Center for Economic Studies Research Report: 2024

Research and Methodology Directorate

Issued December 2025

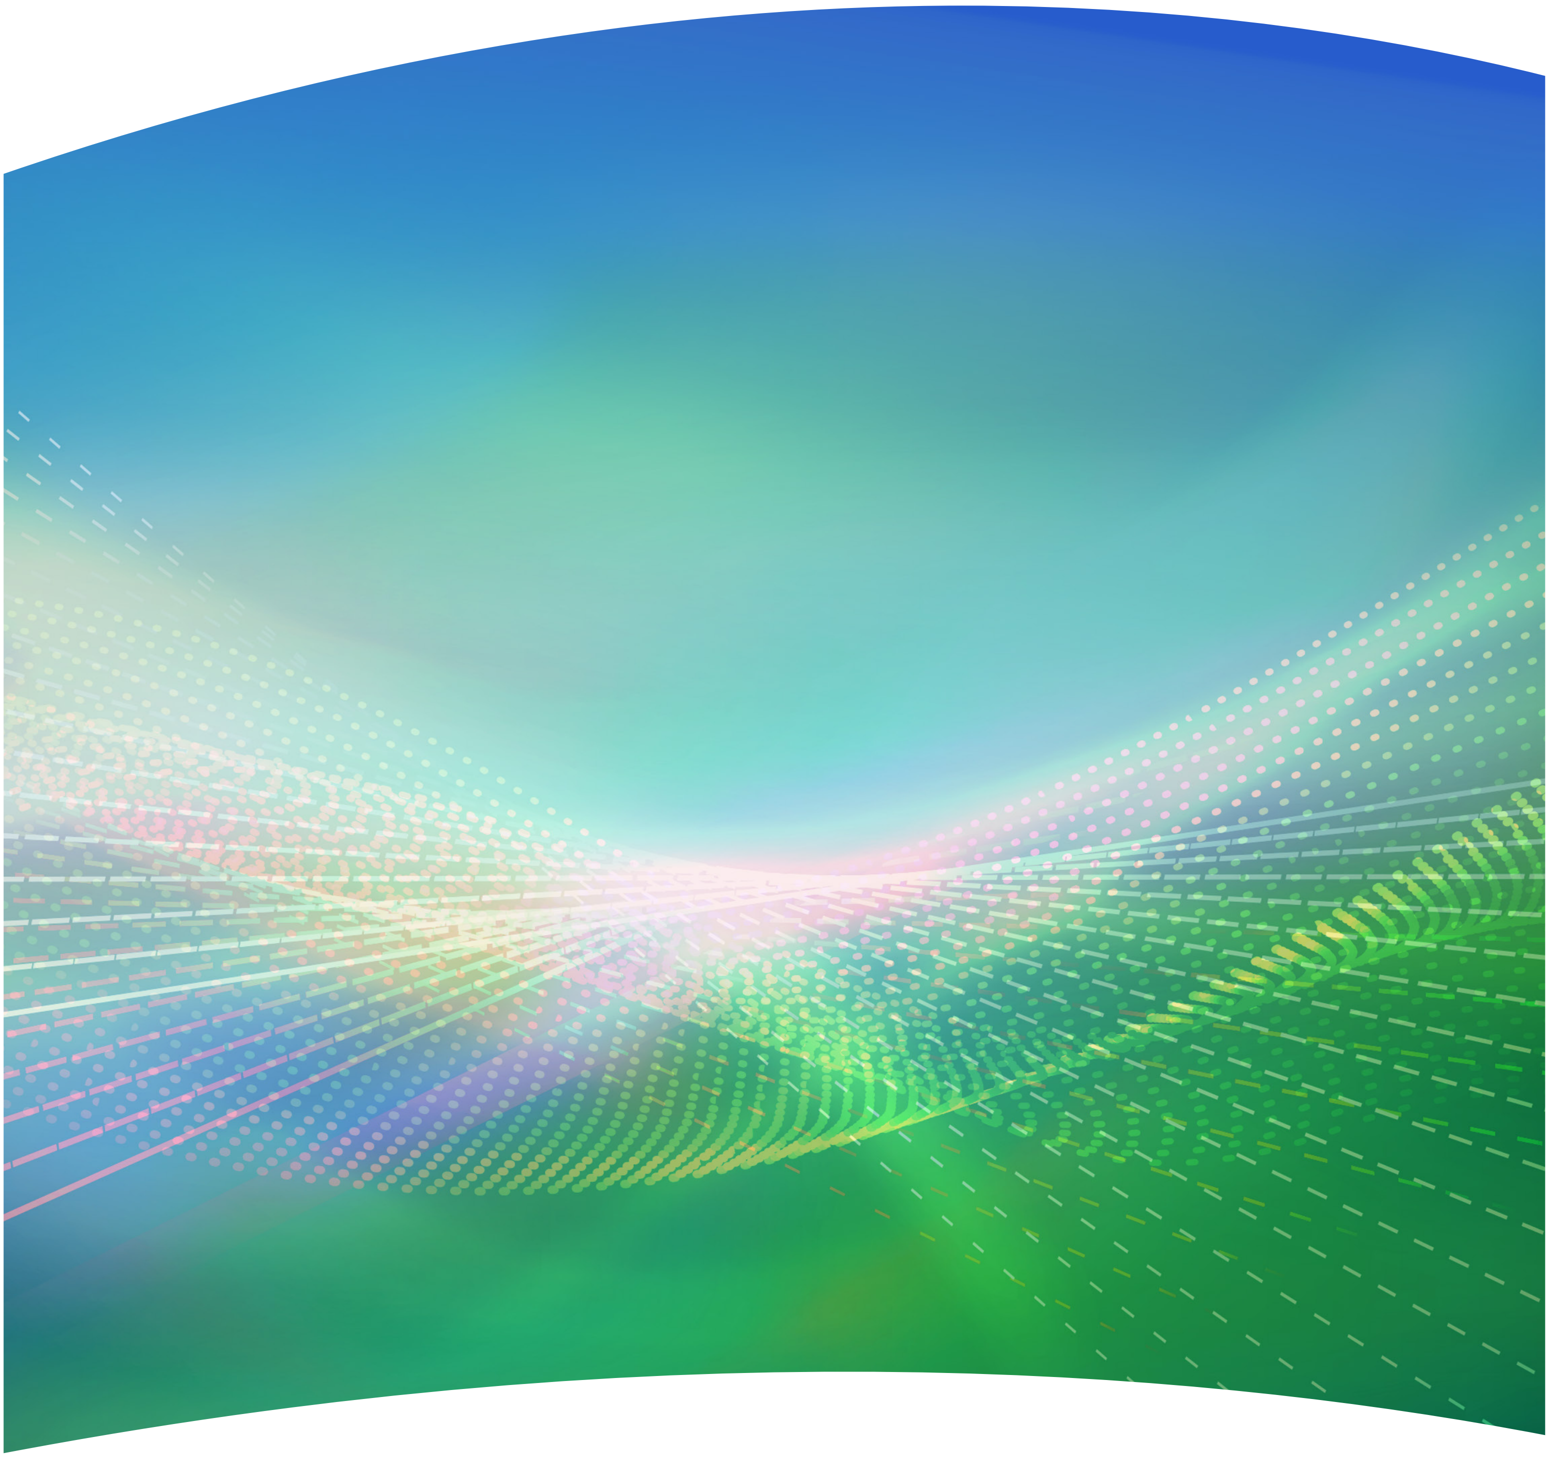

United States Census Bureau

U.S. Department of Commerce
U.S. CENSUS BUREAU
census.gov

Chapter 3.
Measuring Artificial Intelligence Use by Businesses

Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, and John Haltiwanger

INTRODUCTION

Artificial Intelligence (AI) has the potential to radically transform our economy; however, it is not clear how this technology will diffuse across businesses, what complementary investments are required for adoption, nor how employment and productivity at adopting (and nonadopting) businesses will be impacted. As part of the U.S. Census Bureau's mission "to serve as the nation's leading provider of quality data about its people and economy," it is critical we provide accurate and timely measures of this process. This chapter outlines the Center for Economic Studies' (CES) efforts to measure business use of AI through Census Bureau surveys, administrative data, and private data. It synthesizes roughly a dozen research papers detailing the development of these measures, the challenges faced, and lessons learned about the diffusion and impacts of AI. In doing so, this chapter also documents how Census Bureau research and development activities interact with production activities to improve economic measurement. In service to this dual focus on the results and the process used to get these results, this chapter is organized by measurement methodology (process) with a section near the end pulling together common themes (results).

This work has been undertaken by the CES technology team comprised of researchers interested in studying business technology use.[1] Much of our work is done in collaboration with our programmatic partners in the Economic Directorate (especially in the Economic Indicators Division and the Economic Reimbursable Surveys Division) who are responsible for the surveys and some of the products described in this chapter. We also leverage the expertise at other statistical agencies and academic institutions. We seek feedback and provide transparency through peer-reviewed journal publications, presentations, and blog posts (refer to the text box titled "Feedback and Commitment to Transparency").

Given the rapidly changing technology ecosystem, the CES technology team attempts to identify nascent technologies that should be considered for measurement. One such technology is AI, the subject of this chapter, but the team also has measurement projects related to robotics and other advanced technologies, as well as the space economy. There is also an extensive body of research by CES researchers outside the CES technology team who have examined a variety of questions related to technology adoption by businesses.[2] The chapter starts with surveys, then discusses administrative data, and ends with a discussion of private data in future work.

The Census Bureau has long collected information on business technology use through surveys such as the Annual Survey of Manufactures (ASM) and its supplements, the economic census (EC), the Annual Capital Expenditures Survey (ACES) and its supplements, and special surveys such as the Survey of Manufacturing Technology (SMT) (refer to the text box titled "Related Census Bureau Business Surveys"). Apart from the SMT, there has not been a dedicated Census Bureau business survey for measuring technology use. Instead, the CES technology team works closely with the Economic Directorate program areas to develop new content within existing survey instruments.

[1] The team, which extends beyond CES, has included at various times: David Beede, Catherine Buffington, Eric Childress, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Aditya Pande, and Nikolas Zolas. The papers using Business Trends and Outlook Survey include the following Economic Indicator Division experts: Kathryn Bonney, Cory Breaux, and Keith Savage. This chapter draws heavily from the work of all of these (and others).

