



Kevin Hebner, PhD

Managing Director, Global Investment Strategist, TD Epoch

AI Will (Eventually) Turbocharge Productivity and Profits

- Artificial Intelligence (AI) capital expenditure (Capex) is estimated to be \$600 bn this year in the U.S. But when will productivity and profit gains arrive? And how will markets respond if it takes longer than consensus expects?
- We are bullish regarding the long-term prospects for AI. However, this paper demonstrates that technological diffusion and commercial success always take decades, not years.
- Adoption rate vs. return on investment: Recent surveys emphasize a significant gap between experimentation and impact. While many firms have initiated AI pilots, the vast majority see no measurable return. Why? Misalignment with existing workflows and the technology just isn't good enough yet.
- Productivity growth has averaged 1.9% since 1870: Sometimes it is a bit lower (Great Depression) and sometimes a bit higher (WWII, 1986-2007). But growth has never been anything like what techno-optimists are calling for. We forecast a moderate 0.4 ppt impulse to trend productivity growth over the next decade.

- **Technology investing from railroads to the internet:** There are several recurring themes, most notably that tech revolutions always take decades not years. The path from impressive demo to widespread adoption is inherently long and unpredictable.
- **Substitution of capital for labor:** We are richer than our grandparents because of automation. Most wealth creation comes from shifting the production of tasks away from slowly-improving human labor to rapidly-improving machines. However, it's a slow process. Empirically, the elasticity is only 0.5, implying substitution is tedious and time-consuming.
- **Engineering hurdles and the “March of Nines”:** Reaching 90% reliability represents only 20% of the journey. For critical applications like autonomous vehicles (AVs), robotics or radiology, the path to 99.999% reliability requires a decades-long slog through a zillion “edge cases.”
- **Current large language model (LLM) architectures are a “dead end” for true artificial general intelligence (AGI):** Leading AI researchers like Yann LeCun, Fei-Fei Li and Demis Hassabis argue true intelligence requires spatial intelligence and a 3D world model—capabilities that remain years away from commercialization.
- **Implications for investors:** (a) AI capex growth to moderate to a more sustainable trajectory; (b) Emphasize “Quality Tech”—companies with sustainable free cash flow (FCF) and established user bases—rather than speculative startups; (c) Diversify beyond concentrated bets in U.S. tech (although we remain constructive on U.S. equities), including commodities and infrastructure; and (d) AI is set to turbo-charge creative destruction and the innovator’s dilemma. Tech revolutions always induce massive upheaval in corporate leadership.
- **Where could we be wrong?** First, the world of bits transforms more quickly than the world of atoms, so the timeline could be compressed this cycle. Next, rather than “five nines,” some disruptive technologies succeed at first by being “good enough” for massive, lower-stakes markets. Third, fear of missing out (FOMO) plus plentiful liquidity could result in companies over-investing for longer than “rational” models suggest. The U.S.-China superpower race could also incentivize an extended AI capex cycle.
- **Four reasons why this is different from the internet bubble:** 1. Valuations are stretched but not as extreme as 1999 (reflecting the enormous profitability of hyperscalers and how little debt or equity issuance has occurred this time around). 2. Liquidity is plentiful and becoming even more so (vs. aggressive tightening in 1999-2000) 3. Infrastructure capacity (e.g., data centers) is lagging in demand today (vs. 1999 when 97% of new fiber optic cables were dark, not yet being used). 4. Superpower race, with AI central to all three domains of power (economic strength, tech prowess, defense capabilities).

Will we soon be eclipsed by robots with productivity and profits soaring ever higher?

Ever since ChatGPT’s public release in late-2022, we’ve been overwhelmed by headlines about AI’s astonishing progress, with improved models and breathtaking capabilities.

Top executives have described AI as the most significant breakthrough in human history, even surpassing the impact of electricity or fire, and as a technology capable of transforming or saving the world.

This paper examines the economics of previous tech breakthroughs, from railways and electricity through to PCs and the internet. Such innovations always take decades to mature, and evolve along paths that are entirely unpredictable, even for the Thomas Edisons of the world. We are bullish on the long-term potential of AI, but our perspective suggests the optimists’ scenario is acutely implausible. That is, it will take longer than consensus expects for the latest tech wave to turbocharge productivity and profits. It always does.

New tech always takes decades to diffuse across the economy: AI will not be different

AI has enormous long-term potential and could well become as impactful as railways and electricity were in their day. However, such gains only diffuse across the economy after: dozens of complementary innovations are introduced; a vast new infrastructure is built; the long bridge from impressive pilot to commercial product is crossed; workflows are redesigned; and the labor force adopts entirely new skill sets. Creative destruction takes time, but it is coming and always results in once mighty incumbents being displaced by plucky startups.

We are richer than our grandparents because of the history of innovations. Much of this reflects automation, which involves shifting production tasks away from slowly-improving human labor to rapidly-improving machines. This doesn't occur overnight, it takes decades, with most of the value added coming from entirely new products, services and even sectors. No one is ever smart enough to foresee how a tech wave will unfold. Although some, like James Watt, Thomas Edison, Alexander Graham Bell, Guglielmo Marconi, Henry Ford, Steve Jobs, Bill Gates, Jeff Bezos and Elon Musk see a bit further than others.

Roughly 20% of companies were using AI in 2025, but few saw a measurable return

Before diving into the history of productivity and tech waves, this section discusses five surveys that help us understand where we are today, beginning with one by McKinsey.¹ It found that a majority of companies are beginning to use the technology, including agentic AI. However, most are still in the preliminary stages, experimenting and running pilots, with few capturing a significant return to their investment. The study emphasizes that successfully introducing AI requires redesigning workflows, a process that is both challenging and time intensive.

A second survey, this one out of Massachusetts Institute of Technology (MIT), found that 95% of organizations see no measurable return on their investment in AI pilots.² The article stressed that most fail due to misalignment with day-to-day workflows, with the lack of customization constituting a core barrier to success. The sectors that reaped the greatest benefit were tech, media, and finance.

The next survey, of roughly 500 CFOs, is from Duke University, in conjunction with the Fed (**Figure 1**).³ Similar to the results from McKinsey and MIT, it found AI is having negligible impact at the enterprise level.⁴

Workflows

¹ "The state of AI in 2025: Agents, innovation, and transformation," by McKinsey & Company, 2025.

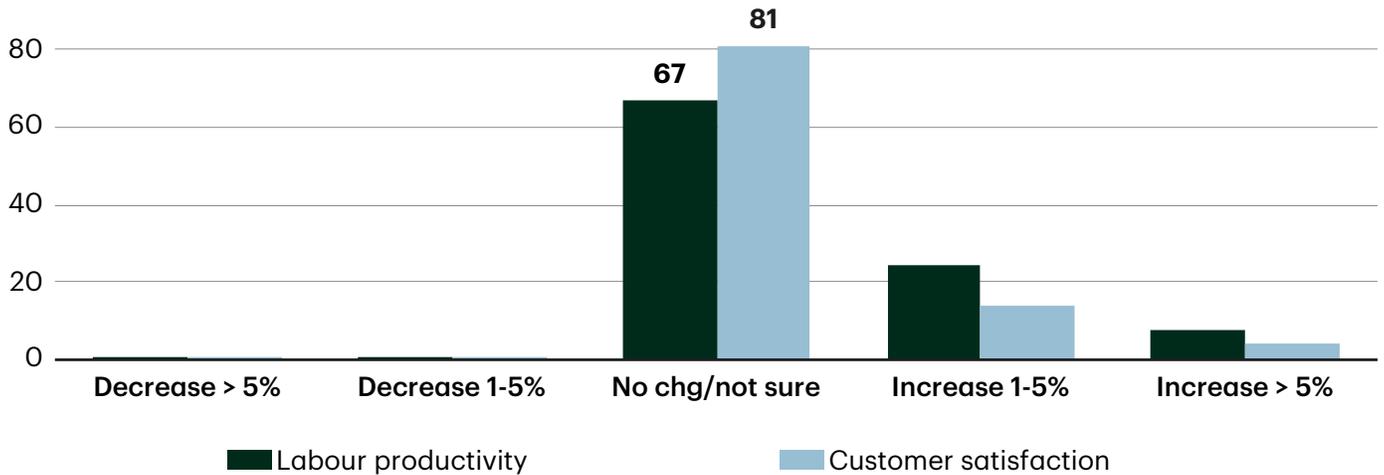
² "The GenAI Divide: State of AI in Business, 2025," A. Challapally et al, MIT, 2025.

³ "CFO Outlook for 2026: Tariffs, Hiring, Prices, and AI Impact," Duke University, 2025.

⁴ "Artificial Intelligence in the Firm," by F. Chen (Harvard) et al, 2026, obtains a similar result. Analyzing a dataset that covers 100,000 engineers at 500 firms they conclude "AI adoption leads to moderate increases in productivity ... However, these productivity gains do not pass through to effects on output, task composition, or employment." That is employee-level improvements do not yet translate into firm-wide gains.

Figure 1: Over the last 12 months, how has your firm's use of AI affected productivity and customer satisfaction?

The vast majority of CFOs are not yet seeing an impact from AI.



