



New work, new world

Hard numbers are finally here. Generative AI could deliver up to \$1 trillion in annual growth by 2032, while potentially disrupting up to 90% of existing jobs. How can leaders navigate an upheaval of this scale and realize the technology's full potential? By investing in people.

Table of contents

8 A \$1 trillion story

Three scenarios that put the final productivity tally in human hands

10 Decoding job disruption

Exposure and friction scores reveal the scope of change

19 2032 on the horizon

How a spike in adoption may drive a strategic evolution

25 Forging a new trust compact

Four recommendations for maximum productivity

30 Final word

31 About the authors

32 Methodology

Executive summary

As generative AI has swept across the global stage, it's been clear something big was about to happen—to the economy, to business and society, to people's livelihoods. But how big would the change be? When would it hit? And would the impact be destructive or productive, embraced or feared, unifying or polarizing?

Answers are beginning to emerge, and it appears gen AI will be very big indeed.

By 2032, according to our research, the technology could inject the US GDP with up to \$1.043 trillion in additional annual value—an economic surge greater than the entire US construction industry.


We'll see a spike in gen AI adoption; while as many as 13% of companies could be leveraging the technology in three to four years, nearly half could embrace it in a decade's time, according to our most bullish scenario.

This remarkable projection not only underscores the power of gen AI to *enhance, augment and automate* many work tasks across all areas of the economy, but also sets the stage for a profound shift in how we approach work, productivity and economic growth.


Annual value of gen AI to US GDP by 2032

\$1.043 trillion

Businesses adopting generative AI



in 3-4 years



in 10 years

Executive summary

To quantify generative AI's impact on productivity and the future of work, we partnered with Oxford Economics to create an economic model. The model is calibrated to reveal three scenarios reflecting low, middle and high levels of business adoption.

While multiple studies have been published on this topic, we wanted to analyze how this powerful technology could impact actual people, actual jobs and actual careers—and by extension, business productivity and the economy. Our model digs into the 18,000 tasks that drive the US economy, and carefully examines the impact generative AI could have on the 1,000 jobs composed of these tasks.

While we focused on the US workforce, the general themes that emerged from the findings can be applied globally. This is because the tasks—and the theoretical ability for generative AI to perform or assist with them—remain the same despite differences in their relative importance and rate of adoption from country to country. [See more details](#) on the methodology.

Our model digs into the 18,000 tasks that drive the US economy and carefully examines the impact generative AI will have on the 1,000 jobs composed of these tasks.

Executive summary


What we learned has profound implications for workers—and for the future of work. Our model shows that in the next 10 years, most jobs (90%) could be disrupted in some way by generative AI. Everyone from administrative assistants to CXOs will be affected.

In some job families, workers may face months of joblessness as they seek to transfer their existing skills to new roles. In total, approximately 9% of the current US workforce may be displaced by generative AI. Based on an analysis of previous economic shifts, it's possible 11% of displaced employees—or almost 1% of the total workforce—may struggle to find new work again.

Left unmanaged, this level of disruption would have severe consequences for not only organizations and the people who work in them but also for productivity itself.

Because achieving the high end of our productivity forecast requires two things: high levels of business adoption of generative AI and low levels of disengaged or permanently displaced employees. And here's where businesses—and the humans who run them—come in.

Addressing both of these factors requires trust—trust between AI makers and AI users, between businesses and policy makers, between employer and employee, and trust in the technology itself. While all these areas are vital, we are primarily focused on the last two, because it's here that businesses can exercise the most control to maximize the productivity potential of generative AI.



of jobs could experience some degree of disruption by gen AI

Executive summary

The time is now, as gen AI begins to ramp up, for leaders to lay the foundation for a new trust compact, ensuring gen AI is a positive force not just for economic productivity but also for workers and society. If they succeed, we could be entering an age of unfathomable wealth and efficiency. If they don't, we could be facing a long period of unrest and strife, glimmers of which can already be witnessed among AI makers themselves.

By embracing the following tenets, organizations can take the crucial first steps to bolster confidence, build trust and open the door to a new age of productivity and prosperity.

1. Take care of your people

Roll out strategic reskilling programs at a pace and scale never seen before.