[2] Interested readers can use the Census Research Exploration and Analysis Tool (CREAT) to search the CES Working Paper series for keywords such as "technology."

FEEDBACK AND COMMITMENT TO TRANSPARENCY

Getting feedback and providing transparency is critical for any scientific endeavor. The research and development activities described in this chapter have been documented via presentations to federal advisory committees, at academic conferences, and other conferences and through publications including Census Bureau blogs, working papers, and peer-reviewed journal articles. A few presentations and blogs are highlighted here; publications are cited in the text and appear in the reference section.

Presentations to Federal Advisory Committees

- Federal Economic Statistics Advisory Committee (FESAC)
    - o "Measuring AI Use by U.S. Businesses," December 13, 2024, Emin Dinlersoz.
    - o "Measuring Business Adoption and Use of Advanced Technologies, Artificial Intelligence, and Data," December 13, 2019, Lucia Foster.
    - o "Direct and Indirect Measures of the Economic Impact of the Digital Economy," December 15, 2017, Nathan Goldschlag.
- Census Scientific Advisory Committee (CSAC)
    - o "Results of New Measures of Technology: The Annual Business Survey," March 27, 2020, Nikolas Zolas.
    - o "Measuring Technology Use by U.S. Businesses," December 7, 2018, Cathy Buffington.

Presentations at Selected Conferences

- "Measuring Business Use of AI," Regional Disparities in AI Conference, NYU Stern School of Business, July 12, 2024, Lucia Foster.
- "Measuring the Diffusion of AI in the U.S.," MIT Computer Science & Artificial Intelligence Laboratory, November 2, 2023, Emin Dinlersoz.
- "Automation and the Workforce: A Firm-Level View From the 2019 Annual Business Survey," NBER Conference on Research in Income and Wealth, March 18, 2022, Daron Acemoglu.
- "Advanced Technologies Adoption and Use by U.S. Firms: Evidence From the Annual Business Surveys," Comparative Analysis of Enterprise Data conference, November 20, 2021, Nikolas Zolas.

Blogs

- "Is AI Use Increasing Among Small Businesses?" Emin Dinlersoz and Nathan Goldschlag, *Research Matters*, December 3, 2024.
- "How Many U.S. Businesses Use AI?" Cory Breaux and Emin Dinlersoz, *America Counts*, November 28, 2023.
- "Three Results From Recent Research on Advanced Technology Use and Automation," David N. Beede and Emin Dinlersoz, *Research Matters*, September 11, 2023.
- "New Technologies, Automation and Productivity Across U.S. Firms," Daron Acemoglu, Gary Anderson, David Beede, Catherine Buffington, Eric Childress, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas, *VoxEU*, August 7, 2023.

RELATED CENSUS BUREAU BUSINESS SURVEYS

Census Bureau business surveys have long collected information on capital and technology as an important input to production (along with labor, energy, materials, and services). This is especially true in the manufacturing sector where measurement issues are not as challenging as in other sectors and where information is collected at the establishment level. However, the economic census is an economy-wide survey and collects some related content; and starting in 1993, the Census Bureau started economy-wide collection of detailed capital expenditures at the firm-level.

Details are below.

Annual Survey of Manufacturers (ASM): Collected during intercensal years from 1949 through 2021 when it was integrated into the Annual Integrated Economic Survey (AIES) in 2023. In economic census years (those ending in “2” or “7”), the ASM was collected as part of the economic census covering the manufacturing sector. Capital information was collected via questions on “machinery and equipment” and “building and structures.”

Computer and Network Use Supplement (CNUS): The 1999 ASM included a supplement on computer and network use. The CNUS asked seven detailed questions about computer network and software use in the plant, one more open-ended question about networks, and two contact information questions (Atrostic and Nguyen, 2001).

