The AI Value Levers: How Innovation-Focused Strategies Outperform—for Firms and Workers

David Schatsky
Technology and Geopolitics Fellow
dschatsky@tenonresearch.com

JUNE 2026

HARVARD Kennedy School
BELFER CENTER
for Science and International Affairs

Introduction

Firms are investing in generative AI at a pace unprecedented in the modern history of information technology (see sidebar). Generative AI, or gen AI, is a branch of artificial intelligence focused on training models—particularly large language models (LLMs)—to generate new content, such as text, code, or data, in response to user prompts.

Predictions about the potential impact of gen AI are proliferating. Among those who focus on its economic consequences, some see the dawn of an era of rapid economic growth and unparalleled material abundance. Others warn of mass unemployment and social upheaval. The aim here is not to offer a prediction. Instead, this paper seeks to understand how companies are using the technology today and how they might amplify its benefits and reduce its potential harms to workers.

There are many ways a company might create value with gen AI. Often these "value levers" can be used in combination. One lever involves reducing *labor intensity*—reassigning work from people to AI. If accompanied by job cuts, this is the approach that could be most disruptive to the labor market, particularly if not combined with other value levers. And, among the companies I interviewed, reducing labor intensity was the most common value lever. My research suggests that focusing too much on reducing labor may be a mistake. It fails to exploit the full potential of the technology. It tends to point companies toward cutting costs, when the bigger opportunity may be innovation-driven growth.

Prior academic research, discussed below, has found that companies that apply technology like AI toward innovation tend to generate stronger returns than those that focus on cost reduction. This suggests that firms' generative AI strategies should aim to engage multiple value levers rather than focusing solely on reducing labor intensity and labor costs. This path can be better for both workers and firms.

How fast are enterprises adopting generative AI?

I compare a 2025 estimate of firm spending on generative AI with spending on older enterprise software categories: software as a service and cloud infrastructure as a service (SaaS and Cloud IaaS).

The comparison is challenging because the gen AI market is less than four years old and comparable market size data for those categories at the same stage of development is not available. Instead, I look at how long it took those markets to attain the size of today's gen AI market.

According to venture capital firm Menlo Ventures, global gen AI spending by enterprises reached $37 billion in 2025, three years after the launch of ChatGPT. Gartner estimates that global spending in 2025 on AI applications, AI infrastructure software, and gen AI models was $581 billion, more than ten times 2025's figure. I use Menlo Ventures' more conservative figure for comparison with Gartner figures for SaaS and Cloud IaaS.

So far, the gen AI market is growing many times faster than the SaaS and Cloud IaaS markets did in their early years. In 2010, ten years after the launch of Salesforce, SaaS revenues were just $9.2 billion ($13.6 billion in 2025 dollars), about a third the size of today's gen AI market. In 2016, ten years after the launch of Amazon Web Services, Cloud IaaS revenues were estimated to be $25.4 billion ($34.14 billion in 2025 dollars), about the same amount of gen AI spending we are seeing after just three years.

The motivation for this research

Many studies have looked at the impact of generative AI on firms. Some are case studies of individual companies and use cases, often provided by technology vendors or consultants. Others are based on surveys whose quantitative findings can attract much attention. A widely discussed 2025 paper by a team at MIT, for example, based largely on survey data, found that 95 percent of companies that invested in generative AI got no return on their investment (Challapally et al. 2025). Other recent studies, discussed below, have reported more positive findings on the impacts of AI generally without separating out the impact of gen AI specifically.

Recent survey-based research by Boston Consulting Group (BCG), Deloitte, and McKinsey found that many companies do create value with AI. The McKinsey report, published in November 2025, found 39 percent of companies have seen positive firm-level impact on earnings before interest and taxes (EBIT), a financial metric that measures a company's core operating profitability (Singla 2025). Most respondents said the impact is modest, though, accounting for less than 5 percent of EBIT. Most are still only experimenting or piloting. And the majority report little or no reduction in cost or increase in revenue within business functions that regularly use AI. The study by BCG, published in September 2025, had similar findings: Only 40 percent of companies are achieving material value, of which only 5 percent received substantial value (Apotheker et al. 2025). In a Deloitte study, published in January 2026, 40 percent saw cost savings and 20 percent achieved revenue improvement. The report does not describe the size of the impact. As in the other studies, most respondents reported no value yet (Rowan et al. 2026).

Because the BCG, Deloitte, and McKinsey studies asked about AI generally, including traditional machine learning, they do not isolate the impacts of generative AI specifically. BCG did segment some responses by AI type but found that most of the applications respondents mentioned were based on what it called "predictive AI" rather than generative AI.

An October 2025 study published by the Wharton School and GBK Collective, a marketing strategy consultancy, did focus specifically on generative AI. It found that positive returns were widespread among its respondents: Seventy-four percent of respondents reported positive return on investment (ROI), including 35 percent who described returns as significantly positive. Larger companies—those with $2 billion or more in revenue—were less likely to report positive results. Among these firms, 57 percent reported positive ROI, including 26 percent who characterized the returns as significantly positive (Korst et al. 2025). This data warrants careful interpretation. As one of the authors of the study wrote to me in an email, "I think that the ROI data are best seen as reflecting the 'vibe.' As such, they are probably closer to a perceptual variable (reflecting people's thoughts and feelings) than to an objective financial measure."[1]

Besides these studies, several controlled experiments have measured the impact of using generative AI for a certain task or job. These are useful for understanding how gen AI might influence productivity, quality, and worker experience in narrow, usually artificial situations. Surveys are useful for taking the pulse of a phenomenon like the adoption of generative AI. All these approaches make a contribution. But they do not paint a picture of what is really happening with gen AI inside companies.

I conducted thirty interviews with executives at companies in diverse sectors including consulting, financial services, pharmaceuticals, tech, and travel. We discussed twenty-two generative AI applications executives believe have created a positive impact on their business, focusing on how the applications benefit each

1 Stefano Puntoni, Faculty Co-Director of Wharton Human-AI Research, University of Pennsylvania, personal communication, February 26, 2026.

organization and affect its workers. My aim was to understand qualitatively what firms were achieving with their use of gen AI, rather than to derive quantitative benchmarks like return on investment. Therefore, we did not discuss financial metrics such as the cost to build and operate the solutions, the costs they might have saved, or the incremental revenues they may have enabled. My analysis is based on these interviews and a review of academic and industry research. Given the relatively small number of interviews and the fact that executives were drawn from my professional network rather than selected at random, what I found may not be representative of broader gen AI adoption patterns. But it does provide a useful picture of the impact of gen AI at the level of individual firms.

