HB Wealth

Technological Revolutions, Fast and Slow

By: Josh Rowe, CFA, PhD, Managing Director of Research & Family Office Investment Strategy, HB Wealth

This article is the second installment in our four-part series examining artificial intelligence through a wider economic lens. In the first essay, we explored the history of artificial intelligence and the recurring cycles of enthusiasm and disappointment that have shaped the field's development. If you did not read the first piece, [we encourage you to start there](#).

Across this series, we examine: (1) the history of AI and machine learning, (2) the economics of innovation and how automation spreads through the economy, (3) the implications for labor markets and the possibility of worker displacement, and (4) the potential impact of generative AI on financial markets and portfolio positioning.

Executive Summary

Technological revolutions rarely translate into immediate productivity gains. Economists have long observed that transformative technologies often appear across the economy before their benefits show up in the data. This lag reflects the time required for complementary investments, organizational changes, and new skills to develop around technology.

Generative Artificial Intelligence (AI) shows many characteristics of a general-purpose technology like electricity, the internal combustion engine, or digital computing. Such technologies reshape entire economies, but their impact is not truly felt until supporting infrastructure, new institutions, and novel workflows evolve to put them to commercial use.

Early evidence suggests that generative AI is meaningfully improving productivity at the task level. Just a few years after this technology exploded into public consciousness, we are already seeing faster output in knowledge work in activities such as writing, coding, and customer service. However, these gains do not automatically translate into either a firm-level or economy-wide productivity pickup. Bottlenecks in energy and power requirements and regulatory uncertainty conspire to slow the pace of adoption, at the same time as human users and organizations must adapt their activities to take fuller advantage of AI's potential.

Generative AI, however, shows signs of diffusing more quickly than past platform technologies. Its natural-language interface and integration with existing business software look to significantly compress the learning curve. Major complementary infrastructure (chips and datacenters) is being built at breakneck speed. Even so, widespread economic transformation will unfold gradually rather than immediately, as some bulls forecast.

For investors and business leaders, the challenge is less determining whether AI will revolutionize the economy, but understanding who will capture the value it creates, and how that value ultimately shows up in corporate earnings and markets.

About the Author

Josh Rowe, Managing Director of Research at HB Wealth, wrote a PhD thesis in the history and economics of technology, focusing on computer automation of office work in the 20th century. He has studied the history of AI, venture capital's funding of technological innovation, and the impact of technological change on financial markets—both as a resident of the ivory tower and as an investor. This surprising moment in history is the first time that he can say with any confidence that the years he spent in libraries and databases working on a doctoral dissertation might be of any practical use. He used AI in organizing and editing these essays, but the ideas (right and wrong) here are his own.

Introduction: Solow's Paradox and the Productivity J-Curve

In 1987, the Nobel laureate economist Robert Solow made an observation that would reverberate for decades among watchers of the ICT revolution. "You can see the computer age everywhere," he wrote in the New York Times, "but in the productivity statistics." It was a throwaway line, really—eleven words tucked into a discussion of manufacturing competitiveness. But those eleven words captured something that had been nagging at economists for years, and the remark lent its name to an eponymous paradox that remains central to understanding how transformative technologies actually transform.

By 1987, the personal computer had been around for more than a decade. IBM had shipped its first PC in 1981. Apple had unveiled the Macintosh in 1984. Businesses across America were buying computers by the millions, installing them on desks, training workers to use spreadsheets and word processors. The investment was enormous. The much-prophesied productivity gains were…where, exactly?

Solow's puzzle wasn't entirely concerned with computers. It was about the gap between what we see technology doing and what we can measure about its contributions to economic output. Understanding the lag, why it exists, how long it persists, and what eventually closes it, is essential for anyone trying to think clearly about generative AI.

What Makes a Technology "General Purpose"

Not all technologies are created equal, at least not in their economic impact. A better mousetrap is a fine thing, but it won't reorganize civilization. A better way to harness energy, such as steam, electricity, and internal combustion, can reshape where people live, how they work, and what becomes possible. Information, in this respect, is like energy; it is a medium that is processed by every industry to produce real, useful economic value.

Economists have a term for technologies in this latter category: general-purpose technologies, or GPTs (no relation to Generative Pre-trained Transformers). The label was formalized in the mid-1990s by Timothy Bresnahan and Manuel Trajtenberg, two economists who were trying to explain why some innovations have effects that ripple far beyond their immediate applications while others, however ingenious, remain confined to narrow domains.

A general-purpose technology, in their framework, has three defining characteristics. First, it is pervasive—applicable across many sectors of the economy rather than to a single industry. Second, it improves over time, growing cheaper or more capable in ways that continually expand its range of uses. Third, and most crucially, it exhibits what Bresnahan and Trajtenberg termed "innovational complementarities": that is, it makes other inventions more valuable, and other inventions make it more valuable in return.

Think about electricity. When Edison's Pearl Street Station began supplying power to lower Manhattan in 1882, electricity was essentially a novelty—a cleaner, safer alternative to gas lighting. But electricity wasn't just better illumination. It was a platform on which inventors could build: electric motors, refrigeration, radio, elevators, air conditioning, and assembly lines. Each of these inventions expanded the value of electrification, and electrification expanded the value of each of them. The complementarities compounded. Technological products exist in webs, and these webs display network effects, as with internet platforms.

The same pattern appears with steam power, the internal combustion engine, the semiconductor, and, by most economists' reckoning, the internet. Each began as a solution to a specific problem. Each turned out to be a platform for solutions to problems no one had yet imagined. For decades, steam engines were used for pumping water out of coal mines, not for powering factories or driving locomotives. Many platform technologies' ultimate range of applications were wildly underestimated in their own time. Early critics of the telephone, for instance, dismissed it as vastly inferior to the then-pervasive telegraph. A luxury device or "electrical toy," according to Western Union executives, not a tool for doing business. And certainly not a computer that everyone would one day possess in their pockets. Western Union officials rejected Alexander Graham Bell's offer to purchase its patents for $100,000. The historical record is littered with such bad forecasts[1]:

- "Computers in the future may weigh no more than 1.5 tons." – *Popular Mechanics*, 1949.
- "I think there is a world market for maybe five computers." – Thomas Watson, chairman of IBM, 1943.
- "There is no reason anyone would want a computer in their home." – Ken Olson, president and founder of Digital Equipment Corp., 1977.
- "The wireless music box has no imaginable commercial value. Who would pay for a message sent to nobody in particular?" – RCA's David Sarnoff rejecting the opportunity to invest in consumer radios.
- "Who the hell wants to hear actors talk?" – H.M. Warner, Warner Brothers, 1927.
- The internet's impact on the economy will be "no greater than the fax machine's." – Paul Krugman, 1998.

The common mistake here is to misjudge the path of technical improvement, often along the dimensions of cost and miniaturization, and to fail to recognize the development of the suite of technological complements that make an innovation useful beyond its initial scope. Looking back and estimating the economic consequences, each technology did register in the productivity statistics, but that process took far longer than contemporary observers expected.

Does generative AI qualify as a general-purpose technology? The case is strong, though the evolution of a robust "application layer" will ultimately render the verdict. The "pervasive" criterion seems clearly met: summarizing documents, drafting emails, writing code, analyzing images, translating languages, generating designs—these are tasks that exist across virtually every industry. Knowledge work, broadly defined, accounts for something like 40 percent of the U.S. labor force. A technology that can assist with reading, writing, and reasoning touches a vast swath of economic activity.