Management and Organizational Practices Survey (MOPS): The 2010 and 2015 MOPS included questions on digitization and use of predictive analytics as a precursor to the more AI-focused questions on the 2021 MOPS described in this chapter (two questions in 2010 and six questions in 2015), as described in Brynjolfsson and McElheran (2016).

Industrial Robotic Equipment Experimental Product: The 2018–2021 ASM included three special questions on robotics: number of robots in use, robots purchased, and the amount spent on robotic equipment (Brynjolfsson et al., 2023). Published as an experimental series at <www.census.gov/library/publications/2024/econ/2021-asm-robotic-equipment.html>.

Survey of Manufacturing Technology (SMT): The SMT was conducted in 1988, 1991, and 1993 with partial funding from the Department of Defense. Collected information on use (or planned use) of 17 advanced technologies including computer aided design, pick and place robots, and programmable controllers (Dunne [1994] and Dinlersoz and Wolf [2018]). Results published in a series of reports.

Economic Census (EC): The economic census is conducted in years ending in “2” and “7.” In addition to the questions from ASM noted above, various EC collections have asked about technology. As an example, concerning technology and automation, the EC captured the introduction of self-service pumps at gas stations (Basker et al., 2017) and more recently the use of self-service in restaurants and self-checkout in supermarkets, drugstores/pharmacies, and general merchandise stores (Basker et al., 2019).

Annual Capital Expenditures Survey (ACES): Collected firm-level information on capital expenditures 1993–2022 and was integrated into the AIES in 2023. Refer to Becker et al. (2006) for a detailed discussion of the ACES. Robotics were added to the ACES in 2018. There was also an occasional supplement to the ACES (the Information & Communication Technology Survey).

This decentralized approach, developing technology use content within existing surveys (rather than having one dedicated survey), has distinct advantages and disadvantages. One benefit of the decentralized approach is the synergies between new and existing survey content that might not have otherwise been collected on a standalone survey. Examples of these synergies, described in greater detail below, include management practices, business innovation strategies, and expectations for firm performance. One cost of the decentralized approach is consistency is harder to achieve since definitions, scope, and sample characteristics may vary.[3] For this reason, we provide background on each data source before discussing the results from new technology content.

In proposing additional content on a survey, CES considers three broad metrics: appropriateness, consistency, and optimality. That is, whether the content is (1) appropriate based upon the Census Bureau's mission and our role in the larger federal statistical system, (2) consistent with the specific survey's structure and goals, and (3) optimal in terms of weighing the benefits of an additional collection to fill an information gap against the costs of additional respondent burden. Importantly, our focus on consistency ensures that newly developed content fits with the existing structure of the survey and, for surveys with rotating special topics, aligns additional technology questions with those special topics.

Before describing patterns of technology use, it is worth noting challenges common to any attempt to measure production inputs using surveys. The first challenge, which is particularly salient for measuring technology use, is whether the business has the information required to respond. A second challenge, specific to measuring technology, is determining which technologies warrant measurement. Survey collection is a costly investment. As such, we aim to focus on technologies that will be used broadly or have a significant impact on the economy. A cautionary example that highlights how hard it is to forecast which technologies will become important is the adoption of self-service technologies in supermarkets on an EC; while it seemed that this technology would become ubiquitous, at least one major supermarket chain leapfrogged over this technology while others de-adopted it (refer to the text box titled "Example of Challenges of Measuring Technological Change"). This approach also runs the risk of missing out on the measurement of a technology initially deemed not to be important but that then takes off. Another challenge is developing language to describe frontier technology that will be universally understood by respondents, both by businesses using the technologies and those that are not. Misunderstandings can lead to nonresponse that is not random and the reporting of incidental or insignificant use of the technology (in this case, AI) as "no use" or "unknown" (to the respondent). Finally, adapting measurement approaches becomes increasingly difficult as the speed of technology evolution increases, as has been the case in recent years for AI.

BUSINESS SURVEYS: ABS, MOPS, BTOS

Three business surveys have collected information about business use of AI: the Annual Business Survey (ABS), the Management and Organizational Practices Survey (MOPS), and the Business Trends and Outlook Survey (BTOS). We discuss the surveys in roughly chronological order based on AI content development, but there is overlap due to repeat collections for a given survey instrument. These repeat collections are critical for longitudinal analyses that allow for a deeper understanding of the causes and consequences of technology adoption. Importantly, the survey content evolves over time to capture changes in technology. For example, the MOPS captured successive waves of data-dependent technology starting with data-driven decision-making, then adding predictive analytics, and most recently adding AI. CES worked closely with the program areas and

[3] Refer to Seamans (2021) for additional discussions of the benefits of a standalone technology survey.

academic experts in developing the questions, analyzing the data, and interpreting the results. In addition to their own expertise, the program areas also ensure that the questions undergo cognitive testing and comply with all Office of Management and Budget requirements.

To preview an important theme of this chapter, there are important tradeoffs between different types of survey instruments. Large annual surveys provide more scope for detailed and complex questions. They also allow for measurement of other characteristics of firms within the same survey—an especially important feature for relating technology use to firm characteristics and outcomes—and tangible and intangible assets complementary to technology use not typically available in administrative data. However, they can be slower to collect and publish. Smaller, timelier surveys can measure rapid changes in adoption, providing almost real-time information on technology use. These require fewer and simpler survey questions so some subtleties may be missed. Given that many technology diffusion patterns follow an S-shaped curve of adoption, we need both large, complex surveys to provide less frequent detailed snapshots of use and more nimble surveys to capture the rapid adoption dynamics—especially at the early stages of the diffusion. Administrative data can supplement both goals, depending on the type of data being used, but provides less control over the object of measurement.

EXAMPLE OF CHALLENGES OF MEASURING TECHNOLOGICAL CHANGE

Automation can mean a machine entirely taking over tasks formerly accomplished by workers, but some tasks can only be partially automated and still require some human input. This is true in many service industries where manual tasks are resistant to full automation since they require either complex interactions or onsite production. Sometimes the adoption of a new technology allows the human input to be shifted away from paid employees to customers. This shift introduces mismeasurement in productivity because the shift is from a measured form of labor (employee) to an unmeasured form of labor (customer). Basker et al. (2017) study this phenomenon over the rise of self-service gas stations (and note as other examples, automatic teller machines and self-checkout at supermarkets) using successive economic censuses (specifically, the Census of Retail Trade) from 1977 to 1992.

Building on this experience, CES proposed adding questions to the 2017 Economic Census to capture the rising adoption of self-checkout technologies at supermarkets (and other retail and service industries) asking if the establishment offered “a dedicated self-checkout lane for customers,” with checkbox responses of “yes” and “no.” Basker et al. (2019) highlight the challenges that this collection faced in light of continually evolving technology and, in some cases, de-adoption. For supermarkets, this was highlighted on the one hand by the introduction of Amazon Go, which leapfrogged over self-checkout and, on the other hand, by the de-adoption of self-checkout by some supermarket chains and locations.

