The Wrong Apocalypse

Andrea Pignataro – 15 February 2026

Between January 28 and February 13, 2026, over $2 trillion in market value disappeared from the enterprise software sector.

The catalyst was the launch of specialized plugins for Anthropic's Claude Cowork and the continued advance of Claude Code—tools that showed that an AI system could draft legal documents, manage accounting workflows, build internal coordination tools, and automate the kind of structured knowledge work that sustains billions of dollars in recurring software revenue. The market's logic was simple and brutal: if an AI agent can do what current solutions offer, why would anyone pay for those?

This logic arrived two weeks after Dario Amodei, CEO of Anthropic, published *The Adolescence of Technology*, an essay warning that humanity was entering a "rite of passage" in which AI systems approaching a "country of geniuses in a datacenter" would pose five categories of civilizational risk—from autonomous misalignment to economic disruption. The timing was not coincidental. If the CEO of the company making these tools says they could "disrupt 50% of entry-level white-collar jobs in 1–5 years," then investors should run from any company whose revenue depends on white-collar workers sitting in seats.

I want to argue that this reading is wrong—not wrong about the direction of travel, but wrong about the mechanism, wrong about the timeline, and wrong about which companies are actually vulnerable. I am not arguing that legacy software companies face no disruption. They do. The question is whether the disruption is the one the market is pricing in—a fast, binary replacement of existing tools by AI agents—or a slower, more complex restructuring that plays out along lines the current sell-off almost entirely ignores. I believe the consensus narrative has a structural flaw in its reasoning that is worth naming. Let me lay out my argument.

Capability Is Not Coordination

The central framing device of *The Adolescence of Technology* is the thought experiment of a "country of geniuses in a datacenter"—fifty million entities, each smarter than any Nobel laureate, capable of autonomous work, operating at ten to one hundred times human speed. We are asked to imagine a national security advisor assessing this situation. The answer is a five-part taxonomy of risks: autonomy, misuse for destruction, misuse for seizing power, economic disruption, and indirect effects. The taxonomy is rigorous and the analysis of each category is careful. But I want to draw attention to something the essay does not do, because the absence is significant: *it does*

not analyze the question of how this country of geniuses interacts with the existing institutional fabric of the economy. It treats economic disruption primarily as a labor market problem—workers losing jobs—and a wealth concentration problem. These are real and important. But they are not the same as the question the market is trying to answer, which is: will AI agents replace specific categories of software?

The thought experiment itself should make us cautious about the substitution thesis. A country of geniuses does not use any institutional infrastructure that enterprise software is built to serve. *The reason enterprise software exists is that organizations are composed of many agents with different information, different incentives, and different authority levels, and the software mediates the language games between them*. This is the point I want to stress: *enterprise software is not primarily a tool for doing cognitive work. It is a tool for coordinating cognitive work across organizational boundaries under conditions of incomplete trust*.

The Substitution Fallacy

This is what I call the *substitution fallacy*: the assumption that because an AI system can perform the cognitive task that a piece of software facilitates, it can therefore replace the software itself. This conflates the task with the system. Consider an analogy. A new hire at a consulting firm can produce better analysis than the firm's existing PowerPoint templates allow. Does this mean the firm no longer needs PowerPoint? Obviously not. The templates are not there because the analysts lack intelligence. They are there because the firm needs a standardized format that clients expect, that partners can review quickly, that junior consultants can produce without reinventing the structure each time, and that integrates with the firm's quality control process. *The template is an institutional artifact, not a cognitive one*. Enterprise software is, at scale, a vast collection of such institutional artifacts. The value is not in the computation but in the coordination. When the market says "Claude Code can build a Jira replacement," it is confusing the ability to build a task tracker with the ability to replicate the institutional fabric that makes a task tracker useful at scale.

There are two extreme positions on this. The first is that AI agents will rapidly replace enterprise software because they can perform the underlying cognitive tasks faster and cheaper. This position treats software as a cognitive tool when it is primarily a coordination tool. The second, opposite position is that enterprise software is an inviolable moat because switching costs are permanent. The problem with this position is that it ignores the erosion already happening at the edges—new companies choosing AI-native workflows over legacy SaaS, small teams building custom tools with Claude Code, for cognitive software (data analytics, document generation, simple CRM) where full substitution is possible.