[1] Economist Sergiu Hart's personal webpage.

The “improvement over time” criterion also looks solid. The leap from GPT-3 to GPT-4 to the current generation of models has been rapid and continuing. Costs per token have fallen dramatically. Capabilities that seemed impressive two years ago, passing the bar exam, writing competent Python code, are now table stakes. If anything, the pace of improvement has surprised even insiders.

But the third criterion, innovational complementarities, is where things get interesting, and where the history of previous GPTs becomes instructive. Because the productivity impact of a general-purpose technology depends less on the technology itself than on the ecosystem of complementary innovations, investments, and organizational changes that grow up around it. And the evolution of those complements inevitably takes time.

The Dynamo and the Computer

Paul David, the Stanford economic historian, spent much of his career thinking about why productivity gains from new technologies arrive so much later than the technologies themselves. His most celebrated paper, published in 1990, drew a parallel between Solow’s computer paradox and an earlier puzzle: why did it take so long (almost a half century) for electric power to boost manufacturing productivity, even though the technology was manifestly more practical than steam.[2]

The answer, David argued, was not about the technology. It was about everything else.

When electric motors first appeared in factories in the 1880s and 1890s, manufacturers used them the way they had used steam engines: as a centralized power source driving overhead line shafts, pulleys, and belts that transmitted power to individual machines. This was a perfectly rational response. Factory buildings were designed around a central steam engine. Workers were trained for steam power.

---

[2] Notably, technological superiority does not guarantee success. David is equally famous for a paper entitled “Clio and the Economics of QWERTY,” in which he elucidates two counterintuitive properties of technological evolution: path dependence and “lock-in.” QWERTY was, by design, a slower keyboard arrangement than alphabetical or other ergonomic options. Due to the problem of typewriter keys sticking when the best typists tapped away at full speed, QWERTY emerged as a standard to throttle back their speed. Because it was associated with a then-dominant platform, it ended up setting the standard for all future keyboard input. Currently dominant platforms—not necessarily the best future technologies—often set the standards that dictate the path of future development. Thus, they become “locked in.” The same phenomenon may be observed in the victory of VHS videocassette format over its rival Betamax, which many users viewed as better tech, thanks to cheaper and more flexible licensing arrangements. Certain bugs in the code of VisiCalc, an early spreadsheet software, reappeared in subsequent programs like Lotus 1-2-3 and Microsoft Excel.

Replacing the steam engine with an electric motor, while keeping everything else the same, offered modest efficiency gains but nothing revolutionary. Better to run an existing factory on steam or convert piecemeal to electricity than tear it down. Businesses with high fixed costs would prefer to slowly depreciate assets rather than restructuring wholesale.

The transformative potential of electricity lies in its scalability. Electric motors could be distributed—one motor per machine, each drawing power independently from the electrical grid. This meant factories no longer needed to be organized around a singular power source. Machines could be arranged according to the logic of production rather than the physics of power transmission. Natural lighting could replace the dim canyons required by belt-driven layouts. Smaller firms could afford power without investing in their own steam plants. Manufacturing could be reconceived from the ground up.

But that reorganization required new factory buildings, new machine tools designed for individual motors, new skills, and new management practices. It required, in other words, a broad-based rethinking of how manufacturing was done. And that rethinking could not happen overnight. Existing factories had decades of useful life left. Managers trained in steam-era methods remained in charge. The knowledge of how to design production flows for distributed power had to be developed, codified, and diffused.

David estimated that the productivity surge from electrification didn't arrive until the 1920s—roughly forty years after the first commercial power stations began operation, and a full generation after electric motors became clearly superior to steam on narrow technical grounds. The technology was ready long before the economy was.

The story of digital computers follows a similar script. By the late 1970s, the technical advantages of organizing business on computers, whether mainframe, mini, or micro, were patent. Machines that could process information millions of times faster than humans, never tired, and never made arithmetic errors—surely that would show up in improved output per worker? But it didn't, not in the aggregate statistics, not for more than a decade.

Why? Because computers, like electric motors, require complements. They needed software that did more than replicate paper processes. They needed workers who could use them effectively. They needed managers who understood how to design workflows around new capabilities. They needed networks that let machines talk to each other, standards that let data move between systems, and business models that exploited the possibilities of digital information. Early adopters frequently bought clunky and expensive installations like IBM's System 360 because, in the 1960s, enterprise capex on computers (much like AI today) was the thing to do. Then they left them sitting idle or lightly used, except for rudimentary automation of inventory and payroll. IBM, like Microsoft after it, won share not because of a technically dominant product, but because its army of sales and service professionals could help buyers integrate machines gradually into their core workflows, usually in a very high-touch, hands-on manner.

For these reasons, it is often not immediately optimal for businesses to invest in the latest-and-greatest gadgets, nor to disrupt themselves in a rush to become “AI companies.” Lawncare.ai probably does not need to exist—or at least not yet. Implementation bottlenecks within organizations must be broken down, but likely utterly demolished. This phenomenon leads to very real innovator’s dilemmas—to invoke the conceptual vocabulary of Harvard’s late business thinker Clayton Christensen. The old way of doing business works very well for many incumbents and indeed may be critical to serving their core customers. The new possibilities frontier is more likely to be explored by startups, addressing for small, enthusiastic audiences. Only iteratively and carefully will older workflows be replaced in large organizations.

Christensen’s framework is in some sense an intellectual heir to that of Joseph Schumpeter.[3] Schumpeter, an Austrian economist writing mostly before World War II, is seen as the prophet of “disruptive innovation.” He argued that waves of what he called “creative destruction,” often in the form of new technologies and business models that at first slowly, then violently displace incumbents, constitute the engine powering continual growth in capitalist economies. But disruption typically emanates from the edges, not the center of the economy.

At the time of Solow’s 1987 remark, much of the physical and conceptual infrastructure needed to grease the integration of computers was still missing or immature. His productivity surge manqué came later, in the mid-1990s, when the internet began to knit everything together, when enterprise software finally matured, when a new generation of managers who’d grown up with PCs took the reins. What appeared paradoxical in 1987 looked, in retrospect, like a predictable lag.

[3] Schumpeter was a colorful character. He is famous for many things, not least the declaration that his life’s ambition was to simultaneously become the greatest economist in the world, the greatest horseman in all of Austria, and the greatest lover in Vienna. He claimed to have achieved two of those goals but did not specify which.

The Shape of the S-Curve

The pattern David identified—technology first, complements later, productivity last—has repeated often enough that economists now treat it as something approaching a law of technological diffusion. New capabilities propagate through economies not in a straight line but along an S-curve: slow initially, then accelerating rapidly, then leveling off as saturation approaches. The graph of this phenomenon, called a logistic curve, may be familiar to those who looked at epidemiologists' models of how the COVID-19 pandemic spread. So-called SIR models and technological diffusion models indeed are specified very similarly. The same mechanisms, exponentially increasing "infection" leading to saturation and stabilization, are at work.

The economist Carlota Perez, building on decades of research into historical episodes of technical change, has developed perhaps the most influential model for understanding how adoption and diffusion unfold in a series of cycles. In her telling, major technological revolutions pass through two broad phases, separated by a "turning point" that often involves a financial market eruption.