Annual Business Survey

The first collection of AI use by businesses occurred on the Annual Business Survey (ABS). The ABS is a firm-level survey representing all nonfarm private employer businesses, conducted in partnership with the National Center for Science and Engineering Statistics (NCSES) at the National Science Foundation (refer to text box titled “Partnering with Other Statistical Agencies”).

The ABS sample size is about 300,000 firms in most years, but about 850,000 in EC years (those ending in “2” and “7”). The core ABS questions collect detailed information on business owners and on research and development activities. In addition, the ABS hosts revolving special interest modules covering topics such as firm financing, technology, innovation, and globalization.

The CES technology team worked with two teams of experts to develop two distinct technology modules for the ABS. In building these teams, the CES technology team benefited from the topical working meeting hosted by the American Economic Association (AEA) Committee on Economic Statistics (AEAStat) that brought together economists working on technology and attending the 2018 annual AEA meeting. First, the technology team worked with Erik Brynjolfsson (Stanford University) and Kristina McElheran (University of Toronto) to develop three questions related to advanced technologies for the inaugural (2018) ABS. Second, the team worked with Gary Anderson (NCSES), Daron Acemoglu (MIT), and Pascual Restrepo (Yale University) to develop questions related to automation for the 2019 ABS. We discuss each in turn below.

PARTNERING WITH OTHER STATISTICAL AGENCIES

Part of the role of chief economist of the Census Bureau is to look for opportunities to collaborate with other statistical agencies to improve or develop new measures of the U.S. economy and its people (for example, Chapter 2 of the "2014 CES Annual Report" describes the BLS and Census Bureau collaboration producing measures of productivity dispersion).

In this case, the CES technology team contributed to the existing partnership between the Economic Reimbursable Surveys Division (ERD) and the National Center for Science and Engineering Statistics (NCSES) that enabled the development, collection, and dissemination of the Annual Business Survey (ABS) by working closely with both in developing the content for a module on Automation. The resulting paper includes NCSES researcher, Gary Anderson, as a coauthor.

NCSES also conducted their own research on the ABS technology modules, especially for the Digitization module, including these papers that Timothy Wojan wrote or cowrote:

- Han, Luyi, Timothy R. Wojan, and Stephan J. Goetz, "Experimenting in the Cloud: The Digital Divide's Impact on Innovation," *Telecommunications Policy*, 47, 2023, 102578.
- Wojan, Timothy R., "Digitalization, Cloud Computing, and Innovation in U.S. Businesses," NCSES Working Paper 22-213, 2022.
- Wojan, Timothy R., "The Geography of Self-Reported Innovation: Results From the 2017 Annual Business Survey," NCSES Working Paper 22-206, 2022b.

Technology and Intellectual Property Module: Digitization (2018)

The 2018 ABS included a module on "Technology and Intellectual Property" containing seven questions. The first four questions are on intellectual property and the last three questions pertain to technology use. The technology questions start by asking about the precursors or necessary complements to AI use. The ABS asks about data availability and computing power. The first question concerns the intensity of data use over seven business functions (including personnel, supply chain, and marketing), with check boxes as follows: the data is not collected, not used (none), used up to 50 percent, used more than 50 percent, used all the time, and don't know. The second question asks about the intensity of cloud computing use across nine possible functions including security, data analysis, and customer relationship management. The third question asks the intensity of use for ten advanced technologies; Zolas et al. (2020) identify five of these as AI-related technologies: automated guided vehicles, machine learning, machine vision, natural language processing, and voice recognition.[4] The reference

[4] Zolas et al. (2020) is a more general "first-look" at using the 2018 ABS to understand advanced technology adoption across firm size, age, and sector. McElheran et al. (2025), discussed below, focuses particularly on AI adoption using the same data.

period for these questions is 2017; these questions are repeated in the 2021 ABS (with reference period 2020).

In interpreting the statistics from the 2018 ABS, it is important to note that the questions involve "AI related technologies" perhaps reflecting broader use than narrower questions asked in later surveys discussed below (e.g., in the 2019 ABS and BTOS) that specifically ask about AI use. For example, early voice recognition used rule-based and other statistical methods not consistent with what would be considered AI today. On the other hand, even these earlier methods could be interpreted as precursors to the use of what is interpreted as AI today.

While this technology module is relatively brief, its content was developed in consideration with other elements especially the other modules. Of particular interest from the intellectual property module are questions about patents (pending and owned) and the business' assessment of the importance of intellectual property. The innovation module contains 16 questions; of particular relevance to the technology module are questions concerning growth-oriented innovation strategy, product innovation, and process innovation. The firm financing module contains 17 questions; of particular interest are questions on the reasons for owning the business and two questions on the source and size of capital used to start or acquire the business. The analysis utilized all of these questions in relating firm characteristics to AI use.

Results

McElheran et al. (2024) focus on the adoption and diffusion of AI-related technology, characteristics of firms that adopt these technologies, and identifying correlates of adoption of these technologies. They leverage the rich set of business owner characteristics (demographics and previous experience) and business characteristics (industry, location, innovation, intellectual property, and financing) available in the ABS. They further supplement this information with employment and revenue drawn from the Longitudinal Business Database (LBD), a longitudinal establishment-level panel of all non-farm employer businesses constructed using Census Bureau's Business Register (Chow et al., 2021). Combining these data sources allows them to capture the entire life cycle of firms.

McElheran et al. (2024) find fewer than 6 percent of firms use any of the five AI-related technologies in 2017.[5] However, use is concentrated in larger firms. They find the employment-weighted adoption rate, which measures how many workers are "exposed" to AI by employment at an adopting firm, is 18 percent. Adoption is more prevalent in sectors such as Manufacturing, Information, Health Care Services, and Professional and Technical Services.

While the overall adoption rate for AI is low, the impact of AI may still be significant if adoption is elevated among the economic actors that drive growth. To investigate this possibility, McElheran et al. (2024) conduct empirical exercises focused on young firms, which play a disproportionate role in driving reallocation, job creation, and business dynamism (Decker et al., 2016). They use the LBD to identify about 75,000 young firms (5 years and younger) in their sample. They find that, conditional on firm size and industry, young firms are more likely to use AI.

Using ABS characteristics, they find that AI-using young firms tend to have younger, more educated owners with prior business experience whose primary motivations are community-oriented or idea-oriented (rather than lifestyle-oriented, for example, work-life balance). In general, young AI users are more likely to be high-growth, to undertake process and product innovations, have a growth-oriented innovation strategy, obtain venture capital funding, and start with larger capitalization. They are also substantially more likely to patent. Lastly, McElheran et al. (2024) find that AI-use often occurs alongside other "enabling technologies"

[5] The studies of the 2018 ABS and 2019 ABS referenced in this chapter defined use as the percent of respondents who answered YES to questions as a fraction of those who answered YES or NO—explicitly excluding "Do Not Know" responses. This contrasts with the studies using BTOS where use is a percent of all respondents including those that said they "Do Not Know." A discussion below highlights the importance of this distinction.

such as digital collection and other "high-potential technologies" including robotics and cloud computing, consistent with the emphasis on complementarities in the general-purpose technology adoption literature.