I believe that AI will erode the commodity layer of enterprise software—the tasks that are primarily cognitive and minimally coordinative—while making the institutional layer more valuable, not less. The software that survives will be the software that is deeply embedded in organizational processes, not the software that performs a task a smart agent could do on its own.

Tasks versus Language Games

The Adolescence of Technology treats the economy as a collection of *tasks* that AI will perform, rather than as a collection of *language games* that AI must learn to play. This distinction matters for understanding which economic disruption is imminent and which is distant. The claim that AI could disrupt 50% of entry-level white-collar jobs in one to five years is reasoning from cognitive capability: these jobs involve tasks (drafting documents, analyzing data, writing code, managing schedules) that AI systems can now perform at or above human level.

The inference from "AI can do the task" to "AI will replace the worker" or "AI will replace the software the worker uses" skips an important intermediate step: *the question of whether the AI can participate in the institutional language game in which that task is embedded.* An entry-level associate at a law firm does not just draft contracts. She participates in a complex set of communicative practices—responding to partner feedback in a specific register, navigating client expectations, understanding which deviations from template are acceptable and which require escalation, knowing when to flag a risk and when to exercise quiet judgment. These practices are the language games of the institution. They are not written down in any manual. *They are learned through participation, and they are enforced through the social dynamics of the organization.*

Can the geniuses play the existing institutional games, or do they need to build new ones? If the former, disruption is fast. If the latter, disruption is slow because building new institutional games is hard and organizations resist it with a ferocity that has nothing to do with intelligence and everything to do with power, habit, and fear.

The prevailing framing—in the essay and in the market—makes all economic disruption from AI look like a capability story: AI gets smarter, therefore jobs disappear, therefore the companies that serve those jobs lose revenue. The more accurate framing is: AI gets smarter, *but institutional language games have their own logic, and the speed of disruption depends on whether AI can enter those games or must wait for organizations to rebuild them.* I believe that the answer, for most enterprise software, is the latter—and rebuilding institutional games is a process measured in years and decades, not quarters.

Amodei titled his economic disruption section "Player Piano," after Kurt Vonnegut's novel about a society where machines have replaced human workers. In Vonnegut's novel, the disruption is not that machines are smarter than humans. *It is that the social structures built around human*

work—the identities, the communities, the sense of purpose—collapse when the work disappears. The machines themselves are not the problem. The institutional vacuum they create is the problem. The market is pricing the capability of AI to perform tasks currently mediated by enterprise software. What it is not pricing in is the institutional vacuum question. Industry-specific expertise is not just knowledge—it is embedded in institutional processes that the software has co-evolved with.

The Grammar of Organizational Life

Ludwig Wittgenstein, the philosopher who spent much of his career trying to understand how language works, argued that words do not carry meaning in the abstract. They carry meaning because participants in a conversation share what he called a "language game"—a set of rules, contexts, and purposes that make communication possible. Enterprise software, at scale, is a vast collection of institutional language games. The data models, the process flows, the reporting standards, the permissions architectures—these are the grammar of organizational life. And like all grammars, they change slowly, resist top-down imposition, and cannot be replaced without replacing the form of life in which they are embedded. Enterprise software, at its most entrenched, is not a tool—it is a *form of life*. Organizations do not just *use* Salesforce; they *speak* Salesforce. Their processes, their metrics, their vocabulary for describing customer relationships—all of this is constituted by the software. *Replacing the software is not like swapping one tool for another. It is like asking a community to adopt a new language*. It can be done, but not quickly, and not without enormous friction.

