



Investing in Light of AI

What Changes and What Endures?

AI agents have crossed a meaningful capability threshold for knowledge work.

Executive Summary

Improved reasoning models, ability to turn those models into autonomous workers ("agents"), and greater-than-10x annual declines in inference costs have made AI-driven knowledge work economically viable at scale. Software engineering is the leading proof point, as demonstrated by frontier labs that generate the vast majority of their code with AI. General-purpose agents are rapidly expanding into legal, finance, customer support, and other knowledge-intensive domains. For investors, particularly in private markets where illiquidity demands long-term conviction, this shift introduces additional new risk considerations to existing portfolio exposures as well as future investments. As it relates to a relatively large swath of the market, these risks include potential disruption of business models and redistribution of industry profit pools, erosion of switching costs and workflow lock-in as AI agents more easily select best-in-breed tools, and meaningful price compression for commodity software functionality.

The near-term landscape will be defined by a race to adopt AI comprehensively, not just in engineering and product organizations, but across most service and administrative functions. Middle-market companies may hold a structural advantage: they possess an existing customer base and revenue engine but retain enough organizational nimbleness to retool faster than large incumbents burdened by legacy process and change management. Products themselves will evolve from static human interfaces toward customizable, generative experiences, and ultimately, headless architectures that serve AI agents as first-class consumers. Companies that fail to create space in the P&L for this transformation, or that lack leadership attuned to AI-native thinking, risk falling behind in a period where the cost of feature delivery is collapsing and the pace of competitive iteration is accelerating.



At the same time, the shift clarifies which **moats endure and strengthen**. **Businesses that automate decisions on proprietary, multi-participant or "federated" datasets** (e.g., fraud detection networks or risk-scoring platforms) that create value from outcomes, not interfaces. The benefits compound with network effects that new GenAI-powered entrants cannot replicate with algorithms alone. **Multi-sided network platforms**, where the assembly of supply and demand creates reinforcing structural advantages, will deepen their defensibility as intelligence becomes cheaper, which in turn will be used to power decisions for network participants. **Regulatory barriers** further protect incumbents that have invested in certification, compliance infrastructure, and human oversight requirements. CEOs of middle market businesses should focus on competitive advantages that not only endure but strengthen as the cost of intelligence declines.

Recent Developments in AI

Over the past 15 months, a combination of improved models, agent harnesses, and steep cost declines has enabled AI to execute valuable knowledge work. Reasoning models can now coherently solve long-duration, multi-step problems. Advances in models to call on information outside of models (function calling or tool use) and multi-modal understanding (e.g., images and audio) allow these models to reliably interact with code, real-world documents, APIs, and visual interfaces. These capabilities are now controllable in "agent harnesses" that unleash models' core intelligence into file systems, databases, 3rd party products, and operational workflows. The prominent Claude Code platform was originally built for software engineering, but has quickly taken off as a **general purpose "knowledge worker"**. Cost has continued to decline >10x per year, with increasing speed and quality ¹, finally making these "knowledge workers" economically viable at scale.

The proof points are no longer theoretical. Andrej Karpathy, founding member of OpenAI, reported in December 2025 that his personal workflow flipped from 80% manual coding to 80% agent-directed in a matter of weeks after Claude Opus 4.5 crossed key capability thresholds, calling it "the biggest change in roughly two decades of programming." AI generated 90% of Claude Code itself, and OpenAI builds internal products with zero manually-written code ². In January 2026, Anthropic launched Claude Cowork, a user-friendly version of Claude Code, including "plugins" that extend the agent's capabilities to general knowledge work (e.g., legal, finance, HR, etc.). Nearly \$1T in software market value has been wiped out since the start of the year as investors are pricing in that the same capabilities transforming code could generalize across enterprise software categories.

“The biggest change in roughly two decades of programming.”