The first phase, what Perez calls "installation," is dominated by financial capital. Investors, entrepreneurs, and speculators pour money into the new technology, erecting infrastructure, funding experiments, and driving adoption in leading-edge applications. This is the era of exuberance, of grand visions, of "this changes everything." It is also the era of overbuilding, asset bubbles, and spectacular busts. The canal manias, railway booms, and dot-com bubbles all occurred during installation periods. Financial euphoria and panic are both symptoms and causes of major technological revolutions, in a self-reinforcing, reflexive dynamic.[4]

---

[4] Investors George Soros and Bill Janeway have both described how bubbles form from self-propagating feedback loops between investor sentiment and fundamental industry change. Capital markets don't merely observe change in business—they cause it, leading to further loops that can become amplified by investor herding into booms and busts. The seminal model of financial booms was articulated by Charles Kindleberger and Hyman Minsky, and looks remarkably similar in its periodization to Perez' technology cycles. Janeway, like Perez, sees these feedbacks as essential to driving not just the pace, but the direction of innovation. See Soros' *The Alchemy of Finance* and Janeway's *Doing Capitalism in the Innovation Economy*.

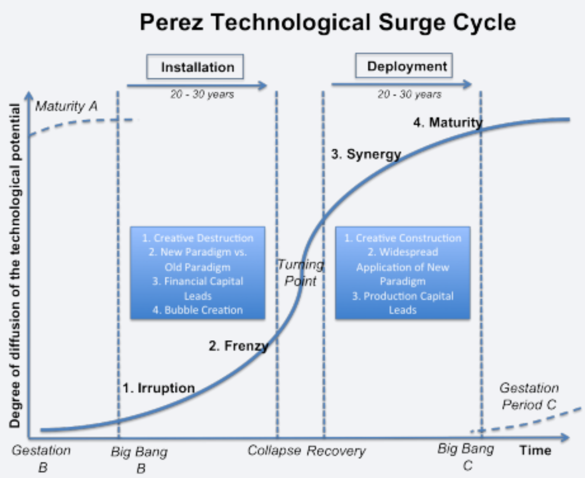

Gartner Hype Cycle

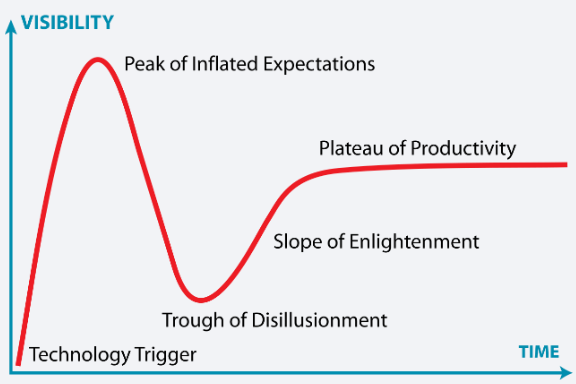

Two maps of technological adoption cycles. Sources: [Brian Manning] and [Wikipedia]

The turning point arrives when the financial excesses correct, expectations reset, and attention shifts from building new infrastructure to actually using it profitably. The second phase, ”deployment,” is led by production capital: the steady, incremental investment that spreads the technology through the broader economy. This is when authentic, lasting productivity gains materialize, when the new paradigm becomes normalized, when what seemed revolutionary becomes simply how things are done. This is the point where growth rates settle down, and valuations mature.

If we situate generative AI on Perez’s map, we appear to be somewhere in the late installation period—probably in what she would call the “frenzy” phase. The signs are recognizable: enormous capital expenditure on chips and data centers; energy and land constraints emerging as binding limits; a proliferation of pilots, proofs of concept, and early products; valuations that bake in extrapolative assumptions of widespread adoption. The key question is how long the installation phases lasts before the turning point, and how smoothly the deployment proceeds afterward.

Perez's framework offers a helpful corrective to both excessive optimism and excessive skepticism. To the optimists, it says: yes, this technology is real, but the productivity payoff requires organizational and institutional changes that require time to germinate. To the skeptics, it advises: the lag is normal, not a sign that the technology is overhyped. The dynamo took forty years. The computer took twenty. Patient observers saw the revolution coming before growth statistics confirmed it.

What the Firm Sees vs. What the Statistician Measures

One source of the productivity paradox is simply measurement. National income accounts were designed to track the output of factories making physical goods, not knowledge workers producing information. When a lawyer uses AI to draft a contract in two hours instead of six, how is this increase in efficiency recorded? If the lawyer bills for a third of the hours, the measured output of legal services, revenue, actually falls. If the lawyer keeps the same billing rate but handles more clients, the gain might surface, but slowly and with considerable noise. If the quality of the contract improves, fewer errors, better tailored to the client's needs, that's a welfare gain for the client, but it may never register in productivity statistics at all.

Measurement holes are more than a minor technical issue. A substantial fraction of the economy consists of services where output is difficult to quantify, and quality improvements harder still. Healthcare, education, professional services, government—in all these sectors, the relationship between inputs and measured product is fuzzy at best. A technology that makes knowledge work better, faster, or more accurately may confer enormous benefits that official statistics simply miss.

Economists possess various techniques for adjusting productivity measures to capture quality improvements, but these adjustments are imperfect, contested, and fail to track changes in real time. If you want to know whether a new technology is generating real value in real firms, the macro data will often be the last place to look.

What does show up more promptly is evidence at the task level and the firm level. Here the early returns on generative AI are striking, even if the aggregate statistics remain unmoved. Consider a 2023 study by Shakked Noy and Whitney Zhang, two MIT economists who ran a randomized controlled trial with several hundred college-educated professionals. Participants were given realistic writing tasks, drafting press releases, short reports, analysis memos, and randomly assigned to either use ChatGPT or work without it. The AI-assisted group completed tasks 40 percent faster, on average, and the output was rated higher quality by evaluators who didn't know which submissions were AI-assisted.

Even more striking: the productivity gain was largest for workers who started out as below-average performers. The best writers got better. But the weaker writers improved dramatically—their output quality rose while their completion time fell. AI acted as a kind of equalizer, compressing the distribution of performance.

A larger study by Erik Brynjolfsson, Danielle Li, and Lindsey Raymond examined customer service agents at a software company who were given access to an AI assistant. Again, productivity rose significantly—about 14 percent on average—and again, the gains were concentrated among less experienced and less skilled workers. Novice agents using AI matched the performance of agents with months more experience. The technology didn't just make work faster; it rendered the less skilled more capable.

A different sort of study was conducted by researchers at Harvard Business School and Boston Consulting Group, who ran an experiment with actual BCG consultants. Some were given access to GPT-4 for a set of realistic consulting tasks; others worked without it. The AI-assisted consultants completed 12 percent more tasks, did so 25 percent faster, and produced work judged 40 percent higher quality. Once again, the gains were largest at the bottom of the performance distribution.

These are not trivial effects. They are large enough to matter at the firm level, and they are consistent across multiple studies in different contexts. Will they generalize across all organizations? It's early days. Hints of real productivity windfall exist in micro and anecdotal evidence, even if the economic drivers today are upstream in manufacturing and infrastructure.

But there is a catch: task-level gains don't automatically become firm-level gains, and firm-level gains don't automatically become economy-wide gains. The same BCG study found that when consultants applied AI to tasks outside its zone of competence—tasks demanding judgment, experience, synthesis, or creativity beyond what the model could reliably deliver—performance actually declined. Consultants leaned on the AI when they shouldn't have, trusting outputs they should have questioned. The technology created value at one frontier and destroyed it at another.

