

The Intelligence Economy

Artificial intelligence, productivity, and the next capital cycle

Summary

At Pender, our equity team focuses on identifying opportunities at the intersection of **sector and capital cycles, structural change, and market inefficiencies**. Our core thematic areas include **Enterprise Software, Artificial Intelligence, the Energy Transition, and the evolving geopolitical landscape**.

This white paper examines artificial intelligence as a major technological and economic shift. It is complemented by a companion podcast featuring **Toufic Boubez**, Venture Partner at Pender Ventures and a serial entrepreneur with more than two decades of experience in machine learning, cloud architecture, and enterprise software.

Artificial intelligence is one of the most consequential technological development in decades, advancing at a pace that is rapid, compounding, and difficult to fully comprehend. Industries are being restructured, entire categories of labor face obsolescence, and new sectors are emerging in their place. McKinsey estimates that AI could contribute between \$17.1 trillion and \$25.6 trillion annually to the global economy.¹

We believe this not as a source of apprehension, but as one of the most compelling investment opportunities of our time.

Before proceeding, it is useful to define a few key terms:



Artificial Intelligence (AI):

A branch of computer science focused on enabling machines to perform tasks that typically require human cognition, including learning, reasoning, problem-solving, and perception. Unlike traditional software, AI systems leverage large datasets to identify patterns, make probabilistic predictions, and adapt dynamically to new inputs.

Agentic AI: an artificial intelligence system that can accomplish a specific goal with limited supervision. It consists of AI agents, machine learning models that mimic human decision-making to solve problems in real time.²



Machine Learning (ML):

A subset of AI that involves training algorithms on data to make predictions or decisions without being explicitly programmed for specific outcomes. It encompasses a broad range of statistical and computational techniques.

Deep learning: a subset of machine learning driven by multilayered neural networks whose design is inspired by the structure of the human brain. Deep learning models power most state-of-the-art AI today.



Large Language Models (LLMs):

A class of deep learning models trained on vast amounts of data. These systems generate and interpret natural language by predicting token sequences based on learned statistical relationships. Applications such as ChatGPT and Claude fall within this category.



Hyperscale:

A distributed computing environment and architecture that is designed to provide extreme scalability to accommodate workloads of massive scale. The related term “hyperscaler” refers to hyperscale data centres, which are significantly larger than traditional on-premises data centres.³

Historical Context

AI as a General-Purpose Technology

To the untrained eye, the current moment can appear abrupt. A sudden, highly disruptive breakthrough arriving without precedent. This perception is understandable, but misleading. History, studied carefully, is among the most powerful tools an investor can have, and the future, in many respects, has a past. Rather than enabling precise prediction, history provides a framework for understanding causality, context, and the mechanisms through which change unfolds. By examining how prior technological revolutions evolved: from initial innovation, through adoption and eventual economic integration, we can situate AI within a broader continuum and develop a clearer perspective on where we may be in its development trajectory. The objective is not certainty, but better judgment.

Over time, economic development has exhibited a recurring, cyclical pattern. Roughly every half century, a new general-purpose technology emerges (steam power, electricity, the automobile, or information technology etc.) triggering a surge of capital investment. This influx of capital often leads to periods of speculative excess and asset bubbles, followed by correction. Yet, in the aftermath, these technologies enter a “golden age” characterized by sustained productivity growth and widespread economic transformation.

Each major economic era has been catalyzed by such a technology. Importantly, these innovations do not merely augment existing processes; they reorganize entire economic systems. They redefine labor markets, enable new industries, and reshape global supply chains. This pattern can be observed across several historical episodes:



Steam Power and Rail Infrastructure

(early–mid 19th century)

Early steam engines enabled the expansion of rail networks, culminating in the railway boom of the 1840s. This period of speculative overinvestment was followed by a financial collapse, after which rail infrastructure became foundational to industrial growth.



Electricity and Heavy Industry

(late 19th–early 20th century)


Initial electrification spurred large-scale infrastructure investment and utility expansion. Following financial instability in the 1890s, electricity entered a period of broad industrial adoption, driving significant productivity gains.



Automobiles and Mass Production

(early–mid 20th century)

The emergence of the automobile led to intense capital investment and industry fragmentation, followed by consolidation and the rise of mass production systems that defined modern manufacturing.