2. Innovate or stagnate

Plan for how you'll operate and create value in a decade's time.

3. Build confidence with transparency

Show how AI outcomes will serve the broader goals of business and society.

4. Put AI gains to good use

The wealth generated from AI also needs to benefit the workforce and the world.

A \$1 trillion productivity story—with a catch

Three scenarios that put the final
productivity tally in human hands

A \$1 trillion productivity story—with a catch


Generative AI is a breathtaking feat of technology that could have a substantial impact on society and the economy. But the way it plays out will be grounded in deeply human factors. Will we resist or welcome it? Adapt or stay unchanged?

The extent of that impact will be determined by the rate of business adoption and how quickly people adapt to working in new ways. These factors have historically slowed other productivity innovations. For example, microprocessors arrived in the early 1970s, but it took two decades for personal computers to be widely adopted and for productivity gains to materialize.

This is why we worked with Oxford Economics to develop three adoption scenarios.

If business adoption of gen AI is low, the annual productivity boost in the US could grow to 1.7% by 2032. Even this bearish estimate represents a significant upswing, given that the long-term annual average US growth hovers around the 2% mark. And if adoption comes in at the high end, that figure could soar to 3.5% by 2032 (see Figure 1). From the perspective of total economic output, this means US GDP would see a boost between \$477 billion and \$1 trillion in 10 years' time.

Three scenarios for gen AI impact on US GDP



Generative AI could inject anywhere from \$477 billion to \$1 trillion into the US economy by 2032, depending on the level of business adoption.

Source: Oxford Economics and Cognizant
Figure 1

Decoding job disruption

Exposure and friction scores
reveal the scope of change

Decoding job disruption

Another key factor that will either bolster or deflate gen AI-related productivity gains is disruption to jobs and job markets. The more disruption workers face, the more difficult it will be to adapt to new ways of working. To understand what that disruption will look like, we analyzed almost 1,000 professions and their associated tasks and calculated a “theoretical maximum exposure score” for each one (referred to in this report as the “exposure score”).


A lower score means a lower percent of the job’s high-value tasks will be automated, so people in that profession will be relatively unaffected. For example, those who work on a heating, ventilation and air-conditioning maintenance crew have an exposure score of 5% because generative AI will neither assist with nor automate most of their job tasks. Contrast this with the 62% exposure score for computer science teams—individuals who do these jobs are looking at major change.

Explained: Exposure score

This score reflects the degree to which an occupation will be affected by generative AI. The exposure score takes into account how many of the job’s tasks could be fully automated by gen AI; how many could be assisted; and the relative importance of these tasks.

Through most of the report, unless otherwise indicated, we use the “theoretical maximum exposure score,” which assumes immediate and unanimous adoption of gen AI, vs. the “predicted exposure score,” which is calculated using the central adoption scenario.

Example: Family medicine physician




- Gen AI could automate many lower-value admin tasks, such as collecting and maintaining patient information.
- It could also assist with higher-value diagnostic tasks.
- But it would be less effective for an array of very high-value tasks, such as successfully helping a patient with chronic illness to make life-saving lifestyle changes.
- **Resulting exposure score: 33.3% by 2032**

Decoding job disruption

By 2032, few jobs could go untouched by generative AI


Some impact:

Exposure scores of at least 5%




Greatly impacted:

Exposure scores of at least 25%



Little impact:

Exposure scores of 5% or lower



Most jobs will see some change from gen AI, and over half could be greatly impacted.

Source: Oxford Economics and Cognizant
Figure 2

Decoding job disruption

In the past, technology advances and automation largely impacted manual labor and process-centric knowledge work. This time around, it's different. The script has been flipped; recall our HVAC technician whose job will see little change. Instead, gen AI will change what's generally called knowledge work—everyone from entry-level number crunchers to seasoned heads of business units to C-suite executives will see their job evolve over the next decade.

In fact, our research shows CEOs could see an exposure score of more than 25% by 2032, as they begin using gen AI for everything from reviewing reports and analyzing operations to conducting competitive assessments and making strategic decisions. Even areas long considered the purview of senior execs—cost reduction, program improvement and policy change—will increase significantly in “automatability.”