This is the challenge of what the researchers called "the jagged technological frontier." The utility of AI systems is uneven. Models that excel at summarization may stumble at reasoning. Models that generate plausible prose may hallucinate facts with serene confidence. The skill of the user lies in discerning where the frontier is: when to trust but verify, when to delegate, and when to do the work yourself.

This quality of judgment takes time to develop. Institutionalizing it takes longer still—building repeatable workflows, quality controls, and training programs that reliably capture productivity gains while avoiding traps. Anyone who has taught high school or college-age students in the age of generative AI will have pondered whether the next generation of workers are AI-native cyborgs with uncanny technical facility or slavish automatons bereft of critical thinking skills. In the interim, many organizations are rationally cautious, adopting AI in circumscribed pilots rather than wholesale transformation. Learning how to wield these new tools where the stakes are limited.

The release of AI coding agents in late 2025 has sharpened these dynamics considerably. Tools like Anthropic's Claude Code don't merely assist with discrete tasks; they can architect, implement, and debug entire software projects with minimal human direction. For some users, the experience has been disorienting in its intensity. Some companies have begun tracking employees' "interactions per day" with coding agents, on the assumption that more interactions signal greater productivity. An ongoing UC Berkeley study of a two-hundred-person organization found something counterintuitive: even as workers offload tasks to AI agents, they are simultaneously working longer hours. The researchers labeled one cause of this phenomenon "task expansion"; when nontechnical colleagues use AI to vibe code rough prototypes, engineers inherit the cleanup work, adding to already growing workloads. The line between roles that were once clearly delineated has blurred. A recent survey found that more than 40 percent of C-suite executives reported that AI tools save them at least eight hours per week. Among non-managers, 67 percent said AI had saved them fewer than two hours—if it saved them any time at all.

The gap between experimentation and value capture is borne out in many contemporary firms' experience. A late-2025 study by Boston Consulting Group found that 60 percent of companies report seeing little to no material value from their AI investments—minimal revenue growth or cost savings despite substantial outlays. Only 5 percent of firms globally have successfully integrated AI into core workflows in ways that deliver significant financial and operational impact. Corroborating research from MIT Sloan Management Review suggests that up to 95 percent of generative AI pilots fail to advance to production or deliver a return on investment.

The reasons are by now familiar: organizations treat AI as a standalone tool rather than redesigning core business processes around it; individual productivity effects at the task level fail to translate into faster company-wide delivery; and the time required for employees to realize meaningful performance improvements, at least eleven weeks of consistent use, by one estimate, exceeds the patience of managers hunting for quick wins. Perhaps most striking, data from S&P Global indicates that 42 percent of companies abandoned the majority of their AI initiatives in 2025, a sharp increase from 17 percent the year before. It's not that AI doesn't work; it is evidence that making it work, at scale, reliably, inside the messy reality of existing organizations, is very hard. These findings counsel more sober expectations about the timeline to material (sticky rather than experimental) revenue than some model builders, hyperscalers, and their investors may have reckoned. Technology has leapt ahead; complements are still catching up.

Why Complements Take Time

If the history of electrification and computing teaches anything, it's that the most arduous part of deploying a general-purpose technology is not creating the technology itself. It's building everything else around it that makes the technology generically useful.

Start with infrastructure. For electricity, the complement was a continental network of power generation, transmission, and distribution—plus the factory buildings, machinery, and electrical equipment designed to exploit it. For computing, it was networks, storage, software, and the "data plumbing" (often fiber-optic cable that, once laid, lay dormant for years) that ultimately enabled information to flow freely across systems. For generative AI, this infrastructure buildout is well underway: data centers, specialized chips, power capacity, and cooling systems.

Constructing all this supporting infrastructure is expensive, slow, and constrained by physical and financial reality. You can't vibe code a power plant into existence or accelerate the permitting process for a new substation just because an AI model is ready to consume more electricity. Physical energy systems and grid capacity are tangible, real-world brakes on breakneck model growth. A trillion parameters, a billion tokens – these sound like abstract improvements governed by Moore's Law. But in many cases, they scale linearly or even sub-linearly in real-world energy costs. Where cloud software licenses can be created ex nihilo, generative AI models exhibit high fixed and variable costs. Project and debt finance tend to be stickier wickets than a Series A round. Throwing money at a problem works to an extent, but transmission, interconnection, and more efficient training and reasoning algorithms are thorny questions that will require adaptation and ingenuity.

Next, consider human capital. Every major technological transition has required new skills. And not just technical skills, but managerial and organizational ones. The productivity fillip from electrification was only seen when a new generation of factory managers, trained to think in terms of production flows rather than power transmission, took charge. Computing's killer apps arrived only after a generation that had grown up with PCs and networks reached positions of authority. Skills diffuse more slowly than algorithms. Knowledge is embodied in people, and people learn gradually and retire slowly. Whole subfields of organizational behavior, "human computer interaction," and UI/UX design are dedicated to easing this skills transition.

Then there are the organizational complements: workflows, processes, reporting structures, incentive systems, and quality controls. A firm that simply bolts AI onto existing processes may capture some gains, faster document drafting, quicker data retrieval, but the larger gains come from AI-native workstreams around the machines can do differently. That means rethinking task bundles: which activities should humans do, which should AI do, and which require a hybrid collaboration? It means redefining jobs and changing how work is monitored, evaluated, and rewarded. The rise of such augmented intelligence, variously termed "cyborg" or "centaur" modes of interaction, will have a huge role in determining how organizations capture value. This process is very much in its early innings.

This type of organizational change is exceptionally hard. It will necessitate experimentation, and experiments often fail. It requires coordination across departments and functions (sales and IT, products and finance, for example). It will depend on buy-in from workers who may reasonably suspect that "reorganizing around AI" is a euphemism for "reorganizing you out of a job." Even when the potential is clear in principle, achieving ROI in practice requires leadership, patience, and a willingness to tolerate disruption. Surveys consistently find that only a small fraction of firms—perhaps as few as 1%, according to McKinsey—have reached so-called "AI maturity," meaning their data, processes, and governance are ready for scaled deployment.

Lastly, there are institutional complements: regulation, liability rules, professional standards, and the broader social infrastructure that governs how technologies are used. These matter more for some applications than others. An AI that helps draft marketing copy operates in a relatively permissive environment. An AI that assists with medical diagnoses operates under layers of regulatory scrutiny, liability risk, and professional ethics. In sectors like healthcare, finance, and critical infrastructure, where the consequences of error are severe and accountability is tightly regulated, adoption will lag the advancing frontier by years, perhaps decades. Though some technologists might decry big-government Luddism, this caution is appropriate for systems designed to protect against harm.

In these heavily regulated industries, and ones more sensitive to error than the "move fast and break things" ethos of Silicon Valley, widespread adoption may be permanently retarded by safety and reliability thresholds that prove stickier than technologists project. At minimum, these sectors, which remain less tech-enabled than the enterprise, require error traceability and transparency—a major challenge for foundation models. Such "last miles" can hold back deployment in workflows that otherwise seem like they should be automated – managing data from patients' electronic medical records, for example. Investor Sarah Guo of Conviction, a VC firm investing in AI applications, has coined the phrase "Minimum Viable Quality" as distinct from "Minimum Viable Product." She argues that, in highly regulated domains, entrepreneurs and investors may be underestimating the challenge of surmounting this quality hurdle.

Where We Sit on the Curve

So, where does this leave us with generative AI? How far along the S-curve have we progressed?