90%

Claude Code
written by AI

0 lines

human-written
code

2026

Claude Cowork
expands to general
knowledge work

~\$1T in software market value repriced
As AI generalized beyond code

[1] Epoch AI. (March 12, 2025). LLM inference price trends. <https://epoch.ai/data-insights/llm-inference-price-trends>

[2] OpenAI. (2026, February 11). Harness engineering: Leveraging Codex in an agent-first world. OpenAI. <https://openai.com/index/harness-engineering/>

Potential Near-Term Business Implications

The Near-Term Race

Managers at most companies will be faced with a key decision of responding to potential productivity improvements vis-a-vis their own current product development roadmaps. In the near term, we believe productivity gains from product and technology spend will likely be reinvested to meet and exceed customer expectations. Falling behind for the sake of maximizing near-term profitability is likely not a winning strategy. As the very definition of product evolves (detailed below), companies will have to invest to meet this rapidly changing customer expectation. Software companies that have simple workflow tools could be easily replicated. These products, in their current formulation, will likely face meaningful price pressure and therefore will likely race to add feature sets that could help mitigate against that.

Potential Tailwinds for Middle-Market Companies

Large incumbents will be challenged by change management and may struggle with the “curse of scale” — they have more people, processes, and legacy systems to retrain and retool. At the other end of the spectrum, early-stage and smaller companies often lack an established customer base to build from.

In contrast, **middle-market companies may be best positioned**. They have meaningful customer relationships but remain agile enough to adapt quickly, develop new products, and will compete now on an effectively more level playing field.

The Evolution of Product

At the outset, most companies are unlikely to build these GenAI first software capabilities in-house. They will need enterprise-grade security, cost efficiency, and performance delivery at scale. That said, these shifts are likely to put pressure on the pricing and margins of companies that rely on traditional user interfaces and workflow products. Competition may also come from open-source solutions that are easy to deploy, customize, and scale over time with an ecosystem to support it (which the industry has seen many times before).

In the near term, human-based user interfaces will become more customizable and predictable. Over time, however, companies will need to redesign their core operations to serve machines (the AI agents), not just people. That means moving beyond human-only interfaces toward machine-based connections, with core functionality designed directly for agents. In practice, this could make many platforms “headless” (without a traditional User Interface) at their core — built to serve agents as seamlessly as they serve human users.

Generative AI enables fundamentally better products through new experiences and interfaces. Some business problems (e.g., software IDEs, simple workflow, analytics/insights dashboards, information gathering/research platforms) that are currently solved by traditional software products will be better served by GenAI first products will likely see significant disruption. Additionally, AI agents will become the customers of these products and they will more easily switch, adapt, and select best-in-breed tools vs. humans with legacy habits.

We do not believe that being a “system of record” is, by itself, a durable competitive moat. Today’s core platforms (e.g., HRIS and CRM systems) are likely to make their data available through Model Context Protocol (MCP) servers so that agents and algorithms can access it directly. Those agents will then use the data to inform decisions and automate workflows on behalf of customers. In many ways, this is the risk markets are already pricing in: that familiar user interfaces fade into the background. The underlying systems will still exist, but their primary role may shift to supplying data to intelligent agents rather than serving as the main point of interaction. **The critical question becomes: who ultimately captures the value from the decisions being powered by that data?**

Creating “Space in the P&L”

In the near term, companies will need to adopt AI quickly and comprehensively. This goes beyond embedding AI in the core practices of engineering and product teams to improve development speed and quality. It also requires applying AI across service-heavy and administrative functions (e.g., customer support, sales support, legal, HR admin, and other back-office roles) where automation can reduce reliance on lower-cost labor. Realizing these efficiency gains is not optional. The cost savings created by AI will be essential to protect margins and, importantly, to free up room in the P&L to reinvest more aggressively in other areas.

Organizational Readiness

Some companies, particularly larger organizations, are likely to experiment with leaner teams built around AI-native talent. This shift will also require reassessing the senior leadership team through a different lens. Beyond assembling a strong executive team, companies will need leaders who are genuinely aligned with and capable of driving broad, AI-led transformation across the organization.

