February 21, 2025

THE AMERICAN ECONOMIC ASSOCIATION

Committee on Economic Statistics

Measuring the Economic Effects of Artificial Intelligence

Summary of an AEAStat Working Session held January 3, 2025, at the
2025 Allied Social Science Association meetings in San Francisco, CA

Background: It is widely expected that artificial intelligence (AI) will have significant effects on employment, wages, productivity, prices, and profits in the decades to come. Although representative data show U.S. businesses to be in early stages of adopting AI technologies for the production of goods and services, a burgeoning empirical literature using data of various types is examining empirical evidence as to AI's early effects. Research to date has used multiple types of data to gain insights into economic effects of AI, including firm-level data, large-scale representative business surveys, nontraditional labor-market data, and data on the task- and skill-composition of occupations. While use of multiple data types yields many valuable insights, it is also the case that, as Agarwal, Gans, and Goldfarb (2019: 16) have pointed out, characterizing broad-based effects of AI requires "consistent measures of AI, productivity, intangible capital, and growth across sectors, regions, and contexts."

Objective: To discuss priorities in advancing our ability to measure economic effects of AI, the AEA's Committee on Economic Statistics (AEAStat) organized a working session at the 2025 ASSA meetings in San Francisco. The session brought together 40 economists and other experts from academia, industry, and federal statistical agencies for a facilitated discussion of data innovations that will be needed to accurately measure the effects of AI on employment, wages, productivity, prices, and profits in the years to come. The working session focused on measuring economic effects of AI applications in the business context; individuals' non-work use of AI is also of clear interest but was beyond the scope of what could be covered in a 90-minute working session. The session was chaired by AEAStat Chair Karen Dynan and AEAStat member Erik Brynjolfsson and was conducted under the Chatham House rule.

AEA Contact: Martha Starr, PhD, Washington, DC, Representative of the AEA and Director of Government Relations.
martha.starr@aeapubs.org. 703-283-0054.

This document first briefly reviews types of data that research to date has used to measure economic effects of AI. It then summarizes some of the key issues that were discussed in the working session. In general, session participants identified a broad range of potential effects of AI that future research will want to be able to measure, including not only effects on employment, wages, productivity, prices, and profits, but also such varied issues as increased efficiency in innovation, changes in the functioning and stability of capital markets, and how effectively the market for post-secondary education is responding to changing demand for skills. Discussion of data innovations that are needed now to ensure that AI's economic effects can be accurately measured going forward included: refining how AI adoption and use are measured and finding ways to efficiently increase research access to firm-level and nontraditional labor-market data. Discussion highlighted the priority of continuing to collect high-quality, representative data on AI use in large-scale government business surveys, without which accurate measurement of AI's economy-wide effects will not be possible. It should not be inferred from this summary report that all participants agreed on points raised in the discussion.

Streets of San Francisco

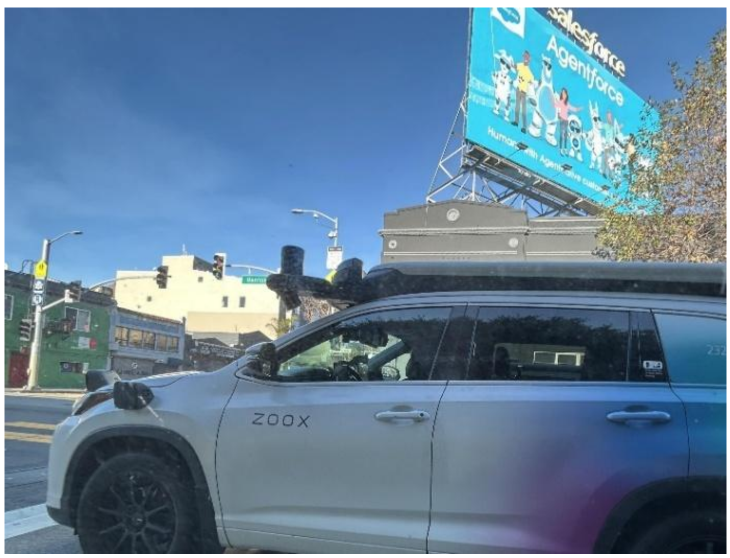

An Amazon Zoox fully autonomous vehicle (with safety driver), stopped under a sign for a new software-as-service product that puts friendly-looking AI agents to work with humans to "drive customer success together."