A realistic answer is that we are early—probably earlier than much market excitement implies. Adoption surveys tell a story of intense experimentation but limited scaled deployment. A majority of large enterprises are running pilots. A much smaller fraction have moved AI into production workflows with measurable business impact. The phenomenon is typical of the early diffusion stage: visible in leading-edge firms and startups, perceptible in certain occupations and industries, but not yet normalized across the economy.

Productivity statistics bear this out. As of mid-2025, U.S. labor productivity growth has been solid but not spectacular—better than the doldrums of 2010–2019, but nothing like the surge that accompanied the internet buildout in the late 1990s. Macro data show no clear AI effect yet, though they also don't capture evidence of the technology doing harm. We occupy the ambiguous zone where believers can find confirmatory anecdotes and skeptics find disconfirming ones simply because the data are still so noisy. Anyone betting on AI to redeem an economy laden with debt, deficits, and elevated interest rates (indeed, rates likely pushed higher by frenetic capital spending), might have to wait for an uncomfortable interval.

What we can say is that the micro evidence is more encouraging than the macro evidence. Controlled studies show real productivity gains at the task level—ranging from modest improvements on routine tasks to gains exceeding 50 percent in favorable conditions. Early production deployments report meaningful improvements in efficiency and quality. The case for AI as a potential productivity enabler is compelling at the firm level, but we are a long way off from calls for a grand economic reset. There has yet to be any discernible momentum in the aggregate productivity statistics. That’s precisely what Paul David would predict at this stage of adoption.

There are additional reasons why task-level productivity gains may take time to surface in economy-wide statistics. Jobs are collections of tasks. Even if AI dramatically accelerates some tasks, others act as bottlenecks. A software developer who can write code twice as fast still must wait for code reviews, attend meetings, coordinate with teammates, and navigate organizational processes. If workflows are not redesigned to integrate AI output, the non-assisted tasks may actually take longer—more code produced means more code to review, more features to test, more documentation to maintain. The gains at the task level get absorbed by friction elsewhere in the system.

Moreover, we may be on the descending portion of what economists call a productivity “J-curve”. When firms adopt transformative general-purpose technologies, measured productivity often initially falls, because resources are diverted to investment, reorganization, and learning that do not immediately show up as measured output. The resistance training and intangible capital being accumulated, new workflows, tacit knowledge, and organizational adaptation, are invisible to standard statistics. Only later, as the complements mature and the investments bear fruit, does measured productivity rise to reflect the underlying transformation. The J-curve framework suggests that today's disappointing macro numbers may actually be consistent with significant long-run gains—the dip before the climb.

Meanwhile, critical stumbling blocks are real and visible. Hallucinations, the tendency of language models to generate plausible-sounding falsehoods, are not a mere curiosity. They are an economic cost, requiring human oversight that blunts efficiency gains. A 2024 analysis by researchers at the Boston Fed estimated that hallucination rates in production settings, while improving, remained high enough to require significant verification effort, especially in high-stakes domains. Until reliability improves, many firms will rationally keep humans in the loop, limiting the efficiency pickup that full automation could deliver.

Data infrastructure remains a bottleneck. Many enterprises have years of accumulated data in formats that AI systems struggle to use—scattered across legacy systems, poorly documented, and inconsistently labeled. The work of integrating, cleaning, structuring, and governing that data is unglamorous but essential. Without it, AI models trained on clean public datasets may perform poorly in the Wild West of private corporate information.

And then there's the matter of trust. Managers who can't explain why an AI made a particular recommendation may be reluctant to delegate important decisions. Workers who fear their expertise is being devalued may resist adopting tools that seem to threaten their status. Customers who've been burned by chatbot hallucinations may prefer to talk to a human. Trust is built slowly, through experience and track record. A technology that has been in widespread use for barely two years hasn't had time to build the institutional trust that is foundational to widespread deployment.

Why This Time Might Be Different (Somewhat)

If the history of GPTs were perfectly deterministic, we might predict a twenty- to forty-year lag between the emergence of generative AI and its full economic impact—the same timeframe associated with the electricity and computing revolutions. But history doesn't repeat so neatly, and there are reasons to suppose the current transition could unfold more rapidly.

The most consequential difference between this technology rollout and past ones is the interface. Prior computing paradigms required users to adapt to the machine—to learn programming languages, master complex software, or, at a minimum, navigate menus and commands designed by engineers. Successively more abstract layers of software were built over many generations to more ergonomically translate from machine code to something human users understood and back: assembly language, compilers, scientific and business programming languages such as FORTRAN and COBOL, PC operating systems, and graphical user interfaces (GUIs). All this is to help ordinary workers interact effectively with the computer system.

Generative AI inverts this hierarchy. The interface is natural language; the same medium humans use to communicate offline. You don't need to learn Python to use ChatGPT. You just need to articulate clear prompts in one of 95+ languages and describe what you want. Previous generations of executives, trained on paper spreadsheets and intra-office memoranda, were slow to adopt Lotus 1-2-3 and email because they were initially unnatural and difficult to use for non-digital natives. Compare that to today—company executives describe being "Claude-pilled" and experiencing the thrill of writing their own software. Software engineers are enthusiastic adopters, but sales, finance, marketing, and customer service departments are finding agentic tools surprisingly intuitive.

This matters enormously for diffusion velocity. The classic barrier to technology adoption is the user's learning curve: how long it takes a new user to become productive. Organizations often move at the speed of the slowest adopters in decision-making positions, not the fastest tech-savvy employees. More value is created when whole teams are using systems – so intra-organization diffusion is epidemic in nature as well. It has been a commonplace in the history of computing that polished, human-friendly user interfaces lagged well behind frontier capabilities.

A technology with a steeper learning curve diffuses more slowly because each potential adopter must invest more time and effort. Technologies with a shallow learning curve can spread much faster. Here, generative AI tools might be poised to buck the trend. Natural language and visual interfaces don't obviate all need for learning-by-doing; users still need to understand what AI can and can't do, how to write effective prompts, how to discriminate valuable information, and accelerating AI slop—but they dramatically lower the barriers to initial trials.

A second factor suggesting more rapid uptake is the computability of these novel tools with existing systems. Generative AI isn't arriving on a greenfield. It's being integrated into the instruments that knowledge workers already use: email, document editors, spreadsheets, search engines, customer relationship management systems, and coding environments. Microsoft's Copilot is woven into Office 365. Google's Gemini is embedded in Workspace. GitHub Copilot sits inside the development environments where programmers already dwell. This kind of integration means adoption need not force users to change their workflows dramatically or replace their physical devices; the AI comes to them in the form of user-friendly extensions already embedded in daily routines.

Contrast this with previous paradigm shifts. The transition from typewriters to word processors occasioned new hardware, new software, and new skills. The transition from filing cabinets to databases remade how information was organized. The transition to e-commerce saw the buildout of entirely new systems for inventory, fulfillment, and payment. Each of the changes involved significant switching costs. Generative AI, by comparison, can often be deployed as an augmentation layer atop existing systems, which should reduce friction.

A third factor is the competitive structure of the AI ecosystem. Unlike some previous platform shifts, which quickly consolidated around dominant players, the AI landscape in 2025 remains relatively pluralistic. Multiple frontier model providers compete vigorously, exchanging performance leadership positions every few months. Open-source alternatives, frequently Chinese, are capable enough for many applications. Specialized models fine-tuned for particular domains, legal, medical, and financial, are proliferating. This competition drives down prices, improves quality, and creates options for firms that don't want to stake their operations on a single vendor.