Information Technology and the Internet (late 20th century–present)

The development of microprocessors, computing, and the internet culminated in the dot-com bubble and subsequent correction. In the decades that followed, digital technologies became deeply embedded across the global economy, enabling sustained productivity growth and the rise of platform-based business models.

This historical lens brings us to the present: the emergence of artificial intelligence as the next general-purpose technology. Thus far, the AI cycle has been characterized by foundational breakthroughs in deep learning, the rapid development of large language models, and a surge of capital investment across the technology ecosystem. Yet a central question remains unresolved: what constitutes the inflection point at which AI transitions from early adoption to speculative excess, and is that transition already underway?


History suggests that such turning points are only clearly identifiable in retrospect. What appears increasingly clear, however, is that AI is following a familiar pattern, one that has historically preceded profound and enduring economic transformation.

The AI Adoption Curve

One of the defining characteristics of the current AI cycle is the unprecedented speed of adoption. Compared to prior technological waves, dispersion is occurring at a dramatically accelerated pace.

For example, ChatGPT reached 100 million users in approximately 60 days, compared to nearly two years for Instagram. More broadly, generative AI achieved roughly 40% adoption among U.S. users in under two years, whereas personal computers required more than a decade to reach comparable levels of penetration. This compression of adoption timelines reflects both the maturity of existing digital infrastructure and the inherently scalable nature of software-based innovation.

Despite this rapid uptake, the full economic impact of AI has yet to be realized. This apparent paradox can be understood through the lens of the Productivity J-Curve,⁴ which describes an initial phase in which technological adoption suppresses measured productivity. During this period, organizations incur significant costs related to infrastructure investment, workflow reconfiguration, and workforce retraining. Only once these complementary investments are fully integrated do productivity gains begin to materialize at scale.



Source: David Viney | david-viney.me

Historical evidence strongly supports this dynamic. Electrification, for instance, was technologically viable by the late 19th century, yet meaningful productivity gains did not emerge until the 1920s, when factories reorganized production systems around decentralized electric motors. Steam power exhibited an even longer lag, requiring several decades before delivering widespread economic benefits through rail networks and industrial expansion.

Artificial intelligence is likely following a similar trajectory and appears to be in the early stages of navigating the downward slope of this J-curve. However, the duration and severity of this phase may be less pronounced than in previous technological cycles. Unlike earlier general-purpose technologies, AI is being deployed on top of an already mature digital foundation built by the internet and mobile era.


It's also worth highlighting that advances in generative AI may follow a more nonlinear trajectory, as the technology possesses the unique capacity to accelerate its own development through iterative improvement. This dynamic introduces the possibility that both adoption and productivity gains could materialize more rapidly than historical precedent might suggest.

The AI Capital Cycle

Transformational technologies share a defining feature: they attract capital at a scale with their perceived potential. Artificial intelligence is no exception, and by nearly every measure, the current buildout is unprecedented in both magnitude and concentration. The closest historical parallel is the industrial capital expansion of the late 19th-century railway era, when massive upfront investment in physical infrastructure created short-term economic dislocation while laying the foundation for long-term structural growth. The AI capital cycle bears a striking resemblance to that moment.

Today, a small number of hyperscale technology firms are driving this investment cycle. In 2026 alone, the four largest (Amazon, Alphabet, Microsoft, and Meta) are projected to collectively deploy more than US\$600 billion in capital expenditures.

CapEx-to-revenue ratios, which historically averaged ~10% for large technology firms, have risen above 20% and are expected to peak near 25% in 2026, reflecting the intensity of this infrastructure buildout.



Source: Data center server. Photo by Sumali Ibnu Chamid | Alemedia.id

Importantly, this investment cycle extends beyond public markets. Private capital deployment within the AI ecosystem has also reached unprecedented levels. OpenAI recently raised \$110 billion at an \$840 billion post-money valuation, while Anthropic and xAI have secured \$30 billion and \$20 billion respectively in recent financing rounds. These capital flows reflect both the scale of opportunity and the intensity of competition.