How firms create value with generative AI

To better understand how gen AI is creating value at organizations, I developed a typology of gen AI value-creation levers. I use it to categorize what AI leaders told me about the benefits they were obtaining from the gen AI applications we discussed. Figure 1 shows the levers and the benefits they afford.

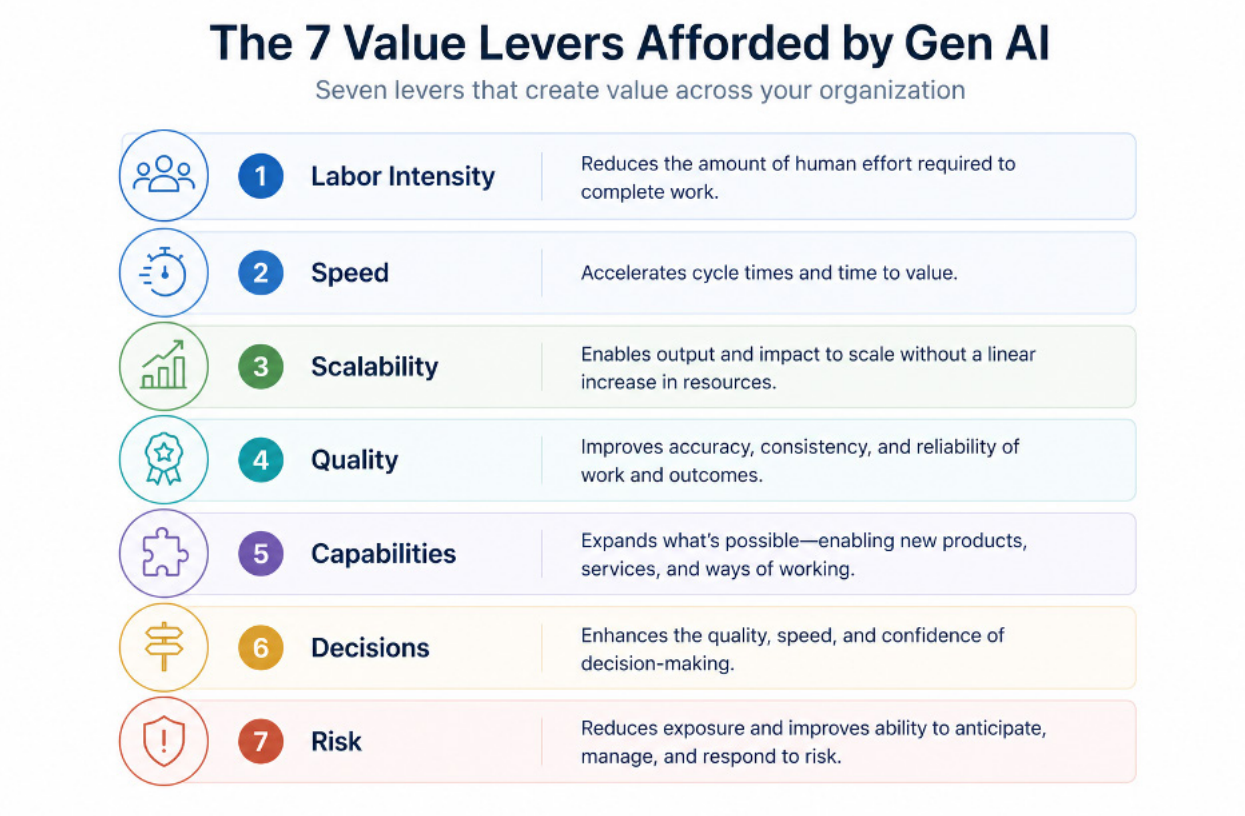

Figure 1: Gen AI Value Levers graphic

Some gen AI applications produce benefits via multiple levers. For instance, some applications that reduce labor intensity also reduce cycle time. Figure 2 illustrates ways the value levers may interact.

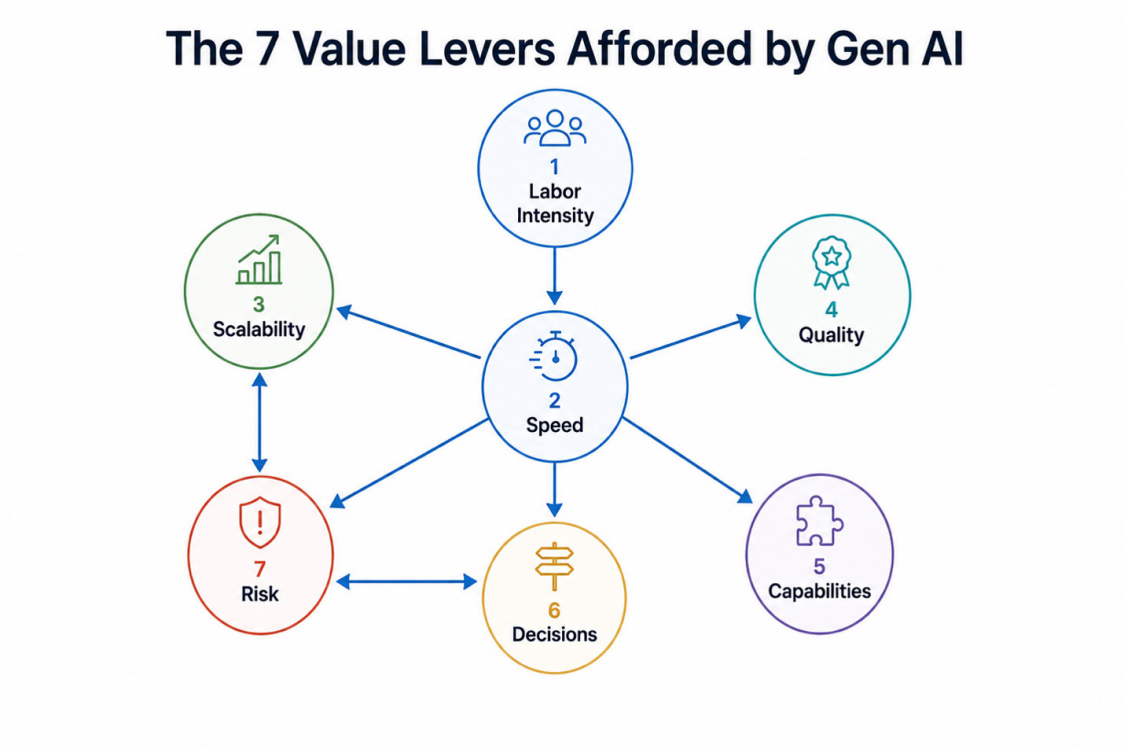

Figure 2: Ways the value levers interact

Figure 3 shows the distribution of value levers used in the gen AI use cases discussed in my interviews.