The result is a swarm of applications, each finding its niche. Rather than wait for a single dominant platform to win, firms can experiment with multiple tools, adopt what's useful, phase in rollout, or switch providers as the market evolves. Multi-parallel experimentation may accelerate learning across the economy.

These dynamics will not repeal the S-curve or eliminate the importance of the coevolution of complementary innovations. Training workers, reforming processes, building data infrastructure, and establishing governance frameworks will all delay greater implementation. The physical infrastructure of AI, data centers, power generation, and semiconductor fabrication, faces constraints that won't be solved by clever software, along with potentially destabilizing debt and depreciation overhangs (more on this later). But the combination of a natural language interface, compatibility with existing tools, and a competitive application layer could plausibly compress the lag from decades to years.

If electricity's economic significance would not be felt for forty years after Edison and Tesla, and digital computing's postdated Moore and Noyce by twenty, perhaps generative AI may fully arrive in ten. That's still a long time if you're wagering in the stock market, but it's fast by the standards of technological revolutions.

The Measured and the Mismeasured: The AI Productivity Puzzle

A further reason to treat GDP growth point estimates with caution is that we may not recognize the productivity gains even when they arrive. The measurement problem discussed above is especially pertinent with generative AI.

First, look at what generative AI actually does in many of its current applications. It drafts documents that humans then edit. It generates code that developers review and modify. It summarizes information, suggests options, and answers questions. In each case, the AI is contributing to a process whose final output is attributed to a human worker or team. It is embedded. Humans and technological artifacts, researchers find, interact in loops that reinforce each other.5 What value gets attributed to which step in the production process? How do you measure the productivity of a tool that operates in the middle of a workflow? If a consultant uses AI to complete a project faster, but the billing rate stays the same, measured productivity (revenue per hour) is unchanged even though real productivity (output per effort) has risen. If a software developer uses Copilot to write code more efficiently but spends the freed-up time on higher-value design work, the measured output of "lines of code per hour" may not

[5] Several seminal papers on this subject are by the management sociologists Brian Pentland and Martha Feldman, found here and here.

change much even though the value of the developer's contribution has increased. This is at least as important for companies as it is for economic forecasters: countable "value" ascribed to a technology is profoundly important for pricing.

The same ambiguity affects quality improvements. AI-assisted legal research might produce contracts with fewer errors, and that's valuable—but GDP doesn't measure "errors avoided." If AI-enabled customer service resolves issues faster and more accurately, customer satisfaction rises—but satisfaction isn't in the national accounts. The welfare gains from better quality, fewer mistakes, and faster service are real, but they're largely invisible to the statistics we use to track productivity.

Over time, some of these gains will seep into measurable categories. Firms that use AI effectively will outcompete those that don't, and market share will shift. New products and services that AI makes possible will appear in the data as new output. Price indexes will eventually adjust to reflect quality improvements, though always with a lag. But for now, and for perhaps another decade, we should expect the official statistics to understate the true productivity contribution of AI—just as they understated the contribution of electrification in the 1920s and computing in the 1990s.

The fact that consumer surplus is undermeasured by official statistics should not license magical thinking. 10% per annum real GDP growth runs up against all kinds of economic and physical barriers. Even more outlandish forecasts pepper the techno-optimist narrative bubble. This is just not how economic growth works. Even the most transformative technologies like steam power do not suddenly unleash step-change growth; too much of the enabling ecosystem remains to be built; too much knowledge and technique remain to be learned; too much of the benefit is unmeasured.

A technology (like social networking, for example) that never shows up in the statistics ever probably isn't generating the true value its promoters claim. But the difficulty of quantifying well-being is a reason for humility about using macro data to assess a technology that is still early in its diffusion and that operates in sectors where output is inherently difficult to measure.

The Investor's Dilemma

For investors, the S-curve model presents a difficulty. Both the horizontal and vertical axes are unspecified. The horizontal axis represents time, but we have no way of knowing until later which section of the curve the present moment represents. The vertical axis may be conceptualized in terms of penetration or "percent adoption" among the susceptible audience, with the top being 100%. 100 per cent of what, exactly?

If generative AI is a genuine GPT, the economic value it creates could be enormous—comparable to electricity or computing, transforming industry after industry over decades. But if history is any guide, most of that value may accrue to users of the technology rather than producers of it, and much of it will arrive after the current wave of capital spending has cycled through boom and bust. Thus, how to make money from adoption depends greater on the terms by which a technology is adopted. In short, who pays?

The railroads reshaped the American economy. They enabled continent-spanning markets, accelerated industrialization, and created consumer surplus that dwarfs any reasonable estimate of the profits earned by railroad companies themselves. But railroad investors, famously, fared poorly. The stock market history of American railroads is a chronicle of overbuilding, rate wars, bankruptcies, and consolidations. Even in a network industry that lends itself to natural monopolies, the competitive dynamics, the capital intensity, and government regulation of railroads made it difficult to capture significant economic rents. The technology was indeed transformative; investment returns were mediocre.

A similar pattern holds for electricity, for automobiles, for airlines, for telecoms. The industries that deployed transformative technologies created vast economic value. The investors who financed that deployment often captured only a fraction of it, after losses to competition, obsolescence, and miscalculation.

The question for AI investors is not whether the technology is real, but how the value it creates will be distributed. Will it flow primarily to the model builders, who control the most sophisticated AI systems? To the cloud providers, who supply the compute infrastructure? To the chip makers, who fabricate the specialized hardware? To the application developers, who build the tools that enterprises actually use? To the enterprises themselves, who capture the productivity gains or penetrate new markets? Or to consumers and workers, in the form of lower prices, better products, and new economic opportunities?

The best current answer is that we do not yet know. Platform technologies often exhibit winner-take-most dynamics, with network effects and switching costs concentrating value among a few dominant players. The history of digital technology is full of such examples: Microsoft in PC operating systems, Google in search, Apple in smartphones, and Amazon in e-commerce. If AI follows this pattern, the leading model providers or infrastructure players could extract enormous rents. The frenzied AI sentiment in mega-cap public tech companies and foundation model private fundraising are highly leveraged to this outcome.

The Consequences of Abundance and Substitution

New markets and new technologies introduce recursive feedback loops that can be hard to predict. Today's big bets are being placed on infrastructure developers, chip suppliers, and model builders. They are being placed against sellers of application software. Let's imagine, as one widely read essay does, that these trends don't mean-revert; instead, they accelerate. What if the replacement of human software engineers, of human accountants, of human consultants proves initially successful? What if it spreads more widely? Companies would report improved margins and improved earnings. This would encourage more competitors to automate greater portions of their white-collar workforce in an arms race. It would similarly prompt AI tool builders to double down on spending, investing greater and greater sums into data centers, chips, and model training. In this manner, layoffs and digital capex are two faces of the same phenomenon. One dystopian vision describes a self-reinforcing cycle where structural unemployment of the greater part of the knowledge workforce leads to a collapse in consumer spending, weakening aggregate demand, but perversely further pushing firms and employers to substitute AI for people to cut costs. A spiral ensues.