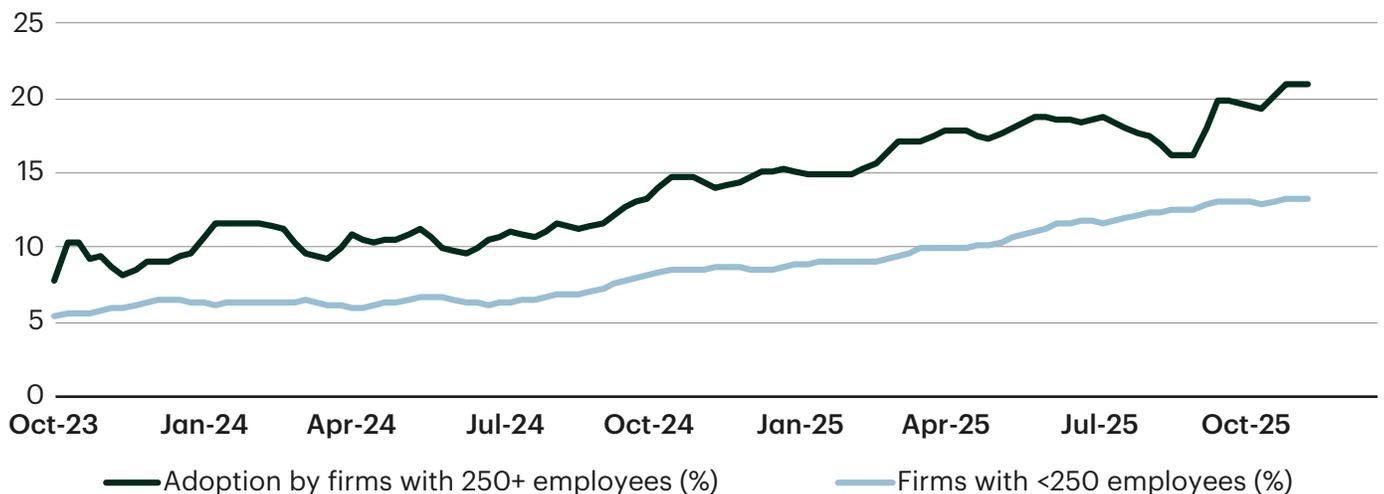
Source: CFO Outlook for 2026: Tariffs, Hiring, Prices and AI Impact. Duke University, FRB Richmond and FRB Atlanta, The CFO Survey – Q4 2025 (November 11 – December 1, 2025)

Fourth, a University of Chicago survey asked a large panel of economists to evaluate the statement, “Adoption of AI will lead to a substantial increase in the growth rate of real per capita income over the next ten years.”⁵ The most common response was “uncertain,” with participants emphasizing: (a) there’s considerable uncertainty about both the timeframe and the net impulse to growth, and (b) AI is likely to prove more positive for corporate profits than for personal income, exacerbating the K-shaped economy.

Finally, the U.S. Census Bureau surveys 1.2 million firms every two weeks (**Figure 2**). Note that the share of enterprises adopting AI is rising, but the pace of increase is much lower than many overly enthusiastic headlines would have you believe.

Figure 2: AI adoption rate expected over next six months (% of firms).

Has increased by 10 ppts over last two years.
At this rate 80% of large companies will be using AI by 2035-2040.



Source: Bloomberg, U.S. Census Bureau. As of 12/31/2025

⁵ “AI and growth,” University of Chicago, 2025.

What does history tell us about the likely impact of AI on productivity?

The previous section demonstrated that, while much is going on below the surface, AI progress at the enterprise level has been modest, at least so far. Regardless, techno-optimists remain adherents of the “explosive growth” theory, which predicts a golden age is imminent with productivity growth at least doubling from the historical trend.

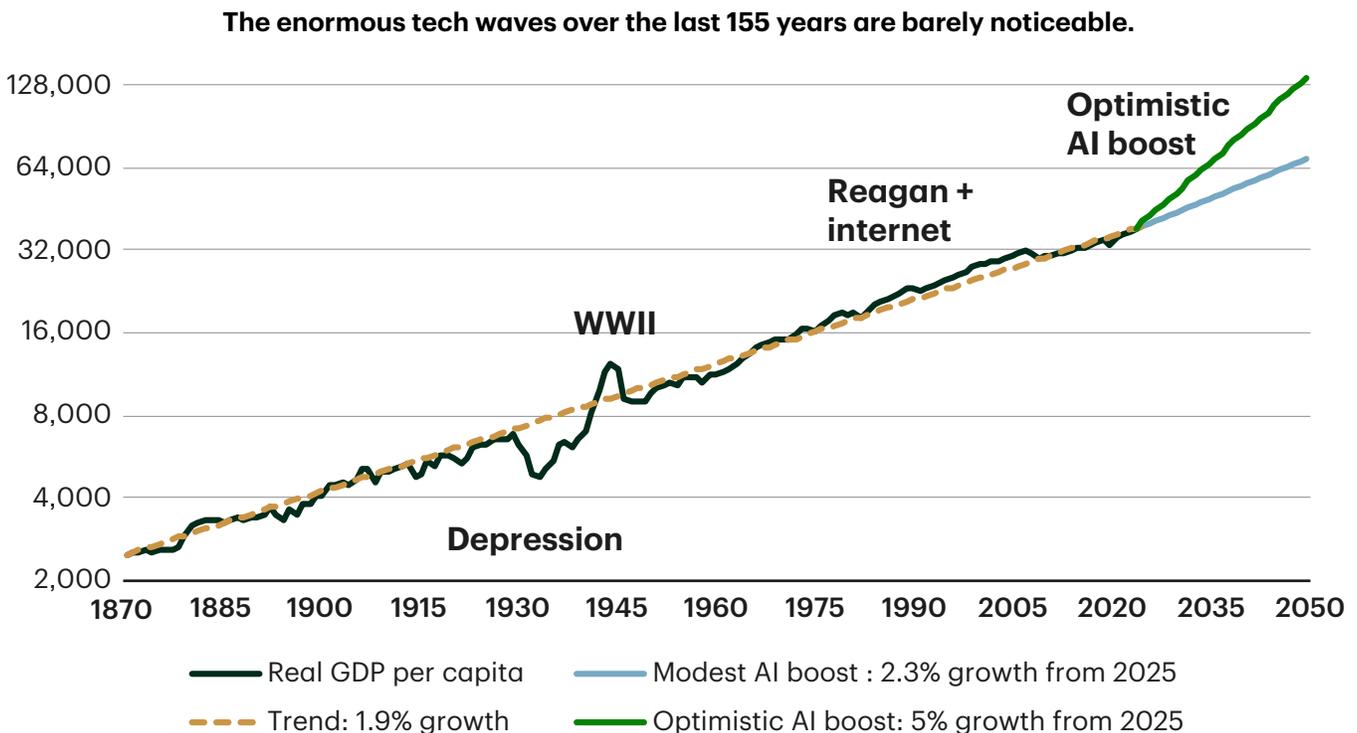
We push back hard on this view and believe the impulse from AI is likely to be in the range of 0.2 to 0.6 ppts, with a median forecast of 0.4 ppts.⁶ We also emphasize that, while 0.4 ppts of extra growth might sound small, if it is maintained for a decade or more, it represents enormous value creation, probably the biggest global wealth impetus in history.

A key reason for our prediction of a 0.4 ppt impulse is the history of productivity growth (**Figure 3**). Since 1870 it has averaged 1.9% with the odd episode

below (the Great Depression) and the occasional period above (WWII, 1986-2007). Today it is bang on 1.9% and there has never in history been anything similar to what techno-optimists are calling for. To illustrate, we’ve added to Figure 3 an “Optimistic AI boost” scenario (3.1 ppt impulse, bringing productivity growth up to 5%, beginning from 2025).

The period since 1870 has witnessed more tech waves than ever before in history, yet there is nothing remotely similar to what AI optimists believe is in the cards. As Sir John Templeton cautioned, the four most dangerous words in investing are “This time is different.” It is conceivable that explosive growth could happen, but it hasn’t before in our species’ 300,000-year history, so I wouldn’t bet my lunch money on it.

Figure 3: Long-term trend for real growth is 1.9%, with small deviations for wars and recessions.



Source: Federal Reserve Bank of Dallas, TD Epoch. As of 03/31/2025

⁶ Among the dozens of projections that have been published, the most credible are from the OECD, ECB, McKinsey & Company, McKinsey Global Institute, Penn Wharton Budget Model and the Peterson Institute, as well as academics such as Daron Acemoglu (MIT), Ricardo Reis (LSE) and Tyler Cowen (GMU). We discount unduly optimistic views, such as those from Anthropic and Epoch AI, as historically implausible.

Technology investing from railroads to the internet: Implications for the AI cycle

We now turn to the history of tech waves, which serves as a vital, though not exact, guide for understanding the AI revolution. Based on our analysis of innovations since 1870, we emphasize five insights.⁷ They are summarized in this section with all the gory details provided in Appendix A (Four takeaways from the last two centuries of tech innovation) and Appendix B (Complementary Innovations – It takes a village).

The first takeaway concerns commercial uncertainty. Achieving sustainable business and financial success is far more difficult than the initial invention, however brilliant that might have been. History is full of brainiacs—from James Watt to Guglielmo Marconi to Bill Gates—who famously misunderstood the eventual scale and evolution of breakthroughs like the steam engine, radio, the PC, and the internet.