Background: Data types used for measuring economic effects of AI [1]

Firm-level data

Firm-level data have been used to study effects of AI on productivity in precise, context-relevant ways. For example, Kanazawa et al. (2022) examined effects of an AI app on the productivity of taxi drivers, measured as cruising time between fares. The app suggested routes along which demand was predicted to be high, which decreased cruising time of low-skill drivers without affecting that of high-skill drivers; as a result, average productivity rose, and the productivity gap between low- and high-skill drivers narrowed. Brynjolfsson et al. (2023) have a similar finding for customer-service workers. After a firm introduced a generative AI-based conversational assistant, the number of problems resolved per hour rose for novice and low-skill workers while that for experienced and high-skill counterparts minimally changed; again, average productivity rose. Toner-Rodgers (2024) had an opposite finding for material scientists. After a research firm introduced AI assistance, the number of new materials discovered by its scientists rose by 44%, but this was due to a doubling in output of top researchers; discoveries by the bottom third of scientists hardly changed.

Firm-level data also provide detailed insights into AI's potential for outperforming human judgment in specific contexts, or whether it is better used as a complement. Cao et al. (2024) built an AI analyst capable of beating 54% of year-end stock price predictions made by human financial analysts over the 2001-2016 period. The AI analyst was stronger when information was high-dimensional, transparent, and voluminous, while humans were competitive when institutional knowledge (such as the nature of intangible assets) was critical to prediction. Agarwal et al. (2023) compared radiology readings done by radiologists alone, an AI app, or radiologists with AI assistance. AI assistance did not improve humans' diagnostic quality on average, even though AI predictions were more accurate than approximately 75% of participants. An issue was that radiologists did not efficiently update their beliefs about the value of the AI prediction – which training could presumably correct.

Downsides of firm-level data are that it is hard to find companies willing to share data, and that data can rarely be shared with other researchers. Firm selection and generalizability are also concerns.[2] For example, if firms that eliminate jobs after adopting AI rarely share data, while those with more nuanced outcomes more often do, the body of firm-level research would suggest less potential for job displacement than there actually is.

Non-traditional labor-market data

Availability of non-traditional data on the labor market – such as online job postings, LinkedIn, resume data, and payroll-processing datasets like ADP – provide rich new opportunities to study AI effects on labor markets, as they cover huge numbers of jobs and businesses, and techniques like machine learning and large language models can be used to analyze text data.

---

[1] This section was circulated to participants before the working session. Valuable insights on trends in AI uptake and use have also been obtained from focused business surveys, like those of the McKinsey Global Institute (2017, 2018, 2019, 2020, 2021, 2022, 2023, 2024). Measurement issues relating to AI uptake and use are discussed in detail in a recent National Academics of Science, Engineering, and Medicine (NASEM) report on Artificial Intelligence and the Future of Work (2024), especially Chapter 6.

[2] Raj and Seamans (2018: 562).

Several studies of AI's economic effects have used data from Burning Glass (now part of Lightcast), which scrapes job postings from 51,000+ job boards and websites of U.S. companies. Acemoglu et al. (2022a) compiled online postings into an establishment-level data set that could be used to study AI effects on hiring. Among establishments initially high in AI-suitable tasks, availability of AI technologies increased hiring into AI-related positions and shifted demand towards certain types of skills – but it also decreased non-AI hiring and hiring overall. Babina et al. (2024) combined Burning Glass data with Cognism resumé data and firm-level Compustat data on sales, employment, and market value. Firms with higher AI investments had higher growth in sales, employment, and market value, which resulted from their higher levels of product innovation (as proxied by increased trademarks, product patents, and updates to product portfolios). This finding is consistent with AI reducing the costs of new product development, which can shift product demand to firms that use new AI-based opportunities to develop new and better products.

Downsides of non-traditional labor-market data are that they are proprietary (limiting access to them) and non-standardized (requiring work to convert the rich raw information to analyzable form). Non-traditional data also overrepresent certain types of businesses (e.g., larger and more digitized businesses) and jobs (e.g., medium-to-higher skill jobs). As AI adoption has so far been concentrated in these segments, the data valuably capture where the action is. But using the data to infer AI effects on the labor market overall would require benchmarking or at least caution in interpreting results.