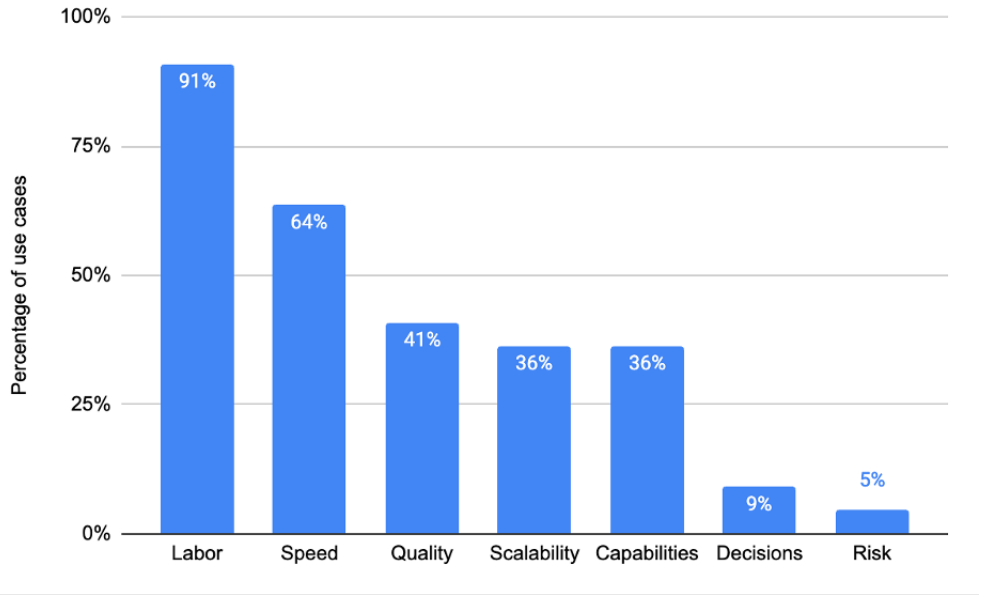

Figure 3: Distribution of value levers. n = 22. Source: Author interviews

Without a broad understanding of how gen AI can create value, business leaders risk overlooking significant opportunities. Indeed, evidence suggests that firms deploying gen AI may be overly focused on reducing labor intensity at the expense of other value levers.

Gen AI is used today mainly as a labor reducer

Most people's initial exposure to generative AI is through a chatbot or copilot. These experiences help create the impression of gen AI as a productivity tool that saves time and allows people to focus on "higher-value tasks." This, plus firms' historical propensity to justify IT investments in terms of cost savings (Gunasekaran et al. 2001; Peffers and Dos Santos 2013), may account for the fact that eliminating labor—and labor costs—appears to be firms' dominant use of gen AI today. **Ninety-one percent** of the use cases discussed in my interviews entailed automating tasks previously performed by workers—in other words, reducing labor intensity. It is by far the most common value lever.

Reducing labor intensity is not the same as reducing labor costs. It reduces the *human effort required per unit of output*. Whether that produces cost savings depends on what happens next: If headcount is reduced, costs fall; if liberated hours are redeployed to other work, the result is a productivity gain—more output from the same wage bill—rather than a cost reduction. Both are valuable, but they are different, and firms and economists tend to account for them differently.

A growing body of evidence suggests firms that adopt AI are beginning to reconfigure their workforce, including cutting staff, reducing hiring for certain skills, and increasing hiring for others. This is evident in surveys of company executives. In the McKinsey survey, depending on the department where AI was used, between 12 percent and 22 percent of respondents said they had reduced headcount as a result. And 25 to 39 percent expect headcount reductions in those departments in the coming year.

We also see evidence of AI-driven headcount reductions in corporate announcements of layoffs. Challenger, Gray and Christmas, an outplacement firm that monitors layoffs, found that U.S.-based employers announced 55,000 layoffs attributed to AI in 2025—5 percent of all layoffs they tracked that year. Most of these occurred in the tech industry, which, besides being an enthusiastic early adopter of AI—particularly for coding—may have overhired during the pandemic, according to *The Wall Street Journal*. So, this may not be a sustained trend. By Challenger's count, as of February 2026, AI has been cited in connection with 12,304 job cut announcements, or 8 percent of job cut plans, trailing reasons such as market conditions, contract losses, and restructurings (Challenger, Gray, and Christmas 2026).

It is not clear whether those cuts were a consequence of labor intensity reductions that firms have already achieved with AI or in anticipation of these reductions. The layoffs announced by Jack Dorsey, CEO of Block, in February 2026 were clearly anticipatory. In justifying the 4,000 layoffs, which reduced the size of the company by nearly half, he wrote on X: "i had two options: cut gradually over months or years as this shift plays out, or be honest about where we are and act on it now."

Companies are taking advantage of gen AI-powered automation to eliminate some jobs and slow hiring for others. These may be jobs that comprise mostly routine tasks that gen AI can do automatically. Recent research has found that job postings for roles like these have already declined significantly (Chen, Srinivasan, and Zakarinia 2025). If technology can do them as well and more cheaply, traditional business logic would dictate that those jobs be replaced by technology.

At the same time, Chen, Srinivasan, and Zakarinia found that gen AI increased both demand for and skill complexity in positions that involve human-AI collaboration. McKinsey's survey found between 5 percent and 15 percent of respondents said they had actually increased headcount in departments where AI is used. The technical skills needed to build and operate AI systems will likely be in high demand for some time as well.

While headlines and studies suggest generative AI is beginning to affect the composition of the workforce, it is difficult to discern impacts at the level of the macroeconomy. Despite the rapid growth of firm spending on generative AI, "the technology is still too nascent for the effects on workforce outcomes to be easily observed in official statistics," according to a March 2026 report by the Federal Reserve Bank of Atlanta (Baslandze 2026).

Ultimately, the impact of generative AI on the labor force will depend on the choices made by firms.

Cut costs or grow?

Firms have long justified their investments in information technology on a variety of grounds, ranging from the financial to the strategic. Even when executives believe that an IT investment will produce strategic or other intangible benefits, the presence of finance department leadership on IT steering committees tends to lead firms to build business cases for IT investments at least partly around cost savings (Peffers and Dos Santos 2013).

