

NBER WORKING PAPER SERIES

FORECASTING THE ECONOMIC EFFECTS OF AI

Ezra Karger
Otto Kuusela
Jason Abaluck
Kevin A. Bryan
Basil Halperin
Todd R. Jones
Connacher Murphy
Philip Trammell
Matt Reynolds
Dan Mayland
Ria Viswanathan
Ananaya Mittal
Rebecca Ceppas de Castro
Josh Rosenberg
Philip Tetlock

Working Paper 35046
<http://www.nber.org/papers/w35046>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2026

This research was funded by a grant from Open Philanthropy (now known as Coefficient Giving). Karger and Kuusela are joint first authors. The views expressed in this paper do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. We thank seminar attendees at the Congressional Budget Office for thoughtful feedback on this project. We also thank Morgane Bascle, Oskari Luoma, Aashish Reddy, Sam Holton, and Victoria Schmidt for outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Ezra Karger, Otto Kuusela, Jason Abaluck, Kevin A. Bryan, Basil Halperin, Todd R. Jones, Connacher Murphy, Philip Trammell, Matt Reynolds, Dan Mayland, Ria Viswanathan, Ananaya Mittal, Rebecca Ceppas de Castro, Josh Rosenberg, and Philip Tetlock. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Forecasting the Economic Effects of AI

Ezra Karger, Otto Kuusela, Jason Abaluck, Kevin A. Bryan, Basil Halperin, Todd R. Jones, Connacher Murphy, Philip Trammell, Matt Reynolds, Dan Mayland, Ria Viswanathan, Ananaya Mittal, Rebecca Ceppas de Castro, Josh Rosenberg, and Philip Tetlock

NBER Working Paper No. 35046

April 2026

JEL No. C83, E27, J21, O33, O47

ABSTRACT

We elicit forecasts of how AI will affect the U.S. economy, comparing the beliefs of five groups: academic economists, employees at AI companies, policy researchers focused on AI, highly accurate forecasters, and the general public. The median respondent in each group expects substantial advances in AI capabilities by 2030, small declines in labor force participation consistent with demographic shifts, and an annual GDP growth rate of 2.5%, which exceeds both the typical medium-run (2.0%) and long-run (1.7%) baseline forecasts from government agencies and private-sector forecasters. Conditional on a “rapid” AI progress scenario, in which AI systems surpass human performance on many cognitive and physical tasks, experts forecast substantial, though not historically unprecedented, economic shifts: annualized GDP growth rising to around 4% and the labor force participation rate falling from its current level of 62% to 55% by 2050, with roughly half of that decline—equivalent to around 10 million lost jobs—attributable to AI. A variance decomposition suggests that expert disagreement about these effects is driven primarily by different beliefs about the economic effects of highly capable AI systems rather than by disagreement about the pace of AI progress. These forecasts map onto notably different policy preferences across groups: experts strongly favor targeted measures such as worker retraining, whereas the general public supports both targeted programs and broader interventions, including a job guarantee and universal basic income.

Ezra Karger
Federal Reserve Bank of Chicago
ezra.karger@gmail.com

Otto Kuusela
Forecasting Research Institute
otto@forecastingresearch.org

Jason Abaluck
Yale University
and NBER
jason.abaluck@yale.edu

Kevin A. Bryan
University of Toronto
kevincure@gmail.com

Basil Halperin
University of Virginia
basilh@virginia.edu

Todd R. Jones
Mississippi State University
trj234@msstate.edu

Connacher Murphy
Forecasting Research Institute
connacher.murphy@gmail.com

Philip Trammell
Stanford University
trammell@stanford.edu

Matt Reynolds
Forecasting Research Institute
matt@forecastingresearch.org

Dan Mayland
Forecasting Research Institute
dan.mayland@forecastingresearch.org

Ria Viswanathan
Forecasting Research Institute
riavishi2004@gmail.com

Ananaya Mittal
Forecasting Research Institute
mittalananaya@gmail.com

Rebecca Ceppas de Castro
Forecasting Research Institute
rebecca@forecastingresearch.org

Josh Rosenberg
Federal Reserve Bank of New York
josh@forecastingresearch.org

Philip Tetlock
University of Pennsylvania
tetlock@wharton.upenn.edu

1. Introduction

1.1 Background

The diffusion of generative artificial intelligence into workplaces, consumer products, and public services has renewed a familiar set of economic questions: Will automation raise productivity enough to change the economy’s long-run growth path? What will happen to work—employment, participation, wages, and occupational structure—as machines become capable of performing an increasingly broad set of cognitive tasks? And if the economic gains from AI are large, will they be broadly shared or concentrated among owners of capital and workers whose skills complement AI?

Despite intense attention, the evidence on AI’s economic effects to date remains mixed. Recent work finds signs of labor-market adjustment concentrated among the parts of the workforce most exposed to generative AI: Brynjolfsson, Chandar, and Chen (2025) document a sizable decline in early-career employment in AI-exposed occupations following the widespread rollout of generative AI, with limited effects on older workers and less-exposed fields. Complementary evidence, however, complicates this interpretation. Humlum and Vestergaard (2025) find similar early-career declines in Danish data, but do not find that the declines are strongly linked to firm-level adoption of generative AI, raising the possibility that measured effects reflect broader changes in demand, task organization, or labor supply. Davis (2026) adds further context to these findings, showing that while employment in highly AI-exposed sectors has lagged the rest of the economy since late 2022, wages in those same sectors have not fallen. Related evidence in Gimbel et al. (2025) likewise points to limited near-term aggregate impacts. Taken together, the emerging empirical record suggests that some groups of workers and occupations may already be experiencing measurable effects from AI, but the aggregate macroeconomic signal remains difficult to isolate from typical fluctuations in a dynamic economy.

At the same time, prominent voices in business and the AI industry warn of transformative upheaval: Jamie Dimon, CEO of JPMorgan Chase, argues that AI “will eliminate jobs” and that “people should stop sticking their head[s] in the sand” (Gerut, 2025); Sam Altman, CEO and cofounder of OpenAI, predicts that “whole classes of jobs” will disappear even as unprecedented wealth is created; and Dario Amodei, CEO and cofounder of Anthropic, suggests that AI could push overall unemployment to 10–20% within the next five years (Altman, 2025; VandeHei and Allen, 2025). Quantitative analyses span a similarly broad range. Arnon (2025), using the Wharton Budget Model, projects AI-driven productivity and GDP gains (in levels) of just 1.5% by 2035 and 3.7% by 2075—modest figures that translate to less than 0.04 percentage points of additional annual productivity growth in the long run. The OECD, by contrast, estimates that AI could add 0.4–1.3 percentage points to annual aggregate labor productivity growth over a ten-year horizon in high-exposure countries (Filippucci et al., 2025). The gap between these assessments is wide, and the policy stakes are correspondingly high.

A central challenge in debates about the economic effects of AI is that such forecasts are, unavoidably, joint forecasts about the capabilities of AI systems and the diffusion of AI-related technology into the economy. In practice, these discussions often combine answers to three distinct questions into one. First, will AI capabilities advance meaningfully, such that

AI systems become capable of independently performing, or assisting with, a large quantity of economically valuable work? Second, if such progress occurs, what will happen to key macroeconomic outcomes, including GDP growth, productivity, labor-force participation, and inequality? And third, given predictions and uncertainty about the effects of AI on the economy, what are the optimal policy responses?

The first question about the progression of AI capabilities is primarily addressed by two research fields: (1) AI benchmarking, where computer scientists develop tasks that track the limits of AI systems’ capabilities over time in work such as Russakovsky et al. (2015), Hendrycks et al. (2020), and Jain et al. (2024); and (2) capability forecasting, where researchers collect predictions about AI capabilities in projects like Grace et al. (2024) and the Longitudinal Expert AI Panel (Murphy et al., 2025). The second question on how AI diffuses into the economy is the core focus of a large and growing economics literature, as discussed above, yet standard models yield ambiguous predictions: the impact of rapid technological change on employment is theoretically indeterminate in the short run and neutral in the long run under canonical frameworks (Aghion, B. F. Jones, and C. I. Jones, 2019; Agrawal, Gans, and Goldfarb, 2018), in part because automating some human tasks often augments the value of others (Agrawal, Gans, and Goldfarb, 2022). Diffusion and adoption introduce further uncertainty, as even the most powerful technologies can take decades to reshape aggregate outcomes (Comin and Hobijn, 2010). On the third question—which policies to optimally pursue—progress is difficult because of unresolved disagreement about the answers to the first two questions, since the optimal policy interventions depend critically on both the pace of technological change and its economic consequences.

1.2 This paper

We elicit expert beliefs on all three of these questions simultaneously, surveying four carefully selected expert groups alongside a general public sample. We draw our expert economist panel from those publishing on the intersection of AI and economics, speakers at major conferences about AI and the economy, and faculty at top-100 economics departments—with a dedicated sub-sample of prominent economists. We recruit AI industry professionals from the organizations behind frontier AI models, and we sample AI policy professionals from U.S.-based think tanks and research institutions. We draw highly accurate forecasters (superforecasters) from a set of individuals with a verified track record of exceptional predictive accuracy.¹ See Section 2.1 for more information about our survey respondents.

First, we collect unconditional (all-things-considered) forecasts of key U.S. economic variables—such as annual GDP growth, total factor productivity (TFP) growth, the labor-force participation rate (LFPR), and wealth inequality—at both near-term (2030) and long-term (2050) horizons. These forecasts reflect each respondent’s current beliefs about the likely trajectory of the economy given their overall views about both AI and other trends and shocks. Second, we elicit forecasts of the same variables conditional on three AI-progress scenarios by 2030—slow, moderate, and rapid—and each respondent’s subjective probability that the world will end up in each of these scenarios. This conditional structure allows us

¹The term ‘superforecaster’ was coined in early research on subjective probability elicitation in a geopolitical context, as described in Mellers et al. (2014).

to separate disagreement about AI capabilities from disagreement about economic impacts, given a particular level of capabilities. Third, we ask respondents to predict the marginal effect of six specific policy proposals on GDP and LFPR under both unconditional and rapid-progress conditions, and to indicate their normative support for each policy. This multi-layered design allows us to trace out how beliefs about AI progress propagate through to economic forecasts and, ultimately, to policy preferences.

The three scenarios for AI progress describe capabilities progress achieved by 2030, across the domains of research, problem-solving, creativity, agency, and robotics. While the full scenarios are described in Appendix H.2.1, to summarize:

1. In the “slow” scenario, AI is a capable assisting technology for humans: writing literature reviews at the level of a capable PhD student, handling half of all freelance software-engineering jobs that would take an experienced human a day to complete, topping up your online grocery cart, and physically being able to unload dishwashers in some homes.
2. In the “moderate” scenario, AI is an effective collaborator across domains: autonomous lab systems can make rapid advances in solar-cell technology; almost all freelance software-engineering jobs requiring 5 days of effort from an experienced human are automatable; robots can do dishes as quickly as humans; robo-taxis can drive anywhere that humans can.
3. In the “rapid” scenario, AI systems surpass humans in most cognitive and physical tasks. Autonomous researchers can collapse years-long research timelines into months or even days. AI systems can surpass all freelance software engineers, customer service agents, paralegals, and clerical workers. Models can write 2025-Pulitzer-caliber books—and negotiate the resulting book contract. Robots can assist in an arbitrary home or factory anywhere in the world.

In addition, forecasters were advised that the scenarios above were intended to describe AI capabilities, not adoption, and that they should consider that regulation, social norms, or integration challenges could delay real-world deployment of systems with these capabilities. They were further advised that reasonable people may disagree with our characterization of what constitutes slow, moderate, or rapid AI progress and may expect slow progress in some AI capabilities alongside moderate or rapid progress in others. Nevertheless, they were asked to select the scenario that, on balance, best represented their views. Finally, capabilities were defined as “achieved” only if they could be done by an AI system as inexpensively and as reliably as humans today.

1.3 Key findings

A detailed analysis of the resulting forecasts yields six key findings:

A majority of survey respondents predicted significant AI progress by 2030. All five groups surveyed anticipate substantial AI capability advances—even if real-world adoption lags—with the average economist assigning a 61.4% probability to moderate or rapid progress.

Despite expecting significant AI progress, unconditional economic forecasts are close to historical trends. Although economists’ unconditional GDP forecasts exceed almost all government and private-sector projections, most do not forecast major departures from recent macroeconomic baselines, citing historical base rates, adoption lags, demographic headwinds, policy responses, potential infrastructure bottlenecks, and longstanding patterns in how general-purpose technologies affect the economy.

Conditional on the rapid scenario, economists expect significant economic shifts, but not the transformative acceleration some have predicted. If the rapid scenario materializes, economists forecast GDP growth of 3.5%, LFPR falling to 55.0%, and the fraction of wealth held by the wealthiest 10% of households rising to 80.0% by 2050—large shifts, but with historical parallels, such as to GDP growth post-WWII, the LFPR before women entered the workforce en masse, or pre-WWII inequality.

Unconditional consensus masks significant uncertainty about rapid scenario outcomes. Economists’ unconditional forecasts are relatively tightly clustered, but under the rapid scenario the range of plausible outcomes expands, suggesting that experts have far less confidence about what would happen if AI proves truly transformative.

Between-group differences are small relative to within-group disagreement, and most disagreement reflects uncertainty about economic effects rather than AI capabilities. In contrast to the view, articulated most clearly by Cunningham (2025), that the primary source of disagreement about AI’s economic effects is disagreement about the pace of AI capability progress, a variance decomposition suggests that expert disagreement about long-run macroeconomic outcomes is primarily driven by different beliefs about the economic effects of highly capable AI systems rather than disagreement about the pace of AI development itself.

Economists and the general public disagree on how to respond to AI’s economic impacts. Economists strongly favor targeted policy interventions such as AI-focused worker retraining (71.8% support) over broad structural interventions like job guarantees (13.7%) or universal basic income (37.4%), whereas the general public supports both targeted and broad interventions.