Large, representative business surveys[3]

Since 2018, the U.S. Census Bureau has included AI-related questions on its Annual Business Survey (ABS), a large, nationally representative survey that covers all private, nonfarm sectors of the economy. The 2018 and 2021 surveys asked businesses about their use of five AI-related technologies to produce goods and services in the previous year.[4] The 2019 and 2023 surveys asked about businesses' use of AI as "part of their processes and methods" for producing goods and services in the previous three years. The 2020 survey asked about the extent to which businesses used AI in innovation in the previous three years, while the 2022 survey asked whether businesses had conducted or funded any R&D on AI in the previous year, and about the extent to which they had used AI in the previous year.

Census staff and academic coauthors have analyzed the 2018 ABS data (covering 850,000 firms) in Zolas et al. (2020) and McElheran et al. (2024), and the 2019 ABS data (covering 300,000 firms) in Acemoglu et al. (2022b, 2023); descriptive statistics for surveys through 2022 are available online. The data show that, at least through 2021 (the reference year for the 2022 ABS), AI use in the production of goods and services was uncommon among U.S. businesses generally – in the low single-digits – but much more common among larger firms. In 2021, for example, only 4.8% of all U.S. businesses reported using AI in the previous year – but the share among businesses with 25,000 employees or more was 44.2%.[5] As employment is concentrated in larger firms, the share of businesses using AI substantially understates the extent to which people work in businesses that are deploying AI.

[3] See Dinlersoz's summary on Census Bureau research on AI, presented at the December 2024 meeting of the Federal Economic Statistics Advisory Committee.

[4] The AI technologies were machine learning, natural language processing, voice recognition, machine vision, and automated guided vehicles.

[5] National Center for Science and Engineering Statistics, Annual Business Survey, Table CET-34: Companies using artificial intelligence, by company size: 2021.

Questions about AI have also been included in the Census Bureau's new Business Outlook and Trends Survey (BTOS), a biweekly experimental survey answered by 150,000+ businesses. In surveys fielded in December 2023-February 2024, businesses were asked about their use of 18 AI-related technologies.[6] Like the ABS, the BTOS data show single-digit AI use for U.S. businesses overall, but much higher use for larger firms.[7] While few AI-adopting businesses reported changes in employment due to AI, more than three quarters said a small number of tasks previously done by employees had been replaced by AI, with a similar share expecting to see this pattern in the next six months as well.[8]

Large federal business surveys have numerous advantages for measuring economic effects of AI, including representative samples, systematic and transparent collection and preparation of the data, and ongoing data collection; as such, they have unique potential for reliably and broadly measuring changes in employment, wages, and sales, as business use of AI diffuses. Researchers can also access confidential micro data files in the Federal Statistical Research Data Centers (FSRDCs) and match them to a treasure trove of other survey and administrative data, such as the Census Bureau's comprehensive Longitudinal Business Data (LBD) (see e.g., Acemoglu et al. 2023).

A downside of large, comprehensive business surveys like the ABS is timeliness, as finished data sets generally take a year or more to prepare after the data are collected. This is not the case for the less complex but also less comprehensive BTOS, which releases results about a month after data collection ends. Even if some ABS questions have been repeated, AI use has not been consistently measured over the 2018-2024 period. Adequate funding to support continued collection of data from large representative samples of U.S. businesses (along with other essential functions of federal statistical agencies) is also a concern.

Data on the task- and skill-composition of occupations

Research has called attention to changing sets of tasks and skills associated with different occupations as a key margin via which new technologies affect employment and productivity (NASEM 2024: 67). An important resource for empirical work has been the U.S. Department of Labor's Occupational Information Network (O*NET) database, which identifies tasks and abilities that make up the work of 19,000 occupations. Several studies have used the O*NET data to identify occupations with relatively high shares of tasks that could potentially be done by an AI application, including machine learning (Brynjolfsson et al. 2018), generative AI (Felten et al. 2023), and large language models (LLMs) (Eloundou et al. 2023). By matching their O*NET findings to occupational data from the American Community Survey, Felten et al. (2023) found that highly educated, highly paid, white-collar occupations may be especially exposed to generative AI. Eloundou et al. (2023) projected that around 80% of the U.S. workforce could have at least 10% of their work tasks affected by introduction of LLMs, with 19% of workers possibly having at least 50% of their tasks impacted.