The authors are careful to note that these are not causal relationships. Moreover, they end the paper with the observation that much more work needs to be done since even with these rich data, which include owner and firm characteristics, the empirical analyses can explain less than a quarter of the variation in AI use across businesses.

Technology and Intellectual Property Module: Automation (2019)

One of the most fundamental questions about AI, or any automation technology, is how it will impact employment. A large literature examines whether AI is a substitute or a complement for employment at adopting businesses and how nonadopting businesses may be impacted. Thus, the 2019 ABS Technology and Intellectual Property module focuses on automation, tasks, and employment. In contrast to the ABS "digitization" module, the ABS "automation" module explicitly included AI as one of the five covered automation technologies (the others are robotics, dedicated equipment, specialized software, and cloud computing). In our description below we focus on AI, but all questions were asked for each of the five technologies separately. The module includes a series of questions on firms' adoption and use of each technology and how adoption impacted the firms' demand for labor. (The module also included questions about the production of these technologies, as described in the text box titled "Businesses That Produce AI").

The first screener question asked about intensity of use of AI. One challenge in measuring business use of AI is that it can be embedded in processes or be incidental in nature, such as email response prompts. The checkbox responses provided a range from did not use, tested but not used, low use, moderate use, high use, and don't know. Respondents who choose did not use, tested, or don't know are told to skip to the last question focused on barriers to adoption.

After this screener question, respondents were asked about the effects of adopting AI on the number of workers employed by the business, the skill level of workers employed by the business, and the scientific, technological, engineering, and mathematical (STEM) skills of workers. The module included questions asking about the effects of adopting AI on the number of production workers, nonproduction workers, supervisory workers, and nonsupervisory workers. Each of these questions includes checkboxes for increased, decreased, did not change, or not applicable.

The module also asked questions about the motivation for, and challenges to, adoption of AI. The motivation checkbox responses include automate tasks performed by labor, upgrade outdated processes or methods, improve quality or reliability of processes or methods, expand the range of goods or services, adopt standards and accreditation, and other. The challenges for adoption include technology is too expensive, technology is not mature, lacked access to required data, required data not reliable, and lacked access to required human capital and talent. The final question concerns factors affecting technology adoption and usage. There are ten checkboxes and respondents can check all that apply; the factors include the technology was too expensive, was not mature, business lacked access to require data, required data was not reliable, lack of access to skilled workforce, laws and regulations, security concerns, lack of access to capital, technology was not applicable to the business, and no challenges to adopting the technology. The 2019 ABS questions are for reference period 2016–2018; many were repeated 4 years later in the 2023 ABS (for reference period 2020–2022).[6]

[6] In sum, CES helped develop questions for the ABS 2018/2021 and ABS 2019/2023. CES did not help develop or analyze the results of AI questions on ABS 2020 and 2022, so they are not discussed in this chapter.

BUSINESSES THAT PRODUCE AI

The focus of this chapter is research related to businesses that use AI. Here we briefly highlight some of the work that technology teams have done on the related research question concerning businesses that produce AI. As in the main text, some of this work is through survey collection and some is through the innovative use of administrative data.

At the request of NCSES researchers, the 2019 ABS automation module asks the full set of questions for both users and producers of advanced technologies. Acemoglu et al. (2025) find that 0.5 percent of firms are AI providers in 2016–2018 and that 2.2 percent of workers are employed at firms selling AI technology. The 2023 automation module does not include questions concerning producers of advanced technologies. We are not aware of research papers focusing on the results for AI producers, but NCSES regularly produces reports on this topic (for example, National Science Foundation [2024]).

Alderucci et al. (2019) identify businesses that produce AI by combining patent administrative data with inventor-firm ties from Longitudinal Employer-Household Dynamics (LEHD) data. Interestingly, they use AI techniques to identify AI patents rather than relying upon patent classification schemes or keyword searches as other researchers have done. This is important, they note, because if AI is a general purpose technology (GPT) these other methods would likely miss some AI producing activity (or could obfuscate it by lumping together AI and non-AI inventions). They build a large dataset of AI-related and non-AI-related patents to identify AI producers.

Once they have built and validated this dataset, they train a machine learning algorithm to identify AI-related patents. Their method yields 52,000 AI-related patents, which is considerably more than earlier studies found, such as the nearly 14,000 AI-related patents found by Cockburn et al. (2019). Linking this data to Census Bureau microdata allows them to examine outcomes such as growth and productivity of AI producing businesses using an event-study approach. They find that AI related innovations are positively associated with firm growth in terms of both employment (25 percent faster employment growth) and revenue (40 percent faster employment growth) relative to otherwise similar firms.

In a related vein, working with experts at Carnegie Mellon University, the CES technology team explored whether it would be possible to ask AI producing companies for information that would provide some sense of the breadth and depth of their business users. This would have provided an indirect method for measuring AI use by businesses. The CES technology team then worked with the Census Bureau programmatic experts for the ABS to determine whether this could be implemented; it was deemed infeasible at the time due to measurement, processing, and disclosure avoidance challenges.

Results

Acemoglu et al. (2025) find that 3.2 percent of firms report using AI. This is the second lowest usage rate among the five automation technologies, just above robotics at 2 percent. Consistent with McElheran et al. (2024), Acemoglu et al. (2025) find that AI use is concentrated in large firms—12.6 percent of workers are at firms that use AI. Along with relatively low adoption rates, they also find lower intensity use of AI (relative to the other automation technologies). One of the benefits of collecting information on multiple advanced technologies is that it enables us to assess complementarities between technologies. Consistent with the notion that computing power is a prerequisite for AI adoption, they find that 86 percent of firms that use AI also use cloud computing (and 90 percent use specialized software).

Adoption of AI across sectors is uneven relative to the other technologies (especially when employment-weighted). Not surprisingly, they find that the Information sector (NAICS 51) has the highest adoption rate, but it is still relatively low at less than 10 percent (about 30 percent when employment-weighted). In comparison, less than 5 percent of firms in Manufacturing (NAICS 31-33) adopted AI (about 20 percent when employment-weighted). The authors then look at dispersion in adoption by firm size and firm age (defining size and age distributions by 6-digit industry). They find that AI adoption rises with firm size; less than 3 percent of firms below the median firm size use AI, while more than 5 percent of the largest firms (top 1 percent) use AI. They interpret this as suggestive that adoption requires large integration costs so that scale is an advantage (and cite Acemoglu et al. [2023] for further evidence).[7] As in McElheran et al. (2024), they also find that for a given firm size, AI adoption tends to fall with age, as younger firms may face fewer organizational barriers to adoption. The most important limiting factor selected by firms that do not use AI is that the technology is not applicable to the business, with 49 percent of nonadopters choosing this (the second most important concern was the technology is too expensive, with 7 percent).