For some people, gen AI-induced change is a reality today. Jobs involving credit analysis, computer programming, web development, database administration and graphic design already have an exposure score of about 50%. By 2032, as the technology advances, exposure scores for some of these jobs will climb to 80%.

Other types of workers will experience a surge in disruption as time goes on, as maturity and trust in gen AI lead to more applications of it (see Figure 3). For example, while customer service representatives currently have an exposure score of 11%, that is projected to surpass 63% by 2032.

Job impact will grow over time

From 2023 to 2032, the percent of jobs with high exposure scores (25% and above) could increase from 8% to 52%, a 550% uptick.




Chart data reflects “theoretical maximum exposure scores.” See methodology for more on how these scores are calculated.

Source: Oxford Economics and Cognizant
Figure 3

The degree to which a CEO’s tasks could be automated could see them with an exposure score of more than



25%
by 2032

Highly impacted jobs may surge



550%
by 2032

Decoding job disruption

As the exposure scores demonstrate, gen AI will affect some jobs more than others. Some workers will need to learn how to work with gen AI as it assists with some job tasks, while others will see their roles taken over by the technology.

To understand how difficult it will be for displaced workers to find new employment, we calculated a “friction score” for various occupational groups. This score represents the ease or difficulty of finding new employment with existing skills. (See methodology for more on how the friction score is calculated.) Interestingly, some of the hardest hit occupational groups (those with high exposure scores) have relatively low friction scores, meaning their journey to new employment is less complex.

Explained: Friction score

High friction score = more challenging to find new work

Low friction score = easier to find new work

Some of the hardest hit occupational groups (those with high exposure scores) have relatively low friction scores, meaning their journey to new employment is less complex.


Decoding job disruption

To better understand this dynamic, we plotted an array of job categories by their exposure scores and friction scores (see Figure 4). In this case, however, we used the “predicted exposure score,” which takes into account a more gradual rate of adoption of generative AI, as reflected in our central scenario. As a result, these predicted exposure scores are lower than the theoretical maximum exposure scores used throughout the report, which assume unanimous and immediate adoption.

Decoding job disruption

A look at job and worker impact in five years

We've mapped jobs that will be most impacted, and those with the most difficulty or ease of finding a new job.



Source: Oxford Economics and Cognizant

Figure 4

Decoding job disruption

Combining the friction scores and the predicted exposure scores provides us with a clear lens into the disruption ahead. For instance, software developers and database administrators (with relatively high predicted exposure scores of 8%) have relatively low friction scores of about 40, meaning their journey to new employment is less complicated, likely due to their in-demand skill set.

But many may face more significant and prolonged disruption from generative AI. Office and administrative roles (such as clerical assistants and receptionists) have a predicted exposure score of 4% but are saddled with a friction score of 80, indicating they will face an uphill battle.

US Census Bureau data, which informs our modeling, indicates this displacement won't be easily solved. Previous periods of change have seen around 11% of displaced workers struggle to find work—taking an average of 39 weeks to find new job opportunities. Disruption at that pace and scale is significant—and if left unchecked could result in trust-shattering levels of unemployment and economic instability for large portions of the workforce.

Disruption at that pace and scale is significant—and if left unchecked could result in trust-shattering levels of unemployment and economic instability for large portions of the workforce.

Decoding job disruption

At the same time, we also foresee an upside to generative AI workforce disruption, such as its positive impact on ongoing labor shortages that exist in key areas of the economy. For example, in healthcare, ongoing staffing challenges are driving up service costs while dragging down service quality. Our exposure scores indicate that emergency physicians, for example, will see up to one-third of their tasks automated over the next decade—a process that will enable organizations to deliver more healthcare services without the need to ramp up hiring.

Another potentially positive impact of gen AI on labor challenges is the way it can be flexibly deployed to accentuate an individual's existing strengths while softening their weaknesses—in effect enabling a broader spectrum of cognitive capabilities into a wider range of roles. Armed with the right tools, for instance, an equity analyst may not need to be a mathematical expert if generative AI handles that part of the job. Instead, they would perhaps lean on a strength such as communication to add value in the workplace. In this way, generative AI has the potential to significantly lower barriers to entry to a large segment of the economy.