While the stock market might initially rejoice, investors would have to revisit many of their assumptions about the structure of the economy. The old saw that the U.S. consumer is the engine of the global economy would soon seem antiquated. What in 2025 looked like a cyclical capex splurge becomes the primary ongoing driver of GDP and stock market growth instead. Perhaps we should be asking whether the entire U.S. economy, built on high-value skilled services, has actually been a "daisy chain of correlated bets" on an obsolete technology—humans' ability to move information around and communicate with each other. In other words, while China's economy is built on advanced manufacturing complex value chains, America's (and its whopping debt load) is built on lawyers, consultants, and bankers. Which model looks more sustainable in the age of agentic AI?

Face-to-face interaction (the province of the attorney, the salesperson, the investment banker) here is revealed to be nothing more than a friction impeding the arrow of pure economic growth rather than a scarce and highly valuable resource. With a "giant extraction layer" removed and substituted for by AI agents, the owners of capital drive ever more efficiently to a future of abundance. No matter that trillions of dollars of contracts link that frictional layer together, that the livelihoods of tens of millions depend on it, that the fundamental impetus of economic growth assumes the viability of Americans' middle-class consumption lifestyle. How could we hope to predict the political or financial consequences of such a shock?

Economists and strategists largely dismiss this "dystopian abundance" scenario as science fiction; for one, it underestimates how difficult it is to remove frictions—it assumes purely rational corporate behavior, pliable regulation, and an infinite well of cheap power at hand to sink into a single economic project. It imagines that the pace of change speeds up for plausible microeconomic reasons, and yet, in order to expedite the thought experiment, it conveniently assumes away all manner of bottlenecks. But it asks a question well worth asking.

The mechanisms for a structural reorientation of the economy are too varied and too interrelated to allow for such grand predictions over short periods. Fundamental lags occur due to economic "irrationality" or the indecisiveness of corporate boardrooms, regulatory pushback, irregular model performance, and critical shortages. But we can observe a directional thrust toward ever-increasing dependence on AI integration for competitive success, of herding, cascading, and path dependence of capital allocation, of the weakening bargaining power of legacy vendors, expert services, and white-collar labor. Rather than counteracted by regulation or fiscal policy, these trends may indeed be supported rather than deterred by government due to national security imperatives—such as concerns that if we don't build the energy and infrastructure to maintain frontier AI leadership, the Chinese will.

Most of economic history is described by mean reversion, at least in the medium run. Big capital markets moves tend to unwind. New technologies' hype tends to dissipate. Extrapolating rapid change accelerating into the future is rarely a safe bet. But new technologies can also create economic dynamics that reinforce their primacy in the economy, drawing more and more economic resources and activity toward them. Think of social networks, smartphones, and streaming. Those calling Netflix's streaming model in 2011 a "fad," and betting on linear broadcasters to capture the value from streaming would be surveying a wrecked portfolio today.

The scenario envisaged above is certainly thought-provoking, but it rather self-consciously runs against the tide of history. When the great economist John Maynard Keynes predicted, in 1930, that technological productivity gains would reduce the human workweek to fifteen hours, his error was to underrate the "elasticity of human wants"—in other words, the more we could produce, the more we wanted. New productive capabilities merely enabled us to explore new ways to put people to work. Worries over imminent displacement may be making the same mistake.

Second, as one economist points out, the number of simultaneous extreme phenomena that would need to occur for a displacement spiral leading to double-digit unemployment stretches credulity. One needs to assume near-total labor substitution, negligible investment absorption, and a complete failure of fiscal response—simultaneously. Simply put, even if capital's share of income rises sharply at labor's expense, capital income doesn't have zero spending velocity; profits get reinvested, distributed, taxed, or spent. For demand to collapse structurally, redistribution mechanisms would need to fail persistently while investment opportunities dry up at the same time.

Finally, and most prosaically: current data don't support imminent collapse. As of early 2026, job postings for software engineers are rising 11% YoY, not falling; new business formation is expanding; and data center construction is boosting employment in other sectors. Recent studies have found that AI-exposed sectors in Europe are actually increasing employment, notwithstanding small productivity gains. Macroeconomists struggle to forecast payroll growth two months forward with any reliability—inferring the forward path of labor destruction with high confidence, an epistemological courage ungrounded in existing evidence. We will have a lot more to say about labor markets in the next part of our essay. For now, we will narrow our focus to the macroeconomic effects of technology shocks.

Investing in a Productivity Boom

It is of course salutary for investors to bear in mind extreme tail scenarios, so-called "black swans" or six-sigma events, but the market reaction is unlikely to reflect a clear consensus for some time. The current structure of the AI market remains fragmented at many links in the value chain; it's not yet obvious who commands disproportionate market power. Multiple frontier models compete closely on performance. Open-source alternatives have proven surprisingly effective. The application layer is dynamic and crowded. Enterprise customers are wary of lock-in and cautiously managing vendor relationships. It is not yet clear that any single player will develop the kind of durable moat that characterized platform winners like Google in search, Amazon in e-commerce, Microsoft in enterprise software, or Facebook in social networking.

Meanwhile, the infrastructure arms race is imposing enormous capital requirements that show little sign of slowing, while new equipment's functional depreciation is incredibly rapid. Fiber-optic cable laid in 2000 is still in use today. Graphics chips from the same time are not. Today's cutting-edge data center may be tomorrow's stranded asset in a few years as models become more efficient, chip architectures evolve, and demand patterns change. The history of technology infrastructure is littered with investments that looked essential at the time and worthless shortly after. The fiber-optic boom of the late 1990s created the backbone of the modern internet, but the investors who financed that buildout largely lost their shirts. Most small investors in the railway boom that covered England in thousands of miles of track within the span of a few years were wiped out entirely.

Investing in AI today looks as precarious as when some of these earlier bubbles were inflating. A fundamental seed lies at the core of market enthusiasm – but markets are prone to extrapolation in the "installation" and "frenzy" phases and disappointment later. The relationship between investor returns and the magnitude of a given technology's social import is weaker than our intuition tells us.

The present moment feels all the more vertiginous when executives most responsible for the AI buildout have begun issuing cautionary notes about all the capital flowing their way. Jensen Huang, NVIDIA's CEO and the primary beneficiary of the chip boom, acknowledged in mid-2024 that "there's a different set of risks" when customers are buying infrastructure ahead of proven demand, and that the AI trade had elements of "speculative investment." Satya Nadella, boss of Microsoft, who has bet hundreds of billions on OpenAI and Azure AI infrastructure, has warned that "there's hype, there's froth" in the market and emphasized that "we have to be grounded in real usage and real value creation." Sam Altman, OpenAI's chief executive and perhaps the most prominent evangelist for AGI (artificial general intelligence, or a future where algorithms are far better at every cognitive task than humans), told an interviewer in late 2024 that "some of this is going to end in tears" and that valuations in parts of the ecosystem were "disconnected from near-term reality." When the principal players absorbing all the capital start sounding notes of caution, outside observers might reasonably take heed.

Most troubling is the self-referential quality of much of the early deal activity. A significant portion of AI revenue, particularly among the hyperscalers and model providers, comes from deals with other participants in the AI ecosystem itself: cloud credits extended to startups, infrastructure partnerships among the major players, and contracts where consideration often takes the form of equity or deferred commitments rather than upfront cash. This pattern rhymes uncomfortably with the fiber-optic buildout of the late 1990s, where carriers booked revenue by swapping capacity with one another, creating an illusion of end-demand that evaporated when investors got cold feet. Numerous large companies filed for bankruptcy amid discoveries that their financial numbers had been massively inflated.