Task automation is an important motivator for businesses that adopt AI. Among firms that use AI, 54 percent report doing so to automate tasks. Using regression analyses, Acemoglu et al. (2025) find that use of advanced technologies is strongly associated with higher productivity. However, when focusing only on AI, they find a positive but statistically insignificant association. This weaker relationship is consistent with AI use still in its early stages; Brynjolfsson et al. (2021) highlight that at early stages of new technology adoption there can be a "J-curve" effect with productivity actually falling for new adopters. Most firms report no change in employment levels due to AI adoption, about 15 percent report an increase in employment, and less than 10 percent of firms report a decrease in employment. More than 40 percent of firms report rising demand for skills with AI adoption and this share is higher than for the other advanced technologies.

Management and Organizational Practices Survey (2021)

A recurring theme in the adoption of new technologies by businesses is the need for complementary changes and investments.[8] These changes can be production-oriented (e.g., reconfiguring the shop floor to handle industrial robotics), process-oriented (e.g., realigning monitoring systems to handle an increasingly remote workforce), or organizational (e.g., shifting towards a more skilled workforce). Thus, management practices may play an important role in the adoption of AI. Nicholas Bloom (Stanford University) and John Van Reenen (London School of Economics), who had experience developing and conducting management practices surveys, reached out to the Census Bureau to propose a partnership to collect this information on what became the MOPS. This partnership expanded, and after collaborating with academic experts Brynjolfsson and McElheran, the Census Bureau collected information on critical components required for AI use on the inaugural MOPS and expanded and refined this information on subsequent waves of data collection to keep pace with rapidly changing technology.

The MOPS is a supplement to the 2010, 2015, and 2021 Annual Survey of Manufactures (ASM) collections, covering a representative sample of about 30,000 manufacturing plants (Bloom et al., 2019). The ASM, with its detailed collection of inputs to production, is a natural testbed for beginning collection

[7] Acemoglu et al. (2023) dig deeper into the finding that AI (and robotics) adoption is concentrated in large firms, determining whether this reflects selection (large firms tend to adopt more) or causal effects (adopting firms become larger). Using the time series from the LBD with the ABS, they find the pattern is due to selection effects.

[8] Refer to Buffington et al. (2017) and McElheran et al. (2025) for a review of this literature.

related to management practices. The 2010 MOPS asks two questions about data availability and use in decision-making, building on the insights of Brynjolfsson and McAfee (2014) who note that relevant data and computing power are increasingly key assets for firms including for the development and use of AI. The 2015 MOPS expands this, adding six questions about data use, decision-making, and the use of predictive analytics (Buffington et al., 2017). Predictive analytics uses statistical models for forecasting and is a digital technology related to, but distinct from, AI. Using the MOPS data, Brynjolfsson and McElheran (2019) find 76 percent of manufacturing plants use predictive analytics and are more productive.

The 2021 MOPS asks AI-specific questions as part of a new section, “Data, Decision Making, and Artificial Intelligence,” which contains nine questions. The section retains questions about data availability, cloud services, and the two precursors to AI: descriptive analytics and predictive analytics. The section also includes four new questions specifically about AI. The first three are prompting questions helping respondents conceptualize AI use to focus their attention on AI use embedded in processes and products (McElheran et al., 2025). The first question focuses on how AI was used across six business functions: production scheduling and monitoring, quality control, environmental or safety compliance, equipment maintenance, supply chain management and logistics, and sales forecasting. Five check box responses are provided for each function capturing the intensity of AI use for each function: none, up to 50 percent, more than 50 percent, nearly all, and function not performed at the establishment. The next two questions gather information on the technical applications of AI (and retrospectively for 2019) with five types provided: machine vision, speech recognition, predictive maintenance, artificial intelligence-enabled industrial robots, and automated guided vehicles (AGV).

The final question asks about challenges businesses face in adopting AI. Respondents are asked to select all that apply from a list that includes: difficulty in identifying business use cases, uncertainty about government regulations or industry standards, level of AI expertise or skills at business, cost, employee attitudes towards AI, and none.

Results

McElheran et al. (2025) study the impact of business use of AI in industrial applications on productivity in manufacturing. They find that 65 percent of plants use predictive analytics and 23 percent of plants use any form of AI (this drops to 13 percent using a more restrictive definition of AI similar to that in the ABS). The most important impediment to adoption is the cost (43 percent), followed by identifying business use case (28 percent), business expertise on AI (12 percent), employee attitudes (9 percent), and uncertainty about industry standards and/or government regulation (1 percent).

They find that some structured management practices are correlated with AI use. Process-oriented structured management practices (setting targets and monitoring) are positively associated with AI adoption, which could reflect the importance of structured management practices for adopting AI or AI inducing businesses to adopt more structured management practices. Personnel structured management practices (decisions about bonuses and reassignments) are negatively associated with AI adoption. AI adoption is also associated with cloud computing and non-IT capital, but not with IT capital.

Evaluating the causal impact of AI use on productivity is complicated by the fact that more productive businesses may select into adoption. The MOPS and ASM data provide a rich set of variables that help control for selection bias. McElheran et al. (2025) find causal evidence of a “micro-level” productivity J-curve—adoption initially causes establishment productivity to fall before rising again. These J-curve patterns vary by age of the establishment, with older establishments facing more negative impacts.

