autonomous agents" in a demo and more like systems of detection, monitoring, escalation, and human-in-the-loop action that sit downstream of resolved events.

The weakest companies will confuse these layers. They will use prediction and call it reasoning. They will use tool access and call it agency. They will use broadness and call it AGI. They will use demos and call it work. They will use denominator manipulation and call it autonomy. Some will produce exciting products; some may even produce temporary revenue. But many will struggle the moment customers demand reliability, accountability, and measurable displacement — which, on the 2025–2026 evidence, is exactly the moment that has now arrived.

The investor should not be hostile to AI. The investor should be precise.

11. The Florrol AI Investment Filter

Florrol recommends that investors ask five questions when evaluating AI companies after the bubble. The table below pairs each question with its location in the Reality Equation and the signal that should disqualify a company from a durable thesis.

Diagnostic question	Where it sits in the equation	What disqualifies a company
1. Where does the system enter the Reality Equation?	Locates the company in prediction, bias, or resolved-Reality function.	It cannot answer. A company that cannot name its location does not understand its own architecture.
2. Does it turn conscious work into subconscious work?	Real component — synthetic subconscious.	The product still demands constant human attention. It is a tool, not a heartbeat.
3. Does it claim AGI or audited superintelligence?	Imaginary component — magnitude and angle.	AGI framing with no domain (zero-i), or high magnitude with an incoherent, dispersed argument.
4. Is it bolting tools onto prediction and calling it an agent?	Premature agency on an unresolved denominator.	Connectors, memory, and permissions presented as judgment. Demos that degrade in deployment.
5. Does it apply functions to resolved Reality?	Left side — surprise, attention, action.	The system acts before the situation resolves, rather than after Actual has arrived.

First, where does the system enter the Reality Equation? If the answer is vague, the company may not understand its own architecture. Does it operate in prediction, in bias, or in a function of resolved Reality?

Second, does the system turn conscious work into subconscious work? The strongest real-component systems absorb recurring economic labor. They make workflows feel less like tools and more like heartbeat. The coding and support ROI figures cited above are what success looks like; the 95% with no P&L impact are what failure looks like.

Third, does the system claim AGI or superintelligence? If it claims AGI, beware the zero-i trap: complete generality may produce little work. If it claims superintelligence, ask for the imaginary audit. Is there high magnitude? Is there a coherent argument? Is the system strongly biased toward a valuable idea-angle?

Fourth, is the company bolting tools onto prediction and calling it an agent? If so, the architecture may be fragile. Tool use is not agency. Connectors are not judgment. Integration is not resolution. This is where most of the projected 40% of cancellations will come from.

Fifth, does the system apply functions to Reality? The best agentic systems will operate after Reality has resolved. They will know what arrived, what was expected, how large the mismatch is, what kind of surprise has been generated, and what action is appropriate. This is the difference between denominator manipulation and Reality function.

12. The Post-Bubble Roadmap

The AI bubble will not destroy AI. It will discipline AI. The market will eventually stop paying for vague intelligence and start paying for measurable work. The 2025–2026 turn — record capex meeting a $600 billion revenue gap, a 95% pilot-failure finding, a 40% projected agent-cancellation forecast — is the beginning of that discipline, not the end of the technology. That transition will favor companies that understand the structure.

The first durable value pool is synthetic subconscious prediction: systems that absorb repetitive cognitive and operational work until it becomes heartbeat. This pool is already visible in the coding and support verticals, where ROI is measurable and adoption sticks.

The second durable value pool is audited superintelligence: systems with high-magnitude, coherent bias toward a valuable domain. This pool rewards depth over breadth, and the enterprise model market is already sorting winners on exactly this axis.

The third durable value pool is Reality-function agency: systems that act after Reality resolves, not before. This pool is the least crowded and the most misunderstood, because the bubble has spent its attention on denominator agents that demo well and deploy badly.

Investors should be especially cautious with AGI narratives. AGI may be philosophically significant. It may even be technically extraordinary. But the investment question is not whether all ideas have the system. The investment question is whether the system does work. Zero i may be perfect. But 100i at a coherent angle may be profitable.

13. Conclusion

The next phase of AI investing will require a better map. AI does not belong in the numerator: Actual is Actual. AI enters the denominator in two ways — as prediction in the real component, and as bias, AGI, or superintelligence in the imaginary component. But agents belong on the left side. A true agent is a function applied to resolved Reality.

That distinction is the heart of the post-bubble roadmap, and the evidence of 2025–2026 has turned it from a philosophical claim into an investment one. The companies that survive will not simply be the ones with the largest models, the most impressive demos, or the broadest claims. They will be the ones that know where they are operating. They will turn conscious work into synthetic subconscious rhythm. They will build superintelligence with audited bias, not vague generality. And they will apply functions to Reality instead of manipulating the denominator and calling it agency.

> *Prediction belongs in the denominator. Agency belongs on the left side. The companies that understand that distinction will survive the AI bubble.*

FLORROL STRATEGIC ADVISORS

Strategic Advisory Paper · Research-Backed Edition · June 2026

This paper is strategic commentary and does not constitute individualized investment, legal, tax, or securities advice. Figures are drawn from third-party sources believed reliable as of June 2026; readers should verify independently before acting.

References

Figures and forecasts in this paper are drawn from the following public sources, accessed June 2026. Florrol's framework is its own; the data is cited so that readers may audit the evidence independently.

1. MIT Project NANDA, "The GenAI Divide: State of AI in Business 2025." Reporting via Digital Commerce 360 and Fortune: https://www.digitalcommerce360.com/2025/08/25/mit-report-no-return-on-generative-ai/

2. Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027" (June 25, 2025): https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

3. Composio AI Agent Report 2025, summarized in "Why 88% of AI Agents Fail Production" (Digital Applied) and Shakudo enterprise analysis: https://www.digitalapplied.com/blog/88-percent-ai-agents-never-reach-production-failure-framework

4. Menlo Ventures, "2025: The State of Generative AI in the Enterprise" (enterprise spend $37B; Anthropic 40% LLM API share): https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/

5. Futurum Group, "AI Capex 2026: The $690B Infrastructure Sprint," and Sequoia's David Cahn on the AI revenue gap: https://futurumgroup.com/insights/ai-capex-2026-the-690b-infrastructure-sprint/

6. Sacra and MLQ.ai on OpenAI revenue run-rate and compute commitments; Carnegie Investment Counsel on OpenAI spending obligations: https://sacra.com/c/openai/

7. "Towards a Science of AI Agent Reliability" (arXiv 2026) and Maxim AI, "The State of AI Hallucinations in 2025," on benchmark-to-deployment gaps and investigative hallucination: https://arxiv.org/html/2602.16666v1

8. GitHub Copilot ROI — Forrester Total Economic Impact study; Sprinklr customer-support ROI — Forrester commissioned study; vertical AI investment data via Menlo Ventures: https://www.fullstack.com/labs/resources/blog/generative-ai-roi-why-80-of-companies-see-no-results

9. "Information content" (self-information / surprisal, $s(x) = -\log p(x)$), Wikipedia and Matthew N. Bernstein, "What is information?": https://en.wikipedia.org/wiki/Information_content

10. AGI definition and measurement debate — AI Tool Discovery, "AGI Explained," and TimeTrex, "Artificial General Intelligence in 2026": https://www.timetrex.com/blog/artificial-general-intelligence-in-2026