This concentration of capital is increasingly reshaping the physical economy. AI development is not purely a software-driven phenomenon; it requires extensive investment in tangible infrastructure, including semiconductors, data centers, real estate, and energy systems. Unlike prior software cycles, the constraints are not only computational but also physical, defined by power availability, land, and grid capacity.

Modern AI training clusters can require between 500 megawatts and one gigawatt of continuous power, placing significant strain on existing utility infrastructure. At the same time, traditional grid interconnection timelines in major U.S. regions, often spanning 36 to 60 months, are fundamentally misaligned with the 12 to 24 month deployment cycles demanded by hyperscalers. This growing mismatch highlights the emergence of energy and infrastructure as critical bottlenecks in the next phase of AI development.

Taken together, the AI capital cycle reflects a transition from a purely digital investment to one that is increasingly constrained, as well as enabled by the physical world.

AI Industrial Stack

Artificial intelligence is not a singular technology. It is a structured, interdependent system that requires coordinated development across a vertically integrated stack spanning physical infrastructure, data, models, and applications. This distinction matters for investors. Unlike prior software-driven cycles, where value creation was largely concentrated at the application layer, AI demands simultaneous investment across every level of the stack. Understanding where capital is being deployed, where constraints are emerging, and where value ultimately accrues requires understanding how these layers interact.

While definitions may vary, the AI stack can be broadly categorized into four core layers: infrastructure, data, models, and applications.

The AI Tech Stack

Applications

AI-infused software applications
AI-powered devices

Platform

AI software platform components
Data
Foundation models

Infrastructure

AI datacenters
Chips
Land, electricity, and connectivity

Source: Microsoft | blogs.microsoft.com

At the foundation lies the **infrastructure layer**, which underpins the entire system. This includes semiconductors, data centres, and cloud platforms, as well as the physical systems required to support them – think power generation, cooling, high-speed networking, etc. As computational intensity increases, physical constraints have become increasingly important. The ability to efficiently move large volumes of data and manage heat now represents a critical bottleneck. In this sense, performance is no longer determined solely by advances in compute, but also by energy availability and thermal management capacity.

Above this sits the **data layer**, which serves as the essential input to all AI systems. Model performance is fundamentally dependent on the quality, scale, and accessibility of underlying data. As a result, modernizing data architecture, through cloud migration, data integration, and real-time processing, has become a prerequisite for effective AI deployment. Significant investment is being directed toward building robust data pipelines and storage systems. AI systems are only as effective as the data on which they are trained.

The **model layer** represents the core intelligence of the stack. This includes machine learning algorithms and, increasingly, large-scale foundation models such as LLMs. Advances at this layer have driven much of the recent progress in AI capabilities. However, these advancements remain dependent on the underlying infrastructure and data layers.

At the top of the stack is the **application layer**, where AI is translated into economic value. This includes enterprise software, consumer applications, and industry-specific solutions that embed AI into real-world workflows. While early investment has been concentrated in infrastructure and models, it is at the application layer where monetization and widespread adoption ultimately occur.

It's important to highlight that as a result of this stack infrastructure, AI development could be constrained by the weakest link across the stack. Each layer has levels of interdependencies, while large-scale adoption depends on seamless integration into operational systems. Further, the foundation of this entire stack rests on semiconductor technology. Semiconductors are materials whose electrical conductivity can be precisely controlled, allowing them to function as both conductors and insulators under different conditions. This property enables the creation of transistors, microscopic switches that regulate the flow of electricity. Modern chips contain billions of these switches, enabling the execution of complex computations at extraordinary speed and scale. As such, semiconductors form the backbone of all digital systems, including those powering artificial intelligence.

The AI industrial stack makes one thing clear: the potential investment opportunity in artificial intelligence is not confined to any single layer. It is distributed across the entire system. For investors, this is both the complexity and the opportunity. Those who understand how the layers interact, where the bottlenecks lie, and where capital is flowing may be better positioned to identify where durable returns are most likely to emerge.

The Productivity Shock

As the AI industrial stack matures and transitions from experimentation to large-scale deployment, it could continue to generate a productivity shock across the global economy. This shift is not merely incremental; it could represent a fundamental change in how work is performed, how knowledge is generated, and how organizations operate. AI is increasing operational velocity, redefining research and development processes, and introducing structural changes to labor markets.