1.4 Prior Work

A growing body of economic research examines how advanced AI may reshape productivity and growth, labor markets, and inequality. While there is near-universal agreement that more capable AI tends to raise productivity, the magnitude and timing of this and other effects are contested. The core disagreement reflects a debate about two previously discussed questions: how fast will AI capabilities progress, and how fast will capable AI systems diffuse through the economy? This paper elicits expert beliefs on both dimensions, as well as on policy responses, and our survey design maps directly onto this debate.

Productivity and Growth

When it comes to productivity and growth, the research largely falls into three schools of thought, each defined by how broadly they believe AI can replace human labor across economic tasks. At the most optimistic end, Trammell and Korinek (2023) review theoretical models which suggest that, if AI can automate both production and research and development, self-reinforcing feedback between capital accumulation and idea generation could produce double-exponential or hyperbolic growth, marking a structural break from the near-linear growth of the past two centuries. Erdil and Besiroglu (2024) synthesize the arguments for “explosive growth,” which refers to economic growth rates significantly higher than recent historical trends, and assign roughly 50% odds of this occurring by 2100 if AI capable of broadly substituting for human labor is developed, while acknowledging the many potential regulatory, physical, and alignment constraints that create wide uncertainty in this forecast.

A more moderate view anticipates meaningful but smaller bumps in productivity and economic growth from AI. Based on Eloundou et al. (2024) estimates of the tasks GPT-4 could perform with scaffolding, Aghion and Bunel (2024) estimate annual TFP gains of 0.5–1.3 percentage points depending on assumptions about adoption and the share of tasks that are profitably automatable, noting that gains are likely transitory unless AI also automates idea production. Filippucci, Gal, and Schief (2024) embed sector-level AI exposure in a general equilibrium model and find a similar 0.25–0.6 percentage point increase in TFP, while highlighting that Baumol Cost-related effects from sectors with limited AI penetration may constrain aggregate gains.² Beraja and Zorzi (2022) add an important caveat to this discussion by explaining that when private incentives diverge from social welfare, which often happens in the rollout of a new technology, automation may be inefficiently allocated, slowing welfare gains even when measured productivity rises.

The most skeptical view is described by Acemoglu (2024), who uses task-based growth accounting calibrated to data to conclude that what he considers to be realistic near-term AI adoption adds only around 0.07 percentage points per year to TFP growth. His key finding is that most currently exposed tasks are either low-value, not yet reliably automatable, or both. Consistent with this view, Comunale and Manera (2024) confirm that firm-level productivity effects are positive, but that macro magnitude remains highly sensitive to adoption speed and institutional context. This “productivity without prosperity” framing has broader implications: income for many may fall, since even if aggregate output rises, labor’s share of that output may fall, concentrating gains at the top of the income or wealth distribution.

Labor Markets

The empirical labor market literature has struggled to keep pace with a technology diffusing faster than standard data-collection cycles can capture. Brynjolfsson, Chandar, and Chen (2025) document a 13% relative employment decline for early-career workers (ages 22–25) in the most AI-exposed occupations following the widespread rollout of generative AI. This may be evidence of concentrated displacement of entry-level workers; an early indicator of eventual macroeconomic effects. In this analysis, the mechanism appears to resemble a negative labor demand shock rather than simple labor substitution. Wages in exposed roles rose alongside

²For more on Baumol’s Cost Disease, see Baumol (2012).

employment declines, consistent with firms using fewer, more experienced workers rather than replacing all workers with AI. In contrast, Gimbel et al. (2025) find the overall occupational mix (by AI exposure) to be broadly stable, and Humlum and Vestergaard (2025) look for and find Brynjolfsson, Chandar, and Chen’s employment decline in a Danish sample, but show that the decline is uncorrelated with firm-level generative AI adoption, raising questions about potential confounders.

These conflicting early results are consistent with an economics literature that often struggles to find technological displacement in aggregate data over short time horizons. Frey and Osborne (2017), whose framework pre-dates generative AI, remains the standard baseline for classifying occupations as being exposed to automation from computerization; they summarize the literature finding such effects in data from the prior decades. On the other hand, Korinek and Juelfs (2022) provide a theoretical counterpoint to pessimism about displacement. They argue that a utilitarian social planner would phase out work gradually and only for workers with low productivity and weak non-monetary attachment to employment. The implication is that the appropriate response to automation is a managed transition rather than resistance.

Inequality

A recurring concern in the literature is that AI’s productivity gains will be distributed unequally, potentially decoupling aggregate economic growth from broad-based welfare improvement for most people. Acemoglu (2024) argues that rising TFP alongside a falling labor share is the most plausible scenario given task-based substitution patterns. Korinek (2024) extends this argument, modeling a post-AGI economy³ in which productivity growth could exceed 18% annually while workers’ wages collapse unless there is active redistribution. This model motivates a set of interconnected policy responses that the authors propose, ranging from UBI to global AI governance. Korinek and Stiglitz (2021) echoes this inequality concern on a global scale, arguing that AI strengthens superstar-firm dynamics and concentrates rents in the most developed economies, weakening the development pathways that allowed developing countries to grow quite fast in the 20th century. Lastly, Abbott and Bogenschneider (2018) identify a mechanism that could amplify these trends: existing tax systems treat automation-related capital as deductible while taxing labor, creating a potential bias toward over-automation that erodes the labor tax base.

Policy Responses

Several studies emphasize that optimal policy responses to AI (or technological innovation more generally) depend on assumptions about the pace and breadth of AI adoption. Slow, uneven diffusion may imply that targeted worker adjustment programs are optimal; rapid, broad-based automation may require structural reforms to taxation and redistribution at a scale with no peacetime precedent. Comunale and Manera (2024) and the OECD (OECD, 2023) both catalog existing regulatory approaches and stress that optimal interventions vary

³AGI stands for Artificial General Intelligence and refers to a world where AI systems can perform tasks and bundles of tasks at a level that exceeds the vast majority of current workers, especially for white-collar workers whose jobs are not particularly physical in nature.

sharply with assumptions about adoption speed and displacement risk.

On the labor-market adjustment side, Hyman, Kovak, and Leive (2024) use a regression discontinuity in the U.S. Trade Adjustment Assistance program to show that this type of program shortens unemployment spells by 1.26 quarters on average, raises cumulative earnings by over \$18,000, and is self-financing once one accounts for fiscal externalities from affected workers. This result speaks directly to the debate about modernized unemployment insurance. Korinek and Juelfs (2022) argue that broader institutional reform is ultimately needed, shifting social insurance from work-conditioned benefits to unconditional income support as automation capital replaces labor income.

On the tax and redistribution side, Abbott and Bogenschneider (2018) propose neutralizing the current potential bias towards capital (and against labor) by taxing capital income at the same rate as labor income, instead of at an implicitly lower rate. Bastani and Waldenström (2024) provide a more systematic optimal-tax approach, concluding that capital taxation becomes increasingly important as the labor tax base erodes, though they note practical constraints on feasibility in open economies. At the far end of the spectrum, Korinek (2024) and Anthropic (2025) advance comprehensive policy frameworks organized by scenario severity, from modest workforce training grants to sovereign wealth funds and new revenue structures, an approach that directly motivates this survey’s conditional policy questions.

Prior Surveys

Several prior surveys have elicited expert beliefs about AI capabilities and economic impacts. Grace et al. (2024) survey 2,778 AI researchers on their predictions about the pace of AI progress, finding a median estimate that AI will outperform humans at all tasks by 2047, with substantial disagreement across respondents. The Clark Center’s US Economic Experts Panel (Clark Center Forum, 2025) has polled leading economists on whether AI adoption will substantially increase growth rates or unemployment over the next decade, with most economists expressing uncertainty about the magnitude though agreeing on the direction of productivity effects. Murphy et al. (2025) establish the Longitudinal Expert AI Panel (LEAP), a three-year panel survey tracking the views of computer scientists, economists, AI industry professionals, and policy researchers through monthly surveys, providing a dynamic alternative to one-time cross-sectional surveys. (We compare our results directly to LEAP’s in Appendix C.) Our survey differs from these prior studies in three key ways: (1) we elicit forecasts of specific quantitative outcomes rather than qualitative agreement or disagreement; (2) we condition on explicit AI-progress scenarios, thereby allowing us to separate beliefs about capabilities from beliefs about economic effects; (3) and finally, we simultaneously collect forecasts and normative policy preferences, allowing us to link empirical expectations to policy support.

The Source of Disagreement and Our Contribution

A key open question that our survey is designed to address is whether disagreement about future economic outcomes is primarily about AI capabilities progress or about economic mechanisms conditional on that progress. Cunningham (2025) argues it is predominantly the former: forecasters largely agree on the economic logic but diverge on whether transforma-

tive AI capabilities will actually arrive. This view is consistent with the observation that Acemoglu’s skepticism and Trammell and Korinek’s optimism are in part reconcilable as they agree on the theoretical mechanisms, but apply to very different scenarios for AI capabilities and adoption. Here, we examine forecasts based on scenarios that span a wide range of AI development. Contra Cunningham, we find that disagreement centered on whether new AI capabilities will have an economic impact, rather than disagreement over whether such capabilities will arise.

2. Methods

This section describes the strategy for recruiting survey participants, the survey instrument, and the data processing procedures used in this study. The survey was launched in October 2025 and concluded in February 2026.

2.1 Participant Recruitment

We surveyed a diverse panel of experts and non-experts. Our sampling strategy targeted five participant groups: (i) economists, (ii) AI industry professionals, (iii) AI policy professionals, (iv) superforecasters, and (v) members of the general public. These groups were selected to capture complementary forms of expertise relevant to forecasting technological progress, economic outcomes of that progress, and policy responses. For ease of presentation, in most analyses, we group AI industry and policy professionals under ‘AI experts’.

In total, we contacted 4,866 experts, 54 superforecasters, and 1,395 members of the general public. Of the invited, 90 economists, 30 AI industry professionals, 27 AI policy professionals, 38 superforecasters and 401 members of the general public completed the survey. After filtering our expert respondents to ensure they meet our population criteria, our final sample contains 69 economists, 27 AI industry professionals, and 25 AI policy professionals. For details on participant recruitment, see Appendix A.

Economists, AI industry professionals and AI policy professionals were compensated at \$100/hour for a minimum of five and a maximum of ten (self-reported) hours. Superforecasters were paid \$60/hour, with the same time limits on compensation. Participants from the general public were initially paid \$30 for completing the survey, but this was increased to \$40 for later batches of participants. In addition to these payments, we incentivized participants to give insightful rationales by promising to award ten \$500 prizes among the expert participants, two \$500 prizes among superforecasters, and twenty \$100 prizes among the general public to participants with the highest-quality rationales.

Economists

Our primary population of interest was economists, who are best positioned to understand and provide quantitative forecasts of macroeconomic outcomes. We further subdivided this target population into three sub-populations: (i) economists working on AI-related topics, (ii) economists working on economic growth and technological changes more broadly, and (iii) well-known economists, such as Nobel prize winners. In the final sample, these groups




Figure 1: *Summary statistics of the economist sample.* The bars on the left show the distribution for participants, and the bars on the right show the distribution in our sampling frame (both participants and non-participants). The figure is based on publicly available information collected by human assistants and an AI-based system. Tier A and B publications are defined in Table 58.

have 24, 43, and 2 respondents, respectively. Figure 1 shows summary statistics about our economist respondents. Our respondents were younger, less experienced, more male, and more European than the sampling frame. They also had fewer high-quality publications and were less likely to have high levels of citations. That being said, we still have significant coverage across the less represented groups in our participant set. For example, 45% of our respondents had at least 16 years of experience, 38% were age 45 and older, 17% were female, 35% were from North America, 37% had at least one Tier A publication (a so-called ‘top-5’ publication in economics), and 58% had at least 1,000 citations. In our main results, we reweight the participants on experience and geography to match the sampling frame. This has minimal effects on our results (see Appendix G).

Economists working on AI were identified using three primary sampling pools: (i) a literature-based pool of authors publishing on the economics of AI, identified through Research Papers in Economics (RePEc) using relevant Journal of Economic Literature (JEL)

codes and AI-related keywords, (ii) an event-based pool of speakers and participants at major academic and policy conferences focused on AI and economic outcomes, and (iii) an institution-based pool targeting top-100 economics institutions according to RePEc. In the literature pool, invitations were extended in descending order of citation counts adjusted by paper age. In other pools, we sampled randomly.

Economists working on growth and technological change were similarly identified through (i) publications indexed in RePEc using JEL codes related to technological innovation and economic growth, and (ii) an institution-based pool. In the literature pool, authors were ranked by age-adjusted citation counts, and invitations were sent in descending order.

Well-known economists were identified using a combination of Nobel Prize recipients, RePEc author rankings, and participation in the Clark Center U.S. Economic Experts Panel. Invitations were extended to all individuals meeting these criteria.

AI Industry Professionals

AI industry professionals were sampled from companies developing or applying frontier AI models. We constructed sampling pools using three sources: (i) institutions associated with frontier models ranked by training compute, (ii) institutions associated with organizations producing top-performing models on public evaluation leaderboards, and (iii) highly funded AI startups identified via fundraising databases. Within each institution, we randomly sampled research and engineering staff. Sampling was stratified to ensure maximum coverage across chosen organizations.

AI Policy Professionals

AI policy experts were sampled from U.S.-based think tanks, research institutions, and government-affiliated organizations engaged in AI governance and technology policy. Participants were identified via institutional staff directories and professional networking platforms, focusing on researchers and policy practitioners working directly on AI governance.

Superforecasters

Superforecasters were recruited through our existing connections. All individuals in this pool have a demonstrated track record of forecasting accuracy.