Positive impacts of generative AI on the workforce are its ability to address ongoing labor shortages in key areas of the economy such as healthcare, and its ability to accentuate an individual's existing strengths while softening their weaknesses.

2032 on the horizon

How a spike in adoption may
drive strategic evolution

2032 on the horizon

Gen AI-driven change will not happen all at once, nor will it follow an evenly paced trajectory. Rather, our analysis indicates it will follow a similar s-curve as other tech advances, such as enterprise-grade software: a gradual rise, to a dramatic spike, to a plateau in which the technology becomes refined and pervasive. This progression will unfold for gen AI as businesses adapt, workers reskill and stakeholders adjust to the societal impact of the technology (see Figures 5 and 6).

Fast uptake in the next decade

Adoption of generative AI could move quickly in the decade ahead.

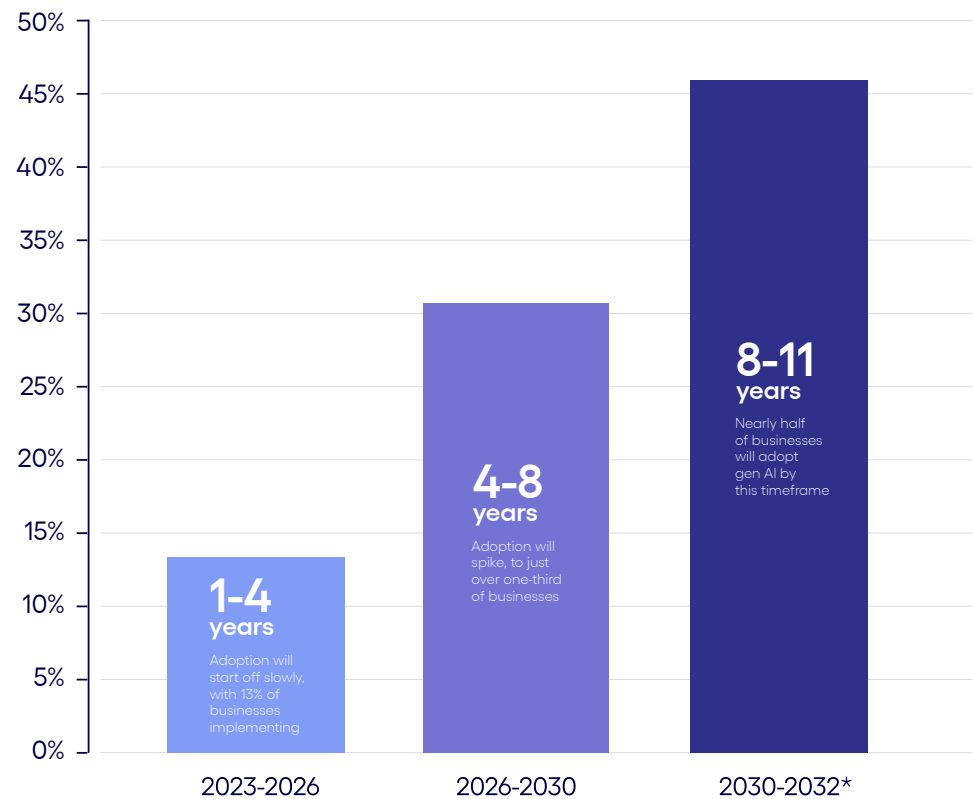



Chart data reflects the highest possible adoption rates, as reflected in our high adoption scenario.
*To ensure we captured the full adoption rate through 2032, we included 2033 data in the calculation for this figure.

Source: Oxford Economics and Cognizant
Figure 5

2032 on the horizon

The 30-year adoption view

After the 15-year mark, the steep climb slows, but adoption continues to grow for at least 20 more years.



Source: Oxford Economics and Cognizant
Figure 6

The difference between gen AI and its predecessors is one of pace, scale and the unique societal challenges that must be overcome. In our analysis, there are three stages to this growth story over the next 10 years—each with its own drivers, challenges and opportunities. These stages can and should guide an organization's approach to making the major institutional changes that will pave the way for a positive and productive embrace of generative AI.

2023 → 2026

Experimenting and preparing

This is the critical time for organizations to commit to major overhauls of their business and operating models and to establish a foundation of trust with their employees and the world at large.