Cost savings alone may not be sufficient to justify some IT investments. Discussing the rationale for investing in cloud computing, McKinsey partner James Kaplan put it this way in a 2023 podcast: "You can't make the case on IT cost by itself. The reason you go to cloud is because of the business value it enables, through agility, scalability, innovation, and flexibility, and that's organizationally a harder business case to make" (Fusaro and Rahilly 2023).

For many firms investing in gen AI, cost-cutting is a prominent objective. Deloitte found that twice as many companies had reduced costs with AI as had increased revenues: 40 percent vs. 20 percent. McKinsey's survey found that the number of companies reporting significant cost savings was up to thirteen percentage points higher than those reporting comparable revenue gains.

Cutting costs with gen AI is often assumed to imply cutting staff, but this is not necessarily so. Gen AI affords potential cost reduction mechanisms that do not necessarily involve cutting staff. For instance:

Improving quality can reduce the cost of rework.	**Reducing cycle time can reduce carrying costs.**	**Increasing scale can reduce the cost of growth.**	**Reducing risk may involve reducing fraud or compliance failure.**	**Improving decision quality decisions.**

Figure 4: Value levers that can reduce cost without cutting staff

While cost-cutting seems more popular, research has repeatedly found that using IT investments to drive growth is likely to hold more potential.

A 2024 study of around 2,000 companies from 2010 to 2018 found that company investments in AI were strongly correlated with increases in sales, employment, and stock market valuation. And the growth appeared to stem from product innovation and increased product offerings rather than from cost cutting (Babina et al. 2024). While the analysis covered a period that preceded the release of ChatGPT, there is no reason to doubt that a similar dynamic applies to gen AI, especially since it offers both efficiency and growth value-creation levers. Indeed, the paper by the Federal Reserve Bank of Atlanta mentioned above reported a similar pattern among firms using gen AI today: They tended to increase their revenue per worker through innovation and market expansion rather than through cost-cutting and staff reduction (Baslandze 2026).[2]

In the next section, I will discuss how some firms are using gen AI value levers to innovate and grow.

The value levers can support innovation and growth

Firms can use the gen AI value levers to innovate and drive growth. For instance, reducing cycle time may enable companies to close deals faster, accelerating revenue. It could enable firms to offer premium, expedited service to time-sensitive customers. Firms that use gen AI to deliver higher-quality services could win more business or charge premium prices. Building new capabilities and better products with features enabled by gen AI could drive growth as well (Figure 5).

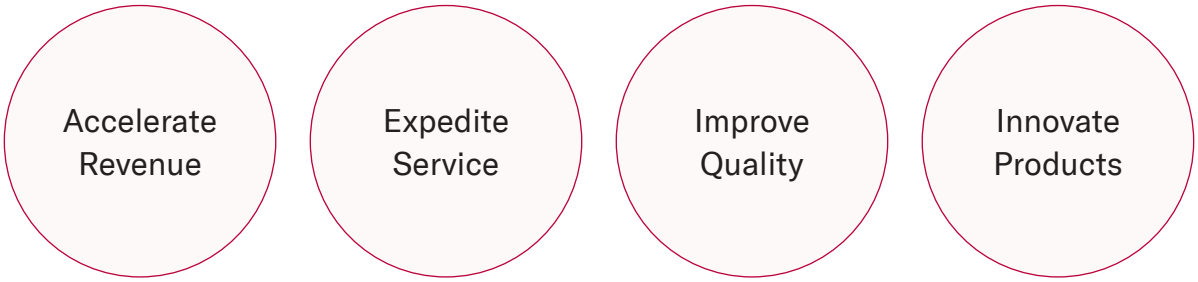

Figure 5: Growth-oriented gen AI strategies

Below, I present examples of how firms are using these value levers.

Reducing cycle time

Gen AI is being used to accelerate tasks that are bottlenecks in time-sensitive processes. This may involve reducing labor intensity, but reducing cycle time is a key part of the value. Interviewees mentioned this lever in **64 percent** of the use cases discussed.

In mergers and acquisitions, anything that enables analysts to complete valuations or due diligence more quickly is valuable. Executives at Deloitte and Ernst & Young (EY), both large professional services firms, told me they were using gen AI to support various phases of the M&A lifecycle, and particularly valued how the technology enabled them to deliver insights to clients faster.

2 Even before the current gen AI era, a 2012 study found that IT-enabled revenue growth tends to yield higher firm profitability than IT-enabled cost reduction (Mithas et al. 2012).

Speed is highly valued in software development as well. The faster software is developed and deployed, the sooner a business can derive value from it. Chris Price, chief information officer of Navan, a travel management service, says AI-powered coding tools have reduced story points—a measure of the effort required to complete coding tasks—by as much as two-thirds in some cases. This means they are achieving more in less time.

IT infrastructure vendor HPE has stated that a gen AI application reduced its financial reporting cycle time by 40 percent, enabling it to have more productive discussions about operational performance, implying the company can act faster to address issues and capitalize on opportunities.

Some work is inherently iterative—writing, software development, and design all involve successive rounds of refinement. In these fields, by accelerating the feedback cycle, generative AI can help raise the quality of the final output.

Speed, in other words, is not just an efficiency gain—it can drive better decisions, competitive differentiation, quality, and revenue.

Improving quality

The tendency of language models to hallucinate—confidently generating erroneous information—earned them an early reputation for unreliability. But with improvements in the models and workflows that include accuracy checks, generative AI implementations are now credited with improving quality in a wide variety of applications. Quality is a value driver in about **41 percent** of the use cases discussed in my interviews. In Wharton's survey of generative AI users, 87 percent of middle management respondents and around 90 percent of VP-level respondents said gen AI enhanced employees' skills and enabled higher-quality outputs.

Examples discussed in my interviews include providing superior customer service, reducing errors during the process of drafting clinical documents in the pharma industry, handling of insurance claims more accurately, screening resumes more accurately, and even enhancing the performance of less experienced employees. Higher quality can support premium pricing, client retention, and competitive differentiation—growth drivers rather than cost reducers.

Increasing scalability

Scalability is the ability to increase outputs without a proportionate increase in inputs. Gen AI-based applications can make a business process more scalable by reducing labor intensity. The value comes not from reducing headcount and associated costs but from enabling lower-cost growth and helping to manage variability in demand. Scalable systems can also help small firms punch above their weight. Scalability was a value driver in about **36 percent** of use cases mentioned in my interviews.