Figure 3-1.
Collection Overview of BTOS

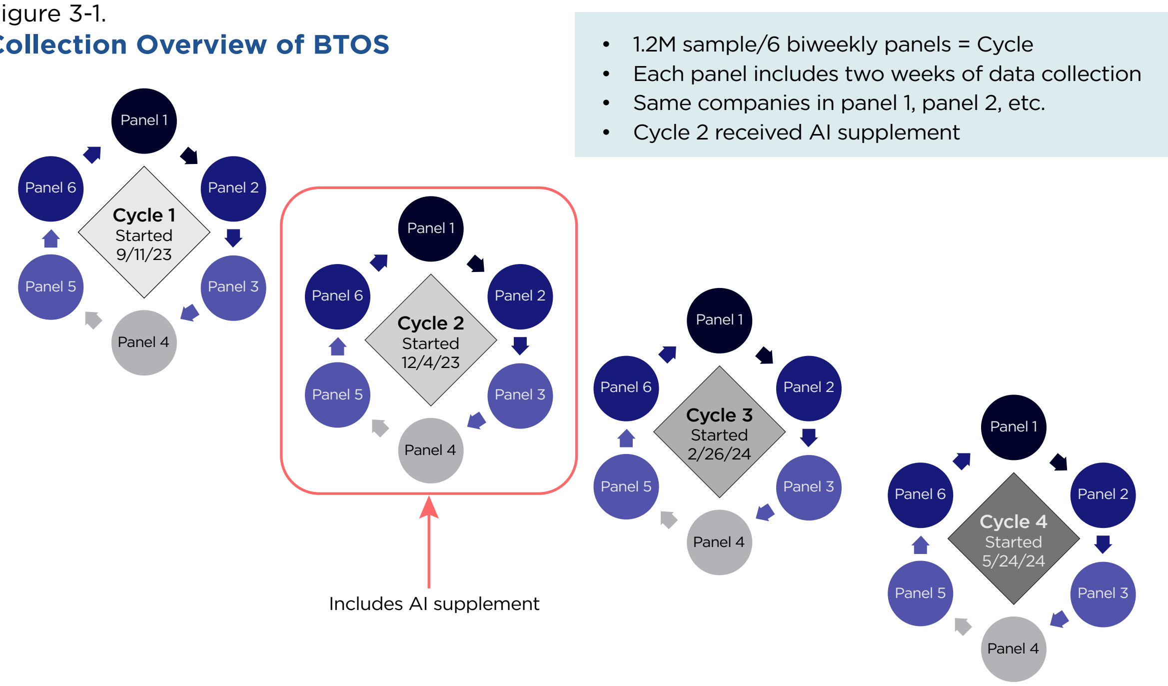

Business Trends and Outlook Survey

Shortly after the CES technology team completed its analyses of the ABS automation questions (Acemoglu et al., 2024), the first generation of large language models (LLMs) were released. Commentators predicted that these models would dramatically increase the diffusion and impact of AI (Hatzius et al., 2023; McKinsey, 2023). Due to the nature of annual collections, and duration of post-collection activities, the reference period of published ABS results are often 2 years lagged. To capture a real-time picture of how AI adoption was changing, the tech team partnered with the Economic Indicators Division (EID) to develop questions for the Business Trends and Outlook Survey (BTOS). Since AI use is still (at the start of the timing of the questions in BTOS on AI use) at early stages, the motivation has been to track the potentially rapid diffusion and adoption of the use of AI.

The BTOS is a biweekly qualitative survey intended to provide timely measures of current and near-term future business trends. The BTOS sample over an entire year covers about 1.2 million employer businesses. Each year there are four to five panels where supplementary questions can be added to a panel. Within each panel, as depicted in Figure 3-1, there are six cycles of collections with each cycle being sent to about 200,000 businesses. About 26 mostly qualitative questions are used to collect information about current and future trends in core economic phenomenon (Buffington et al., 2023). These questions concern current operations (overall performance, demand, revenue, input prices, and output prices) and current conditions (supply chain and hiring). Results from the BTOS collections are published within 2 weeks of collection at the national, sectoral, and state levels. Businesses with establishments operating in more than one NAICS sector are excluded from sectoral totals; they are instead tabulated as unclassified and counted only once in national totals with similar treatment for multilocation (state) firms, as discussed at <www.census.gov/hfp/btos/methodology>.

The BTOS complements the ABS as a firm-level survey, covering the private nonfarm economy but with a shorter dissemination cycle. Moreover, both surveys allow for the introduction of additional content as needed. Whereas the ABS refers to special collections as modules (e.g., the digitization and automation modules described above), the BTOS refers to special collections as supplements.

Two AI questions were added to the BTOS in its second year (2023) and retained in its third year (2024). Since these questions appear in all the panels and cycles, they are referred to as "core AI" questions. These core AI questions ask whether the business used AI in producing goods or services in the last 2 weeks and whether the business thinks they will be using AI in the next 6 months. This pair of questions is intended to provide a near real-time view into the current and expected diffusion of AI.

The CES technology team designed the BTOS supplementary AI questions based upon experiences working on the two ABS technology modules. For details on how these were developed and tested, refer to Bonney et al., 2024a. The supplementary AI questions were run over one panel (fall 2023–winter 2024). The supplementary questions can be grouped into four broad categories: types of AI technologies used; whether AI induced task, labor, and capital augmenting/replacing changes; organizational changes made to accommodate AI; and impediments to AI adoption. The first question in the supplement is intended to address the recurring measurement issue of a consistent understanding across respondents of a frontier technology. It asks whether in the last 6 months, the business used any of 17 detailed AI technologies/applications (for example, virtual agent or chatbots) or select "other."

Respondents who check any of the boxes except "none" are asked five follow-up questions about AI use during the last 6 months. The questions concern whether AI is used to perform tasks previously done by employees (yes, no, don't know), and if so, whether this represents a small, moderate, or large number of tasks, and whether AI impacted the business' total employment (increased, decreased, did not change). A similar question asks whether the business used AI to perform operations previously performed by existing equipment or software in producing goods and services (yes, no, don't know). Another question asks about other types of changes the business made to use AI (for example, trained current staff to use AI, purchased cloud services or cloud storage, and developed new workflows).

The supplement then asks the same set of questions but with a forward-looking time horizon (i.e., the next 6 months). The supplement ends with a question similar to the one in the 2021 MOPS about barriers to adoption. Respondents who selected that they do not intend to use AI in the next 6 months were asked to select all that apply from a list that includes: AI is too expensive, not mature enough, lack of knowledge about AI, concerns about privacy/security, concerns about bias, lack of skilled workforce, lack of required data, laws and regulations that restrict/prohibit use of AI, previous use of AI did not meet expectations, and other.

Figure 3-2.
Current and Expected Business Use of AI

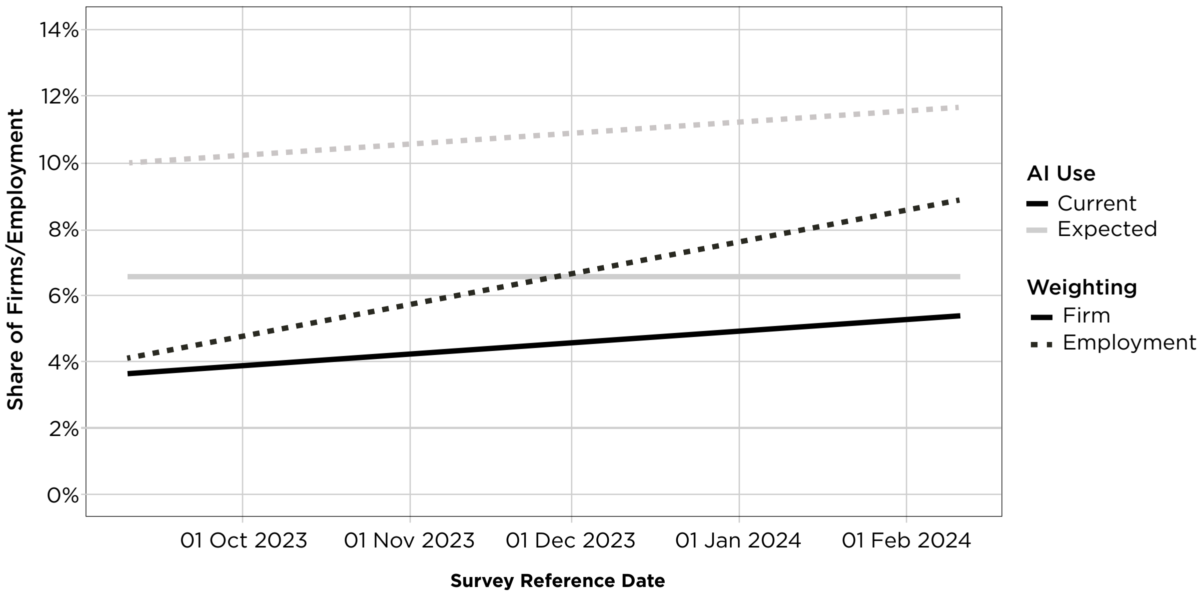

Note: First and last reference periods in the graph ended on September 10, 2023, and February 11, 2024, respectively. "Current AI Use" refers to a business responding "yes" to using AI in producing goods or services during the two-week period ending with the date indicated on the graph. "Expected AI Use" refers to a business responding "yes" to thinking it will be using AI during the 6 months following the date indicated on the graph. Firm-weighted shares should be interpreted as the share of U.S. firms using or expecting to use AI. Employment-weighted shares should be interpreted as the share of U.S. employment working at firms who use or expect to use AI. Trend lines shown on the graph are based on simple (OLS) linear regression.