The stage we're in today is characterized by experimentation and cautious implementation. Organizations are tentatively deploying gen AI for tasks that are manageable in scope and low in risk—like image creation, text generation for reports and emails, and code suggestions for developers. For example, our data shows that marketing managers have a theoretical maximum exposure score of 17% as generative AI helps execute tasks such as content generation and market analysis.

While our most bullish scenario foresees business adoption of only 13% in this timeframe, the importance of this phase cannot be underestimated. This is the critical time for organizations to explore and commit to major overhauls of their business and operating models, and to begin establishing a foundation of trust with their employees and the world at large. Only with a strong bedrock will organizations be able to drive productivity and lead their markets for decades to come.



Low-risk
experimentation



Gradual
adoption



Caution and
uncertainty

13%

Business adoption
("high" scenario)

2032 on the horizon

2026 → 2030

Confident adoption

This stage will see a notable rise in gen AI's role across a range of professions. Exposure scores for general managers and operations managers will grow from 18% today to 52.7% by 2032, with the bulk of the growth occurring in this time period.

In the next phase, we expect accelerated adoption. Our top-end prediction reveals adoption leaps from 13% to 31% in just four to eight years. Looking at historical examples—such as the gradual infiltration of enterprise software into heavily regulated industries—we can expect adoption to grow based on greater regulatory clarity, more robust skills management, stronger public demand and a greater understanding of how the technology resolves business challenges.

This stage will see a notable rise in gen AI's role across a range of professions. For example, exposure scores for general managers and

operations managers will grow from 18% today to 52.7% by 2032, with the bulk of the growth occurring in this time period.


Moreover, the generative AI regulatory environment—currently a mere patchwork—will come into focus. During this stage, gen AI will move beyond mere task automation to become a core component of [business transformation strategies](#). Forward-thinking companies will have developed robust technological infrastructures and a culture that not only accepts but also encourages continuous innovation.



Regulatory clarity



Lower barriers to entry



Reshaped business and operating models

31%

Business adoption (“high” scenario)

2032 on the horizon

2030 → 2032

Embedded organizational collaboration

With a solid foundation of trust, gen AI will assist roles that have traditionally been the exclusive domain of human intuition and experience. For example, our exposure score for legislators hits 37% by 2032.

As gen AI continues to mature, we will see an even greater acceleration in adoption. In our most optimistic prediction, we anticipate 46% of businesses will have meaningfully adopted the technology, heralding a new era of partnership between gen AI and decision makers.

With a solid foundation of trust established through consistent application and proven reliability, gen AI will assist roles that have traditionally been the exclusive domain of human intuition and experience. For example, our exposure score for legislators hits 37% by 2032.

In this phase, the technology is not merely a tool but a collaborator that offers nuanced analysis and strategic insights. This partnership is predicated on the reliability and sophistication of AI systems, which have evolved to understand and anticipate the complex needs of businesses. The data, strategy and culture of innovation fostered in the earlier phases enable generative AI to contribute meaningfully to high-level decision-making, working alongside human leaders to navigate the complexities of modern business landscapes.



Widespread
use cases



Proven
reliability



Meaningful role in
decision making and
strategic insights

46%

Business adoption
("high" scenario)

Forging a new trust compact

Four recommendations for
maximum productivity

Forging a new trust compact

While the timeframe of this research spans over a decade, leaders have little time to spare as they prepare to get the most out of generative AI.

While generative AI promises to boost productivity, its full potential isn't a given. The technology's uncertainties may slow its adoption and impact. Additionally, its influence on the labor market could affect global prosperity and trust. Navigating these aspects is key to realizing AI's benefits.

We've established four areas where we believe businesses can lead the charge by forging the levels of trust needed to optimize productivity so businesses, workers and societies can thrive in the age of generative AI.

1. Take care of your people

The integration of generative AI into business settings is poised to **transform job roles** and bring about some inevitable restructuring, highlighting the need for adaptability and continuous learning. However, the fear of layoffs should neither be downplayed nor ignored. Addressing this concern is crucial for fostering trust between employer and employee, and requires proactive and robust measures.