The second theme is the long struggle for viability and sustainable FCF. Most pioneers, including Thomas Edison, Alexander Graham Bell and Henry Ford, faced decades of financial hardship and skepticism before reaching profitability. The constant need for risk capital incentivizes “overpromotion” and extravagant claims, which historically fueled exuberance in railways, radio, and the 1990s internet. Elon Musk and Sam Altman are just the most recent in a long line of showmen.

Next, the path from impressive demo to widespread adoption is inherently unpredictable and torturously long. For example, it took nearly 50 years from the invention of the light bulb until a majority of U.S. homes were electrified. Most recently, it took the

internet well over a decade for real businesses to emerge. This offers a valuable lesson if we make the analogy that ChatGPT (released Nov 2022) is to AI as Netscape Navigator (launched Dec 1994) was to the Internet.⁸ At this point in the internet boom (early 1998), Google had not yet been incorporated, Mark Zuckerberg and Daniel Elk (founder of Spotify) were in middle school, and the creators of Airbnb were in high school. So, 2026 is just exceedingly early in the AI wave.

Fourth, market dominance and “winner-takes-most” dynamics are not new (think of Western Union in telegraph and AT&T’s long held monopoly). However, first-mover advantage appears more fragile in the age of digital tech, as seen with displaced firms like Netscape (browsers), Yahoo! (search), Napster (music), Friendster (social media), and Nokia (smartphones). Today’s tech behemoths are likely more vulnerable than consensus recognizes.

The final takeaway concerns complementary innovations. The long lag between an impressive tech demo and commercial success is largely due to the necessity of complementary innovations. A single breakthrough is never enough to create a viable, sustainable industry; it requires an ecosystem of secondary inventions. Examples include Bessemer steel for railroads, alternating current for electricity, and pneumatic tires for automobiles. For modern robotics, it will require advances in multimodal world models, spatial intelligence, high-torque actuators, tactile sensing, solid-state batteries and semantic connectivity. Regardless of Elon Musk’s promises, this will take at least a decade, and probably more.

Innovation

⁷ Sources include: “Engines that move markets: Technology investing from railroads to the internet and beyond,” by A. Nairn, 2018 and “How the internet happened: From Netscape to the iPhone,” by B. McCullough, 2018, as well as “The rise and fall of American growth: The U.S. standard of living since the Civil War,” by R. Gordon, Northwestern, 2016 and “Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity,” by D. Acemoğlu (MIT) et al, 2023.

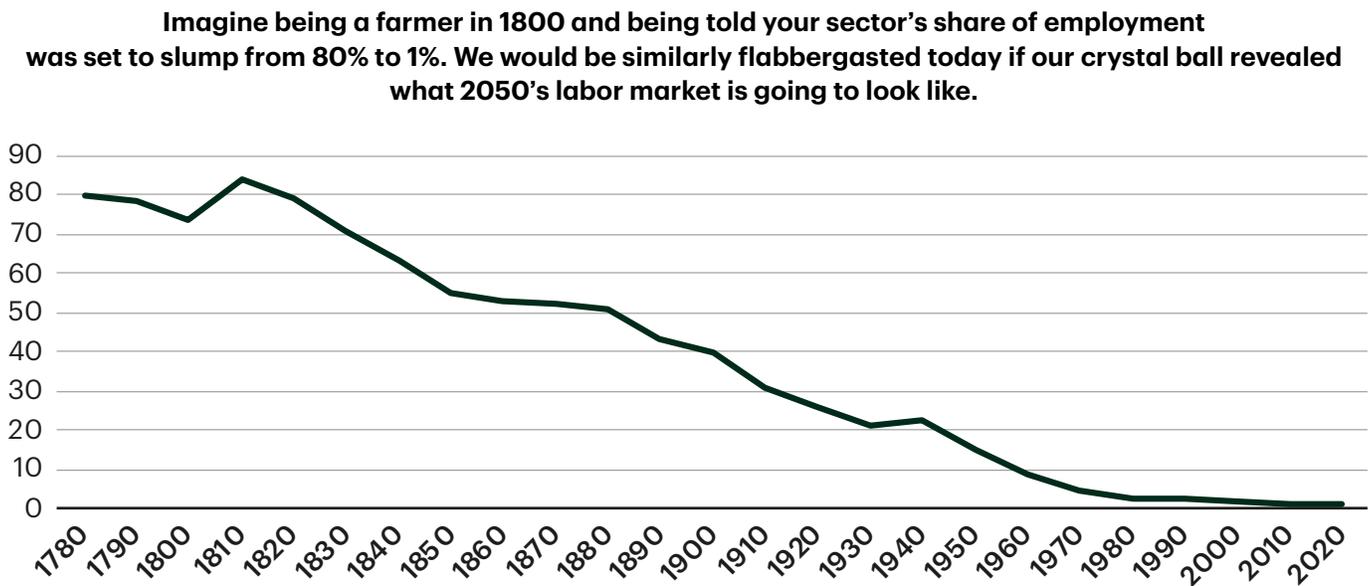
⁸ See Gavin Baker’s “Nvidia v. Google, Scaling Laws, and the Economics of AI,” from the “Invest like the best” podcast, 2025.

AI is just the latest form of automation: A process that has been ongoing for centuries

Having surveyed the last two centuries of tech innovation, let's examine the economics of progress, which largely consists of substituting capital for labor. Such automation has occurred relatively steadily since at least 1780, with AI representing the latest fashion.

Farmers once threshed grain by hand; now a single combine harvester replaces dozens of workers and agriculture productivity has soared (**Figure 4**). Similarly, telephone operators, typists, and travel agents were once ubiquitous; today, software handles most of these tasks. In auto plants, welding and painting have moved from human workers to industrial robots. More topically, LLMs are increasingly able to write code to replace some of the tasks of software engineers. Further, AI agents are progressively replacing customer service representatives (CSRs) and maybe, one day, even radiologists.

Figure 4: Agriculture's share of U.S. employment (%). It's not because we're eating less.



Source: Bureau of the Census, Historical Statistics of the U.S. As of 01/31/2020

We are richer than our grandparents because of automation: But it's a slow process

Agriculture illustrates that most wealth creation over the centuries has come from shifting the production of tasks away from slowly-improving human labor to rapidly-improving machines. Alternatively put, productivity growth has been driven by improvements in the efficiency with which capital performs tasks across a variety of sectors. Since 1950, capital productivity growth has averaged 5% per year, which is roughly ten times the productivity growth rate for labor.

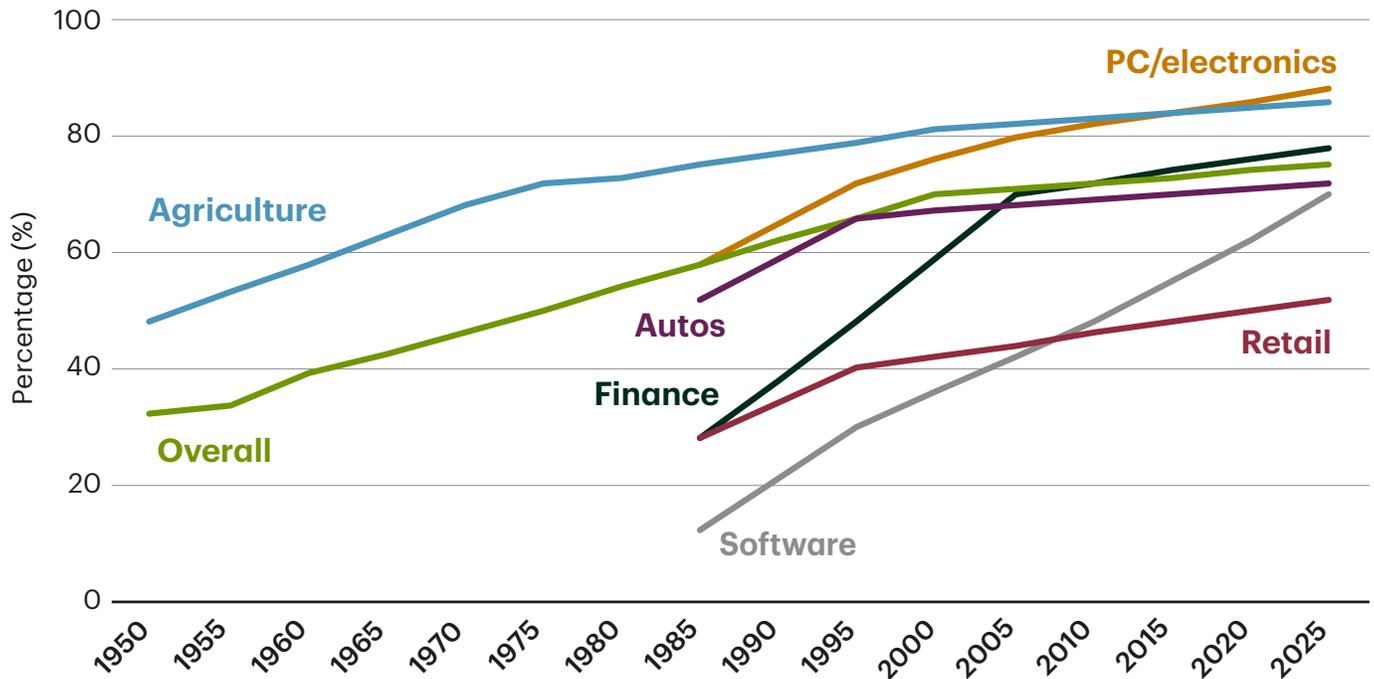
A recent study examined the automation process for six sectors and the aggregate U.S. economy (**Figure 5**).⁹ It found that the pace of substitution is slow, which can be attributed to various bottlenecks

or weak links. To economists, the pace is measured by the elasticity of substitution, with a value of 1 implying it is rapid (capital and labor are close substitutes, so a 10% decrease in the cost of capital results in companies reducing the labor input by 10%). Empirically though, the elasticity is around 0.5, meaning substitution is relatively slow. That is, capital and labor are complements, with the labor input reduced by only 5% as substantial labor input is required to realize the productivity gains from the new and improved capital.