Forecasters are denoted “superforecasters” if they (1) were in the top 2% of the accuracy distribution in a given year of the Intelligence Advanced Research Projects Activity (IARPA) Aggregative Contingent Estimation (ACE) tournament (IARPA, 2011; Mellers et al., 2014) or (2) they were a highly accurate performer on Good Judgment Open, an online continuous geopolitical forecasting tournament. Good Judgment Inc., which runs Good Judgment Open, then adds these top forecasters to the “superforecaster” pool. Most superforecasters come from the first selection criterion. Mellers et al. (2015) find persistent performance of these superforecasters across several years of geopolitical forecasting.

General Public

We included a general public sample to compare expert beliefs against the broader public’s expectations. Members of the general public were recruited through CloudResearch Connect (Hartman et al., 2023).

2.2 Survey Instrument

The survey instrument elicited probabilistic forecasts of AI progress, economic growth, labor-market outcomes, inequality, and policy effects over medium- and long-term horizons (by 2030 and 2050). Respondents provided both unconditional forecasts and conditional forecasts under three AI development scenarios—slow, moderate, and rapid progress, defined using concrete capability benchmarks.

For forecasts of economic outcomes, participants reported 50th-percentile forecasts and, for selected questions, 10th- and 90th-percentile forecasts. The survey also collected qualitative rationales to capture underlying reasoning and mechanisms. Participants received background information, historical data, and detailed resolution criteria for all forecasting questions. The survey instrument was administered online and required 8 hours for completion, on average. The survey questions are shown in Appendix H and the forecasting interface in Figure 80.

2.3 Data Processing

Coherence Checks

To ensure forecasts were both logically coherent and representative of participants’ intended forecasts, we applied a series of consistency checks to each participant’s forecasts. These include verifying that scenario probabilities summed to one and that conditional forecasts were not inconsistent with the unconditional forecasts for a given economic outcome. We reached out to participants to give them the opportunity to update their forecasts if their forecasts were flagged by these checks. In all analyses, we present results using these updated forecasts. If participants did not update their forecasts, we removed their forecasts for some checks, while we left them in place for others. For a full list of these checks, see Appendix F. The appendix also describes the impact of the coherence check process on aggregate forecasts: differences between pre- and post-intervention results are small.

Reweightings

We assign each participant in the economist group a weight to correct for non-response bias in participant recruitment. This ensures that the sample we derive our results from is representative of those we invited. We reweighted on years of relevant expertise and continent (North America, Europe, other), since our sample was biased towards European and more junior participants compared to the sampling frame (see Figure 1). These weights are used in all results presented below. For a full list of variables we considered for reweighting, see Appendix G. The appendix also compares key results with and without reweighting. The procedure does not systematically make results more conservative or extreme: reweighted GDP growth aggregate forecasts are more conservative (small-to-moderate differences) and

LFPR forecasts more extreme (moderate-to-large differences). Even the larger differences, observed for longer-horizon LFPR forecasts, are well within aggregate uncertainty bounds.

For the remaining groups, sample sizes are smaller and reweighting has not been applied; results for these groups should accordingly be interpreted as characterizing the recruited samples rather than their broader populations.

Aggregation

We use two approaches for aggregating forecasts. For our main results, we calculate a weighted (using the participant weights) median for each percentile separately. We see these as our main results, since this aggregates the noisy estimates of the underlying 50th, 10th, and 90th percentiles to arrive at a more accurate estimate.

To understand forecaster disagreement, we additionally performed an alternative aggregation procedure for questions where the additional 10th and 90th percentile forecasts were available. We first fit a distribution to an individual participant’s quantile forecasts (separately for unconditional forecasts and forecasts conditional on the rapid scenario) and then calculated a weighted average over the participants. This has the benefit of showing disagreement among participants. Section 4.3 contains more details about this procedure. For a comparison of these two aggregation approaches and evidence that the former may lead to more accurate forecasts, see Lichtendahl, Grushka-Cockayne, and Winkler (2013).

Rationale Analysis

We used an LLM to extract the key drivers mentioned in rationales accompanying forecasts for the four main results (GDP, TFP, LFPR and wealth inequality). We supplemented these LLM-curated drivers with a manual analysis of the rationales, checking for inconsistencies and adjusting the identified drivers to fully capture those offered in the rationales. We then used a second LLM agent to tag each occurrence of the identified drivers, and then analyzed the frequency with which different drivers were mentioned by various groupings of participants.

The rationales accompanying the policy results received similar treatment. An initial partial human review identified the likely main drivers and checked for systemic inconsistencies or question misinterpretations; this was supplemented by LLM analyses to confirm and expand upon these drivers, roughly quantify the frequency of driver mentions across participant groups, and extract rationales that best exemplified each driver.

See Appendix D.2 and Appendix I for more details on rationale analysis.

3. Results

3.1 AI Progress

In the survey, we elicited participants’ beliefs about the pace and scope of AI capability progress by the end of 2030 using three descriptive scenarios: slow, moderate, and rapid progress. These scenarios, described in full in Appendix H.2.1, are intended as broad benchmarks rather than precise descriptions of how the world will unfold. Participants were asked to judge which scenario a panel of experts would judge as the best overall match to the

true state of AI capabilities, recognizing that progress may be uneven across domains. Across all groups, we find that a majority of respondents place most of their probability weight on the moderate or rapid progress scenarios, expecting AI capabilities to advance significantly by 2030 even if real-world adoption lags.

Figure 2 summarizes the scenario descriptions and participants’ probability forecasts. On average, economists assigned the highest probability to the moderate scenario (47.4%), with the remaining probability split between the slow (38.6%) and rapid (14.0%) scenarios. The rapid scenario was judged to be the least likely according to all groups, including AI policy and industry professionals (here grouped together under ‘AI experts’), superforecasters, and the general public. Superforecasters placed more weight on slow progress when compared to economists, assigning 45.0% to the slow scenario, and were also the most skeptical about the rapid scenario, assigning it 12.6% on average. In contrast, the general public’s probabilities were more evenly distributed across scenarios (40.8% slow, 41.0% moderate, 18.1% rapid).

Figure 3 summarizes the distribution of scenario probability forecasts across expertise groups. The slow scenario category demonstrates the highest variance among respondents, most pronounced among AI experts, whose interquartile range (IQR) stretches from 16.5–60.0%. Variance was slightly lower for the moderate scenario probabilities, with economists having the widest interquartile range, stretching from 31.5–64.5%. Variance across respondent groups was lowest for the rapid scenario, with interquartile ranges across groups spanning from 10 percentage points (superforecasters) to 20 percentage points (the general public).

3.2 Growth

Gross Domestic Product (GDP)

Figure 4 compares historical measures of GDP growth (shown in black) to respondents’ forecasts of annualized GDP growth over the five years leading to 2030 and 2050. Participants provided unconditional forecasts, reflecting their overall views, as well as forecasts conditional on the three AI progress scenarios described throughout the survey. Across all groups, unconditional forecasts for 2030 were clustered tightly around the current rates, spanning from 2.4% to 2.5%. By 2050, unconditional medians remain in roughly the same range, ranging from the economists’ prediction of 2.5% to AI experts’ prediction of 3.0%. However, disagreement between groups, as well as uncertainty within groups, slightly increases (as reflected in the shaded regions). These forecasts suggest that most respondents do not expect strong trends of acceleration in GDP growth absent major shifts.

In contrast, the rapid scenario is associated with significantly higher expected GDP growth for both time horizons. For 2025–2029, the median economist predicted annual GDP growth to be around 3.3% (with median 10th and 90th percentile forecasts of 1.2% and 5.5%), and superforecasters showed a similar median (3.7%, 2.0–6.0%). AI experts’ rapid scenario median is slightly higher in this period, at 3.7%. The divergence between groups is wider by the 2045–2049 period, with annual GDP growth reaching 3.5% for economists (1.0–7.0%), 4.0% for superforecasters (1.0–7.0%), 5.3% for AI experts (2.3–9.3%), and 4.5% for the general public (2.5–7.1%). The moderate scenario falls between unconditional and rapid.

Figure 5 summarizes uncertainty using pooled distributions, shown for the unconditional

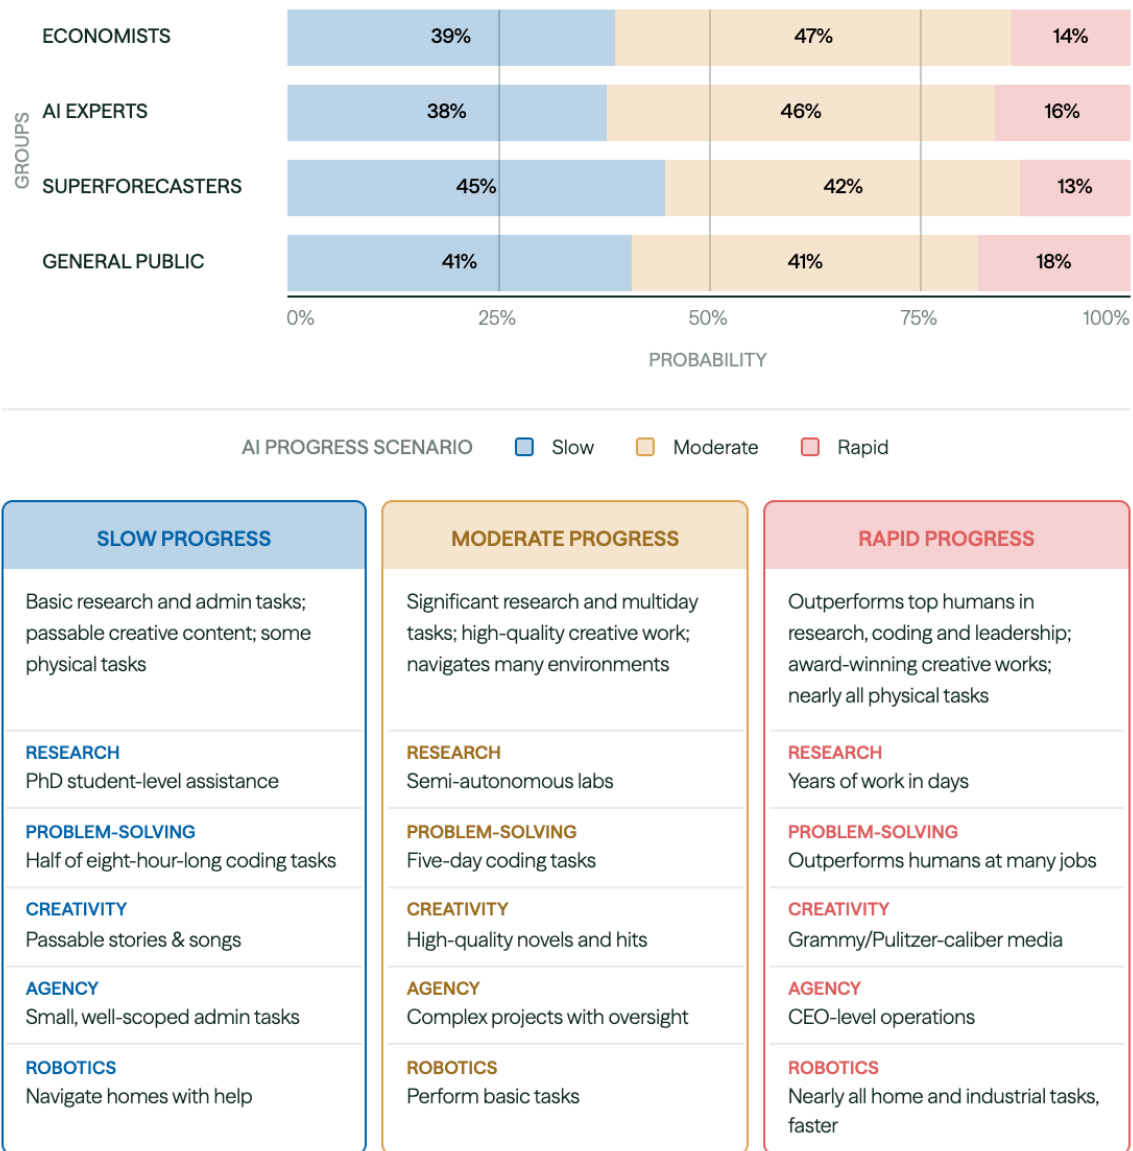


Figure 2: *Average probability of AI progress scenarios by 2030.* The average probability assigned by respondents in each group to the likelihood of a given AI progress scenario most closely describing the real world in 2030. The full scenario definitions can be found in Appendix [H.2.1](#).




Figure 3: *Distribution of AI progress scenario forecasts.* Vertical lines show the 25th, 50th, and 75th percentiles of the scenario forecasts. The 25th and 75th percentiles are labeled. The distribution is shown in two ways: as a histogram and as a density estimate.




Figure 4: *Forecasts for five-year annualized change in the Gross Domestic Product (GDP).* Lines show medians of 50th percentile forecasts across participants. Shaded regions span from the median 10th to the median 90th percentile forecast. The results for economists are reweighted to adjust for non-response bias (see Section 2.3). See Appendix H.4.1 for question details and the source of the historical data.

and rapid scenarios, for which participants provided 10th and 90th percentile predictions in addition to their best-guess 50th percentile predictions. Variance increases at the longer time horizon (see Table 10), consistent with greater uncertainty about growth outcomes, and with this broadening happening especially under the rapid scenario. For example, economists’ pooled distribution for GDP growth in 2030 under the rapid scenario spans 1.21%–6.420% (10th–90th percentile), compared to 0.73%–4.56% for the unconditional scenario. Similar patterns hold for other groups. Notably, the 90th percentiles of the 2030 and 2050 pooled distributions for AI experts reach the 10% winsorization cap (a bound applied to the distribution tails to prevent outlier forecasts from disproportionately skewing the pooled variance) and some experts forecast growth well above our cap. While superforecasters’ medians once more fall close to economists’, they assign higher weight to >10% growth in the rapid scenario, even in 2030.