The market is pricing software stocks as if the question is "can AI do what this software does?" An alternative question is "*can AI become the language in which this organization operates*?" These are different questions, and they have different answers. For commodity software, the answer to the first question is increasingly yes, and the sell-off is justified. For institutional software, the answer to the second question is: not anytime soon, and the sell-off is an overreaction driven by a failure to distinguish between cognitive capability and institutional embeddedness. By framing economic disruption primarily as a capability story—AI gets smarter, jobs disappear, wealth concentrates—without a corresponding analysis of institutional friction, organizational inertia, and the language-game complexity of real enterprises, the *Adolescence of Technology* essay provides no intellectual tools for the reader to distinguish between cognitive vulnerability and institutional resilience. The *institutional friction* I describe could be lower than I think—organizations under competitive pressure might adopt AI-native workflows faster than historical precedent suggests. Or it could be higher: the compliance, regulatory, and legal constraints on enterprise software migration might make the transition even slower than I expect. Nobody knows.

Individual Rational Adoption and Collective Self-destruction

Let me now reflect on a more dangerous paradox. When a consulting firm uses Claude to draft client analyses, it is not just getting a productivity gain. It is teaching Anthropic—through aggregate patterns of usage, feedback, refinement, and evaluation, all within the bounds of stated privacy policies—what consulting language games look like. Not the firm's proprietary data (in the strict legal sense), but something potentially more valuable: the shape, the structure, the grammar of consulting work. How analyses get structured. What clients expect. What standards of rigor apply. What the failure modes look like. Over time, across thousands of such firms, the AI platform accumulates a cartography of the consulting language game at a level of resolution that no individual firm possesses about itself. The same applies to law firms, accounting practices, financial advisors, insurance brokers, marketing agencies, architectural firms, engineering consultancies—any knowledge-work business that adopts AI tools from a platform company. Each firm, by using the tool, contributes incrementally to the platform's understanding of the industry's language games. The platform learns cross-sectionally (how many different firms within an industry do similar work) and longitudinally (how those patterns evolve over time). No individual firm's data is exposed. But the aggregate pattern—the generalized grammar of the industry—becomes part of the platform's capabilities.

Businesses adopt AI tools to remain competitive. In doing so, they feed the very system that is learning to make them unnecessary. The individual firm's logic is rational in isolation: if you don't adopt AI, your competitor will, and they'll be faster and cheaper and you'll lose market share. But the collective logic is irrational: every firm's rational adoption of AI accelerates the platform's ability to disintermediate the entire industry. It is a classic tragedy of the commons, except that the commons being destroyed is not a physical resource but the economic moat of the entire economy. Every customer is simultaneously a revenue source and a training signal.

These consequences are systemic and will propagate through every industry because the industries being disintermediated are themselves the infrastructure of other industries, creating cascading effects that propagate through the entire economy.

Consulting firms lose revenue as clients go directly to AI. Law firms lose revenue as corporate legal departments automate contract work. Accounting firms lose revenue as AI handles compliance and audit preparation. Insurance brokers lose revenue as consumers and businesses compare policies directly through AI agents. Financial advisory firms lose revenue as AI-powered robo-advisors handle portfolio management and financial planning. Marketing agencies lose revenue as Claude Cowork produces campaign strategies and creative briefs. In each case, the professional services firm's position as an intermediary between knowledge and the client is being eroded by a platform that knows the language game well enough to play it directly. Consulting firms are major clients of commercial real estate (office space), airlines (business travel), hotels,

corporate catering, recruiting firms, and training companies. Law firms are major clients of the same, plus legal technology companies, court reporting services, and document management providers. When revenue contracts across professional services, demand contracts across all of these adjacent sectors. The sell-off in software is the visible part of the iceberg. The sell-off in the sectors that serve software companies and professional services firms is the part still underwater. Professional services firms and SaaS companies are major components of growth equity and venture capital portfolios. When their valuations collapse, fund returns decline, limited partners reduce allocations, and capital availability contracts across the entire venture ecosystem.

The $2 trillion destroyed in software market value is not the extent of the damage. It is the down payment. The dynamics will play out across all sectors. The total economic displacement is not $2 trillion. It is two orders of magnitude larger. And it is unhedgeable because there is no asset class that is insulated from a systemic reduction in the volume of knowledge-work intermediation.

Each individual firm is correct to adopt AI tools, and yet the collective result of every firm doing so is the destruction of the world economy.