⁹ "Past Automation and Future AI: How Weak Links Tame the Growth Explosion," by C. Jones (Stanford) et al, 2026.

Figure 5: Share of tasks (%) that have been automated since 1950.

Automation proceeds slowly and was happening way before AI. Currently the pace of automation is highest in software, electronics and finance, but slowest in retail, autos and agriculture.



Source: "Past Automation and Future AI: How Weak Links Tame the Growth Explosion," by C. Jones (Stanford) et al, 2026. As of 12/31/2025.

Bottlenecks or weak links explain why the diffusion of new technologies takes so much time. This is true even when the study’s model, calibrated on the historical elasticity of substitution, is primed for big AI-induced benefits: By 2040, output is only 4% higher than the base-case (0.3 ppts per year), and by 2060 the gain is still only 19% (0.6 ppts per year). That is,

even when advances in AI mean capital is improving rapidly, and many tasks can be increasingly automated, output gains are still constrained by the gradual pace of substitution away from slowly-improving labor. Whew, that was tough, let’s move onto something a little easier.

The march of nines: The long road from a cool demo to five-nines reliability ¹⁰

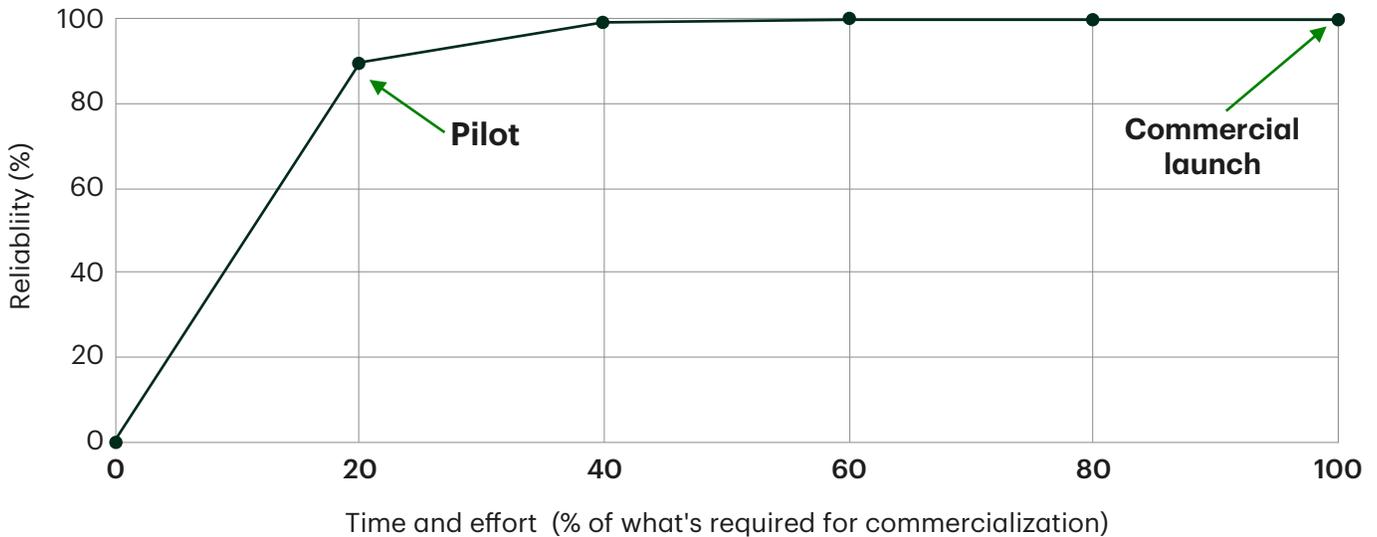
Having discussed the perspective of economists on tech diffusion, we now turn to more of an engineer’s viewpoint. It’s not enough if an AI agent books the right flight or makes the correct hotel reservation 90% of the time. With a 10% error rate, it’s more work to verify and fix. Even though today’s agentic models possess ever more sophisticated tool-calling capabilities, we’re still years away from autonomous systems that can reliably book your next vacation to Kyoto or plan your wedding in Spain.

Ninety percent reliability might be acceptable for marketing copy or routine emails, provided the cost of failure is comparatively low. However, for many purposes, the standard that society sets is much higher, 99.999% (**Figure 6**). Autonomous vehicles (AVs) may already be safer than humans, but the long tail of errors represents complex “edge cases” that challenge the best engineers. For AVs, each successive “nine” of reliability requires a constant amount of work, meaning the slog to near-perfection takes five times longer than the initial breakthrough.

¹⁰ See the podcast, “Andrej Karpathy — AGI is still a decade away,” 2025. “March of nines” is similar to the systems perspective, applicable to many domains, that “linear progress requires exponential resources.”

Figure 6: March of nines – 90% reliability (pilot launch) gets you 20% of the way to a commercial product that is 99.999% effective.

This principal is true of AVs and drones, industrial and household robots, radiology and health care applications, legal and financial services, cloud applications, and nuclear power plants. The world is fat-tailed, and mistakes can be extremely costly.



Source: TD Global Investment Solutions (TDGIS)

LLMs are a “dead end”: World modeling

The previous two sections explained why economists and engineers expect relatively slow diffusion of AI. We now turn to three leading AI researchers who, while extremely bullish on the medium- to long-term, agree that current AI architecture is a “dead end” for achieving true AGI.¹¹ Fei-Fei Li emphasizes that the next frontier of AI requires systems that can see, reason, and act within the 3D physical world. In her view, world modeling is about spatial intelligence. Language is one component, but a world model must also include visual awareness, physical actions, and the natural laws of the world.

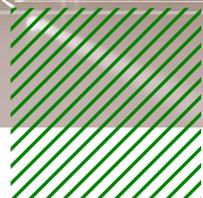
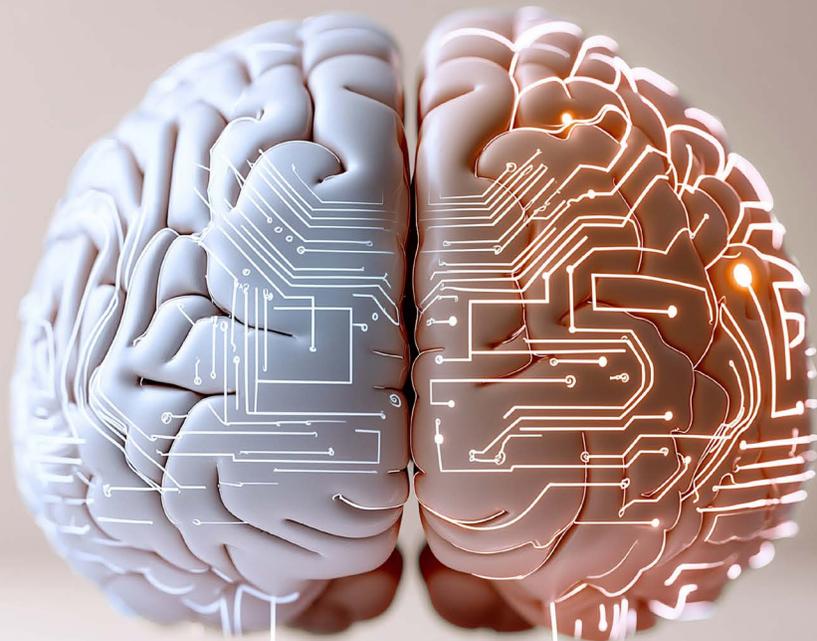
One challenge in creating a world model is obtaining data. Training LLMs is relatively easy because language data is all over the internet. Video data is plentiful, and helpful, but the real world is multimodal. It’s spatial, has fundamental 3D information, including geometry, physics, and dynamics. This type of data is not easily obtainable, but is critical to the success of robots, drones, and AVs. This suggests a long journey to commercialize physical AI.

Demis Hassabis stresses that LLMs can generate brilliant text, images, and even code, but AI lacks a world model. It just predicts patterns, it doesn’t understand. Real insights require an internal model of how the world works, continuous learning, selective memory, stronger reasoning capabilities and long-term planning.

Yann LeCun is especially outspoken, arguing “the AI industry is completely LLM-pilled,” with everyone in Silicon Valley working on the same thing.¹² The key to progress is building machines that understand the physical world as well as a cat or toddler. A world model must be a representation of how reality actually works. Just as a baby learns by watching a ball fall, AI should learn from sensory inputs like video to understand gravity, causality, and object permanence. This is an extremely high bar and raises an almost philosophical issue on the nature of intelligence, which we’ll now discuss.