We compare the median of economists’ 50th percentile GDP forecasts to other forecasts in Figure 6.⁴ Our economist sample’s estimate for 2025–2029 of 2.5% is higher than eight of

⁴These forecasts are: FOMC (Federal Reserve (FOMC), 2025), CBO (Congressional Budget Office, 2026), OMB (Office of Management & Budget (OMB), 2025), IMF (International Monetary Fund, 2025), SPF (Survey of Professional Forecasters, 2025; Survey of Professional Forecasters, 2026), OECD (OECD, 2025), Goldman Sachs (Goldman Sachs, 2025), The Conference Board (The Conference Board, 2026), and Deloitte



Figure 5: *Distribution of forecasts for five-year annualized change in Gross Domestic Product (GDP)*. Distribution is pooled across participants to summarize the full distribution of participant beliefs. Tail mass outside of figure bounds shown as ball-and-stick at 0% and 10%, with numbers in boxes indicating the proportion of the pooled distribution that lies below 0% or above 10%. Interior points show 10th/50th/90th percentiles of the distribution. See Appendix H.4.1 for question details.

the nine other forecasts we consider, which range from 1.9% to 2.1%. The only forecast that is higher is OMB’s of 2.8%. When we consider 2045–2049, this paper’s estimate of 2.5% is higher than both other estimates (1.3% and 1.6%).

Compounded over decades, small differences in growth rates can produce large differences in prosperity; see Figure 7. Consider the 3.5% growth rate in the rapid forecast possesses a doubling time of 20 years versus 28 years for the 2.5% growth in the unconditional forecast. Extrapolating to 2050, the rapid scenario produces a real GDP of \$54.7T, 25% larger than \$43.7T in the unconditional scenario.⁵ This is roughly equivalent to the GDP difference between 2016 and today.

Towards the end of the forecast period, growth approaches the post-World War 2 economic boom. From 1951-1973, growth averaged 4.06%, close to the rapid scenario; because population growth today is lower than during the postwar era, this suggests the rapid scenario could meaningfully exceed the per capita growth experienced from 1951-1973.

Total Factor Productivity

As a separate measure of economic growth, participants were asked to provide forecasts for the annualized change in total factor productivity (TFP) over 5 years. Figure 8 compares participant forecasts to historical measures of this outcome. Median forecasts are homogenous across groups, clustering tightly for the unconditional scenario around 1.0% (superforecasters and the general public)–1.2% (economists and AI experts) for 2030 and 1.0%–1.7% for 2050, with economists at 1.5%.⁶ Under the rapid scenario, forecasts roughly double relative to the unconditional baseline, growing to 1.9%–2.0% for 2030 and 2.2%–2.5% for 2050.

Key finding: Despite expecting significant AI progress, most unconditional economic forecasts are close to historical trends

The results above reveal an apparent tension: economists assign a 61.4% probability to moderate or rapid AI progress by 2030 (see Figure 2), yet their unconditional GDP and TFP forecasts do not substantially depart from recent baselines. The median unconditional GDP forecast of 2.5% for both 2030 and 2050 is only marginally above the 2021–2025 baseline of 2.39%, and TFP forecasts of 1.2% for 2030 and 1.5% for 2050 similarly represent noticeable, but incremental, gains relative to the 2025 baseline of 0.97%. While it is true that these forecasts are not quite as conservative as they may first appear when considered in the context of the medium- and long-run projections we report in Figure 6, the implied AI-driven productivity acceleration is modest by historical standards: the IT boom of the late 1990s arguably led to an additional 0.65 p.p. jump in annual TFP in the late 1990s (Jorgenson and Stiroh, 2017), and the economists’ forecasts of TFP growth in our study imply an acceleration in growth from AI that is roughly half that size by 2030 and on par with it by 2050.

(Deloitte, 2025). We focus on the baseline estimates; some forecasts consider other cases, such as Deloitte’s “Downside” and “Upside” scenarios and OECD’s “Energy Transition” scenarios.

⁵These estimates assume growth rates in 2030-2045 are linearly interpolated between the 2025–2029 and 2045–2049 forecasts, as in Figure 7.

⁶For comparison, the average CBO forecast of TFP between 2025 and 2029 is 0.82%, and between 2044 and 2049 it is 1.1% (Congressional Budget Office, 2026).




Figure 6: *Comparison of economist GDP growth forecasts to other forecasts.* In this figure, we compare the weighted median forecast for economists, which is indicated by the red dashed line, to nine other forecasts for 2025–2029 (left panel) and to two other forecasts for 2046–2050 (right panel). We use the baseline (or analogous) estimates. Where yearly estimates are given, we average these estimates. Where a single estimate is given for a yearly range, we use this estimate. Exceptions are: 1) for 2025–2029 Federal Reserve, we use the “Longer run” value for 2029; 2) for 2025–2029 Conference Board, we use the 2028–2032 value for 2028 and 2029; 3) for 2025–2029 Deloitte, we linearly interpolate values for 2027–2029 based on the 2026 and 2030 values; 4) for 2025–2029 (2045–2049) OECD, we use the value for the range of 2025–2030 (2045–2050). The type of forecast (real GDP or potential real GDP) is indicated.




Figure 7: *Projected GDP trajectory under different scenarios.* Based on economists’ GDP growth forecasts; growth rates between 2030–2045 are linearly interpolated from 2025–2029 and 2045–2049 forecasts. Shaded areas are determined using 10th and 90th percentile forecasts for the unconditional and rapid-scenario forecasts. Note the logarithmic scale on the y-axis.

Economists’ written rationales, however, lend insight into this tension. The most frequently cited reason why transformative technology would translate into only modest economic growth was uneven and time-lagged diffusion, with economists drawing on analogies to electrification, automobiles, and personal computers to argue that multi-decade lags routinely separate a general-purpose technology’s arrival from its measurable productivity impact. Geopolitical, structural, and demographic headwinds—including trade wars, climate change, an aging population, and declining immigration—that could offset AI-driven gains were also cited, as were constraints on energy, chips, and data center construction. Collectively, these factors were deemed likely to cap the pace at which AI capabilities could be deployed regardless of how quickly they advance. Under the rapid scenario, the possibility that large-scale workforce exit, societal unrest, or existential risk could become drags on GDP was also considered. A deeper analysis of these rationales can be found in Appendix D.2.1 (GDP growth) and Appendix D.2.3 (TFP growth).




Figure 8: *Forecasts for five-year annualized change in total factor productivity (TFP)*. Lines show medians of 50th percentile forecasts across participants. Because we elicited only 50th percentile predictions for TFP growth, this figure does not show uncertainty. The results for economists are reweighted to adjust for non-response bias (see Section 2.3). See Appendix H.4.3 for question details and the source of the historical data.

3.3 Labor Markets

Overall Impact

Our main metric for understanding labor market impacts is the Labor Force Participation Rate (LFPR). It measures the fraction of the adult population participating in the labor force, either through employment or active jobseeking. The LFPR has historically been relatively stable, mostly shaped by demographic trends, and is much less cyclical than the unemployment rate. It captures discouragement and long-term exits from the labor force.

Forecasts for the LFPR are shown in Figure 9. Economists expect the decreasing trend of the last two decades to continue. By the beginning of 2030, the median economist forecasts LFPR of 61.0%, a decrease from the January 2025 rate of 62.6%.⁷ This forecast has already been affected by economists' expectations about AI capabilities: conditional on the slow scenario, the median forecast is 61.5%, 0.5 p.p. higher. However, when asked to assume the rapid scenario, their forecast drops to 59.3%. If this were to occur, it would be the

⁷For comparison, the average forecast LFPR between 2025 and 2029 for CBO (using the Census Through 2020 Plus CBO Projection) is 62.4% and for Deloitte (after converting unemployment rate and employment-to-population rates to LFPR) is 62.0% (Congressional Budget Office, 2026; Deloitte, 2025).




Figure 9: *Forecasts for the labor force participation rate (LFPR)*. Lines show medians of 50th percentile forecasts across participants. Shaded regions span from the median 10th to the median 90th percentile forecast. The results for economists are reweighted to adjust for non-response bias (see Section 2.3). See Appendix H.4.4 for question details and the source of the historical data.

first time since the 1970s that a monthly LFPR reading fell below 60% in the U.S., and not since the 1960s has the annual rate fallen below 60%. The potential of AI to influence the LFPR—a difference between the slow and rapid scenarios of 2.2 p.p.—rivals the dip seen during the Covid-19 pandemic (from 63.3% in February 2020 to 60.1% in April 2020), but would reflect a more permanent impact on the labor force. By the beginning of 2050, economists expect a LFPR of 58.3% in the unconditional scenario⁸ and 55.0% when assuming the rapid scenario, with roughly half of that decline—equivalent to around 10 million lost jobs—likely attributable to AI rather than demographics and other non-AI trends.⁹

Directionally, all groups agree: faster progress implies lower LFPR values. The general public is slightly more optimistic than economists, with unconditional median forecasts of 62.0% and 60.0% for 2030 and 2050, respectively. Superforecasters are slightly more pessimistic in the long term: unconditionally, they give a median forecast of 58.2% for 2050 and 54.6% assuming the rapid scenario.

In the written rationales economists offered to explain their LFPR forecasts, the most frequently mentioned themes were AI-driven job substitution and loss, reallocation and reskilling, demographic changes, and historical LFPR baselines. A deeper analysis of the rationales can be found in Appendix D.2.4.

Key finding: Unconditional consensus masks significant uncertainty about rapid scenario outcomes

The median LFPR forecasts reported above may give a misleading impression of expert agreement. Across groups, unconditional forecasts cluster in a narrow band—economists at 61.0% for 2030 and 58.3% for 2050, with other groups nearby—and the decline from the current 62.6% baseline looks orderly. But when we examine the distributions underlying these forecasts, particularly under the rapid scenario, the range of plausible outcomes expands.

Figure 10 shows this uncertainty using pooled distributions for the unconditional and rapid scenarios, for which participants provided 10th and 90th percentile predictions in addition to their best-guess 50th percentile predictions. The pooled distribution represents

⁸For comparison, the average forecast LFPR between 2046 and 2050 for CBO is 62.1% (Congressional Budget Office, 2026).

⁹Both estimates are derived by comparing slow and rapid scenario LFPR forecasts. Under the rapid scenario, economists forecast a 7 p.p. decrease in the LFPR from the February 2026 reading of 62% to the 2050 median prediction of 55.0%. Under the slow scenario, where AI-driven displacement considerations are likely to be minimal, the decrease is forecasted to be 2.8 p.p., from 62% to 59.2%. This implies that $(7 - 2.8)/7 \approx 60\%$ of the projected rapid-scenario decline is attributable to AI rather than demographics or other non-AI trends. The AI-attributable share of the decline is thus roughly 4.2 p.p. (the difference between the rapid-scenario forecast of 55.0% and the slow-scenario forecast of 59.2%). Applying this to the current U.S. civilian non-institutional population aged 16 and over of approximately 270 million yields roughly 11 million fewer labor force participants. We round to 10 million to reflect the imprecision inherent in this calculation, including uncertainty about the size of the 2050 population and the simplifying assumption that the slow scenario captures all non-AI sources of LFPR decline. We also note that worlds with rapid AI progress likely differ from worlds with slow progress in ways that extend beyond AI capabilities themselves, so this gap should be understood as the estimated effect of moving between scenarios that are focused on AI rather than a cleanly identified causal effect of AI capabilities alone. That said, the U.S. civilian population is projected to grow through 2050, which would increase the number of displaced workers, making this estimate conservative.




Figure 10: *Distribution of forecasts for labor force participation rate (LFPR)*. Distribution is pooled across participants to summarize the full distribution of participant beliefs. Tail mass outside of figure bounds shown as ball-and-stick at 45% and 70%, with numbers in boxes indicating the proportion of the pooled distribution that lies below 45% or above 70%. Interior points show 10th/50th/90th percentiles of the distribution. See Appendix H.4.4 for question details.

both within-forecaster uncertainty and between-forecaster disagreement. The variance of the distribution (shown in detail in Table 10) increased at the longer time horizon and in the rapid scenario, consistent with the pattern observed for GDP growth. While economists’ pooled distribution for LFPR in 2030 spans 56.9–65.16% (10th–90th percentile) in the unconditional scenario, it widens to 53.09–64.4% in the rapid scenario. Variance increases even more in the 2050 rapid scenario, spanning 44.77–64.69%, with a significant amount of probability mass (14.3%) below the 45% winsorization floor (note that the 10th percentile falls below the winsorization floor), indicating that, in aggregate, economists predict a non-trivial chance of extreme low-LFPR outcomes. An LFPR at or near the median economist’s 10th percentile forecast for 2050 in the rapid scenario would indicate tens of millions had left the workforce and reflect an unprecedented shift in the structure of the economy.

This heightened uncertainty in the rapid scenario is also visible in GDP forecasts, where the 90th percentile of the pooled economist distribution under the rapid 2050 scenario reaches 8.43%, suggesting that experts have narrow priors for a world in which AI augments the economy incrementally, but far less confidence about what will happen if the technology proves truly transformative.

Impacts by Sector

Economists’ median forecasts for the sizes of different sectors as shares of the labor force are shown in Figure 11. In the unconditional scenario, economists expect the share of business and analytical (“white-collar”) roles to continue to rise slowly, reaching 21.0% in 2030 and 22.0% in 2050 compared to a 2025 baseline of 20.4%. The share of care and service workers is forecast to increase more rapidly, reaching 48.0% in 2030 and 52.5% in 2050 in the unconditional scenario, while skilled trade and industrial (“blue-collar”) occupations would fall to 12.5% in 2030 and 11.0% in 2050.

In the rapid scenario, the growth in white-collar occupations would stall, with their share remaining flat at around 20.0%–21.0% in 2030 and 2050. Care and service occupations are forecast to see larger increases compared to the unconditional scenario (reaching 57.1% in 2050), while blue-collar occupations would sharply decline, with a median share of 8.0% by 2050—a historical low.