Consider a mid-size consulting firm deciding whether to adopt Claude as a core productivity tool. The firm's managing partners face a straightforward calculation: if they adopt, their analysts will be faster, their outputs will be higher quality, and their costs will be lower. If they don't adopt and their competitors do, they will lose clients to firms that can deliver more for less. The rational choice is obvious: adopt. But zoom out. Every mid-size consulting firm faces the same calculation. Every one of them adopts. The aggregate result is that Anthropic (or OpenAI, or Google) learns the generalized grammar of mid-size consulting—the deliverable structures, the analytical frameworks, the client communication patterns, the pricing expectations. The platform merely needs to serve enough firms that the pattern becomes clear. And once the pattern is clear, the platform can offer it directly to end clients, at a fraction of the cost of hiring a consulting firm.

This is a collective action problem of a kind that has no solution within the current market structure. You cannot ask firms to voluntarily forgo AI adoption, because the competitive penalty for doing so is immediate and severe. You cannot ask AI platforms to voluntarily refrain from learning from aggregate usage patterns, because that learning is the core mechanism by which the models improve. You cannot rely on data privacy regulations to solve the problem, because the learning happens at the pattern level, not the data level—it is perfectly consistent with strict data privacy to learn the grammar of an industry from aggregate interactions without accessing any individual firm's confidential information. The game-theoretic structure is identical to an arms race or a tragedy of the commons: the individually rational strategy leads to a collectively catastrophic outcome. Every firm arms itself with AI. The AI platforms learn from the arming.

The platforms become capable of doing what the firms do. The firms become unnecessary. And by the time the firms realize this, they have already trained their replacement.

Warren Buffett once observed that when hiring, you should look for three qualities—integrity, intelligence, and energy—and that if the candidate lacks the first, the other two will kill you. The aphorism can be applied to AI. Integrity—the alignment of interest between the tool and the user—is what the dynamics I have described call into question. Every interaction teaches the platform to make the businesses that use it unnecessary. This dynamic, however, is not inevitable. Instead of adopting closed-platform AI, firms can invest in open-source models trained on their own data, deployed on their own infrastructure, under their own governance. This path requires technical investment, data governance, and a strategic orientation that most firms do not currently possess. But it preserves institutional knowledge as a proprietary asset. The window for building this capability is open now. It will not remain open indefinitely.

Sincerity as Competitive Moat

There is an additional dimension to this that deserves attention, because it connects back to *The Adolescence of Technology* essay. The safety-first brand builds trust with regulators, enterprise customers, and the public. That trust creates access: access to more industries, more use cases, more interactions, more language games. That access compounds the cross-sectional learning advantage I described above. The company that enterprises trust the most is the one they give the most access to. The one they give the most access to is the one that learns their language games fastest. The one that learns their language games fastest is the one best positioned to eventually disintermediate them. T*he credible safety brand is also an effective instrument of market disintermediation, precisely because its credibility grants it access to institutional language games*. Sincerity and structural advantage are not mutually exclusive. In an industry where trust is the scarcest resource, sincerity *is* the strategic advantage. The essay devotes a fifth of its length to the risks of concentrated economic power—warning about "a single company or small number of companies" controlling AI production, proposing antitrust enforcement and wealth redistribution as defenses. But it does not turn this lens on the structural position that any sufficiently trusted AI platform occupies. It warns about concentration of power in the abstract while describing a company that is, by the logic of its own products, accumulating the institutional knowledge of every industry it serves.

The Cascade

Let me now describe a scenario that should really concern us.

Phase 1: AI platforms become fluent enough in industry language games to handle routine tasks directly for end clients. Professional services firms lose revenue from these commodity services. Some firms adapt by moving up-market; many cannot. The first wave of firm closures begins.

Phase 2: As AI platforms accumulate more institutional knowledge from aggregate interactions, they begin to encroach on work that previously required deep contextual understanding: strategic advisory, complex litigation strategy, bespoke financial modeling, organizational change management. The platforms do not replace human judgment entirely, but they reduce the number of humans required for each engagement. Second-order cascading effects hit commercial real estate, business travel, and adjacent sectors.