¹¹ “From Words to Worlds: Spatial Intelligence is AI’s Next Frontier” by Fei-Fei Li, Stanford, 2025, as well as “The Future of intelligence” by Demis Hassabis, CEO of Google DeepMind, 2025 and “The information bottleneck” by Yann LeCun, NYU, 2025.

¹² “An A.I. Pioneer Warns the Tech ‘Herd’ is Marching Into a Dead End,” by C. Metz, NYT, 2026.



Moravec’s paradox: Hard problems are easy, and the easy problems are hard

Hans Moravec wrote in 1988 “It is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility.”¹³

What’s easy for humans is hard for computers, and vice versa. That is, tasks that humans find difficult—such as playing chess, solving calculus problems, writing code, or analyzing vast amounts of financial data—are comparatively easy for computers to master. However, robots find it exceedingly difficult to perform tasks that a toddler finds easy, such as walking across a cluttered room or picking up a glass of water. Not to mention adult tasks such as folding laundry, unloading a dishwasher, or making a sandwich.¹⁴

Moravec believed the reason lies in evolutionary history. Survival skills like vision, movement, and physical coordination have been “optimized” by natural selection over millions of years. The

underlying neural machinery is vast, complex, and highly efficient, even if we aren’t consciously aware of the effort involved. On the other hand, tasks like chess, calculus, abstract reasoning, coding, and protein folding are relatively recent. Being “new” to our brains, they feel difficult, but they actually require much simpler logical structures that are easy to encode into software.

Why does this matter? It seems natural to assume that physical AI such as AVs, drones and robotics will happen quickly, because they only have to master tasks that appear simple to us, things we learn quickly as children. However, Moravec suggests this will take considerably longer than consensus believes. Tech optimists expect AGI to be achieved by 2030 or so, while most serious AI researchers believe 2040 is more likely. Regardless, there is a wide band around these estimates, illustrating the diversity of opinions, even regarding what AGI means.

¹³ “Mind Children,” by Hans Moravec, CMU, 1988.

¹⁴ “A Gap In AI Adoption? Moravec And The AI Productivity Paradox,” by A. Susarla (Michigan State), 2026, argues that Moravec’s paradox explains why the C-suite, which deals more with analytical issues that AI excels at, is constantly disappointed with AI launches that fail to produce firm-wide results.

Why AI is so slow: Radiology, AVs, and robotics as examples

The previous section was quite abstract, so we now turn to a more practical discussion, three concrete examples where AI progress has taken decades. In 2016, Geoffrey Hinton warned that all radiologists would be replaced by AI in 5 to 10 years.¹⁵ Not only didn't that happen, but the number of radiologists in the U.S. has increased by 11% and there is a growing shortage. Why? First, AI isn't perfect at reading scans (still far from the "Five Nines" of reliability). Second, that is only one of a radiologist's tasks. The job's broader purpose is to diagnose disease and help patients. And on these dimensions, AI is unlikely to replace radiologists for decades.¹⁶

The second example concerns AVs, which represent an overnight success that was 40 years in the making.¹⁷ Tremendous strides have been made since 1985, and it is likely that AVs will represent a majority of miles driven in America by 2040-2050. However, it has been a long road involving numerous breakthroughs and still requires complementary innovations such as spatial intelligence (cameras and light detection and ranging (LiDAR)), solid-state batteries, and semantic connectivity (to talk to other cars and traffic signals). Further, road safety nirvana requires better training on edge cases (such as heavy snow, construction sites, and power outages), improved reasoning capabilities, and a more transparent and consistent regulatory environment.¹⁸

Our third example, Robotics, may well become the biggest industry ever but is unlikely to develop nearly as rapidly as boosters suggest. We would definitely take the over on Elon Musk's pledge that robotics will make work optional and money irrelevant in 10 to 20 years.

The robotics industry has been developing for decades, from 1961 when General Motors (GM) installed the first industrial robot (for die-casting) and 1977 when NASA designed, manufactured, and deployed the Sojourner rover on Mars. However, the rollout of industrial robots has been disappointingly slow, with few employed outside of auto plants and warehouses. Further, general-purpose humanoid robots remain in the research stage and are years away from mass adoption. Morgan Stanley forecasts little growth over the next decade but expects 134 million units to be sold in 2040 (valued at \$1.2 trillion), soaring to 1,019 million units in 2050 (\$4.7 trillion).¹⁹

The "march of nines" is especially daunting for robotics, which requires a host of complementary innovations including: multimodal world models, high-torque actuators (the "muscles" of the robot—represent 60% of a humanoid's material cost), tactile sensing (digital skin), spatial intelligence (cameras, LiDAR), solid-state batteries (higher energy density, better thermal stability), and semantic connectivity (to talk to smart elevators, security systems, and other robots in real-time). The robots are coming, just not on the time scale Elon Musk is promising.

¹⁵ Hinton is often hailed as the "Godfather of AI," won the Nobel prize in Physics (2024) and teaches at U Toronto. See "Radiology resident thumbs nose at Nobel Prize winner who predicted specialty would become obsolete," *Radiology Business*, 2024 and "The Growing Nationwide Radiologist Shortage," by S. Mirak et al, *Radiological Society of North America*, 2025.

¹⁶ For Jensen Huang's discussion of the "task" vs "purpose" perspective see "NVIDIA's Jensen Huang on Reasoning Models, Robotics, and Refuting the "AI Bubble" Narrative," *No Priors* podcast, 2026.

¹⁷ Five milestones on the AV journey:

1985: CMU's AV, funded by DARPA, achieved the first road-following demonstration using Lidar and computer vision, reaching a speed of 31 km/h.

1995: The "No Hands Across America" tour. CMU's AV completed the first autonomous U.S. coast-to-coast journey from Pittsburgh to San Diego. The vehicle was driven autonomously for 98.2% of the trip at an average speed of 102.7 km/h.

2005: Sebastian Thrun (widely known as the father of the AV) led the Stanford Racing Team to victory in the DARPA Grand Challenge. It was the first robot car to complete the 132-mile desert course, doing so in under seven hours and winning a \$2 million prize.

2018: Waymo launches its robotaxi service in Phoenix, AZ.

2025: Waymo's robotaxis completed 14 mn trips, driving 150+ mn miles, up 3x from 2024.

¹⁸ "We Don't Know If AVs Are Safer Than Human Drivers," by David Zipper (MIT), *Bloomberg*, 2026.

¹⁹ Humanoids: Investment implications of embodied AI," by Adam Jones et al, *Morgan Stanley Research*, 2025.

Implications for investors

1

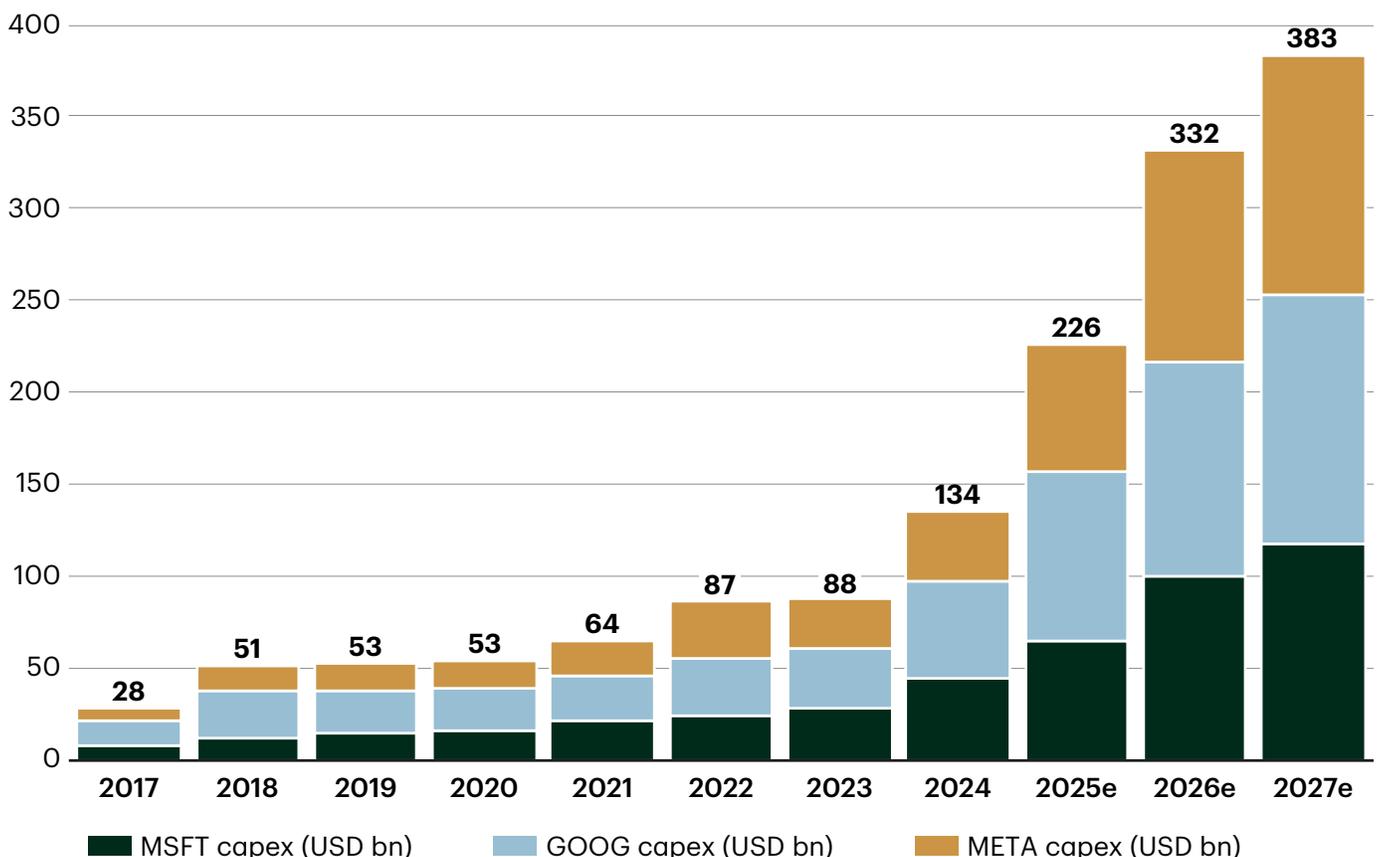
AI capex growth to moderate to a more sustainable trajectory