Aprio, an accounting and business consultancy, is using generative AI to help it scale. The firm grew organically by 16 percent last year while also acquiring more than a dozen companies; it has aggressive acquisition goals for the coming year as well. Chief digital officer Brent McDaniel says the firm's AI program, in which the firm is investing $300 million over five years, is continually reducing the labor intensity of client service, liberating about 20,000 additional person-hours every four months. The firm intends to add capacity and grow without having to add people.

Navan, the travel management company, has achieved greater scalability and staffing flexibility from its gen AI customer service platform. When customers seeking assistance use online chat, the system automatically translates between the languages of the customer and the agent if needed. According to CIO Chris Price, this enables the company to staff a shift with agents around the world, without concern for foreign language skills or the number of senior agents on the shift (since, with the support of gen AI, junior agents perform at nearly the same level).

Regeneration.VC, a small venture capital firm that invests in environmentally sustainable businesses, developed a process leveraging gen AI to help it screen the hundreds of pitch decks it receives each month. The tool extracts or infers the dozens of criteria the firm uses to decide whether to take a closer look at a deal. According to partner Martijn Lopes Cardozo, the tool expands the pool of deals its small staff can consider, increasing the possibility of landing on a winning investment opportunity.

In each of these cases, gen AI does not necessarily reduce what firms spend on labor; it expands what is possible with the labor they have. Whether companies like Navan may take advantage of gen AI to shift their talent model toward lower-skilled staff remains to be seen.

Innovating with new capabilities

Some companies are using gen AI to enhance their products or to expand the capabilities of the business **(36 percent of respondents).** This gen AI value lever can drive differentiation and growth.

Gong, a vendor of "revenue intelligence" software, added a host of AI-powered features to its product after the introduction of ChatGPT. The company says its gen AI features analyze customer interactions; create briefs on in-progress deals that summarize customers' pain points; suggest effective tactics to salespeople; and automate routine tasks like following up with customers by email. Product manager Inbal Ben Yehuda says these features help to make sales teams more effective and to close more deals; most customers quickly adopted them. In this example, generative AI creates value for Gong by helping it innovate its product. And it creates value for Gong's customers by reducing labor intensity, increasing scalability, and improving decision quality.

Adobe integrated generative AI across its product set to provide new features that have resonated with customers. Company executives have attributed accelerating revenue growth to the adoption of gen AI-powered capabilities such as Generative Fill in Photoshop, with which users can add, remove, or modify image content using text prompts.

Many startup companies are taking advantage of gen AI to offer novel products. Outset.ai helps automate and accelerate market and user experience research. The Outset platform helps design questionnaires and conducts interviews with subjects in a naturalistic dialogue as a chatbot session. Once the interviews are completed, the software analyzes the responses and presents a summary with key insights highlighted. Researchers can then query data through a chatbot interface. The product allows researchers to conduct an arbitrary number of interviews automatically, even simultaneously, cutting research cycle time. Because the tool lowers the cost and time required to do research, CEO Aaron Cannon says, the company's customers have tended to undertake more research than they had before.

These examples illustrate how gen AI can create value not by reducing the cost of existing work, but by enabling products and services that did not previously exist.

Making better decisions

Gen AI tools can reduce the cost and the time required for research and analysis. Consequently, these tools enable analysts to identify and interrogate more sources than possible before, potentially supporting better decision-making **(9 percent of respondents).** Anyone who has used a gen AI-powered chatbot has likely experienced how it simplifies and accelerates the research process. Firms use customized tools, powered by gen AI, to reduce the cost and time of many types of research and analysis, from market research to due diligence to legal contract review.

Deloitte's Erik Dilger believes gen AI-powered due diligence can improve the quality of deal outcomes by making it possible to do more thorough research, thus reducing surprises. In Deloitte's survey, 53 percent of respondents said they were using AI to enhance decision-making. In Wharton's study, respondents ranked insights generation and decision-making through data analysis sixth among the top ten benefits they were seeking with gen AI, up from eighth place in 2023.

Better decisions create value by improving outcomes rather than by reducing inputs—a distinction that matters for how firms think about measuring the return on their gen AI investments.

Reducing risk

Better decision quality is linked to another potential value driver of generative AI: reducing risk. For example, Dilger, the only executive I spoke with who mentioned risk management as a gen AI benefit, believes gen AI can help Deloitte identify and manage risks during M&A diligence and execution. Other examples can be found among tech startups. Startup Safetykit uses gen AI to help its online platform customers like Upwork and Eventbrite automate content moderation and fraud detection. Cyberthreat detection and response vendor Trellix uses gen AI to autonomously triage and investigate security alerts. In these examples, generative AI reduces the cost of threat and fraud detection to allow more comprehensive protection. Risk reduction, like the other value levers discussed here, demonstrates that the case for gen AI need not rest primarily on labor cost savings.

Of course, generative AI is also a new source of risks. While gen AI can empower cybersecurity professionals with new defensive tools, it can also afford cybercriminals new tools and expose new attack surfaces to hackers, as Anthropic's Mythos model has demonstrated.

Value recognized but not measured

Some companies have deployed generative AI applications that they believe are valuable but have not been able to quantify that value in conventional terms. Some are measuring the potential for business value ahead of seeing it accrue to the bottom line.

For instance, a large professional services firm developed a chatbot for internal use that automates common tasks, from getting IT support to building firm-branded PowerPoint presentations. The company was pleased with how intensively staff were using the tool but had yet to measure actual productivity benefits.

Rather than trying to calculate ROI from their gen AI investments, some companies believed they had no choice but to embrace gen AI, due to pressure from clients, competitors, or even their boards.

Predii, a software vendor serving the automotive service industry, built a gen AI-powered search function that helps technicians identify and order the right parts for repair jobs. This is a nontrivial task given that an auto manufacturer might have thousands of replacement parts in its catalog with each model, year, and set of options requiring different parts. The company spent twelve months training the models to achieve the desired accuracy. Predii made this investment because their customers were asking for AI features—sometimes under pressure from their boards to implement AI. This represents a defensive investment whose financial returns may be difficult to measure.

A large professional services firm uses gen AI to help estimate the value of an asset that a corporate client is considering spinning off. They have been able to reduce the time required to conduct a valuation from two weeks to a day or two. The accelerated valuation exercise does not directly generate much additional revenue, but it increases the likelihood that the firm will win the lucrative follow-on consulting work if the client decides to proceed with a deal to sell off the asset. In this way, the firm believes it is monetizing the speed that gen AI enables.