Directionally, other groups largely agree with these aggregate economist forecasts, although there are some areas of disagreement. Under the rapid scenario, by 2050 superforecasters and AI experts expect a decline in white-collar jobs to 16.0% and 18.0%, respectively. AI experts predict a much smaller increase in the share of care and service occupations, growing to only 49.8% by 2050 in the rapid scenario, while the general public is the only group expecting a decline in this sector in the rapid scenario, forecasting a share of 42.0%. The results for other groups are shown in Appendix D.2.

Impacts by Occupation

Each survey participant was randomly shown 10 of the 43 International Standard Classification of Occupations (ISCO-08) sub-major occupation groups and asked to rank them by predicted percent change in employment from 2025 to 2030, indicating whether each would grow or




Figure 11: *Economists' forecasts for different groups of occupations as shares of the labor force.* Areas show the median 50th percentile forecast. The 'Other' category is derived by subtracting the sum of the other sectors' median forecasts from 100%. It consists primary of public sector and agricultural workers. Labeled historical values correspond to the beginning of 2025. See Appendix H.4.7 for question details and the source of the historical data.

decline relative to a 0% marker, both unconditionally and conditional on the rapid scenario.¹⁰

Figure 12 plots the percent of respondents who expect there to be growth in employment in a given occupation between the beginning of 2025 and the beginning of 2030. In the unconditional scenario, occupations economists expect will see the strongest employment growth are personal service workers, personal care workers, health professionals, and military occupations—roles characterized by in-person physical presence, human interaction, or security functions that are difficult to automate. At the bottom of the ranking, economists expect declines in general and keyboard clerks, other clerical support workers, stationary plant and machine operators, and assemblers—routine cognitive and manual roles thought to be particularly vulnerable to technological displacement. Point estimates for blue-collar occupations generally fell below the 50% threshold, indicating that a majority of economists placed these occupations in the job-loss category. This pattern was more pronounced than for white-collar occupations, which clustered closer to or above 50%, and for service occupations, which spanned a wide range but included some of the survey’s most optimistic forecasts.

Conditioning on the rapid scenario did not significantly alter these rankings. While most occupation groups shifted modestly leftward—that is, fewer respondents predicted positive growth—the differences between the two scenarios were not statistically distinguishable (also note the relatively small sample size). It is worth noting that the unconditional scenario should not be read as a counterfactual with no LLM exposure; respondents were presumably already incorporating LLM effects into their unconditional predictions.

In Appendix D.2.8, we compare the fraction of economists who predict each occupation will experience growth with the measure of AI exposure from Eloundou et al. (2024).¹¹¹² We do not observe a relationship between our results and AI exposure. We also compare the AI exposure measure to the difference in the fraction predicting positive growth between the rapid and unconditional scenarios, and again find no relationship.

3.4 Economic Inequality

We measure economic inequality as the fraction of wealth held by the 10% wealthiest households. This metric has had a moderate upward trend since the 1980s, reaching 71.2% in 2023. Like the LFPR, this metric has been relatively stable historically, ranging from 62.7% (in 1985) to 73.2% (in 2013-2014).

¹⁰There were 10 blocks, each with 10 occupations, and each person was randomly assigned to see one block. Two blocks mistakenly contained the same occupation twice. In these cases, we consider only one job if the respondent answered consistently for the duplicates and do not consider either if they answered inconsistently.

¹¹When comparing results, it is important to note that (Eloundou et al., 2024) focus on LLM-exposure, while our survey respondents predicted job growth. While one component of job growth is LLM exposure, there are, of course, many other things that the survey respondent may be considering.

¹²Eloundou et al. (2024) quantify how exposed the occupations are to LLMs by obtaining ratings from humans (and ChatGPT) on how exposed a job’s tasks (from O*NET) are. They present several versions of AI exposure; we focus on the version using human raters and Direct Exposure plus 0.5x LLM+ Exposed, where Direct Exposure is defined as “using the described LLM via ChatGPT or the OpenAI playground can decrease the time required to complete the DWA or task by at least half (50%)” and LLM+ Exposed is defined as “access to the described LLM alone would not reduce the time required to complete the activity/task by at least half, but additional software could be developed on top of the LLM that could reduce the time it takes to complete the specific activity/task with quality by at least half. Among these systems, we count access to




Figure 12: *The fraction of economists predicting a positive change in employment for each occupation between the beginning of 2025 and the beginning of 2030. Gray points correspond to the unconditional scenario, while red points correspond to the rapid scenario. Dashed lines indicate statistically non-significant results at the 5% level.*




Figure 13: *Forecasts for wealth held by the 10% wealthiest households.* Lines show medians of 50th percentile forecasts across participants. Shaded regions span from the median 10th to the median 90th percentile forecast. The results for economists are reweighted to adjust for non-response bias (see Section 2.3). See Appendix H.4.9 for question details and the source of the historical data.

These historical values, as well as forecasts by different groups, are shown in Figure 13. Economists expect wealth inequality to increase, with a median forecast of 73.2% in 2030 and 75.0% in 2050. In the rapid scenario, the increase is faster: 75.0% in 2030 and 80.0% in 2050. Our secondary measure of inequality, the labor share of income, leads to similar conclusions: inequality is increasing, and faster AI progress is predicted to lead to greater increases in inequality. These results appear in Appendix D.2.10.

Other participant groups agree with economists in directional terms. However, superforecasters give notably more conservative forecasts, especially for longer-term outcomes. For 2050, their unconditional forecast is 74.5%, and the forecast conditional on the rapid scenario is 75.0%.

Although these forecasts indicate inequality is set to increase, economists also predict that median household income will continue to increase (see Appendix Figure 52). Since we are measuring median household income in real terms (adjusted for inflation), this reflects an increase in the purchasing power of the median household. In addition, according to these forecasts, faster AI progress will lead to larger increases. Setting aside that we are measuring different concepts with these two measures (income vs. wealth), a possible interpretation of

image generation systems.”

these results is that while both the top-10% of households and the median household will see gains as a result of faster AI progress, these gains may tend to disproportionately accumulate to the top-10%.

In the written rationales economists offered to explain their wealth inequality forecasts, the most frequently mentioned themes were the shift from labor to capital, historical wealth inequality trends, redistribution and tax policy, ownership of AI infrastructure and IP, and uncertainty around 2050 outcomes. A deeper analysis of these rationales can be found in Appendix [D.2.9](#).

Key finding: If the rapid scenario materializes, economists expect significant economic shifts, but not the transformative acceleration some have predicted.

The wealth inequality forecasts under the rapid scenario describe a U.S. economy that would be substantially more unequal than today—but not unrecognizably so. The median economist forecast of 80.0% of national wealth held by the top 10% in 2050 would represent the highest concentration since the late 1930s, and yet that is a level the U.S. has reached before, albeit under very different technological and institutional conditions. This result fits a broader pattern we observe under the rapid scenario in which economists forecast large shifts—GDP growth of 3.5%, LFPR falling to 55.0%, wealth inequality at 80.0%—but shifts that still have historical parallels, such as GDP growth post-WWII, or the LFPR before women entered the workforce en masse, or pre-WWII inequality.

These forecasted impacts stand in marked contrast to the “transformative” economic impacts proposed by some technologists (Amodei, [2024](#)) and highlighted as possible by some economists (Brynjolfsson, Korinek, and Agrawal, [2025](#)). Even in this rapid scenario where AI models can surpass humans at most cognitive tasks and robots can surpass humans at most physical tasks, economists do not forecast—and no other group forecasts—anything like the tenfold increase in economic growth to around 30% discussed in the literature (Davidson, [2021](#)).

3.5 Other Outcomes

In addition to the main survey questions we report on above, we elicited forecasts on eleven additional measures, the results of which are summarized below. The full results can be found in Appendix [D.2](#).

Overall, economists project that unemployment will remain remarkably stable and that real median household incomes will continue to grow, but that the nature of work will change due to increasing use of AI. Labor’s share of economic output is expected to drop.

Key Results

Labor Productivity. Economists forecast only modest increases in the annualized change in labor productivity. For 2030, their median forecasts are 2.0% for the unconditional scenario, 2.0% for the slow, 2.5% for the moderate, and 3.2% for the rapid scenario—this compared to a 2025 baseline of 1.94%. For 2050, these expectations shift slightly to 2.5% (unconditional), 2.0% (slow), 3.0% (moderate), and 4.0% (rapid). AI experts, however, are significantly more

optimistic about the rapid scenario in 2050, forecasting a 5.0% annualized change in 2050.

Unemployment Rate. Economists expect the overall unemployment rate to remain relatively stable, even under the rapid scenario. For 2030, they forecast 5.0% for the unconditional, 5.0% for the slow, 5.0% for the moderate, and 6.0% for the rapid. For 2050, their estimates are 5.0% (unconditional), 5.0% (slow), 6.0% (moderate), and 6.0% (rapid). AI policy and industry professionals, however, project higher long-term unemployment under the rapid scenario: for 2050, they forecast 8.0%.

Youth Unemployment Rate. Even for 20-to-24-year-old workers, economists expect the unemployment rate to remain relatively stable. They forecast 2030 youth unemployment rates of 9.5% for the unconditional scenario, 9.0% for the slow, 10.0% for the moderate, and 11.0% for the rapid—predictions that fall well within the historical range. By 2050, their forecasts drop slightly across the board to 9.0% (unconditional), 8.2% (slow), 9.8% (moderate), and 10.0% (rapid). As with the overall employment rate, AI policy and industry professionals were more pessimistic under the rapid scenario, predicting 11.4% for 2050.

Labor Share. Economists project a slight downward trend in the share of economic output going to workers, particularly if AI advances quickly. For 2030, they forecast 54.3% for unconditional scenario, 55.0% for the slow, 54.0% for the moderate, and 52.0% for the rapid, compared to a 2025 baseline of 55.48%. For 2050, this drops to 50.0% (unconditional), 52.0% (slow), 50.0% (moderate), and 45.0% (rapid). AI experts, however, expect a drastic collapse in the labor share under the rapid 2050 scenario, forecasting just 40.0%.

Median Household Income. Economists forecast steady real income growth. For 2030, their median estimates are \$83,967 for the unconditional scenario, \$83,046 for the slow, \$85,000 for the moderate, and \$87,000 for the rapid, compared to a 2023 baseline of \$80,610. By 2050, their forecasts rise to \$93,175 (unconditional), \$91,864 (slow), \$95,142 (moderate), and \$100,000 (rapid). AI experts are markedly more pessimistic than economists in the rapid 2050 scenario, expecting a median household income of \$95,000.

Life Satisfaction. Economists do not expect AI progress to drastically alter average life satisfaction in the U.S. For 2030, forecasts on the Cantril ladder (Cantril, 1965) are 6.6 for the unconditional scenario, 6.7 for the slow, 6.6 for the moderate, and 6.5 for the rapid, compared to a 2024 baseline of 6.72. In 2050, forecasts are 6.7 (unconditional), 6.7 (slow), 6.5 (moderate), and 6.5 (rapid). Forecasts across all groups and scenarios are remarkably clustered, generally hovering between 6.4 and 6.8 with very little deviation. Superforecasters are an exception: for 2050, they forecast an average life satisfaction of 7.0, independent of the AI scenario.

Work Hours Assisted by Generative AI. Economists foresee significant adoption of generative AI. In 2030, they estimate the percentage of assisted work hours will be 10.1% for the unconditional scenario, 8.0% for the slow, 12.9% for the moderate, and 24.2% for the rapid, compared to a 2024 estimate of 3.35%. By 2050, they expect this to surge to 40.0%

(unconditional), 25.0% (slow), 44.7% (moderate), and 62.0% (rapid). AI experts also project high utilization in the rapid scenario, hitting 60.0% in 2050. By comparison, superforecasters are much more conservative, predicting just 33.0% in 2050 under the rapid scenario.

AI Electricity Consumption. Economists predict growing energy demands for AI. Their 2030 median forecasts for the share of U.S. electricity consumption used by AI are 4.0% for the unconditional scenario, 2.3% for the slow, 4.9% for the moderate, and 7.4% for the rapid, compared to a 2024 baseline estimate of 1%. For 2050, these jump to 8.0% (unconditional), 5.0% (slow), 8.3% (moderate), and 15.0% (rapid). AI experts and superforecasters anticipate somewhat higher electricity usage in 2050 under the rapid scenario, both forecasting 19.5%.

3.6 Policy Responses

We asked survey participants to predict the marginal impact of six policy proposals on GDP growth and the LFPR, under both the unconditional and rapid scenarios, and for 2030 and 2050. Participants were instructed to estimate the effect of each policy in isolation, setting aside any consideration of conditions, political or otherwise, under which each policy might be adopted. We also asked respondents to indicate their support for implementing each policy without conditioning on any specific AI progress scenario. Finally, we asked respondents to estimate the probability that each policy (or a similar policy, as judged by an economist panel) would be implemented by the U.S. by the end of 2026. We present a summarized version of each policy, along with our key findings, below. Findings reported are for impacts on GDP growth and the LFPR in 2030. The full policy descriptions, results, and rationale analyses can be found in Appendix D.3.

1. Retraining Support: *Offers displaced workers in high-automation-risk industries up to \$25,000 per year (for up to two years) in training credits, career counseling, and relocation assistance, funded by a small payroll tax.*
2. Modernized Unemployment Insurance: *Expands unemployment benefits to 75% of prior salary for up to 18 months for automation-displaced workers, with added wage loss insurance and streamlined administration, funded by higher employer payroll taxes.*
3. Universal Basic Income: *Gives every American adult \$1,000 per month unconditionally, funded by a 15% VAT on all goods and services.*
4. Manhattan Project for AI: *Deploys 0.4% of U.S. GDP annually¹³ in federal spending to accelerate AI research and infrastructure development, funded by a 0.7% VAT.*
5. Compute Tax: *Taxes heavy AI electricity users \$50 per MWh above a set threshold and redistributes the revenue to consumers as stimulus checks.*

¹³This would have corresponded to approximately 120 billion dollars in 2025. For context, Epoch estimates that the combined capital expenditures at Alphabet, Amazon, Meta, Microsoft, and Oracle in 2025 were 448.2 billion dollars: <https://epoch.ai/data-insights/hyperscaler-capex-trend>

6. Job Guarantee Program: *Guarantees a federally funded job paying at least \$15 per hour (indexed to inflation) with full benefits to any adult who wants one, funded by a 0.5% VAT.*

Key Results

The Manhattan Project for AI had the highest projected GDP impact, and the Job Guarantee Program had the highest projected LFPR impact. A comparison of median marginal impacts across policies is shown in Figure 14 for GDP growth and in Figure 15 for the LFPR. However, economists expressed exceptionally low support for the Job Guarantee Program (13.7%), and only middling support for the Manhattan Project for AI (55.8%), suggesting that broader economic and societal considerations, beyond headline GDP and LFPR projections, are the primary drivers of economists’ policy preferences.