---

[6] Machine learning, natural language processing, virtual agents or chatbots, speech/voice recognition using AI, AI-based recommendation systems, large language models, AI-based text analytics, AI-based data analytics, neural networks, augmented reality, AI-based decision-making systems, deep learning, image/pattern recognition, machine/computer vision, robotics process automation, biometrics, AI-based marketing automation, and "other."

[7] See Bonney et al. (2024a, 2024b).

[8] Census Bureau statistics on business formation can also be used to study start-ups intending to develop AI technologies or produce goods or services that use, integrate, or rely on AI. See Dinlersoz et al. (2024).

A downside of task-related approaches for studying economic effects of AI is that occupation data are generally collected in household surveys, which have limited information on people's employers, while business surveys generally do not collect occupation-level data.[9] New possibilities of linking micro data from household- and employer-level data could be helpful for tracing through employment effects of AI that operate via the task composition of work.

Summary of Working Session Discussion

The working session discussion was organized around three questions circulated in advance to session participants, concerning: key research questions on the economic effects of AI, data innovations that could be most valuable for answering these questions, and concrete steps that can be taken now to promote the creation of such data.

Question 1: What are the key research questions that empirical studies of economy-wide effects of AI adoption and use should address?

In general, session participants viewed three sets of issues as topping the list of broad-based economic effects of AI that economic research should investigate:

1. Potential for AI to cause significant changes in labor demand, wages, and/or job displacement
2. Potential for AI to cause broad-based gains in productivity and product innovation
3. Potential for AI to cause widespread changes in the skill- and task-composition of occupations

Specific research areas that were discussed with respect to the above issues and others are summarized below.

Effects of AI on innovation processes

- As digitization has created vast new supplies of information, we have extraordinary new opportunities to study the whole chain of changes that could result from AI-based innovations -- from the production of AI-based innovations (as shown in, e.g., AI-related patents, publications, start-ups, and the use of AI to innovate in other sectors); to the uptake of AI-based innovations by businesses (existing and startup); to effects of AI-based innovations on productivity within firms, including both physical measures like output per worker and rates of introduction of new products and improvements in product quality.
- The explosion of data availability and new tools to analyze it give us unprecedented chances to examine the time profile of AI's effects on aggregate labor productivity. Until well into the computer age, as Robert Solow famously remarked, computers could be seen everywhere except in the productivity statistics. If this "productivity paradox" repeats itself in the case of AI, better data will provide new opportunities to explain why major technological innovations would affect aggregate labor productivity with a lag.

[9] An exception is the BLS's Occupational Requirements Survey.

IO/markets/antitrust

- In relatively competitive product markets, technological change that decreases production costs gets passed through to consumers in the form of lower prices. But many of the markets that could benefit from AI-driven productivity gains are characterized by scale economies and production of complementary products (e.g., digital products intensive in design and development) – so robust product-market competition cannot be assumed. With competition expected to figure centrally in explaining how AI-driven cost savings get distributed between firms (as higher profits) and consumers (as lower prices), plans to measure variations in competition across product markets should figure into plans for measuring AI's economic effects.

Education and the changing demand for skills

- As educators, economists in academia should be tracking how AI use affects the skills and knowledge that students will need to successfully launch their work lives and thrive in an AI-enhanced workplace.
- There is strong demand across the board – by federal and state policymakers, higher-education decision-makers, media, and people generally – for real-time information on how AI-driven automation of work processes is changing demand for skills.
- A key policy question is how effectively the post-secondary education industry is responding to AI-driven labor-market changes. Are institutions that educate the bulk of the workforce – non-elite colleges and universities, community colleges, training institutes – adding courses, degrees, and certificate programs that will help students succeed in the changing work world (and what in fact are these)? Or are they just making curricular changes that look like they will help, without actually imparting the right skills?