Numerous investors have argued there is a bubble in LLMs, emphasizing that some (mostly private) companies are spending tens of billion of dollars, are unlikely to turn FCF positive this decade, but are still able to fundraise at astronomical valuations. This is especially concerning given the limitations of LLMs discussed above (“dead end” for achieving true AGI) and the prevalence of open-source labs. It seems reasonable to expect that many of the new startups will lose money and are accidents waiting to happen. That was the pattern with railways, electricity, autos, the internet, and all other episodes of tech disruption.

Regardless, the tech cycle will endure. “AI bubble” is too binary a term, as public equity valuations are only moderately stretched and, at the epicenter of AI capex, are the world’s most profitable companies. This is the critical difference from previous tech cycles, including the 1990s. Still, what happens if – as occurred with railways in the 19th century – investors recognize their returns won’t be realized quite as rapidly as they had forecast? We expect this to happen during the next few quarters and believe capex growth will then moderate from its current torrid pace to a more sustainable path (Figure 7).²⁰

Figure 7: Hyperscaler capex (USD bn) -- Up 6x since 2020.

We expect investors to soon realize that the productivity and profit gains from AI will arrive slower than previously anticipated. This will result in a lower and more sustainable capex growth trajectory.



Source: Bloomberg Finance, L.P. As of 01/31/2026

²⁰ Even if AI capex growth in the U.S. moderates it won't turn negative. This is because AI is the key domain in which the U.S.-China contest for global supremacy is being waged. More specifically, AI is core to all three domains of power – economic strength, tech prowess, and defense capabilities.

2

Emphasize quality tech

By quality we mean companies generating sustainable FCF, with solid margins, and a return on invested capital (ROIC) that is greater than its weighted average cost of capital (WACC). This includes the hyperscalers which have real businesses, with large existing user bases. Aside from defensive characteristics, having lots of customers unlocks the flywheel that is at the center of every great tech company. The users make the algorithm better, which improves the product, attracts more customers, and the machine just spins. It's not quite spinning yet in AI, but you can squint and see it.²¹ Additionally, the infrastructure layer is usually a safer place to be at this stage of a tech revolution. There will be tons of opportunities in applications, but it is still too speculative with enormous unknowns.

3

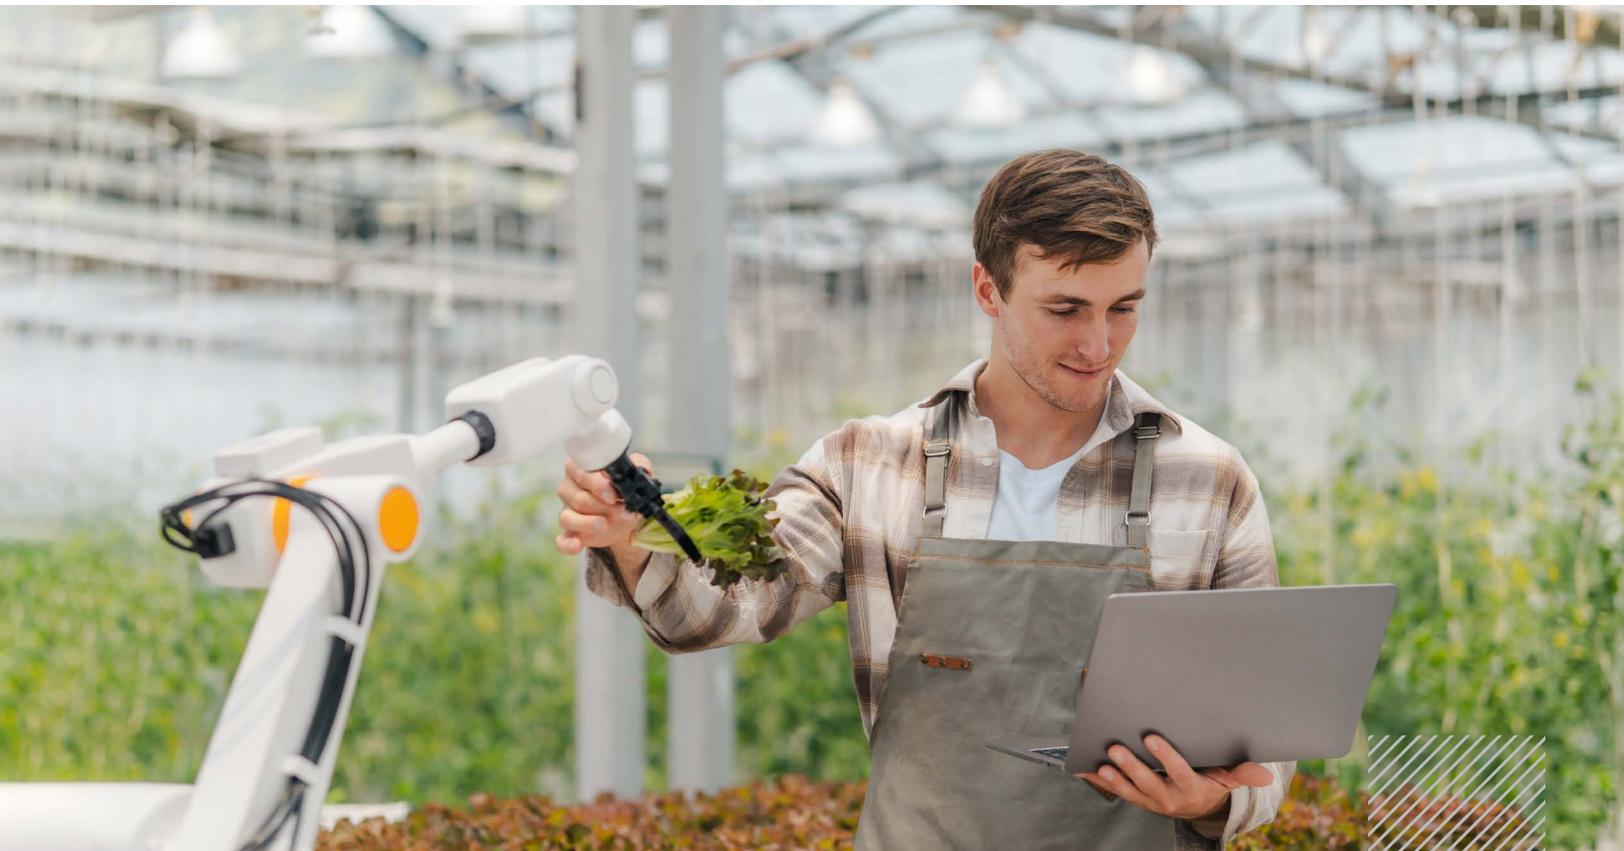
Diversify beyond the U.S.

We remain constructive on U.S. equities, but many investors hold overly concentrated portfolios in U.S. tech. Further, the U.S. valuation premium often doesn't appear justified by fundamentals, and we believe the USD is over-valued.²² We especially like global champions outside the U.S., including China, which meet our definition of quality.

4

Creative destruction and the innovator's dilemma: Turbo-charged by AI

Tech revolutions are always associated with massive churn in corporate leadership. To illustrate, of the top twenty-five global tech companies in 2000, only four of them are still members of that list (**Figure 8**). What will this list look like in 2040? This is difficult to imagine as today's champions (NVDA, AAPL, MSFT, AMZN, GOOG) always appear bulletproof and insurmountable. Regardless, history strongly suggests that a majority of firms on today's leaderboard will fall off within the next decade or two.



²¹ These themes are explored in "Gavin Baker and David George on Positional Strategy in AI," a16z, 2025.

²² See "The Dollar is Our Currency, but It's Your Problem: Rebalancing Requires a Much Weaker USD," 2025.

Figure 8: Largest global tech companies.

Only four of the top twenty-five in 2000 are still there today. Expect a vastly different list in 2040.