Rob Roy, a strategic communications consultancy, has also used gen AI to help it complete projects faster, enabling it to serve more clients in the same period without growing staff. And it is able to take on urgent, time-sensitive projects it might have had to pass on before. Founder Josh Reynolds believes the selective, targeted use of gen AI in his small agency allows it to better manage crunch-time workloads and service large clients while maintaining high levels of strategic value.

Gen AI has helped pave the way to a new business model for Envisioning, a global innovation consulting boutique. Recognizing that gen AI could commoditize the research they do for clients, they have shifted their model: A gen AI-powered research process slashes the time they spend on research, while gen AI coding automates their data visualizations. Today, Envisioning gives some clients the research at no charge, billing them instead for live ideation sessions where executives discuss the findings. Digital technology had a similar impact on the music industry: Some artists who once earned primarily from record sales now derive much of their income from live performance.

Over time, firms will impose greater financial discipline on generative AI investments. And some applications that individual professionals find useful may face scrutiny if their return on investment remains difficult to demonstrate. This is a familiar pattern in enterprise technology adoption, where tools that improve individual productivity do not always translate into measurable firm-level value. If generative AI follows this pattern, we may see firms prune some of their gen AI applications. However, history suggests that widely valued tools rarely disappear entirely. Instead, they often become embedded in broader platforms, justified through indirect benefits such as talent retention or quality improvements, or absorbed into the baseline expectations of knowledge work.

New ways of working

Besides changing labor demand, AI is also changing how people work. In fields as varied as legal services and software development, generative AI is leading some workers to shift their focus from creating to orchestrating, reviewing, and editing. To make the most of generative AI, firms will likely have to change the structure of teams and workflows. What Erik Brynjolfsson and Lorin Hitt argued more than twenty-five years ago remains true today:

> A significant component of the value of information technology is its ability to enable complementary organizational investments such as business processes and work practices; second, these investments, in turn, lead to productivity increases by reducing costs and, more importantly, by enabling firms to increase output quality in the form of new products or in improvements in intangible aspects of existing products like convenience, timeliness, quality, and variety (Brynjolfsson and Hitt 2000).

Consider software development. Deloitte Canada is exploring how to restructure software product teams to take advantage of the productivity boost afforded by gen AI. Rather than staffing a team with a business analyst, a project manager, and a product manager, for instance, Sameer Bohra, who is working on gen AI innovation at the firm, envisions those three roles converging into one. Gen AI might automatically capture and document user requirements from the transcripts of meetings with business users. It could also automate the creation of status updates and release notes and their transmission to stakeholders.

Firms could also consolidate their development teams. Rather than comprising specialists in the various programming languages and tools a project might require, a single developer, aided by AI coding tools, may be able to leverage their expertise in one language to work effectively with all those tools. This might significantly reduce the cost of software development, including overhead of communication and coordination across teams, which would also reduce misunderstanding and speed up development.

A typical software team of the future may be smaller than those of today. This does not imply that mass unemployment among software developers is inevitable. Boosting the productivity of software developers might even increase demand for them. As Peter Swartz, chief science officer and co-founder of Altana, a supply-chain intelligence platform provider that uses gen AI, told me, the demand for software always outstrips the supply—there is always a backlog of feature requests and bugs to fix.

In the pharma industry, gen AI authoring tools are dramatically speeding up the creation of clinical documents and regulatory filings. Novo Nordisk took advantage of the time savings to in-source work that had been outsourced to specialized agencies. Gilead Sciences is developing its own gen AI-powered authoring system. They expect their medical writers to spend the time they save collaborating with other teams. One team member is responsible for training the medical writers in the soft skills their reconceived roles will require.

Experimentation may be required

Readers should view the examples above as experiments. Researchers and firms are still learning that there are better and worse ways to engage with generative AI. These can affect productivity, critical thinking, and skill development. So it is crucial to advance our understanding of effective ways of using this technology and to apply it effectively.

Take productivity. Experiments have shown that the use of gen AI can increase or decrease productivity, depending on how it is used (Simkute et al. 2024). There are questions about the impact of gen AI on critical thinking too; certain AI-supported workstyles can lead workers to reduce the effort they put into critical thinking (Lee et al. 2025). Some researchers predict that the use of gen AI could produce “cognitive surrender”—when workers stop not only doing the work but also evaluating whether the AI is doing the work well (Kim et al. 2026). Recent research has found the “behavioral signature” of cognitive surrender (Shaw and Nave 2026). On the other hand, with proper educational design, gen AI can also enhance learning (Kestin et al. 2025).

Gen AI can help workers build skills. But it might instead lead to the deskilling of workers (Kim et al. 2026). One observer warned of the risk of a "deskilling spiral": "AI handles entry-level work → fewer humans build foundational skills → remaining humans over-trust AI outputs → errors go undetected → the humans who would have caught those errors were never developed. This is not a static equilibrium; it compounds over time."

Whether AI will create value over the long term therefore depends heavily on how work is reorganized around it.

Avoid indiscriminate headcount reductions

Considering the multiple ways generative AI can create value, companies need to be thoughtful about their workforce strategies. We have seen that innovation tends to be a stronger driver of firm performance than cost cutting; that gen AI-powered productivity gains can make some workers such as programmers more valuable rather than less; and that gen AI can narrow the performance gap between less-skilled workers and their higher-skilled colleagues. Blanket job cuts in anticipation of gen AI efficiency gains are unlikely to be an optimal strategy and could inflict unnecessary harm on firms by eliminating valuable knowledge and workers (Davenport and Srinivasan 2026).

As I have shown here, the broad spectrum of value levers affords diverse opportunities for creating value with gen AI. These will often require designing new roles for workers and reconfiguring teams and workflows. Some roles will become obsolete and workers who held them will need to find new ones. But other workers may bring valuable skills and knowledge about the firm and its market to newly defined roles. Firms' investments in gen AI should account for the need to discover new ways of working with AI, to redesign staffing models, and to train workers for new opportunities.

Policy implications

The analysis in this paper speaks primarily to business leaders rather than to policymakers. Its central finding—that companies are overconcentrated on one value-creation lever, labor intensity reduction, while underexploiting others, particularly those oriented toward growth and product innovation—is fundamentally a strategic observation. The appropriate response is a change in how companies think about and deploy gen AI, not necessarily a change in law or regulation.

That said, the analysis has some policy implications, though creating and implementing effective policy would be challenging. I highlight two below.