Estimates of the impact vary widely. MIT economist Daron Acemoglu offers a relatively conservative projection, suggesting that AI may contribute less than 0.1% to annual productivity growth over the next decade. In contrast, Goldman Sachs estimates that widespread adoption of generative AI could increase U.S. labor productivity growth by up to 1.5% annually, potentially adding approximately 7% to global GDP over time.

On a company specific scale however, productivity gains can already be observed. In software engineering, AI-assisted development tools are resulting in massive efficiency gains across processes. Developers using tools such as GitHub Copilot have been shown to complete tasks up to 55% faster, with the largest gains among less experienced engineers.

Industry research suggests that overall coding productivity improvements in the range of 20% to 30% are becoming increasingly common. Notably, these gains are often reinvested into improving code quality, reducing technical debt, and enhancing system architecture, rather than simply accelerating output.

Customer service functions are experiencing similar improvements. Generative AI-powered assistants can analyze customer interactions in real time, recommend responses, and automate post-interaction workflows. Early studies indicate productivity increases of ~14% among support agents, with AI handling routine inquiries and enabling human workers to focus on more complex, higher-value interactions.

Beyond firm-level efficiency, the broader question surrounding AI adoption is how uptake will impact the labor market. As AI systems increasingly substitute for certain forms of cognitive labor, it is still unclear if this transition will result in net job displacement or follow historical patterns of role expansion and general job creation.

Based on what we have observed thus far, the latter appears more likely. The World Economic Forum estimates that while AI and automation could displace ~92 million jobs globally by 2030, they may also create around 170 million new roles, a net positive employment effect. This is consistent with the historical pattern of prior technological revolutions, where short-term labor displacement ultimately gave way to broader economic expansion and the emergence of entirely new categories of work. The industrial revolution did not eliminate employment; it transformed it. There is reason to believe AI may follow the same trajectory.


Investment Implications

Artificial intelligence presents a dynamic familiar from prior technological cycles: near-term uncertainty alongside long-term structural opportunity. The challenge for investors is not simply identifying beneficiaries, but distinguishing between short-term disruption and durable value creation.

We do not believe AI will eliminate economic value; it will redistribute it. Value is likely to shift away from thin application-layer software toward areas characterized by scarcity and defensibility: proprietary data, intellectual property, industry expertise, and physical infrastructure such as energy and data centers. Yet this view is not uniformly reflected in current market pricing. In some segments of the software ecosystem, valuations imply significant long-term impairment, with certain companies effectively assigned minimal terminal value. Where we believe the market is over-discounting disruption risk, we see opportunities for asymmetric outcomes.

This environment underscores the importance of active management. By maintaining a deep understanding of companies, their business models, competitive advantages, and end-market exposure, we are able to selectively increase exposure where we believe market expectations have diverged meaningfully from underlying fundamentals.


Ultimately, our approach is to use this period of volatility to build positions in high-quality businesses where we believe long-term value is being mispriced. As the AI ecosystem continues to evolve, we will remain focused on identifying where durable economic advantages are forming, and positioning the portfolio to benefit from the structural shifts that define the emerging intelligence economy.




References

- ¹ McKinsey & Company. (n.d.). The economic potential of Generative AI: The Next Productivity Frontier. The economic potential of generative AI: The next productivity frontier. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>
- ² Stryker, C. (2026, March 18). What is Agentic Ai?. IBM. <https://www.ibm.com/think/topics/agentic-ai>
- ³ Powell, P., & Smalley, I. (2025, November 17). What is hyperscale?. IBM. <https://www.ibm.com/think/topics/hyperscale>
- ⁴ Brynjolfsson, E., Rock, D., & Syverson, C. (n.d.). General Purpose Technologies (GPTS) such as AI enable and require significant. MIT INITIATIVE ON THE DIGITAL ECONOMY RESEARCH BRIEF. https://ide.mit.edu/sites/default/files/publications/2019-04JCurvebrief.final2_.pdf

To learn more and to find out who your local Pender contact is, **please visit our website.**

 penderfund.com

 [penderfund-capital-management](https://www.linkedin.com/company/penderfund-capital-management)

penderfund.com/disclaimer for Pender's standard disclaimer.

PENDER
PenderFund Capital Management Ltd.