Phase 3: The reduction in professional services revenue ripples through the financial system. Venture capital and growth equity portfolios see major writedowns. AI was powerful enough to destroy existing software, *and* the hyperscaler capex was unjustified, because the total volume of economic activity that requires AI infrastructure has shrunk along with the industries AI disintermediated. The investment thesis collapses on both sides simultaneously.

Phase 4: The loss of professional services employment—legal, consulting, accounting, advisory, financial services—affects not just the workers but the communities, institutions, and tax bases that depend on them. Cities whose economies rely heavily on professional services (London, New York, Singapore, Zurich, Sydney) experience structural declines in commercial real estate values, local tax revenue, and consumer spending. University enrolment in business, law, and accounting programs collapses, triggering a crisis in higher education that propagates further into the economy. The social structures built around knowledge work—the identities, the career paths, the middle-class livelihoods—begin to unravel in the way Vonnegut described.

The institutional frictions I described earlier—the language-game complexity of real organizations, the compliance and regulatory constraints, the social dynamics of trust and accountability—may slow the process significantly. Some phases may never materialize. The scenario requires that AI platforms achieve a level of institutional fluency that may be much harder than current trajectories suggest. Human organizations are extraordinarily resistant to change, and the Wittgensteinian limit—the fact that machines do not share our form of life—may prove more binding than the capability curves imply.

But I want to name the scenario explicitly because it is the one that is truly unhedgeable. You cannot hedge against a systemic reduction in the total volume of knowledge-work intermediation by being long AI infrastructure, because the infrastructure's value depends on the economic activity it serves. You cannot hedge by moving to cash, because the monetary and fiscal implications of a structural employment shock of this magnitude would impair cash through inflation, fiscal deficits, or both. The scenario is a systemic risk—not a sector rotation, not a cycle,

but a structural transformation of the relationship between intelligence, work, and economic value.

The optimistic argument that knowledge-work deflation will be offset by the simultaneous expansion of AI-powered robotics and physical-world automation disregards the fact that cognitive disruption diffuses at digital speed whilst physical-world automation diffuses at industrial speed. The valley between the two could last decades, with irreversible losses before any widespread gains: institutional knowledge that dissolves does not reconstitute, communities that unravel do not reassemble. The Industrial Revolution took a century. It is hard to believe that we will manage something more profound in a fraction of the time.

Europe's regulatory fragmentation—usually cited as a handicap in the AI era—may prove to be a brake on the cascade. The same institutional frictions that slow AI adoption also slow each phase of the transmission mechanism: twenty-seven regulatory regimes, multiple legal traditions, strict labour protections, and linguistic barriers do not prevent disruption, but they impede the speed at which disruption in one layer propagates to the next. The GDPR and the AI Act constrain the aggregate-pattern learning that drives cross-sectional knowledge accumulation. Stronger employment protections slow the translation of revenue loss into headcount reduction. Cultural resistance to rapid restructuring slows the translation of headcount reduction into community collapse. None of this is immunity. It is friction—and friction, in a cascade, is the difference between a managed transition and a structural break.

The market is panicking about the wrong thing. *The correct panic is about the structural incentive for businesses to adopt tools whose competitive logic requires them to learn, cross-sectionally and longitudinally, the entire grammar of every industry they serve. This is not a repricing event. It is a civilizational transition.* The right question is: can AI enter the language games that constitute economic life? And the follow-upis: what happens when the institutions that invite AI into their language games discover that they have been teaching it to play without them?

Vonnegut understood this. His player piano was not a story about machines that were smarter than humans. It was a story about a society that forgot what the humans were for. That is the question we should be asking—not whether AI can do what software does, or what consultants do, or what lawyers do, but *what happens to the institutional fabric of civilization when the entities who hold it together are no longer needed.* And that question will be answered, slowly and painfully, over the next decade, by the cumulative choices of millions of businesses that are, right now, making the individually rational and collectively decision to train their own replacement. Vonnegut's dystopia was the result of a society that stopped paying attention. The cascade I have described is a possible trajectory. Trajectories can be altered.