No business can guarantee layoffs will not occur. However, all organizations have it in their power to roll out a new generation of reskilling programs at a level never seen before. Rather than being a tactical, optional add-on to an employee's work life, these programs should be an essential part of the workday.

Reskilling programs could take a variety of forms. In some cases, businesses could partner with higher education institutions to continuously revamp curricula in selected skill areas. In others, organizations might collaborate with policy makers, government officials, regulators and even across industries to create shared “academy”

systems that would not only teach gen AI skills, but also establish new job tracks for people in roles with high exposure or friction scores.

The more robust the reskilling program is, the more it will become a marker of competitive advantage and brand reputation, which will serve to attract and retain key talent. Reskilling could ultimately be so important to a business's standing that it becomes an area of shareholder scrutiny, with investments and outcomes highlighted in the company's public reporting.

By reskilling even those slated to leave the company, businesses will grow the overall productivity pie by specifically equipping people to succeed and move freely between roles in the new world. They'll also bolster their brand as a company that's working with—not against—the vast socioeconomic impacts of generative AI. Future regulations may even require businesses to provide support for displaced workers when it can be proved a job loss was directly caused by generative AI-powered automation.

The fear of layoffs should neither be downplayed nor ignored. Addressing this concern is crucial for fostering trust between employer and employee.

2. Innovate or stagnate

As much as some employees might put off thinking about generative AI and the changes it will require, they won't want the same from their employers. Employee trust will hinge on whether their employers “get it” and are preparing for the next wave of change.

To shore up employee trust, and fully participate in all generative AI has to offer, businesses need to think about how the technology [will change how they do business](#) and even the business they're in. Rather than engaging in quarter-by-quarter thinking, they need to reimagine how they will operate and create value in 10 years' time.

Gen AI offers organizations new ways to earn revenue, improve operational efficiency, innovate products and services and, ultimately, redefine their businesses. For example, infrastructure services companies in the banking sector may find it less lucrative to provide the plumbing for global finance than to enable access to data for client large language models. Or life sciences companies may carve out an efficient and cost-effective path to direct consumer engagement using gen AI interfaces.

This type of evolution will require both [incremental adoption of new technologies](#) and a fundamental rethinking of business practices, corporate structures and the nature of work itself.

From the boardroom to the business unit, organizations should form cross-functional teams that continuously unravel process glitches and revenue roadblocks, and pave new paths into the future. New decision-making frameworks will be needed to quickly assess the best way forward—whether it's acquiring a burgeoning startup or forming new partnerships.

To execute newly-hatched business models, organizations will need a robust technological infrastructure designed to accommodate the rapid pace of gen AI development. Additionally, new governance policies, organizational structures and operating processes will be required to deal with massively accelerated, dynamic and autonomous operations. The companies that prosper will be those that make clear to all stakeholders that they understand the transformative potential of gen AI and make it a central pillar of their strategic planning.

Employee trust will hinge on whether it's clear their employers “get it” and are preparing for the next wave of change.

3. Build confidence with transparency

Many technologies have roiled the workforce, and all have seen some level of mistrust. But no technology has provoked as much fear and misgiving as generative AI. Whether it's the technology's black-box logic or its potential for bias and error, generative AI is simply not universally trusted.

Building trust in generative AI requires a clear and demonstrable commitment to mitigating any detrimental effects of the technology on people, society and the business itself. Organizations must be transparent about how gen AI systems are developed and deployed, the values they're designed to uphold, and the measures in place to ensure they do so. This transparency must show in concrete terms how gen AI decisions are made, and how outcomes serve the broader goals of the business and well-being of customers.

There are many ways to establish confidence in the decisions and outputs of a gen AI system. For instance, trust metrics can be incorporated into the application by providing visibility into the data the model used to make its decisions. Further, the system could produce a confidence rating in the decision itself. As design of these systems matures, their inner workings will become less mysterious, and overall confidence in AI-generated content and decisions will grow.

Organizations must also **build in safeguards** that minimize the risk of gen AI bias, error and ethically problematic decisions. **Using an array** of gen AI agents, surrogate predictive machine-learning models, explainable decision models, and human oversight and intervention, businesses can ensure their systems provide non-biased, ethically responsible responses and decisions.

Instilling trust in the technology will be key to business adoption, which, as our research shows, is necessary to reach the highest levels of productivity.