Human-resource economics and human capital

- As AI is increasingly being used to manage HR functions like hiring and training and is expected to change the skill composition of jobs, it can be expected to significantly affect employer-employee relations. HR applications like employee-training chatbots could decrease job-training costs and improve training effectiveness; this could incentivize firms to retrain workers and shift them across jobs, rather than letting them go when functions they have been filling become obsolete. At the same time, AI automation of work processes could decrease the level of job-specific skills and knowledge that people need to become effective in new positions; this could again facilitate intrafirm mobility, but it could also decrease human capital accumulation, flatten lifecycle earnings curves, and/or decrease job satisfaction. It could be relatively hard for economists to study these types of AI effects, as they mostly do not involve a short-term response to an AI-adoption event. Finding ways to measure them is clearly important for understanding AI effects on trajectories of people's working lives.

Energy demand and environmental impacts

- The energy demands associated with training and running existing AI systems are known to be substantial. If the energy efficiency of AI-related hardware and software systems cannot be improved by enough to keep energy consumption from surging, environmental consequences would represent a significant social cost of AI adoption that should be factored in when measuring its economic effects.

Capital markets

- The potential for AI to disrupt capital markets is just starting to be studied. To give examples of AI adoption in finance from a call for papers for an upcoming conference:
  - In corporate finance, AI has uses in capital allocation, risk management, financial management, and forecasting.
  - AI-driven trading and asset-management strategies are being adopted by hedge funds, mutual funds, private equity, and venture capital, to different degrees.
  - Banks and other credit providers are deploying AI for credit allocation, managing risk, and detecting fraud.
- Venture capital (VC) markets are going big into AI. The latest Pitchbook data show one-half of all VC funding going into AI projects in 2024:Q4 – almost double the share from a year before.
- Evaluating effects of AI on capital markets – including on market volatility, systemic financial risk, and allocation of capital across sectors – is of clear policy interest.

AI, growth, and public finance

- If AI has as much potential to accelerate productivity growth as some analysts believe, could it increase the size of the economy by enough to improve the sustainability of federal government's finances? How would fiscal implications of AI differ if income increases from productivity gains accrue more to capital than they do to labor?
- Conversely, could AI-driven job loss decrease tax revenues, raise social-insurance payouts, and worsen the deficit, at least in the medium term?

Question 2: What data innovations would be most valuable for answering these questions – how might research benefit from, for example, a new wave of "pin factory visits," a biennial panel survey of business CIOs, machine-learning analyses of all publicly traded firms' public-facing business documents, and/or new AI questions on representative business surveys that also collect data on firms' employment, sales, and wage bills?

Large-scale representative business data

- All the representative information we have on businesses' AI use comes from questions fielded on Census Bureau surveys. The idea of adding questions on advanced technologies to ongoing Census surveys first arose in discussions at a previous AEAStat working lunch, which eventually led to the ongoing, highly productive collaborations between Census, academic economists who study economic effects of technological change, and the National Science Foundation's National Center for Science and Engineering Statistics.
- Adequate federal budget support is critical to continuing this work. Fielding AI questions on large-scale, ongoing Census Bureau business surveys is a highly efficient way of collecting high-value data on AI adoption and use; it is a concern that budget restrictions or reductions could put this

fundamentally important means of measuring AI's broad-based economic effects at risk. Continued collaborations with academics also help ensure that Census collects the information most needed for this purpose.

- There is a trade-off between fielding AI questions on slower-moving annual surveys that collect a lot of complementary information (like the ABS) vs. on rapid-response short surveys (like the BTOS). While rapid movement in the AI frontier may seem to favor BTOS's very short delay between data collection and availability, it is also the case that key economic effects of AI – rising productivity, shifts in employment across sectors and occupations, changes in the skill composition of occupations, etc. – are expected to materialize over the medium- to longer-term. So for purposes of measuring broad-based economic effects of AI (vs. tracking trends in AI uptake), the advantages of more comprehensive data are worth the wait.