Tax code reform. A key policy implication concerns the structure of the tax code. If companies are overinvesting in labor intensity reduction relative to other uses of gen AI, one likely reason is that the tax code incentivizes them to do so. As Acemoglu, Autor, and Johnson point out, the current U.S. tax code places a much heavier burden on firms that hire labor than those that invest in algorithms to automate work. This asymmetry predates generative AI. But AI may amplify its displacement effects at a time when, in some circles, including some AI companies and their investors, the idea of replacing workers with AI has gained currency. A response to this asymmetry would be to set marginal taxes for hiring and training workers equal to those for investing in equipment and software. This would eliminate the tax-driven bonus for replacing workers with technology (Acemoglu, Autor, and Johnson 2026).

There is a case for tax reform. But the political environment for it is not currently favorable. The One Big Beautiful Bill Act (United States Congress 2025) permanently reinstated 100 percent bonus depreciation for qualifying capital equipment—including AI hardware—and created a new Section 174A, permanently allowing immediate expensing of domestic research and development costs, which would include software development and AI model building. These provisions enhance the favorable tax treatment of AI technology investment with no counterpart on the workforce investment side, deepening, rather than correcting, the asymmetry Acemoglu, Autor, and Johnson identified.

Workforce transition. The analysis also supports some form of policy attention to workforce transition. For some workers, this may involve preparing for entirely new careers. Others may have opportunities to acquire and apply new skills in or adjacent to their current field. These skills might include fluency in using digital and AI-powered tools, critical thinking, and the ability to orchestrate work and review outputs relevant to their current or similar jobs.

How firms choose to adopt generative AI will determine what they require of their workforces. The examples described above—redefining roles as gen AI is deployed and investing in new skills to allow workers to navigate change without displacement—are promising but unproven. And, as noted, recent research suggests that some ways of organizing work around AI can degrade the experience of workers and the outputs produced. Codifying effective AI-supported work designs will require more research and experience. Government incentives for worker training could be one policy response here. But the track record of government-incentivized training programs is mixed. Support for further research on how to design AI-supported work could be beneficial.

The gap between what the analysis suggests is desirable and what policy can reliably deliver is wide. The findings here make a more compelling case for change in how companies think about and deploy gen AI than for any specific legislative action. That conclusion may seem modest. But identifying the right change agent—which is firms themselves rather than governments—is itself useful.

Value today; transformation tomorrow?

Firms in every industry are creating value with generative AI today. But for most companies, the impact of gen AI so far is relatively limited and localized—to individual workers or departments—and often hard to measure. Few large companies can say that the technology's impact is transformative, at least not yet.

Gen AI is not yet a mature technology. And despite already high spending by firms, adoption is still at an early stage. So, its impact should grow if the economics prove favorable. However, the leading frontier model providers, despite their rapid growth, remain loss-making—a fact that raises questions about the economics of gen AI over the medium and long term.

We can see from the applications already in use that there are numerous strategies for creating value with gen AI. Reducing labor intensity is just one of them. And even when reducing labor intensity, reducing headcount may not be the only or best way to create value. Companies can also capture value by increasing production without growing labor costs, reallocating staff to higher-value tasks, or reducing the cycle time in areas where that helps drive new revenue. A focus on cost reduction is understandable and may be appropriate for some firms. But research has shown that growth is the bigger opportunity. Gen AI can support both objectives.

Because of the versatility of the technology, its impact on workers and on the labor market overall is not preordained. Much depends on how firms choose to use it.

Capturing value with gen AI, as with any information technology used in an enterprise, requires organizational change. In some cases, firms will need to build new team structures, roles, and workflows. It is currently unclear what those changes should be. The experience of employees doing the work may be essential for determining the right approach. So rather than undertaking speculative layoffs, firms should support their workers and encourage them to explore how to produce more value with their labor.

The power of generative AI presents firm leaders with significant opportunities and consequential choices. Some paths can drive down costs and potentially marginalize workers. Others can power innovation, drive growth, and create new opportunities not only for firms and their customers, but for workers as well.

Acknowledgements

I am grateful to the many practitioners who agreed to be interviewed for this research, including Inbal Ben Yehuda (Gong), Ayan Bhattacharya (Axtria), Will Bible (Deloitte), Kabir Bindra (AHEAD), Sameer Bohra (Deloitte), Daniel Byler (ServiceTitan), Aaron Cannon (Outset), Martijn Lopes Cardozo (Regeneration.VC), Anastasia Christianson, Tim Cote (Brown Brothers Harriman), Erik Dilger (Deloitte), Giselle Doolan (Expedia), Patrick Gallic (Tokio Marine), Catherine Havasi (M2 Foundry), Waheed Jowiya (Metamorphic AI), Tilak Kasturi (Predii), Vivek Kulkarni (Deloitte), Chris Lee (Gilead Sciences), Subha Madavan (Pfizer), Adnan Masood (UST), Brent McDaniel (Aprio), Tim Potter (Deloitte), Chris Price (Navan), Brian Ray (Atos), Josh Reynolds (Rob Roy Consulting), John Shumadine (EY), Duncan Stewart (Deloitte), Peter Swartz (Altana), Elliott Torres (Foxbox Digital), Dave Williams (Trustpilot), Sean Wood (Human Pilots), Michell Zappa (Envisioning) as well as others who spoke on background.

I am especially grateful to Philip Burns (Chief of Staff Roundtable), Joe Hurd (The Katama Group), Erik Kiaer (Atamai), Carolyn O'Boyle (Aprio) , Jennifer Reingold (Jennifer Reingold Advisors), and Andrew Ryscavage (Brintin Bio) for facilitating key introductions. I also benefited greatly from the thoughtful feedback of Filippo Gualtiero Blancato, Joshua Entsminger, Nur Laiq, Charlie Pozniak, and the research assistance of João Hofmeister (all of the Harvard Kennedy School's Belfer Center) on earlier drafts; any remaining errors are my own.

Finally, I thank Michael McQuade, whose feedback on earlier drafts was invaluable, and Rebeka Davis Reed, and the Emerging Technologies, Scientific Advancement and Global Policy fellowship at the Belfer Center for their support throughout this project.

References

Acemoglu, Daron, David Autor, and Simon Johnson. “Building Pro-Worker Artificial Intelligence.” *NBER Working Paper Series*, Working Paper No. 34854. National Bureau of Economic Research, 2026. https://doi.org/10.3386/w34854.