4. Put gains to good use

If you asked 100 employees how they thought their employer would use productivity gains realized by generative AI, 99 would likely respond with a skeptical—and perhaps historically accurate—response. But given the sweeping change generative AI promises for work and society, the past cannot serve as prologue.

Employee-employer trust will only be forged when businesses give at least some consideration to how generative AI benefits will be distributed across society—without hampering innovation and investment. We see this as the next front of stakeholder capitalism. In the same way today's net zero programs redistribute profits to take care of the planet, tomorrow's generative AI gains might be shared with employees and society.

While there are far less radical approaches, some businesses might completely transform their corporate credo and dedicate a high percentage of their generative AI-driven financial gains to a social cause, as sustainability-minded businesses do now to combat negative environmental impacts.

Another possible solution is for those who gain the most from gen AI-related advances to

contribute proportionally to a fund dedicated to societal benefits. For instance, investments could be made in educational programs to further support reskilling. Or funds could be earmarked for research into sustainable technologies or healthcare improvements, laying the groundwork for a society that is not only more efficient but also healthier and more resilient.

Another potential use of gen AI-related abundance, and surely a popular one, would be to improve employees' work lives, perhaps by shortening the workday or week. Indeed, many have already begun experimenting with or instituted a four-day workweek.

The options presented here are meant to stimulate thought and discussion among business and government leaders, policymakers, community organizers and philanthropic organizations. Each choice comes with its own set of complexities and would require careful planning and implementation.

But as with reskilling, it will soon be the case that the businesses which think progressively about reinvestment will be seen as trusted brands that are positively engaged in the generative AI age and the productivity it can deliver.

Employee-employer trust will only be forged when businesses give at least some consideration to how generative AI benefits will be distributed across society—without hampering innovation and investment.

Final word

When we started our research, the world was grappling with a mysterious advancement that had catapulted from popular curiosity to top agenda item for business and government leaders.

For all its potential, generative AI—like many AI tools before it—inspires equal measures of intrigue and aversion, attraction and rejection. The truth is, for everyone who welcomes generative AI into their daily lives, there's an equal number who fear the downside of these technologies: their potential to widen prosperity gaps, spread misinformation and reshape the very essence of what we've always understood to be human.

The questions surrounding gen AI run deep. Can we trust complex decisions made with non-human (and possibly inhumane) logic? Would we rely on an algorithm to have our best interests in mind? Would we feel safe interacting with an entity if we're unsure whether it's a person or a sophisticated AI system?

Such questions may not be answerable with an economic model. It will take time, experimentation and new levels of trust forged between all who would be impacted, for better or worse, by generative AI. No one will, or should, welcome an engine of economic growth if its achievements are detrimental to a vast swath of humanity, their sense of purpose and their very foundation of what it means to earn a living.

The world has the power to usher in a new era of productivity that's extraordinarily positive for the economy, workers and society itself. This path requires collaboration between the public and private sectors to agree on and uphold global standards of ethics and governance, deliver reassurance on the perceived threats of digital disruption, and educate and empower all in society to thrive amid new possibilities.

If those making the decisions, policies, systems and programs know what questions to ask—and answer them with the best interests of humanity in mind—then the potential of generative AI won't be a dry calculation but a living reality with benefits for all.

About the authors



Ollie O'Donoghue

Senior Director, Cognizant Research

Ollie O'Donoghue leads Cognizant Research, leveraging over a decade of experience as an industry analyst and consultant. His primary focus is on understanding the impact of new economic and technological trends on businesses and industries.

Throughout his career, Ollie has provided valuable guidance to C-suite decision-makers, helping them navigate the best paths for digital transformation initiatives, changing economic environments and emerging business models. He has also contributed to refining marketing messaging and developing go-to-market strategies for large IT services and software companies.

OliverRory.O'Donoghue@cognizant.com



Duncan Roberts

Senior Manager, Cognizant Research

Duncan Roberts is a Senior Manager in Cognizant Research. He joined the company in 2019 as a digital strategy and transformation consultant in industries ranging from satellite communications to educational assessment. He has advised clients on utilizing technology to meet strategic objectives and discover the art of the possible through innovation.