Service Industry Conundrum

We strongly believe knowledge-based service industries will need to go through a significant transformation within the next several years and will likely see an evolution towards a more agentic first future. In the private markets there have been a few strategies that have emerged to target a transformation of service-based industries. We ultimately think investors need to price in the risk that as these industries transform, especially where the revenue models are on a per-seat or per-hour basis, they will likely face revenue compression and gross margin dollar compression. Even though it is possible that gross margin percentages themselves may go up, investors may bet that the gross margin dollar compression may be offset by a positive multiple arbitrage.

In the near-term as investors buy and transform these companies, they will likely see a lift in both gross margin percentages and dollars but ultimately those will get competed away by revenue compression that will take place over time.

Sustainable Moats that Strengthen


(1) Federated Data Powering Decisions

The platforms most likely to endure will be those that directly power decisions and transactions using proprietary datasets that no single customer could build on their own. For example, a fraud detection platform using machine learning models under the hood that evaluate every transaction on behalf of its clients becomes more valuable as it aggregates data across merchants. The insights drawn from that shared dataset improve decision quality for all participants, creating a powerful network effect.

Incumbents with years of longitudinal data and a track record of effective decisions have a meaningful advantage. Their historical data and proven models reinforce one another, **creating a flywheel: more customers generate more data, better data leads to superior outcomes, and stronger outcomes attract more customers**. These advantages are difficult for competitors to replicate.

Importantly, companies that automate decisions are paid for results — the accuracy, ROI, and quality of those decisions — not for their user interface.

Customers care about outcomes. Switching providers is not just a matter of learning new software; it means taking on risk with every decision the new system makes. This decisioning framework may be wrapped in a traditional human interface but also increasingly in headless interfaces. Investors will have to distinguish between companies that may look the same on the surface but have very important distinguishing characteristics in terms of why they're getting paid.



Returning to the fraud detection example: the true asset is the proprietary dataset built from millions of transactions across thousands of merchants. Each additional transaction improves the model for everyone. While AI may be able to replicate the algorithm, it cannot replicate the proprietary data. Without comparable data scale and history, even a technically superior model will likely produce inferior decisions.

(2) Commercial Network Effects

Sustainable moats are most likely to exist in multi-sided network businesses, where the network itself creates value for every participant. In these models, bringing together supply and demand is the core functionality. As the network grows, it reinforces itself: more demand attracts more supply, more supply attracts more demand, and customer acquisition costs decline over time. More importantly users benefit from the depth and “liquidity” of the network, which becomes hard to recreate for a new entrant. Multi-sided networks are even more powerful as they introduce an additional class of participants to contribute to the flywheel. These advantages are structural to the company’s market position, not dependent on how cheaply software can be built. Marketplaces, payment networks, and data consortia are classic examples.

While a new entrant may be able to replicate the software layer using AI, it cannot easily replicate the network or the data generated by that network. That embedded scale and connectivity are what make the moats deep and these businesses defensible.

(3) Regulatory Moats

Regulation creates structural barriers that are hard for new entrants to overcome. This usually applies to regulated industries such as financial services and healthcare. Companies that serve clients in these industries are often required to gain certification, licenses, approval, strict adherence to requirements, and human oversight. Specific regulatory knowledge that is built into products creates moats that startups generally find more limiting, giving incumbents the advantage to innovate. In addition, regulations and the associated requirements are also constantly evolving, which makes it more suited to players who are serving multiple clients in a given sector.

Looking Ahead

The speed of human knowledge work has historically been the binding constraint on how quickly businesses operate. Software accelerated it, but humans remained the bottleneck. AI agents first broke that constraint in software engineering because they could verify correctness of code and learn from tight feedback loops. The same capabilities are now expanding to general knowledge work.

AI is no longer a speculative technology. It is a structural shift in how work is performed and how software creates value. As agents take on a growing share of knowledge work, the economics of software will change: products will be built more quickly and cheaply, interfaces will matter less than outcomes, and commodity functionality will face sustained price pressure. AI capability is arguably no longer the bottleneck, but rather human factors: skepticism, process adaptation, and organizational ability to change.