Apotheker, Jessica, Vinciane Beauchene, Nicolas de Bellefonds, Patrick Forth, Marc Roman Franke, Michael Grebe, Nina Kataeva, Santeri Kirvelä, Djon Kleine, Romain de Laubier, Vladimir Lukic, Amanda Luther, Mary Martin, Jeff Walters, and Christoph Schweizer. “The Widening AI Value Gap.” Boston Consulting Group, September 2025. https://media-publications.bcg.com/The-Widening-AI-Value-Gap-Sept-2025.pdf.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. “Artificial Intelligence, Firm Growth, and Product Innovation.” *Journal of Financial Economics* 151 (January 2024): 103745. https://doi.org/10.1016/j.jfineco.2023.103745.

Baslandze, Salomé, Zachary Edwards, John R. Graham, et al. *Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives*. Working Paper 2026-4. Federal Reserve Bank of Atlanta Working Paper Series. Federal Reserve Bank of Atlanta, 2026. https://www.atlantafed.org/-/media/Project/Atlanta/FRBA/Documents/research/publication/working-paper/2026/03/25/04-artificial-intelligence-productivity-and-the-workforce-evidence-from-corporate-executives.

Brynjolfsson, Erik, and Lorin M. Hitt. “Beyond Computation: Information Technology, Organizational Transformation and Business Performance.” *Journal of Economic Perspectives* 14, no. 4 (2000): 23–48. https://doi.org/10.1257/jep.14.4.23.

Brynjolfsson, Erik and L. Hitt. “Breaking Boundaries.” *InformationWeek.* September 22, 1997, 54–61.

Carr, Nicholas G. “IT Doesn’t Matter.” *Harvard Business Review* 81, no. 5 (2003): 41–49.

Challapally, Aditya, Chris Pease, Ramesh Raskar, Pradyumna Chari. *The GenAI Divide: The State of AI In Business 2025*. Cambridge, MA: MIT NANDA, 2025.

Challenger, Gray & Christmas. “Challenger Report: February Cuts Plunge, YTD Hiring Falls 56 Percent.” 2026. https://www.challengergray.com/blog/challenger-report-february-cuts-plunge-hiring-falls-56-percent/.

Chen, Wilbur Xinyuan, Suraj Srinivasan, and Saleh Zakerinia. “Displacement or Complementarity? The Labor Market Impact of Generative AI.” Working paper, 2025.

Davenport, T.H., and Laks Srinivasan. “Companies Are Laying Off Workers Because of AI’s Potential—Not Its Performance.” *Harvard Business Review Digital Articles*, 2026, 1–7.

Fusaro, Roberta and Lucia Rahilly, hosts. “Getting Ahead in the Cloud.” *The McKinsey Podcast*, McKinsey & Company, November 30, 2023. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/getting-ahead-in-the-cloud.

Gunasekaran, A., Peter E. D. Love, F. Rahimi, and R. Miele. “A Model for Investment Justification in Information Technology Projects.” *International Journal of Information Management* 21, no. 5 (2001): 349–64. https://doi.org/10.1016/S0268-4012(01)00024-X.

Hitt, Lorin M., and Erik Brynjolfsson. 1996. “Productivity, Business Profitability, and Consumer Surplus: Three Different Measures of Information Technology Value.” *MIS Quarterly* 20, no. 2 (1996): 121–142.

Kestin, Greg, Kelly Miller, Anna Klales, Timothy Milbourne, and Gregorio Ponti. “AI Tutoring Outperforms In-Class Active Learning: An RCT Introducing a Novel Research-Based Design in an Authentic Educational Setting.” *Scientific Reports* 15, no. 1 (2025): 17458. https://doi.org/10.1038/s41598-025-97652-6.

Kim, Tae Woo, Umair Usman, Aaron Garvey, and Adam Duhachek. “From Algorithm Aversion to AI Dependence: Deskilling, Upskilling, and Emerging Addictions in the GenAI Age.” *Consumer Psychology Review* 9, no. 1 (2026): 142–64. https://doi.org/10.1002/arcp.70008.

Korst, Jeremy, Stefano Puntoni, Prasanna Tambe. *Accountable Acceleration: Gen AI Fast-Tracks Into the Enterprise.* Philadelphia: Wharton, University of Pennsylvania, 2025. https://ai.wharton.upenn.edu/wp-content/uploads/2025/10/2025-Wharton-GBK-AI-Adoption-Report_Full-Report.pdf.

Lee, Hao-Ping (Hank), Advait Sarkar, Lev Tankelevitch, et al. "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers." *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–22, April 26, 2025. https://doi.org/10.1145/3706598.3713778.

McAfee, Andrew, and Erik Brynjolfsson. "Investing in the IT That Makes a Competitive Difference." *Harvard Business Review* 86, no. 7 (2008): 98–107.

Mithas, Sunil, Ali Tafti, Indranil R. Bardhan, and Jeanne M. Goh. *Information Technology and Firm Profitability: Mechanisms and Empirical Evidence.* SSRN Working Paper, January 3, 2012. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1000732.

Peffers, Ken, and Brian L. Dos Santos. "Research Opportunities in Information Technology Funding and System Justification." *European Journal of Information Systems* 22, no. 2 (2013): 131–38. https://doi.org/10.1057/ejis.2012.60.

Rowan, Jim, Nitin Mittal, Beena Ammanath, Costi Perricos. "State of AI in the Enterprise: The Untapped Edge." Deloitte, January 2026. https://www.deloitte.com/content/dam/assets-zone3/us/en/docs/services/consulting/2026/state-of-ai-2026.pdf.

Shaw, Steven D., and Gideon Nave. "Thinking—Fast, Slow, and Artificial: How AI Is Reshaping Human Reasoning and the Rise of Cognitive Surrender." Preprint, PsyArXiv, January 12, 2026. https://doi.org/10.31234/osf.io/yk25n_v1.

Simkute, Auste, Lev Tankelevitch, Viktor Kewenig, Ava Elizabeth Scott, Abigail Sellen, and Sean Rintel. "Ironies of Generative AI: Understanding and Mitigating Productivity Loss in Human-AI Interaction." *International Journal of Human-Computer Interaction*, October 15, 2024, 1–22. https://doi.org/10.1080/10447318.2024.2405782.

Singla, Alex, Alexander Sukharevsky, Lareina Yee, Michael Chui, and Tara Balakrishnan. *The State of AI in 2025: Agents, Innovation, and Transformation.* McKinsey & Company, November 2025.

United States Congress. *One Big Beautiful Bill Act*, Public Law 119-21, §§ 168, 174A. 2025.