How "AI use" is measured

- Adding precision to measures of AI adoption and use is urgent.
  - With AI functionalities spreading so quickly and becoming embedded in so many goods and services, all businesses may soon say they "use AI." This makes it urgent to shift away from asking businesses whether they use AI via a yes/no variable only, as "yes" will include everything from firms that have substantially overhauled their business processes to exploit AI to firms whose only AI use is through AI-based features added to widely used digital products (e.g., search engines that return replies generated by Co-Pilot or AI Overview).
  - At one end of this spectrum, Census business surveys ask specifically about use of AI "for producing goods and services." This focus helpfully screens out passive, incidental uses (like search-engine returns). But it could also miss things like automation of core business functions (like customer service or the hiring process), which *could* be important for measuring AI's economic effects.
- Asking businesses about specific AI technologies (e.g., machine learning, virtual agents, etc.) adds information on how AI is being used and helps screen out incidental use from widely used digital products (e.g., search-engine returns). However, rapid advances in the AI frontier make it tricky to identify <u>which</u> technologies to ask about and <u>how</u> to ask about them.
  - AI technologies that have burst on the scene, as "generative AI" and ChatGPT 4.0 did in 2024, may seem like good things to ask about. But if a given technology quickly becomes ubiquitous, differences in its use between time t and time $t+s$ will not help identify its economic effects, as only unusual laggards will be in the non-adopting group. An example from the past is "internet use" by businesses, which rapidly went from minimal to near universal in the 1990s.[10]

[10] Thus, some research at the time distinguished between basic use of the internet (e.g., for email, web browsing, document sharing) and enhanced use (e.g., for e-commerce, maintaining a website, providing digital customer support, etc.). By the end of 2000, 89% of businesses with 100+ employees were basic internet users, but only 13% had enhanced uses. See Forman, Goldfarb, and Greenstein (2005: 399-401).

  - Ideally researchers want to trace out the whole upswing part of the technology-adoption curve, moving from no use, into early use, then into rapidly diffusing use, and converging to steady-state adoption.
  - Constant change and increased differentiation in AI terminology are also problems, as common ways of referring to specific AI technologies at time *t* may be outdated by time *t+s.* To collect information on AI use that is consistent over time, questions that could be used to track adoption should avoid jargon and ask about technology types rather than specific applications (e.g., "machine learning" vs. "deep learning," or "virtual agents" or "chatbots," vs. "sales bots" or "customer-support chatbots").
  - Conferring with industry experts and pretesting questions is essential for identifying technologies with relatively high probabilities of diffusing and potential to have important economic effects, as well as for developing question wordings that will track AI use consistently over time.[11]
- Time could be ripe to develop more robust taxonomies of AI use. The economist's classic "build vs. buy" distinction could be useful here:
  - *"Build"* – Has the firm developed AI-based innovations in-house and incorporated them into its own processes for producing goods and services and/or running core business functions?
    - Adopting AI via "building" is most characteristic of large and/or highly digitized firms. Adoption patterns outside this group are more fragmented, as they lack the in-house infrastructure and expertise needed to start to build.
    - Observing that a firm has started hiring large numbers of AI software engineers is a good indicator of launch of significant AI build.
  - *"Buy, via outsourced AI services"* – Has the firm hired outside experts specialized in creating AI-based business solutions, as a means of incorporating new AI functionalities into its processes for producing goods and services and/or running core business functions?
  - *"Buy, via goods that embed AI functionalities"* – Has the firm bought new or different software, machines, inputs, etc., that embed AI functionalities, with the intention of incorporating new AI functionalities into its processes for producing goods and services and/or running core business functions?
  - Lines between these categories are not crisp. For example, software-as-a-service (SaaS) products increasingly include low-code options that allow companies to develop and deploy their own AI agents (e.g., for sales coaching, customer-service calls, etc.). In effect, the firm can "build" by buying software that has most steps of a software-engineering project built in.[12]
  - Asking detailed questions about build/buy AI use is constrained by the knowledge of respondents within the firm who complete business surveys. But if important parts of AI

[11] See, e.g., Buffington, Miranda, and Seamans (2018) for discussion of the development of Census survey questions on robotics.
[12] Rosenbush (2024).

adoption come from buying rather than building, it will be important to capture uptake via this route to avoid missing upswings in use.