The advantage will accrue to businesses that sit at the center of decisions and transactions, using proprietary operational data powering AI/ML models, operate multi-sided networks, or benefit from regulatory complexity — moats rooted in structure, data and decisions, not just software.

A compelling investment case in this environment will likely center on businesses whose competitive position strengthens as AI scales. As algorithms become increasingly abundant, the most durable advantages are likely to come from proprietary data, network effects, regulatory embeddedness, and platforms that directly power outcomes rather than relying primarily on workflow or interface lock-in.

Disclaimer

These materials are being provided to the recipient solely for informational purposes by GrowthCurve Capital, LP ("GrowthCurve"). These materials do not constitute an offer to sell or a solicitation of an offer to purchase interests in any investment product or investment vehicle or account managed or advised by GrowthCurve or any of its affiliates.

No representations, warranties or undertakings, express or implied, are given as to the accuracy or completeness of the information or statements contained in these materials. The information and statements contained herein are not legal, tax or investment advice and may not be relied on in any manner as such. GrowthCurve is not making any representation or warranty, express or implied, with respect to the fairness, correctness, accuracy, reasonableness, completeness or timeliness of any of the information contained herein, or of the results obtained from the use of such information and GrowthCurve expressly disclaims any responsibility or liability therefor. The information contained herein is in summary form and has been prepared solely for informational purposes.

Certain information contained herein, including financial information, may constitute "forward-looking statements," which can be identified by the use of forward-looking terminology such as "may," "will," "should," "expect," "anticipate," "target," "project," "estimate," "intend," "continue" or "believe," or the negatives thereof or other variations thereon or comparable terminology. Such statements include statements regarding the intent, opinion, belief or current expectations of GrowthCurve. The projections or forward-looking statements included herein are subject to change without notice, are not guarantees of future performance and involve certain known and unknown risks, uncertainties, assumptions and other factors that are difficult to predict, beyond the control of GrowthCurve. Due to the various risks and uncertainties, actual events could differ materially from those reflected or contemplated in such forward-looking statement. There is no assurance that such views are correct or will provide, with the passage of time, to be correct. Actual events, results or performance could differ as a result of factors beyond GrowthCurve's control.

Any statements involving matters of opinion or estimates, whether or not so expressly stated, are set forth as such and not as representations of fact, and no representation is made that such opinions or estimates will be realized. Any estimates or illustrative examples presented herein reflect GrowthCurve's current expectations. The information and expressions of opinion contained in these materials are subject to change without notice. Reliance should not be placed on such information or expressions of opinions contained in these materials, and such information and expressions of opinion should not form the basis for an investment decision. To the fullest extent permitted by law, GrowthCurve and its affiliates and their respective officers, directors, managers, stockholders, partners, members, employees, agents and other representatives accept no liability or responsibility for any damages resulting from information or expressions of opinions provided in these materials, including any forward-looking statements, and disclaim any liability whatsoever (including for negligence) for any loss howsoever arising from any use of these materials or their contents or otherwise arising in connection therewith.

The information provided herein is based on matters as they exist as of the date of preparation, has not been audited, is not complete or all-inclusive, and may not be updated or otherwise revised to correct inaccuracies that become apparent or to reflect information that subsequently becomes available, or circumstances existing or changes occurring after the date hereof. Delivery of these materials as of any future date does not create an expectation that GrowthCurve has updated the information herein. All information provided on and through these materials is provided "as is," with no guarantee of completeness, accuracy, and timeliness, or of the results obtained from the use of this information.

Certain information contained herein, including information concerning economic trends, is based on, or derived from, information provided by unaffiliated third-party sources. GrowthCurve believes that such information is accurate and that the sources from which it has been obtained are reliable. GrowthCurve cannot guarantee the accuracy of such information, however, and has not independently verified and is not responsible for any inaccuracies, omissions and dated information contained in such third-party information or the assumptions on which such information is based. Any analysis or conclusion regarding such information represents subjective views or beliefs of GrowthCurve and does not constitute statements of fact. The information provided herein does not constitute investment advice or an investment recommendation and does not take into account the particular investment objectives or financial circumstances of any specific person who may receive it.