Retraining Support was the consensus favorite among economists. It drew 71.8% support and only 19.9% opposition (with the remainder registering as “unsure”), received modestly positive forecasts on both GDP (+0.1 p.p.) and LFPR (+0.5 p.p.) for the unconditional scenario, and attracted little concern about its funding mechanism—a small payroll tax. Retraining support was also broadly popular with AI industry and policy professionals, superforecasters, and the general public. The fraction of participants in each group supporting the implementation of the different policies is shown in Figure 16.

Support for the six policies varied significantly between groups. The Job Guarantee Program produced the survey’s largest economist-general public divergence: whereas only 13.7% of economists supported it, 57.1% of the general public did. In general, economists’ normative support tilted toward more incremental, targeted policies, while the general public was willing to entertain broader interventions.

Universal Basic Income was predicted to have no impact on GDP across both the unconditional and rapid scenarios, and carried a substantial projected LFPR decline for both scenarios. A plurality of economists (38.2%) opposed it—the second-highest opposition rate for any policy in the survey—citing labor supply disincentives and the drag of a 15% VAT as primary concerns; by contrast, a plurality of the general public (47.9% for versus 30.3% opposed) supported it.

Modernized Unemployment Insurance was forecast to have zero impact on both GDP and LFPR. However, it still drew a robust 62.3% support from economists, with some viewing it as a social stabilizer independent of its predicted macroeconomic effects.

All cohorts assigned low probabilities to real-world implementation of any policy by the end of 2026. Economists’ median estimates ranged from sub-1% (UBI, Job Guarantee) to 15.0% (Manhattan Project for AI), reflecting a pervasive view that near-term political obstacles are considerable, regardless of a policy’s perceived merit.

Many respondents in all cohorts view rapid AI progress as increasing the likelihood of policy enactment in the long term. This is particularly true for safety-net and redistributive policies like UBI, which some respondents—in their rationales—predicted would become inevitable under the rapid scenario. Others indirectly implied the rapid scenario would trigger stabilizing policy responses by arguing near-term action is unlikely *because* AI has not yet caused enough disruption.

Under the rapid scenario, economists expressed heightened uncertainty about whether labor market policies could keep pace with automation. Several noted that workforce interventions might be overwhelmed by structural displacement, and that policies designed for a normal economy could prove inadequate in a world where AI is advancing faster than humans can retrain.




Figure 14: *The median marginal impact of the policies on GDP growth in 2030 (5-year annualized change), according to economists.* Marginal impact is relative to none of the policies being implemented; the legend shows the predicted GDP growth conditional on this baseline. Several policies have zero predicted impact relative to baseline. See Appendix H.5 for question details.

4. Drivers of Disagreement

4.1 Framing the Debate

Debate on the future economic impacts of AI can largely be reduced to two questions:

1. Will AI capabilities progress meaningfully, such that AI systems are capable of completing a large quantity of economically meaningful work?




Figure 15: *The median marginal impact of the policies on the LFPR in 2030, according to economists.* Marginal impact is relative to none of the policies being implemented; the legend shows the predicted LFPR conditional on this baseline. Several policies have zero predicted impact relative to baseline. See Appendix H.5 for question details.

2. If this progress in capabilities occurs, what will happen to important economic indicators?

In this section, we emphasize *disagreement* on these two questions. Specifically, we ask: is disagreement about the path for economic indicators driven by disagreement on AI capabilities progress, or disagreement on the effects of capabilities progress on economic indicators? Cunningham (2025) asks this question and concludes:

The disagreement is about the AI, not about the economics. The primary reason for the disagreement seems to be about the future rate of AI capabilities progress, not about the more directly economic questions such as (1) the current economic impact of AI; (2) the rate of diffusion & adoption over time; (3) the substitutability between AI-produced and human-produced services.

Below, we develop a quantitative approach and reach to a different conclusion. When forecasting the future path of key economic indicators, *disagreement centered on the economic impacts of AI capabilities progress*, rather than the degree to which progress will be made. The rationale analyses in Appendix D.2 dovetail with this conclusion. They suggest the factors economists weigh most heavily—historical base rates, adoption lags, demographics, policy responses, macroeconomic headwinds, and structural views on how economies absorb technology—are shaped more by priors about economics and institutions than by the specific AI capability scenario assumed. These rationales also support our assumption that respondents shared a broadly similar understanding of what each scenario entailed, although we discuss the limits of this assumption at the end of this section. We focus our discussion in this section on economists, but we include similar analyses for all other respondent groups in Appendix B.




Figure 16: *Proportion of respondents in each group supporting implementation of each policy.* Support reflects responses of 'Yes, with at most minor alterations' to the question 'Do you think [policy] should be implemented?' Respondents could also answer 'No' or 'Unsure.' Policy support was elicited without conditioning on any specific AI progress scenario. See Appendix [H.5](#) for question details.

4.2 Disagreement on Capabilities Progress

First, we ask how much disagreement there is on AI capabilities progress; see Table 1. While the average economist assigns a probability of 14.0% (median 10.0%) to the rapid scenario, the bottom quartile of economists gives the rapid scenario a probability of 6.7% or less, while the top quartile gives forecasts of 20.0% or higher. For the slow scenario, the interquartile range is 16.8% to 60.0%, compared to a mean of 38.6% (median 38.3%). Accordingly, average and median forecasts mask substantial disagreement about the progress of AI capabilities.

Table 1: AI Progress Scenario Probabilities: Economists

Group	Scenario	Mean	25th percentile	50th percentile	75th percentile
Economists	Slow	38.6	16.8	38.3	60.0
Economists	Moderate	47.4	31.5	50.0	64.5
Economists	Rapid	14.0	6.7	10.0	20.0

4.3 Total Variance in Possible Outcomes

Measuring total variance. The disagreement on AI capabilities progress documented in the previous section, while extensive, is not necessarily the primary driver of disagreement between forecasters on their unconditional forecasts of economic outcomes. In order to answer our question—“is disagreement about the path for economic indicators driven by disagreement on AI capabilities progress, or disagreement on the effects of capabilities progress on economic indicators?”—we now assess the *total* variance in possible outcomes.

On select questions, we ask respondents to express their uncertainty in the form of quantile forecasts for the unconditional case and the rapid scenario.¹⁴ We fit a distribution to each respondent’s forecasts to fully characterize their beliefs.¹⁵ Given the mean and variance for the unconditional and the rapid scenario distributions, as well as the forecasted probability of the rapid scenario, we can obtain the mean and variance for a bundled slow/moderate scenario.¹⁶ It is possible that the variance of the unconditional outcome distribution is too low—and the probability placed on the rapid scenario too high—to yield a coherent conditional distribution for the bundled slow/moderate scenario.¹⁷ The proportion of incoherent distributions¹⁸ range from 11% to 21% across questions and time horizons for economists. Other groups, notably the

¹⁴Respondents provide forecasts for the 10th, 50th, and 90th percentiles.

¹⁵We fit a normal distribution to outcomes with unbounded support (GDP growth, although strictly speaking it is bounded from below by -100%), a gamma distribution to outcomes that are bounded only above or only below (median household income), and a beta distribution to variables that are bounded both above and below (LFPR, unemployment rate, wealth inequality).

¹⁶We express the unconditional distribution for the outcome as a mixture distribution over the scenarios. The sub-distributions are given by the distribution of the outcome conditional on the rapid or bundled slow/moderate scenario. We use a respondent’s scenario probabilities to define the mixture weights.

¹⁷In this case, the implied variance of the outcome conditional on the slow/moderate scenario will be negative.

¹⁸Incoherent distributions could reflect incoherence in beliefs or misspecification in the distribution fitting procedure.

general public, tend to have higher proportions of incoherent distributions. Superforecasters, on the other hand, tend to have lower rates of incoherence. We summarize the proportion of incoherent distributions in Table 8. The incoherent distributions were removed from this analysis.

With forecaster-level distributions for the slow/moderate scenario, rapid scenario, and unconditional case, we create ‘pooled’ distributions by taking a mixture distribution over forecasters, using our derived weights to govern selection probabilities across forecasters in the unconditional case.¹⁹ In effect, we are considering a hierarchical mixture: we first select a scenario according to the aggregated scenario probabilities. Within each scenario, we then select a forecaster. Lastly, we can then draw an outcome from this forecaster-scenario distribution. Accordingly, these distributions represent the full variation in expert views, rather than a distribution created with calibration in mind. Indeed, this pooled distribution is not necessarily the optimal way to aggregate forecasts in terms of forecasting accuracy. For example, Ranjan and Gneiting (2010) show that combining well-calibrated forecasts in this fashion yields forecasts that are miscalibrated. Similar results are reported by Lichtendahl, Grushka-Cockayne, and Winkler (2013). Nevertheless, approaches that aim to generate calibrated aggregate distributions will collapse disagreement between forecasters and are thus ill-suited to the questions posed above. In summary, our approach captures four sources of variation:

1. Within-scenario variance: how much does the outcome vary *within* a given scenario?
2. Between-scenario variance: how much does the *expected* outcome vary *between* scenarios?
3. Within-forecaster uncertainty: how uncertain is a given forecaster about the outcome?
4. Between-forecaster disagreement: how much do *expected* outcomes differ *between* forecasters?

We quantify these four sources of variance with a hierarchical law of total variance decomposition. We first decompose the total variation in the outcome into between- and within-scenario components, corresponding to the first two items above. Next, we further decompose each of these two components into within- and between-forecaster components, yielding the four components of the total variance. Lastly, on the within-scenario variance branch, we can assess how much the rapid and slow/moderate scenarios contribute to each component.²⁰ See Appendix B for the associated derivations.

Understanding total variance in GDP forecasts: raw data. We now describe in detail a decomposition of 2030 GDP growth forecasts, focusing only on economists. After this discussion, we summarize the results for the other questions. Results for all groups, questions, and time horizons are in Appendix B.

¹⁹We must accordingly adjust the weights in the rapid and slow/moderate scenarios. For example, a forecaster’s weight in the rapid conditional outcome distribution will be proportional to the product of their original weight and their forecasted probability for the rapid scenario.

²⁰The within-scenario variance is the probability-weighted sum of the variance in each scenario, so we can calculate directly each scenario’s contribution to within-scenario variance.

First, before imposing any assumptions, we can explore the quantile forecasts directly. The median of 50th percentile forecasts in the unconditional scenario is 2.5% (IQR: 2.0, 3.0). We can compare these forecasts to those for the three scenarios: the slow scenario median is 2.0% (IQR: 1.5, 2.2), the moderate scenario median is 2.6% (IQR: 2.2, 3.0), and the rapid scenario median is 3.3% (IQR: 2.9, 4.5)

Next, we can consider the quantile forecasts given in the unconditional and rapid scenarios. For the 10th percentile, the unconditional median is 1.0% (IQR: 0.0, 1.2) and the rapid scenario median is 1.2% (IQR: 0.5, 2.6). For the 90th percentile, the unconditional median is 4.0% (IQR: 3.5, 5.0) and the rapid scenario median is 5.5% (IQR: 4.0, 7.0).

While there are noticeable differences in the central tendencies across scenarios, the 10th and 90th percentile forecasts demonstrate that there remains substantial overlap in outcomes across the scenarios; *within-scenario variance is a critical driver of total variance*.

Nevertheless, the above decomposition does not assess how much of the variance within a scenario owes to forecasters' uncertainty, and how much to disagreement between forecasters. The variance decomposition below addresses this question directly.

Understanding variance in GDP forecasts: fitted distributions. We next assume functional forms for forecaster-level outcome distributions, as described above. We pool these distributions across forecasters to assess total variance in the unconditional and rapid scenarios. The result is shown in Figure 5. In the unconditional pooled distribution, the median growth rate is 2.5%, with an interquartile range of 1.63% to 3.44%. In the rapid scenario, the median growth rate is 3.38%, with an interquartile range of 2.34% to 4.72%. Again, we see notable overlap. The rapid scenario median sits between the 50th and 75th percentile of the unconditional distribution, despite economists assigning the rapid scenario a small 14.0% probability on average.

Understanding variance in GDP forecasts: implied distributions for the slow/moderate scenario and a decomposition. We next assume that the forecasts for the bundled slow/moderate scenario would be consistent with the forecasts for the rapid and unconditional scenarios. Again, this assumption yields the mean and variance for each forecaster's conditional outcome distribution in the bundled slow/moderate scenario.

Lastly, we explore the variance of these three distributions. The standard deviation of 2030 GDP growth in the unconditional case is 1.49%. Conditioning on the slow/moderate scenario reduces this by 9% to 1.36%. In contrast, conditioning on the rapid scenario increases the standard deviation by 33% to 1.98%. Thus, the rapid scenario not only raises expected GDP growth but also significantly increases uncertainty about this economic outcome. While moving to the slow/moderate scenario reduces the variance of the outcome distribution, substantial within-scenario variance remains.

4.4 Decomposition Results

Before presenting the decomposition results, we note that Appendix B.1 provides a detailed illustrative example using two hypothetical forecasters, demonstrating how the decomposition separates within-scenario from between-scenario variance and how within-forecaster uncer-

tainty can dominate between-forecaster disagreement even when forecasters hold different beliefs about AI progress.