- Emerging AI use by individuals at work is also of clear interest, as when people use ChatGPT or Co-Pilot to find information more efficiently or draft routine emails or memos.[13] If productivity effects of individual use at work are notable, relative to what happens when firms overhaul business-level processes to incorporate AI functionalities, leaving individual use at work out of business use will misstate productivity effects.
- We have very little consistently measured, longitudinal information on AI adoption that could be used to measure its economic effects. What data and methods could be used to develop novel longitudinal measures of AI adoption and use that could be both pushed back and maintained going forward (perhaps even turning into an economic indicator)? Potentially relevant examples include the sectoral- and firm-level measures of AI uptake developed by Acemoglu et al. (2022a: S306) and Babina et al. (2024: 7), which use Burning Glass/Lightcast (and other) job-description data and go back to 2010.[14] Another resource could be verbatim answers to Current Population Survey questions about people's usual activities or duties at their jobs. Abundant relevant data sources and new ways of analyzing them could offer valuable scope for developing measures of extensive- and intensive-margin AI uptake, which are needed to study AI's economic effects but have not been collected to date.

Firm-level data

- Researchers can spend a lot of time trying to find firms willing to share their data and working out arrangements for data sharing that satisfy both parties; sometimes, at the last stage the firm decides not to proceed. Thus, even if benefits of using firm-level data to study AI effects are very high, costs of getting ahold of it are as well. In some people's experience, firms are becoming even less willing to share data with researchers than they used to be.
- A key issue is that, within the firm, it's boardroom-level people who sign off on data sharing, and their objective functions are fundamentally different from those of academic researchers (and other people within the firm who may be interested in sharing data, like a chief economist or research director).
    - Boardroom people care about iron-clad assurances that:
        - the firm's identity will not be disclosed
        - data access will be completely secure
        - the firm's intellectual property rights are completely protected.

[13] Bick et al. (2024) found one in nine workers reported using generative AI at work every day in an August 2024 (where generative AI was defined as "a type of artificial intelligence that creates text, images, audio, or video in response to prompts," with "ChatGPT, Gemini, and Midjourney" given as examples). See also McKinsey Global Institute (2023:3, 2024: 5-6).
[14] Note that these measures should be thought of as tracking the "build" aspect of AI adoption and use, as they are based on AI-specific hiring.

    - They also too easily imagine scenarios in which researchers' findings are traced back to the firm, to the firm's detriment (e.g., if findings could be construed as out-of-line with the firm's public financial reporting).
- As some decision-makers do agree to data-sharing, they must have been provided with sufficient assurances, and/or saw, if anything, possible upside to research findings getting traced back to the firm. Are there characteristics that successful data-sharing arrangements have that unsuccessful ones lack?
- In general, asymmetric incentives for data sharing seem to put researchers' access to firm-level data below a level that would be socially optimal for comprehensively measuring AI's economic effects. To reduce this problem, are there steps a research organization (like perhaps the AEA) could take to advance research access? For example:
    - Developing and publicizing guidance for data-sharing partnerships. Being able to point to guidance provided by an authoritative research organization could help researchers assure firms that it is possible and productive for them to share data with academic researchers, who appreciate and know how to accommodate their concerns.
    - More ambitiously, setting up a research center that could securely house data sets supplied by firms, with well-defined assurances about data security and provisions that researchers would have to meet to access and use data. This could again allay firms' concerns about data security and would also allow more than one research team to study a given firm's experience (ensuring more robust research). Securing funding and setting up this type of arrangement would be challenges.

Nontraditional labor-market data

- As with firm-level data, it can take a lot of time to hammer out data-sharing arrangements for nontraditional labor-market data, especially if they have not previously been used for research purposes. The rich and vast amounts of data amassed by firms like LinkedIn, Lightcast, Indeed, payroll processors, etc., can take a lot of work to wrangle into analyzable shape, so that even when data have been used for research purposes before, research teams newly using the data have to reinvent wheels that previous researchers already tackled.
- Are there ways to support creation of partially processed data sets that could be made available for research purposes in some way that takes into consideration the proprietary character of the data? For example, perhaps externally supported academic researchers could work with counterparts in the company to build data sets that subsequent researchers could use as starting points for new analyses. Again, this would require resources – but it would expedite progress towards measuring labor-market effects of AI by eliminating inefficiencies in the data-preparation stage of the research process.
- In general, research on economic effects of AI could be faster and more productive if barriers to data access decreased and start-up costs associated with using nontraditional data could be reduced.