Rank	Dec-92	Jan-95	Jan-00	Jan-05	Jan-10	Jan-15	Jan-20	Jan-25	Jan-26
1	IBM	Escom	MSFT	MSFT	MSFT	AAPL	MSFT	AAPL	NVDA
2	MSFT	IBM	NTT Docomo	Vodafone	AAPL	MSFT	AAPL	NVDA	APPL
3	Hitachi	MSFT	CSCO	IBM	IBM	Intel	AMZN	MSFT	MSFT
4	Panasonic	Panasonic	Intel	Intel	AT&T	AT&T	Meta	AMZN	AMZN
5	Intel	Hitachi	Nokia	CSCO	GOOG	Meta	GOOG	META	GOOG
6	HP	Intel	IBM	Dell	CSCO	IBM	Visa	TSLA	AVGO
7	Alcatel-Lucent	HP	ORCL	Telefonica	Telefonica	GOOG	AT&T	GOOG	META
8	Toshiba	Toshiba	Vodafone	AT&T	Vodafone	ORCL	Tencent	AVGO	TSLA
9	EDS	Sony	Nortel	DTE	HP	CSCO	Alibaba	TSMC	TSMC
10	Sony	Sharp	Dell	Nokia	Intel	Samsung	Intel	Visa	Visa
11	Emerson	EDS	Ericsson	QCOM	ORCL	Visa	Mastercard	MA	Tencent
12	Nintendo	Fujitsu	Sony	HP	QCOM	Comcast	TSMC	NFLX	MA
13	Nortel	NEC	HP	eBay	Samsung	QCOM	Samsung	Tencent	Samsung
14	Sharp	KDDI	Yahoo	Viacom	Canon	AMZN	Comcast	ASML	ASML
15	Novell	Vodafone US	QCOM	ORCL	China Mobile	TSMC	CSCO	CRM	NFLX
16	NEC	Emerson	EMC	Orange	Orange	Vodafone	ADBE	ORCL	PLTR
17	Mitsubishi Elec	Kyocera	Softbank	BellSouth	TSMC	Mastercard	NFLX	SAP	AMD
18	Fujitsu	Canon	Motorola	Canon	Nokia	HP	Salesforce	NOW	MU
19	Xerox	ORCL	Fujitsu	Telstra	AMZN	Tencent	NVDA	CSCO	Alibaba
20	Corning	Alcatel-Lucent	PSI Software	Ericsson	DTE	Telefonica	SAP	AMD	ORCL
21	AAPL	Ericsson	AOL	Yahoo!	Visa	SAP	PYPL	IBM	IBM
22	Sega	Xerox	TXN	Samsung	SAP	China Mobile	ASML	ADBE	CSCO
23	Kyocera	Nokia	China Mobile	TXN	Hon Hai	eBay	TXN	Samsung	SAP
24	Vodafone	Compaq	NTT Data	Motorola	Telstra	Baidu	IBM	QCOM	SK Hynix
25	Ericsson	Vodafone LN	Hikari Tsushin	Comcast	Blackberry	EMC	ORCL	TXN	LRCX

Source: Bloomberg, TD Epoch. As of 01/01/2026

Appendix A: Four takeaways from the last two centuries of tech innovation

An appreciation of previous tech waves is critical to understanding how the AI cycle is likely to play out. This appendix discusses four lessons for investors.

First, achieving sustainable commercial success takes years, sometimes decades, and is much more difficult and uncertain than the initial technological triumph. Even the smartest people failed to understand how their big breakthroughs were going to evolve. James Watt invented his steam engine to pump water out from coal mines. Guglielmo Marconi expected his radio would be used for one-to-one interpersonal communication, competing with the telephone.

A few quotes illustrate misperceptions that existed at the time (today's will appear equally ludicrous a generation from now) and the danger in reading history backwards:

- “What use could this company make of an electrical toy?” The CEO of Western Union, William Orton, turning down the offer of Bell’s telephone patents for \$100,000, 1876
- “I think there is a market for about five computers.” Thomas J. Watson Sr, chair IBM, 1943

Second, most innovators struggled for decades, even the lucky few who became fabulously rich. This includes Edison with the incandescent lamp, Bell with the telephone and Marconi with wireless technology (radio). The long road to profitability often resulted in a shortage of capital, multiple bankruptcies, and pressure to sell patents at bargain basement prices. To illustrate, Bell was cash strapped in 1876 and offered to sell his telephone patent to Western Union (the dominant telegraph company) for \$100,000 (\$3.1 million in today’s dollars).

This explains why new technologies and overpromotion go hand in hand. Innovations are greeted with skepticism, but they require risk capital to develop and commercialize. So, pioneers need to project excess confidence and avoid any hint of failure or setback. For example, Edison was well known for making extravagant claims about his discoveries, motivated by a constant need to raise capital. Innovators must always appear bullish, and this behaviour can encourage bubbles (as has occurred repeatedly -- railway companies in the 1850s and 1880s, radio shares during the 1920s, and the internet bubble of the 1990s).

Next, it takes a long time for technologies to evolve from impressive demos to widespread adoption.

James Watt’s steam engine was invented in 1769, but the first steam locomotive demo didn’t occur until 1801, with the initial commercial railway not until 1826. The next few paragraphs provide five additional examples from subsequent tech waves.

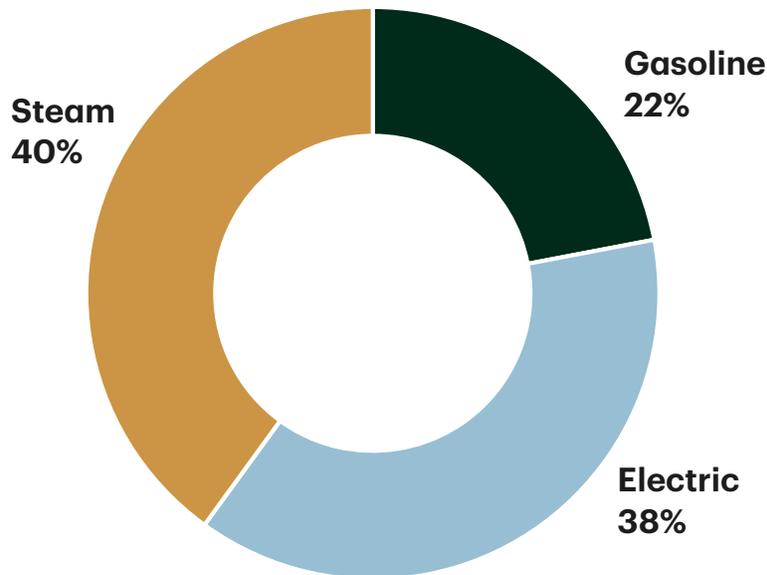
Electricity: Faraday discovered the foundations of modern electricity in 1831, and Joseph Swan demonstrated the principle of an incandescent carbon filament lamp in 1848. However, Edison didn’t receive his historic patent for the incandescent light bulb until 1880, GE didn’t turn a profit until 1895, and it wasn’t until 1924 that a majority of U.S. homes had electricity.

Automobile: In 1860, Lenoir created the first functional, gas-fired engine and in 1876 Nikolaus Otto unveiled the first efficient four-stroke engine. A decade later Karl Benz was awarded a patent for the first practical car that integrated an engine, chassis, and drive system. Still, most cars were powered by steam or electricity in 1900 (**Figure 9**). Henry Ford, after having twice gone bankrupt, introduced the Model T in 1908 but it wasn’t until 1924 that more than 50% of U.S. households owned a car.

Appendix

Figure 9: Automobile technologies in 1900 (U.S. market share, %).

Gasoline didn't become #1 until 1912 after the launch of the mass-produced model T (1908), the appearance of the electric self-starter (1912) and a major oil discovery in Texas (1901). This illustrates two points: the winning technology is rarely clear during the early days of a tech wave; and multiple, complementary innovations are always required to achieve commercial success.



Source: "America adopts the automobile, 1895-1910," by J. Flink, UC Irvine, 1970

Computer: The transistor was invented in 1947 at Bell Labs, while Intel was founded in 1968, Microsoft in 1975, and Apple in 1976. The first killer apps were released shortly thereafter -- VisiCalc in 1979, Lotus 1-2-3 and Word in 1983. In 1984 only 8% of U.S. households owned a computer, a share that didn't hit 50% until 2000.

Internet: The first computer-to-computer link was established between UCLA and Stanford in 1969. The internet's first killer app, Email, arrived in 1971 and the network spread gradually, to include sixty-two academic and military/industrial sites by 1979. A standard internet protocol (TCP/IP) arrived in 1983 and, in the following year, Cisco Systems emerged from Stanford and became the primary provider of networking equipment, such as switches and routers.

The next key development occurred in 1989, with the invention of the World Wide Web, and in 1993 with the release of the Mosaic browser, which sparked the "internet boom." Key company launches included: eBay and Amazon in 1995, Google 1998, Facebook 2004, Spotify 2006, Netflix's streaming service and Airbnb in 2007, and Uber 2009.²³

AI: Its official birth was a conference at Dartmouth in 1956. The first artificial neural network was developed two years later by Cornell's Frank Rosenblatt and the first chatbot was created in 1966 by MIT's Joseph Weizenbaum (it was designed to simulate a psychotherapist). After a long AI winter, it was back above the fold in 1997 when IBM's Deep Blue defeated Garry Kasparov. Next came the AlexNet breakthrough in 2012 when a group from the University of Toronto won an image recognition competition, triggering the modern deep learning boom.