We perform the decomposition, as described in Appendix B.1, and report the results for economists’ forecasts of GDP growth, labor force participation rate, and wealth inequality in Figure 17 below. Specifically for forecasts of GDP growth in 2030, within-scenario (WS) variance—comprising 94.9% of total variance—is the dominant driver of total variance, relative to between-scenario (BS) variance (5.1% of the total). On both paths, within-forecaster uncertainty comprises a larger share of the variance: 83% (78.7% / 94.9%) and 96% (4.9% / 5.1%) of the within- and between-scenario variance, respectively. Forecasters are quite uncertain about the future, and this noise is larger than the variance driven by disagreement.




Figure 17: *Decomposition of the variance in economists’ forecasts of GDP growth, labor force participation rate, and wealth inequality in 2030 and 2050.*

Lastly, we now consider how this analysis relates to the question of whether forecasters disagree on capabilities progress or outcomes conditional on capabilities progress.

1. If forecasters disagree noticeably on capabilities progress, we expect *between-scenario, between-forecaster* variance to be large.
2. If forecasters disagree noticeably on outcomes conditional on capabilities progress, we expect *within-scenario, between forecaster* variance to be large.

For GDP growth, the first component (0.3%) is small in absolute terms and relative to the second component (16.1%). This suggests that disagreement about outcomes conditional

on various levels of progress is a more important driver of total variance in 2030 GDP growth forecasts than disagreement on capabilities progress per se.

This analysis also quantifies the contributions of disagreement versus uncertainty to total variance in outcomes. The *within-forecaster* components all reflect uncertainty, which is outside of this question of *where disagreement* is most prevalent. Hence, we subtract out those components to isolate disagreement. Nevertheless, uncertainty dominates total variance in forecasts, explaining approximately 80% of total variance in Figure 17. Forecasters also highlight uncertainty in their rationales. One forecaster highlights uncertainties around adoption, driving greater variance in the rapid scenario:

Rapid capability progress creates a genuine upside tail via (i) higher TFP, (ii) faster innovation cycles, and (iii) automation of broader task bundles. But it also widens uncertainty because the same world can feature large transitional frictions and misallocation: costly retooling, brittle deployment in complex domains, and the possibility that measured GDP gains are damped by adjustment costs (even if capabilities are “there”).

Table 2: Variance Decomposition: Economists

Outcome	Year	Total Std. Dev.	WS	BS	WS-WF	WS-BF	BS-WF	BS-BF
GDP	2030	1.494	0.949	0.051	0.787	0.161	0.049	0.003
GDP	2050	2.166	0.947	0.053	0.753	0.194	0.050	0.003
LFPR	2030	3.978	0.968	0.032	0.852	0.116	0.030	0.002
LFPR	2050	6.731	0.963	0.037	0.567	0.397	0.035	0.002
Median HH Income	2030	10,568	0.999	0.001	0.928	0.070	0.001	0.000
Median HH Income	2050	16,327	1.000	0.000	0.679	0.321	0.000	0.000
Unemployment Rate	2030	2.534	0.963	0.037	0.810	0.153	0.034	0.002
Unemployment Rate	2050	3.192	0.990	0.010	0.737	0.253	0.009	0.001
Wealth Inequality	2030	3.726	0.939	0.061	0.801	0.138	0.056	0.005
Wealth Inequality	2050	6.424	0.981	0.019	0.790	0.190	0.018	0.001

Note: W = within, B = between, S = scenario, F = forecaster. GDP = Gross Domestic Product, LFPR = Labor Force Participation Rate, HH = household.

We conduct this analysis across all questions, time horizons, and groups. While we report more results in Appendix B, we report results for economists in Table 2. The same patterns hold as above. First, since within-scenario, between forecaster variance dominates between-scenario, between-forecaster variance, we conclude that disagreement on the impacts of AI is primarily driven by disagreement on the conditional impacts of various capability levels, rather than divergence view on AI capabilities progress—although we acknowledge that noise from discretization and scenario ambiguity, as discussed at the end of this section, may account for a small portion of the within-scenario variance. Second, within-forecaster uncertainty dwarfs between-forecaster disagreement. Regardless of realized AI capability progress, forecasters express substantial uncertainty and provide overlapping outcome distributions.

4.5 Limitations of the Decomposition

One caveat is that our scenario descriptions bundled multiple dimensions of AI capability—cognitive tasks such as research and coding alongside physical tasks such as robotics—into single composite scenarios. Respondents were asked to select which scenario, in sum, best represented their views, and were advised that progress might be uneven across domains. Therefore, two respondents who both selected the rapid scenario may have held meaningfully different assumptions about the specific capability profile underlying that label. If respondents interpreted the scenarios differently in this way, some of the within-scenario, between-forecaster variance that our decomposition attributes to disagreement about economic mechanisms could potentially reflect disagreement about the effective capability level assumed within a given scenario. This effect would correspondingly weaken the contrast with Cunningham (2025). In addition to ambiguity, our chosen discretization of progress could introduce some noise in the scenario forecasts. Two respondents might share identical scenario interpretations, but, conditional on rapid progress, one could expect progress to noticeably exceed the threshold between the moderate and rapid scenarios, while another expects progress to fall just above the threshold.

We acknowledge this limitation but believe that its effect was limited. Respondents’ written rationales suggest a broadly consistent pattern: forecasters who deviated from the literal scenario descriptions most commonly assumed faster cognitive AI progress and slower robotics progress, rather than adopting wildly divergent capability profiles. This consistency suggests that scenario ambiguity is unlikely to account for much of the within-scenario variance we observe.

5. Discussion

As we documented in our Section 4 variance decomposition, the primary source of disagreement among economists is likely not about whether AI capabilities will advance significantly—majorities assign meaningful probability to the moderate or rapid scenario—but about how quickly the economy can absorb these potentially transformative capabilities, and what absorption will lead to in terms of economic impacts. Specifically, economists who share similar views on the likelihood of rapid AI progress nevertheless diverge on the likely rate of diffusion, the extent to which new job creation will offset displacement, the degree to which lags will occur between adoption and productivity gains, and how institutional and regulatory responses will shape the transition. We also see this in the analysis in Appendix E.1: while economists who assign above-median probabilities to the rapid scenario tend to have slightly higher unconditional forecasts for GDP growth and lower forecasts for the LFPR, the outcome distributions of these above- and below-median participants overlap significantly.

However, despite disagreement on the magnitude of these effects, the majority of experts agree that their net direction will be to attenuate rather than accelerate AI’s impact on the economy. Indeed, even under the rapid scenario, where AI systems surpass human performance on most cognitive and physical tasks by 2030, experts do not forecast economic outcomes outside the range of historical experience. Instead, their written rationales point repeatedly to diffusion lags, infrastructure bottlenecks, political instability, and demographic

headwinds as mechanisms that will likely prevent even highly capable AI from producing unprecedented economic outcomes in the near term.

This finding stands in marked contrast to warnings, raised by some prominent voices in the AI industry, about rapid economic transformation. Our sample partially captured this view, especially in the AI expert group, which forecast a GDP growth rate of 3.7% in 2030 and 5.3% in 2050 under the rapid scenario. Their 90th percentile forecast is 6.5% in 2030, meaning that they assign a 10% probability to growth equal to or larger than this. These predictions are higher than the economists’ but lower than those of some in the industry (see Cunningham (2025)). One possible explanation is that forecasters in our study over-anchored to historical data, which were prominently displayed in the survey interface (see Figure 80). Additionally, some of the difference (though not all, as described above) could be due to different capability progress beliefs. Our sample and scenario discretization may not capture the most extreme end of rapid progress beliefs, such that even the forecasts conditional on the rapid scenario remain unrepresentative of those in the industry with the most extreme beliefs about capabilities. In Appendix C, we present some early evidence that beliefs about AI timelines may be speeding up across multiple groups: respondents in the Longitudinal Expert AI Panel (LEAP), which ran in February 2026, assigned higher probabilities to the rapid scenario than those in this study (October 2025–February 2026). Even so, we find little difference in economic outcome forecasts between these two samples.

While the aggregate forecasts we find in this study are at the more moderate end of the spectrum of discourse, we note that our survey experts express the most uncertainty under the rapid AI progress scenario where the stakes for policy design are the highest. Experts’ unconditional forecasts—the ones that reflect their actual all-things-considered beliefs—cluster around historical baselines, but their uncertainty under the rapid scenario widens significantly for the LFPR, GDP growth, and inequality. While experts only assign the rapid scenario a 14% probability of occurring, given the magnitude of potential consequences, that number is far from negligible. This matters because, if AI progress is slow, existing institutions making incremental adjustments to current policies may prove adequate. But if progress is rapid, the breadth of the outcome distribution implies that policymakers cannot simply plan for the median outcome; they must contend with tail risks, including the potential for a deep contraction in labor force participation.

On the question of which policies might best mitigate negative impacts from AI, our survey revealed a marked divergence between economists and the general public: economists favor targeted, incremental interventions such as retraining support and modernized unemployment insurance, while the general public expressed support for broader interventions such as job guarantees and universal basic income. This gap is disproportionate to the differences in economic forecasts between the two groups—the general public’s median unconditional GDP growth and LFPR projections only differ modestly from economists’—and yet the general public is nearly four times as likely as economists to support a job guarantee. If AI progress accelerates and labor market disruption becomes more visible, these underlying disagreements may become the central fault lines of policy debates.

References

- Abbott, Ryan and Bret Bogenschneider (2018). “Should Robots Pay Taxes? Tax Policy in the Age of Automation”. In: *Harvard Law & Policy Review* 12.1, pp. 145–175. DOI: [10.2139/ssrn.2932483](https://doi.org/10.2139/ssrn.2932483). URL: <https://heinonline.org/HOL/P?h=hein.journals/harlpolv12&i=147>.
- Acemoglu, Daron (2024). “The Simple Macroeconomics of AI”. In: *Economic Policy* 40.121, pp. 13–58. DOI: [10.1093/epolic/eiae042](https://doi.org/10.1093/epolic/eiae042).
- Aghion, Philippe and Simon Bunel (2024). *AI and Growth: Where Do We Stand?* Federal Reserve Bank of San Francisco. URL: <https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf>.
- Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones (2019). “Artificial Intelligence and Economic Growth”. In: *The Economics of Artificial Intelligence: An Agenda*. Ed. by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago: University of Chicago Press, pp. 237–290. ISBN: 9780226613338. DOI: [10.7208/chicago/9780226613475.003.0009](https://doi.org/10.7208/chicago/9780226613475.003.0009).
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press.
- (2022). *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press.
- Altman, Sam (June 2025). *The Gentle Singularity*. <https://blog.samaltman.com/the-gentle-singularity>. Blog post, published June 10, 2025. Accessed March 30, 2026.
- Amodei, Dario (Oct. 2024). *Machines of Loving Grace: How AI Could Transform the World for the Better*. Essay, [darioamodei.com](https://www.darioamodei.com/essay/machines-of-loving-grace). URL: <https://www.darioamodei.com/essay/machines-of-loving-grace>.
- Anthropic (2025). *Preparing for AI’s Economic Impact: Exploring Policy Responses*. URL: <https://www.anthropic.com/research/economic-policy-responses>.
- Arnon, Alexander (Sept. 2025). *The Projected Impact of Generative AI on Future Productivity Growth*. Tech. rep. Penn Wharton Budget Model. URL: <https://budgetmodel.wharton.upenn.edu/p/2025-09-08-the-projected-impact-of-generative-ai-on-future-productivity-growth>.
- Bastani, Spencer and Daniel Waldenström (2024). *AI, Automation and Taxation*. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4811796.
- Baumol, William J (2012). *The cost disease: Why computers get cheaper and health care doesn’t*. Yale university press.
- Beraja, Martin and Nathan Zorzi (2022). *Inefficient Automation*. w30154. National Bureau of Economic Research. DOI: [10.3386/w30154](https://doi.org/10.3386/w30154). URL: <https://www.nber.org/papers/w30154>.
- Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Stanford Digital Economy Lab. URL: <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>.
- Brynjolfsson, Erik, Anton Korinek, and Ajay Agrawal (2025). *A Research Agenda for the Economics of Transformative AI*. w34256. National Bureau of Economic Research. DOI: [10.3386/w34256](https://doi.org/10.3386/w34256). URL: <https://www.nber.org/papers/w34256>.