Data from firms that supply AI goods and services and/or products that embed AI

- Firms that produce AI-based goods and services and/or widely used digital products with AI features (search engines, social media) clearly amass a lot of granular data on effects of adopting specific AI applications; some presentations in the slate of AI sessions organized for the 2025 ASSA meetings featured detailed insights in this respect. Granular data on use of AI apps that are widely diffusing could be a breaking frontier of economic research. An example is Peng et al.'s study of the availability of GitHub Co-Pilot (a generative-AI coding assistant) on the time it took programmers to complete a task; availability of GitHub decreased average time by 56%, with less experienced and older programmers and programmers who worked more hours per day benefiting especially.

Question 3: What steps can be taken now to ensure that, over the next 10 years, we collect data needed to measure AI's effects on employment, wages, productivity, and profits, in order to inform policies that will positively shape both AI's future course and related economic outcomes?

In addition to data innovations discussed above, the following steps that could be taken now were mentioned:

- As the federal budget could come under intense pressure in the years ahead, economists can help support continued collection of high-quality representative data on businesses' AI adoption and use by explaining its very high value and unique contributions to a broad audience, especially policymakers and the media.
- Ability to measure AI's economic effects depends in part on taxonomies used for categorizing employers by industry (NAICS codes) and workers by occupation (SOC codes) – yet these largely blur the emergence of AI-based economic activity.
  - In the current NAICS scheme (2022), only one code specifically refers to AI.[15] Many or most firms that produce goods and services that have important AI components will be categorized in NAICS codes that include firms that produce similar goods and services without an AI focus.[16]
  - In the current SOC scheme (2018), no code specifically refers to AI.[17] People whose work is strongly AI-related will be categorized in SOC codes that include people doing similar work without an AI focus (e.g., 15-1252 software developers, 15-1251 computer programmers, 15-1221 computer and information research scientists).
  - As both the NAICS and SOC codes are currently being revised,[18] researchers should review current codes and share recommendations with the statistical agencies as to code revisions that could be beneficially made to improve our ability to measure AI-driven shifts across industries

[15] 541715 Artificial intelligence R&D laboratories or services.
[16] e.g., 513210 Software Publishers, 541511 Custom Computer Programming Services, 541512 Custom Computer Systems Design Services.
[17] The SOC code 15-2050 Data Scientists lists machine learning and natural language processing among applications used in the occupation.
[18] Revision of the SOC codes is underway but not currently expected to be completed until 2028. Revision of the NAICS codes started recently with a new version currently expected for 2027.

and occupations over the periods when the new codes are expected to apply (2027-2031 for the 2027 NAICS revision, 2028-2037 for the 2028 SOC revision).[19]

- As mentioned, there is strong interest across the board – by federal and state policymakers, higher-education decision-makers, media, and people generally -- in getting reliable real-time information on changing demand for skills caused by AI-driven economic changes. The abundance of high-frequency nontraditional labor-market data now available could be used (e.g., Indeed, LinkedIn, ADP), perhaps in combination with data on the market for new graduates collected by colleges, community colleges, and certificate programs and by the NSF. Regularly updated, easily interpreted information on skills that are currently strongly demanded, vs. those for which supply exceeds demand, would be highly valued by people trying to navigate new labor-market uncertainties due to AI-driven shifts in labor demand, as well as to federal and state policymakers and education-system decision-makers trying to adjust education and training programs according to new skill-demand patterns. If such an information resource resulted in better decisions about education, training, job changes, geographic mobility, etc., it could help reduce economic costs of adjusting to the "AI shock" over the medium term.[20]
- There is also a considerable amount of state-level administrative data that could be used to study employment effects of AI, including data sets that are being compiled and standardized by the Coleridge Initiative. Linking state-level data with confidential federal statistical and administrative data via the FSRDCs could substantially improve our ability to study AI effects on gross and net flows of workers between jobs, employers, sectors, and geographic areas. The LEHD data infrastructure at Census offers a jobs frame that links employers and employees that can be used for this purpose.

[19] Note, however, that shares of firms producing AI-specific goods and services and shares of workers in AI-specific occupations remain small; for example, according to Lightcast data, only 1.6% of all job vacancies were AI-focused in 2023. Stanford Institute for Human-Centered Artificial Intelligence (2024:223).

[20] See, for example, NASEM panelist Tom Mitchell's live-streamed presentation on this topic at the 2025 ASSA meetings.