Before Cognizant, Duncan worked for one of the largest publishing houses in Europe, playing a leading role in the digital publishing revolution, helping transform their operations end-to-end and launching new innovative products. He holds a Master's in Philosophy and Classics from the University of St. Andrews.

Duncan.Roberts@cognizant.com



Acknowledgments

The authors would like to thank the following for their contributions to this report:

From Cognizant

- Lynne LaCascia, Head of Brand and Creative Strategy, Content, Thought Leadership and Research
- Catrinel Bartolomeu, Director of Storytelling and Content
- Mike Turner, VP, Software and Platform Engineering
- Naveen Sharma, Global Practice Head, AIA Practice Leadership
- Mary Brandel, editor

From Oxford Economics

- Vasilis Douzenis, Lead Economist
- Henry Worthington, Director
- The entire Oxford Economics team

Methodology

Cognizant partnered with Oxford Economics to assess and forecast the economic impact of generative AI technology on the US over the next 10 years. Here is a comprehensive explanation of our approach and rationale.

The project was executed through five main phases of work that enabled us to develop assumed inputs for Oxford's Global Economic Model. These inputs reflect how we anticipate the use of generative AI in US businesses will influence structural drivers of economic activity, notably total factor productivity (TFP) growth. To reflect the uncertainty inherent in such a process, we ran three scenarios to develop a range of outcomes.

We began by using a US Department of Labor database that contains detailed information on the skills and requirements for approximately 1,000 occupations—a database that now spans more than 20 years—to classify the more than 18,000 unique tasks that must be performed in these occupations into categories reflecting their maximum potential for automation through generative AI. This involved the development and testing of a machine learning model, which we used to generate a “theoretical maximum exposure score” for every occupation.

To model possible adoption rates of generative AI, we analyzed precedents from transformative technologies such as steam engines, electricity, automobiles and computers. We also applied a Bass Diffusion Model, a framework commonly used to predict the speed and shape of adoption for new technologies and products entering the market. Using this adoption data, we calculated predicted occupation exposure to generative AI by mapping the theoretical maximum exposure scores to these predicted industry adoption rates.

Because generative AI in the workplace will inevitably have a disruptive effect, it is reasonable to expect transitional costs. To build them into the model, we turned to the Current Population Survey of displaced workers from the US Census Bureau. This enabled us to understand the proportion of workers that will be displaced permanently by generative AI and also to generate a “friction score” to determine the degree to which displaced workers may struggle to regain employment.

These factors, combined with other, more detailed inputs (noted in the full methodology), enabled us to estimate generative AI's TFP and labor market participation impacts. The aggregated impacts of generative AI on productivity and employment were modeled as shocks applied to Oxford's Global Economic Model to estimate the contribution of the technology to US GDP growth above Oxford's baseline forecasts.

For the full methodology, [click here](#).



Cognizant Research

At Cognizant Research, we help leaders make sense of today's volatile socio-economic climate. We bring quality research and insights rooted in Cognizant's deep industry and technology expertise, helping leaders make the decisions that fuel their companies' success. Visit us at www.cognizant.com/us/en/insights.

[Learn more](#) about how we're working with generative AI and the people who use it to engineer modern businesses and improve everyday life.

About Cognizant

Cognizant (Nasdaq: CTSI) engineers modern businesses. We help our clients modernize technology, reimagine processes and transform experiences so they can stay ahead in our fast-changing world. Together, we're improving everyday life. See how at www.cognizant.com or [@cognizant](#).

World Headquarters

300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Toll Free: +1 888 937 3277

European Headquarters

280 Bishopsgate
London
EC2M 4RB
England
Tel: +44 (01) 020 7297 7600

India Operations Headquarters

5/535, Okkiam Thoraipakkam,
Old Mahabalipuram Road,
Chennai 600 096
Tel: 1-800-208-6999
Fax: +91 (01) 44 4209 6060

APAC Headquarters

1 Fusionopolis Link, Level 5
NEXUS@One-North,
North Tower
Singapore 138542
Phone: +65 6812 4000

© Copyright 2024, Cognizant. All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the express written permission of Cognizant. The information contained herein is subject to change without notice. All other trademarks mentioned herein are the property of their respective owners.