- Cantril, Hadley (1965). *The Pattern of Human Concerns*. New Brunswick, NJ: Rutgers University Press.
- Clark Center Forum (2025). *AI and Growth*. URL: <https://kentclarkcenter.org/surveys/ai-and-growth/>.
- Comin, Diego and Bart Hobijn (2010). “An Exploration of Technology Diffusion”. In: *American Economic Review* 100.5, pp. 2031–2059. DOI: [10.1257/aer.100.5.2031](https://doi.org/10.1257/aer.100.5.2031).
- Comunale, Mariarosaria and Andrea Manera (2024). *The Economic Impacts and the Regulation of AI: A Review of the Academic Literature and Policy Actions*. DOI: [10.5089/9798400271663.001](https://doi.org/10.5089/9798400271663.001). URL: <https://papers.ssrn.com/abstract=4774326>.
- Congressional Budget Office (Feb. 2026). *The Long-Term Budget Outlook Data: 2026 to 2056*. Tech. rep. Data available at: <https://www.cbo.gov/system/files/2026-02/57054-2026-02-LTB0-econ.xlsx>. Accessed: March 19, 2026. URL: <https://www.cbo.gov/publication/62044>.
- Crunchbase (2025). *Crunchbase*. URL: <https://www.crunchbase.com/> (visited on 04/08/2025).
- Cunningham, Tom (2025). *Forecasts of AI and Economic Growth*. URL: <https://tecunningham.github.io/posts/2025-10-19-forecasts-of-AI-growth.html>.
- Davidson, Tom (2021). *Could Advanced AI Drive Explosive Economic Growth?* <https://coefficientgiving.org/research/could-advanced-ai-drive-explosive-economic-growth/>. Coefficient, accessed April 1, 2026.
- Davis, Scott (Feb. 2026). *AI is Simultaneously Aiding and Replacing Workers, Wage Data Suggest*. Dallas Fed Economics. Accessed: 2026-03-26. Federal Reserve Bank of Dallas. URL: <https://www.dallasfed.org/research/economics/2026/0224>.
- Deloitte (Dec. 2025). *US Economic Forecast 2026-2030*. Tech. rep. Accessed: March 13, 2026. URL: <https://www.deloitte.com/us/en/insights/topics/economy/us-economic-forecast/united-states-outlook-analysis.html>.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock (2024). “GPTs are GPTs: Labor market impact potential of LLMs”. In: *Science* 384.6702, pp. 1306–1308.
- Epoch AI (2024). *Data on AI Models*. URL: <https://epoch.ai/data/ai-models> (visited on 11/22/2024).
- (2025). *AI Capabilities and Benchmarking*. Online database. Accessed: 2026-03-30. URL: <https://epoch.ai/benchmarks>.
- Erdil, Ege and Tamay Besiroglu (2024). *Explosive Growth from AI Automation: A Review of the Arguments*. DOI: [10.48550/arXiv.2309.11690](https://doi.org/10.48550/arXiv.2309.11690). URL: <http://arxiv.org/abs/2309.11690>.
- Federal Reserve (FOMC) (Dec. 2025). *Summary of Economic Projections*. Tech. rep. Accessed: March 13, 2026. URL: <https://www.federalreserve.gov/monetarypolicy/files/fomcprojtab120251210.pdf>.
- Filippucci, Francesco, Peter Gal, Karin Laengle, and Matthias Schief (2025). *Macroeconomic Productivity Gains from Artificial Intelligence in G7 Economies*. 41. OECD. DOI: [10.1787/a5319ab5-en](https://doi.org/10.1787/a5319ab5-en). URL: <https://doi.org/10.1787/a5319ab5-en>.
- Filippucci, Francesco, Peter Gal, and Matthias Schief (2024). *Miracle or Myth? Assessing the Macroeconomic Productivity Gains from Artificial Intelligence*. OECD. DOI: [10.1787/b524a072-en](https://doi.org/10.1787/b524a072-en). URL: https://www.oecd.org/en/publications/miracle-or-myth-assessing-the-macroeconomic-productivity-gains-from-artificial-intelligence_b524a072-en.html.

- Frey, Carl Benedikt and Michael A. Osborne (2017). “The Future of Employment: How Susceptible Are Jobs to Computerisation?” In: *Technological Forecasting and Social Change* 114, pp. 254–280. DOI: [10.1016/j.techfore.2016.08.019](https://doi.org/10.1016/j.techfore.2016.08.019).
- Gerut, Amanda (2025). *Jamie Dimon on AI: “People Should Stop Sticking Their Head in the Sand”*. URL: <https://fortune.com/2025/10/08/jamie-dimon-head-in-the-sand-ai-will-take-jobs-tip-iceberg/>.
- Gimbel, Martha, Molly Kinder, Joshua Kendall, and Maddie Lee (2025). *Evaluating the Impact of AI on the Labor Market: Current State of Affairs*. URL: <https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs>.
- Goldman Sachs (Oct. 2025). “What Is the US Economy’s Potential Growth Rate?” In: Accessed: March 13, 2026. URL: <https://www.goldmansachs.com/insights/articles/what-is-the-us-economys-potential-growth-rate>.
- Grace, Katja, Harlan Stewart, Julia Fabienne Sandkühler, Stephen Thomas, Ben Weinstein-Raun, Jan Brauner, and R. C. Korzekwa (2024). *Thousands of AI Authors on the Future of AI*. DOI: [10.48550/arXiv.2401.02843](https://doi.org/10.48550/arXiv.2401.02843). URL: <https://arxiv.org/abs/2401.02843>.
- Ham, John C, Julian Wright, and Ziqiu Ye (June 2025). “The Dramatic Rise of the New Society Journals in Economics”. In: *The Economic Journal* 136.673, pp. 259–278. ISSN: 0013-0133. DOI: [10.1093/ej/ueaf044](https://doi.org/10.1093/ej/ueaf044). eprint: <https://academic.oup.com/ej/article-pdf/136/673/259/63522232/ueaf044.pdf>. URL: <https://doi.org/10.1093/ej/ueaf044>.
- Hartman, Rachel, Aaron J. Moss, Shalom N. Jaffe, Cheskie Rosenzweig, Leib Litman, and Jonathan Robinson (2023). “Introducing Connect by CloudResearch: Advancing Online Participant Recruitment in the Digital Age”. In: DOI: [10.31234/osf.io/ksgyr](https://doi.org/10.31234/osf.io/ksgyr). URL: <https://doi.org/10.31234/osf.io/ksgyr>.
- Hendrycks, Dan, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt (2020). “Measuring Massive Multitask Language Understanding”. In: *arXiv*. DOI: [10.48550/arxiv.2009.03300](https://doi.org/10.48550/arxiv.2009.03300).
- Humlum, Anders and Emilie Vestergaard (May 2025). *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. Working Paper 33777. Revised March 2026. National Bureau of Economic Research. DOI: [10.3386/w33777](https://doi.org/10.3386/w33777). URL: <https://www.nber.org/papers/w33777>.
- Hyman, Benjamin G., Brian K. Kovak, and Adam Leive (2024). *Wage Insurance for Displaced Workers*. DOI: [10.3386/w32464](https://doi.org/10.3386/w32464). URL: <https://www.nber.org/papers/w32464>.
- IARPA (2011). *IARPA Aggregative Contingent Estimation (ACE) Program*. URL: <https://www.iarpa.gov/research-programs/ace>.
- International Monetary Fund (Oct. 2025). *World Economic Outlook*. Accessed: March 13, 2026. October 2025 WEO Entire Dataset in Excel, Countries Tab, Gross domestic product (GDP), Constant prices, Percent change variable for US. URL: <https://data.imf.org/en/datasets/IMF.RES:WEO>.
- Jain, Naman, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica (2024). “LiveCodeBench: Holistic and Contamination Free Evaluation of Large Language Models for Code”. In: *arXiv*. DOI: [10.48550/arxiv.2403.07974](https://doi.org/10.48550/arxiv.2403.07974).
- Jorgenson, Dale W and Kevin J Stiroh (2017). “Raising the speed limit: US economic growth in the information age”. In: *Knowledge economy, information technologies and growth*. Routledge, pp. 335–424.

- Kits, Remi Sebastian (Feb. 2, 2026). *Weightipy. A modern, high-performance RIM (iterative raking) weighting library for survey data*. Version 0.4.2. Documentation available at <https://kaitumisuuringute-keskus.github.io/Weightipy/>. URL: <https://github.com/kaitumisuuringute-keskus/Weightipy> (visited on 03/20/2026).
- Korinek, Anton (2024). *Economic Consequences of Frontier AI*. NBER Working Paper No. 32980. Cambridge, MA. DOI: [10.3386/w32980](https://doi.org/10.3386/w32980). URL: <https://www.nber.org/papers/w32980>.
- Korinek, Anton and Megan Juelfs (2022). *Preparing for the (Non-Existent?) Future of Work*. w30172. National Bureau of Economic Research. DOI: [10.3386/w30172](https://doi.org/10.3386/w30172). URL: <https://www.nber.org/papers/w30172>.
- Korinek, Anton and Joseph E. Stiglitz (2021). *Artificial Intelligence, Globalization, and Strategies for Economic Development*. w28453. National Bureau of Economic Research. DOI: [10.3386/w28453](https://doi.org/10.3386/w28453). URL: <https://www.nber.org/papers/w28453>.
- Lichtendahl, Kenneth C., Yael Grushka-Cockayne, and Robert L. Winkler (2013). “Is It Better to Average Probabilities or Quantiles?” In: *Management Science* 59.7, pp. 1594–1611. DOI: [10.1287/mnsc.1120.1667](https://doi.org/10.1287/mnsc.1120.1667).
- LMarena (2024). *Overview Leaderboard*. URL: <https://lmarena.ai/leaderboard> (visited on 11/21/2024).
- McGann, James G. (Jan. 2021). *2020 Global Go To Think Tank Index Report*. Tech. rep. 18. Think Tanks and Civil Societies Program, Lauder Institute, University of Pennsylvania. URL: <https://repository.upenn.edu/entities/publication/9f1730fa-da55-40bd-a1f4-1c2b2346b753>.
- Mellers, Barbara, Eric Stone, Pavel Atanasov, Nick Rohrbaugh, S. Emlen Metz, Lyle Ungar, Michael M. Bishop, Michael Horowitz, Ed Merkle, and Philip E. Tetlock (2015). “Identifying and Cultivating Superforecasters as a Method of Improving Probabilistic Predictions”. In: *Perspectives on Psychological Science* 10.3, pp. 267–281. DOI: [10.1177/1745691615577794](https://doi.org/10.1177/1745691615577794).
- Mellers, Barbara, Lyle Ungar, Jonathan Baron, Jaime Ramos, Burcu Gurcay, Katrina Fincher, Sydney E. Scott, Don Moore, Pavel Atanasov, Samuel A. Swift, Terry Murray, Eric Stone, and Philip E. Tetlock (2014). “Psychological Strategies for Winning a Geopolitical Forecasting Tournament”. In: *Psychological Science* 25.5, pp. 1106–1115. DOI: [10.1177/0956797614524255](https://doi.org/10.1177/0956797614524255).
- Murphy, Connacher, Josh Rosenberg, Jordan Canedy, Zach Jacobs, Nadja Flechner, Rhiannon Britt, Alexa Pan, Charlie Rogers-Smith, Dan Mayland, Cathy Buffington, Simas Kučinskas, Amanda Coston, Hannah Kerner, Emma Pierson, Reihaneh Rabbany, Matthew Salganik, Robert Seamans, Yu Su, Florian Tramèr, Tatsunori Hashimoto, Arvind Narayanan, Philip E. Tetlock, and Ezra Karger (2025). *The Longitudinal Expert AI Panel: Understanding Expert Views on AI Capabilities, Adoption, and Impact*. Working paper 5. Forecasting Research Institute. URL: <https://forecastingresearch.org/s/the-longitudinal-expert-ai-panel.pdf> (visited on 03/18/2026).
- National Center for O*NET Development (Feb. 21, 2024). *European Skills, Competences, Qualifications and Occupations (ESCO) to O*NET-SOC Crosswalk*. Accessed: March 19, 2026. U.S. Department of Labor, Employment and Training Administration. URL: <https://www.onetcenter.org/crosswalks.html>.

- OECD (2023). *OECD Employment Outlook 2023*. Report. Accessed: 2026-03-30. Organisation for Economic Co-operation and Development. URL: https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/07/oecd-employment-outlook-2023_904bcef3/08785bba-en.pdf.
- (Sept. 2025). *OECD global long-run economic scenarios: 2025 update*. Tech. rep. Accessed: March 13, 2026. URL: <https://www.oecd.org/en/topics/sub-issues/economic-outlook/long-run-economic-scenarios-2025-update.html>.
- Office of Management & Budget (OMB) (Sept. 2025). *Mid-Session Review*. Tech. rep. Accessed: March 13, 2026. URL: https://www.whitehouse.gov/wp-content/uploads/2025/09/MSR_2026.pdf.
- Ranjan, Roopesh and Tilmann Gneiting (2010). “Combining Probability Forecasts”. In: *Journal of the Royal Statistical Society: Series B* 72.1, pp. 71–91.
- RePEc (2025). *IDEAS*. URL: <https://ideas.repec.org/> (visited on 03/07/2025).
- Russakovsky, Olga, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei (2015). “ImageNet Large Scale Visual Recognition Challenge”. In: *International Journal of Computer Vision* 115, pp. 211–252. DOI: [10.1007/s11263-015-0816-y](https://doi.org/10.1007/s11263-015-0816-y).
- Survey of Professional Forecasters (Nov. 2025). *Fourth Quarter 2025 Survey of Professional Forecasters*. Tech. rep. Accessed: March 13, 2026. Federal Reserve Bank of Philadelphia. URL: <https://www.philadelphiafed.org/-/media/FRBP/Assets/Surveys-And-Data/survey-of-professional-forecasters/2025/spfQ425.pdf>.
- (Mar. 2026). *First Quarter 2026 Survey of Professional Forecasters*. Tech. rep. Accessed: March 13, 2026. Federal Reserve Bank of Philadelphia. URL: <https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/spf-q1-2026>.
- The Conference Board (Mar. 2026). *Global Forecast Update*. Tech. rep. Accessed: March 13, 2026. URL: <https://www.conference-board.org/publications/global-economic-forecast-update>.
- Trammell, Philip and Anton Korinek (2023). *Economic Growth under Transformative AI*. w31815. National Bureau of Economic Research. DOI: [10.3386/w31815](https://doi.org/10.3386/w31815). URL: <https://www.nber.org/papers/w31815>.
- U.S. Bureau of Labor Statistics (Apr. 2024a). *May 2023 National Occupational Employment and Wage Estimates*. National XLSX file. Accessed: March 19, 2026. URL: <https://www.bls.gov/oes/tables.htm>.
- (Apr. 2024b). *May 2023 OEWS Occupation Definitions [XLSX file]*. Occupational Employment and Wage Statistics (OEWS). Accessed: March 19, 2026. URL: https://www.bls.gov/oes/2023/may/occupation_definitions_m2023.xlsx.
- VandeHei, Jim and Mike Allen (May 2025). “Behind the Curtain: A white-collar bloodbath”. In: *Axios*. URL: <https://www.axios.com/2025/05/28/ai-jobs-white-collar-unemployment-anthropic> (visited on 03/30/2026).