A few years later Google DeepMind's AlphaGo defeated Lee Sedol, proving that AI could use intuition and strategic thinking rather than just brute force. The following year, 2017, Google researchers introduced the transformer architecture in a paper "Attention Is All You Need," with this model becoming the bedrock for modern LLMs. Next, in late-2022, OpenAI released GPT-3.5, bringing generative artificial intelligence (GAI) to the mainstream with the fastest-growing consumer application in history.

²³ An even longer list would include rising stars that burst, including Netscape, AOL, AskJeeves, Pets.com, Webvan, WorldCom, Friendster, and Myspace.

Fourth, winner-takes-most is not a recent development, with historical episodes initially featuring just one or two successful companies. This includes Western Union in telegraph, Bell's AT&T in telephones, Edison's GE in lighting, Ford and GM in autos, IBM for computers, Microsoft with software, Intel in semis, and Google for search. Yet the eventual success of these companies was far from a foregone conclusion: for every successful disruptor, there were dozens of others who tried and failed.

Over any extended period of time though, fierce competitors and new technologies arrive, so monopoly-like returns eventually dissipate. Past exceptions include Western Union (dominant from 1866 until 1929), AT&T (leader for almost a century, until its breakup in 1984) and GE (maintained a leading position long after 1880). This often reflects a first-mover advantage.

However, this advantage appears more fragile in the age of digital tech, as seen with displaced first movers such as: Commodore (PCs), Netscape (browsers), AOL ("walled garden"), Yahoo! (search), Webvan (grocery delivery), Napster (music), Friendster (social media), and Nokia (smartphones). The digital economy features a faster pace of innovation, and your competitor's product is just one click away. This suggests today's tech behemoths are more vulnerable than consensus recognizes.

Appendix B: Complementary innovations – It takes a village

Edison famously said, “Genius is 1% inspiration, 99% perspiration.” That’s true, but it wasn’t just his perspiration. A major reason it takes so long for technologies to evolve from impressive demos to widespread commercial adoption is that the journey requires not just one major breakthrough but dozens of complementary innovations. This section provides a partial list for six tech waves to illustrate their importance:

Rail: High pressure steam engines, advanced boilers, bituminous coal, Bessemer steel, steel rails, standard gauge, air brakes, standard time zones, automatic couplers, refrigerated cars.

Electricity: Modern steam turbines, industrial scale power stations, transformers, Alternating Current, insulated wiring, electric meters, electric motor, two-prong wall outlets, electric iron, vacuum cleaner, washing machine, fridge, toaster.

Autos: Carburetor, electric starter, differential, radiator, vanadium steel, moving assembly line, pneumatic tires, fuel injection, leaded gasoline, gas stations.

PC: Microprocessors, operating systems, graphical user interface, the Mouse, killer apps (VisiCalc), hard drives, modems, gaming software, GPUs and eventually the internet.

Internet: Standard internet protocols, packet switching, world wide web, graphical browsers, fiber optic cables, routers and switches, search engines, smart phones, payment systems, advertising-based business models, shopping, social media, music and video streaming, ride sharing, cloud services.

Robots: Multimodal world models, continuous learning, stronger reasoning and long-term planning capabilities, high-torque actuators (the “muscles” of the robot—represent 60% of a humanoid’s material cost), tactile sensing (digital skin), spatial intelligence (cameras, LiDAR), solid-state batteries (higher energy density, better thermal stability), semantic connectivity (to talk to smart elevators, security systems, and other robots in real-time).

Learn more about TD Global Investment Solutions



www.tdgis.com



The information contained herein is for information purposes only. The information has been drawn from sources believed to be reliable. Graphs and charts are used for illustrative purposes only and do not reflect future values or future performance of any investment. The information does not provide financial, legal, tax or investment advice. Particular investment, tax or trading strategies should be evaluated relative to each individual’s objectives and risk tolerance.

This material is not an offer to any person in any jurisdiction where unlawful or unauthorized. These materials have not been reviewed by and are not registered with any securities or other regulatory authority in jurisdictions where we operate.

Any general discussion or opinions contained within these materials regarding securities or market conditions represent our view or the view of the source cited. Unless otherwise indicated, such view is as of the date noted and is subject to change. Information about the portfolio holdings, asset allocation or diversification is historical and is subject to change.

This document may contain forward-looking statements (“FLS”). FLS reflect current expectations and projections about future events and/or outcomes based on data currently available. Such expectations and projections may be incorrect in the future as events which were not anticipated or considered in their formulation may occur and lead to results that differ materially from those expressed or implied. FLS are not guarantees of future performance and reliance on FLS should be avoided.

No investment strategy or risk management technique can guarantee returns or eliminate risk in any market environment. Any characteristics, guidelines, constraints or other information provided for this material is representative of the investment strategy and is provided for illustrative purposes only. They may change at any time and may differ for a specific account. Each client account is individually managed; actual holdings

will vary for each client and there is no guarantee that a particular client's account will have the same characteristics as described herein. Any information about the holdings, asset allocation, or sector diversification is historical and is not an indication of future performance or any future portfolio composition, which will vary. Portfolio holdings are representative of the strategy, are subject to change at any time and are not a recommendation to buy or sell a security. The securities identified and described do not represent all of the securities purchased, sold or recommended for the portfolio. It should not be assumed that an investment in these securities or sectors was or will be profitable.

The securities disclosed may or may not be a current investment in any strategy. Any reference to a specific company listed herein does not constitute a recommendation to buy, sell or hold securities of such company, nor does it constitute a recommendation to invest directly in any such company. Any projections, targets, or estimates in this presentation are forward-looking statements and are based on our internal research, analysis, and assumptions. There can be no assurances that such projections, targets, or estimates will occur and the actual results may be materially different. Additional information about our assumptions is available upon request. Other events which were not taken into account in formulating such projections, targets, or estimates may occur and may significantly affect the returns or performance.

For Jurisdictions Outside of Canada and the United States: This information is only intended for use in jurisdictions where its distribution or availability is consistent with local laws or regulations. TD Asset Management Inc. and Epoch Investment Partners, Inc. products are also available through a network of affiliated and unaffiliated distributors. Please contact our distribution partners to find out more.

Australia: Epoch Investment Partners, Inc. (ARBN: 636409320) holds an Australian Financial Services License (AFS License No: 5308587) and is incorporated in Delaware, USA (liability of members is limited). To the extent any statements contained in this document constitute financial product advice, those statements are issued by Epoch Investment Partners, Inc. The information contained herein is intended for wholesale clients and investors only as defined in the Corporations Act of 2001.

Japan: For Japanese residents, please note that if you have received this document from The Toronto-Dominion Bank entities based outside Japan, it is being provided to qualified financial institutions ("QFI") only under a relevant exemption to the Financial Instruments and Exchange Act.

If you have received this document from TD Securities (Japan) Co., Ltd., it is being provided only to institutional investors. TD Securities (Japan) Co., Ltd. is regulated by the Financial Services Agency of Japan and is distributing this document in Japan as a Type 1 and Type 2 Financial Instruments Business Operator registered with the Kanto Local Finance Bureau under registration number, Kinsho 2992, and a member of Japan Securities Dealers Association.

South Africa: Epoch Investment Partners, Inc. is a licensed Financial Services Provider (license number 46621) with the Financial Sector Conduct Authority. TD Global Investment Solutions represents TD Asset Management Inc. ("TDAM") and Epoch Investment Partners, Inc. ("TD Epoch"). TDAM and TD Epoch are affiliates and wholly owned subsidiaries of The Toronto-Dominion Bank.

New Zealand: This document is not an offer of financial products requiring disclosure under the Financial Markets Conduct Act 2013 (New Zealand). This document is only for "wholesale investors" within the meaning of clause 3(2) of Schedule 1 of the Act and any offer of financial products will only be made to "wholesale investors". Neither TDAM, TD Epoch, TD Global Investment Solutions nor The Toronto-Dominion Bank are a registered bank in New Zealand.

Additional Disclosure

On September 10, 2025, the Chicago Mercantile Exchange ("CME") made findings involving The Toronto-Dominion Bank ("TD") related to a transaction effected on the Chicago Board of Trade ("CBOT") in June 2023. The CME determined that the transaction was executed in order to transfer a futures position between accounts with common beneficial ownership in violation of CBOT Rules 534 and 538.C. Without admitting or denying any violations or factual findings, TD consented to the entry of findings by the CME and paid a penalty of \$25,000. None of the conduct regarding this matter involved Epoch Investment Partners, Inc. or its employees. A copy of the CME findings can be found here: <https://www.cmegroup.com/content/dam/cmegroup/notices/disciplinary/2025/09/cbot-24-1744-bc-toronto-dominion-bank-nda-posting.pdf>

TD Global Investment Solutions represents TD Asset Management Inc. ("TDAM") and Epoch Investment Partners, Inc. ("TD Epoch"). TDAM operates in Canada and TD Epoch operates in the U.S. Both entities are affiliates and wholly owned subsidiaries of The Toronto-Dominion Bank.

®The TD logo and other TD trademarks are the property of The Toronto-Dominion Bank or its subsidiaries.