Whether today's circular arrangements constitute the same species of accounting legerdemain or are simply hallmarks of an ecosystem in the early stages of maturity is a matter of much investor dispute; what is less debatable is that round-trip deals create tight interdependencies among multiple companies' valuations, leaving the entire complex particularly vulnerable to contagion when sentiment inevitably shifts. When Company A's revenue depends on Company B's spending, and Company B's valuation depends on contracts with Company A, a wobble anywhere can propagate everywhere.

We will examine these investor implications in more detail in a later section. For now, the takeaway is simply that the AI capital market today has outrun the reasonable forecasting ability of anyone familiar with past technology cycles. "Can OpenAI go bankrupt?" is the thrust of recent editorials in The Economist magazine, and, conspicuously, by venture enthusiast Sebastian Mallaby in *The New York Times*. The answer is, of course, "yes," but the most interesting questions concern what happens next.

Patience is a Virtue

The recurrent cycles of technological adoption identified by economists tell us that much of the AI story remains unwritten and warn us against undue certainty. Excitement is rational, but we hardly need to point out the absurdity of investors believing that AI will bring about a critical disjuncture in history (AGI perhaps), akin to the industrial revolution, and *simultaneously* believing that they can today forecast the economic winners. Profound change in the way businesses process information is doubtless ahead of us. Humility is the watchword for anyone today trying to anticipate what this means.

For firm managers, the implication is that AI adoption is a process, not an event. Pilots are fine; full integration will require much more comprehensive organizational evolution. The firms that capture the largest productivity windfall will be those that treat AI implementation as a strategic investment in complements—data infrastructure, workflow redesign, training, governance—rather than a bolt-on to existing processes. The historical record strongly suggests that the spoils go to those who reorganize when it makes competitive sense, not those who automate around the edges. Sometimes being a fast follower is preferable to being an early adopter. To take only one example—Apple's business faces challenges on many fronts, but demurring to pour a trillion dollars into cloud infrastructure may come to be seen as prudence rather than "fumbling the future."

For policymakers, the implication is that the productivity effects of AI are coming, but not on any schedule that neatly fits electoral cycles. Investments in education, infrastructure, and institutional adaptation should pay dividends over decades. The temptation to take short-term positions for or against the "AI revolution" should be tempered by the recognition that the technology's significance will not be perceived or understood for a generation. Policy should focus on future-proofing the economy and the labor force, on enabling the critical complements from power to interconnection to accounting and legal rules, and on ensuring AI safety guardrails.

For investors, the implication is that great fortunes are to be won and lost surfing the waves of market sentiment. But these are, at heart, leveraged beta trades. Alpha—at least the durable kind—comes when investors know something about the shape of future fundamental demand and place bets that pay out in the form of cash flows, not paper-trading profits. It is very early to assess the shape of the opportunity for this sort of endeavor. Much of the value, we suspect, will ultimately accrue at the capillaries of the AI ecosystem, not the heart. AI tools that transform mundane businesses like gas stations and staffing companies, rather than shiny data centers, may be the economic story of 2035.

This generation of AI may fail to achieve its proponents' forecasts of superintelligent abundance. And yet despite our cautions, it will not fizzle. This is not tulipomania. All our historical analogies (railroads, electricity, computing, the Internet) delivered on their transformative promise, eventually. Hoped-for productivity improvements arrived—less in a big bang than in a steady, year-over-year drumbeat. The economy gradually reorganized to take advantage of the latent technological potential. It will do so again –perhaps too slowly for investors betting on a revolution, but more swiftly than in past episodes. How can we look for early signals of who will win, and how should we position portfolios across major asset classes? This is the subject of our next essay.

If you have any questions, please reach out to your client service team, visit us at hbwealth.com, or call 404.264.1400.

Notes

1. Robert Solow's remark appeared in his review of Manufacturing Matters by Stephen Cohen and John Zysman, published in the New York Times Book Review, July 12, 1987. The phrase "everywhere but in the productivity statistics" entered the economics lexicon almost immediately.

2. The formal treatment of general purpose technologies appears in Timothy F. Bresnahan and Manuel Trajtenberg, "General Purpose Technologies: 'Engines of Growth'?" Journal of Econometrics 65, no. 1 (1995): 83–108. A more accessible overview appears in the Brookings Institution's analysis of AI's economic effects: https://www.brookings.edu/articles/the-effects-of-ai-on-firms-and-workers

3. Paul David's canonical treatment of the productivity paradox is "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," American Economic Review Papers and Proceedings 80, no. 2 (1990): 355–361. The paper remains essential reading for understanding technology diffusion lags.

4. Carlota Perez's framework is fully developed in Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages (Edward Elgar, 2002). Her taxonomy of installation/frenzy/turning point/deployment has become standard vocabulary among technology historians and venture capitalists alike.
5. The Noy and Zhang study is Shakked Noy and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," working paper, MIT, March 2023. Available at economics.mit.edu.
6. The Brynjolfsson, Li, and Raymond study examines customer service agents and was published as Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," NBER Working Paper 31161 (2023), subsequently appearing in the Quarterly Journal of Economics (2025).
7. The BCG/Harvard study is Fabrizio Dell'Acqua et al., "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Harvard Business School Working Paper 24-013 (2023).
8. Estimates of AI maturity among enterprises vary by survey methodology but consistently find that only 10–20 percent of large firms have moved beyond pilots to scaled deployment. See McKinsey's annual "State of AI" surveys and Deloitte's enterprise AI reports for representative findings: McKinsey state of AI 2025, https://www.deloitte.com/content/dam/Deloitte/cr/Documents/consulting/2024/the-state-generative-ai-enterprise.pdf
9. On the measurement challenges for productivity in service sectors and knowledge work, the literature is vast. A useful starting point is the work of Chad Syverson, particularly "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown," Journal of Economic Perspectives 31, no. 2 (2017): 165–186.
10. The comparison between railroad investor returns and the broader economic value created by railroads is a recurring theme in the economic history of transportation. See, among others, Robert Fogel's work on the social savings from railroads and the extensive literature on nineteenth-century railroad finance.

Important Disclosures

This article may not be copied, reproduced, or distributed without HB Wealth's prior written consent.

All information is as of the date above unless otherwise disclosed. The information is provided for informational purposes only and should not be considered a recommendation to purchase or sell any financial instrument, product, or service sponsored by HB Wealth or its affiliates or agents. The information does not represent legal, tax, accounting, or investment advice; recipients should consult their respective advisors regarding such matters. This material may not be suitable for all investors. Neither HB Wealth nor any affiliates make any representation or warranty as to the accuracy or merit of this analysis for individual use. This information contains forward-looking statements, predictions, and forecasts ("forward-looking statements") concerning our belief and opinions in respect to the future. Forward-looking statements involve risks and uncertainties, and undue reliance should not be placed on them. There can be no assurance that forward-looking statements will prove to be accurate, and actual results and future events could differ materially from those anticipated in such statements. Certain information herein is based on third-party sources believed to be reliable, but which have not been independently verified. Past performance is not a guarantee or indicator of future results; inherent in any investment is the risk of loss. Specific investments described herein do not represent all investment decisions made by the above date. The reader should not assume that investment decisions identified and discussed were or will be profitable. Specific investment advice references provided herein are for illustrative purposes only and are not necessarily representative of investments that will be made in the future. Investors are advised to consult with their investment professional about their specific financial needs and goals before making any investment decision. HB may hold positions (long or short) in the companies mentioned in this paper.