Source: U.S. Census Bureau, Business Trends and Outlook Survey (BTOS) and Bonney et al. (2024a), Figure 1.

Results

Starting with the core AI questions, Bonney et al. (2024a) find, consistent with ABS results, relatively low AI use that rises over time. Figure 3-2 demonstrates current AI use and expected usage by firm and employment-weighted over the entire collection through February 2024 (Bonney et al., 2024a; Figure 1). The series have been smoothed to reveal general patterns using linear trends fitted through the use rates (Bonney et al., 2024a). Current usage (last 2 weeks) rose from 3.7 percent in September 2023 to 5.4 percent in February 2024. Expected future usage (next 6 months) rises from 6.3 percent to 6.6 percent over that same period.[9] The employment-weighted results, representing the fraction of workers at AI-using firms, are higher but with similar upward trends: current use rises from 4.5 percent to nearly 9 percent, and expected future use rises from 10 percent to 12 percent.

Not surprisingly, there is tremendous heterogeneity underlying the national numbers. Figure 3-3 demonstrates results by firm size pooled across the supplement collection (Bonney et al, 2024a; Figure 4a). In contrast with the ABS results, where usage rises roughly linearly in size, Bonney et al. (2024a) find a U-shape pattern of higher AI usage for the smallest and largest firms. This holds true for both current and future usage. Figure 3-3 plots current AI usage rates by firm size class for current AI use and expected AI use. The higher use of AI by small firms, relative to the ABS results, may reflect the fact

[9] Bonney et al. (2024a) report usage percentages based on those with positive responses to usage relative to all respondents while the earlier discussed ABS analyses exclude "Don't Know" from their calculations. If "Don't Know" is excluded from the BTOS calculations, the current use increases to 6.0 percent in September 2024, which is a near doubling of the 3.2 percent use rate in the 2019 ABS.

Figure 3-3.
Business Use of AI by Firm Size

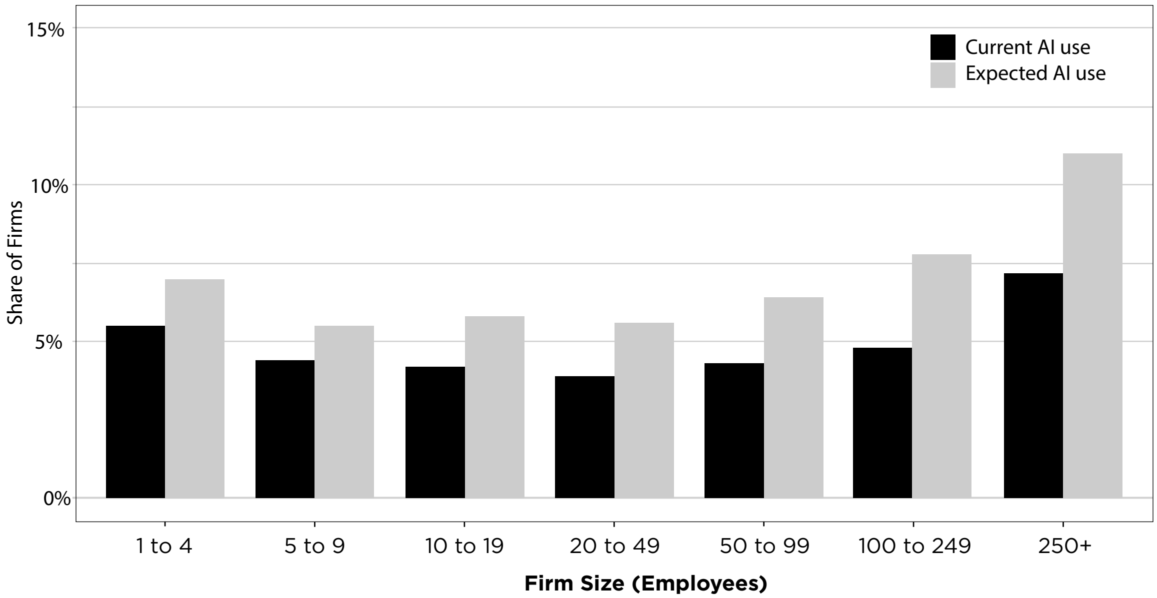

Note: Respondents pooled across 6 two-week panels, with the first and last reference periods ending on December 3, 2023, and February 11, 2024, respectively. "Current AI Use" refers to a business responding "yes" to using AI in producing goods or services during the 2-week period prior to data collection. "Expected AI Use" refers to a business responding "yes" to thinking it will be using AI during the 6 months following data collection. Firm-weighted shares should be interpreted as, within each size class, the share of U.S. firms using or expecting to use AI.

Source: Business Trends and Outlook Survey (BTOS) and Bonney et al. (2024a), Figure 4a.

that newer vintages of AI may be more amenable to use by small firms than prior vintages of AI that were highly reliant on large-scale data resources.

Figure 3-4 demonstrates AI use by sector for current usage and future expected usage pooled across the supplement collection (Bonney et al., 2024a); Figure 2a). Bonney et al. (2024a) find that the sectors with the highest rates of current and expected AI use are Information (NAICS 51) and Professional, Scientific, and Technical Services (NAICS 54). This is consistent with the ABS; however, unlike the ABS, the BTOS does not have high usage rates in the Manufacturing sector (NAICS 31-33) nor in the Health Care Services sector (NAICS 62).

It is possible to gain additional insights about business AI use through other core content on the BTOS. Bonney et al. (2024a) focus on current and expected future outcomes for overall performance and for employment. Correlating AI use with responses to other items on the BTOS, they find that firms using AI, or expecting to in the future, have both better current performance and better expected future performance. The differences are especially stark for expectations: 48 percent of AI users expect excellent or above average firm performance in the next 6 months as compared to 31 percent of non-AI using businesses. Given the interest in AI as either a substitute or complement for labor, the results for current and expected employment changes are of interest, but it must be noted that the changes captured are only at the extensive margin (increase, decrease, or no change). With this caveat in mind, they find 29 percent of AI using businesses expect an employment increase in the next 6 months as compared to 16 percent of non-AI using businesses.

Figure 3-4.
Business Use of AI by Sector

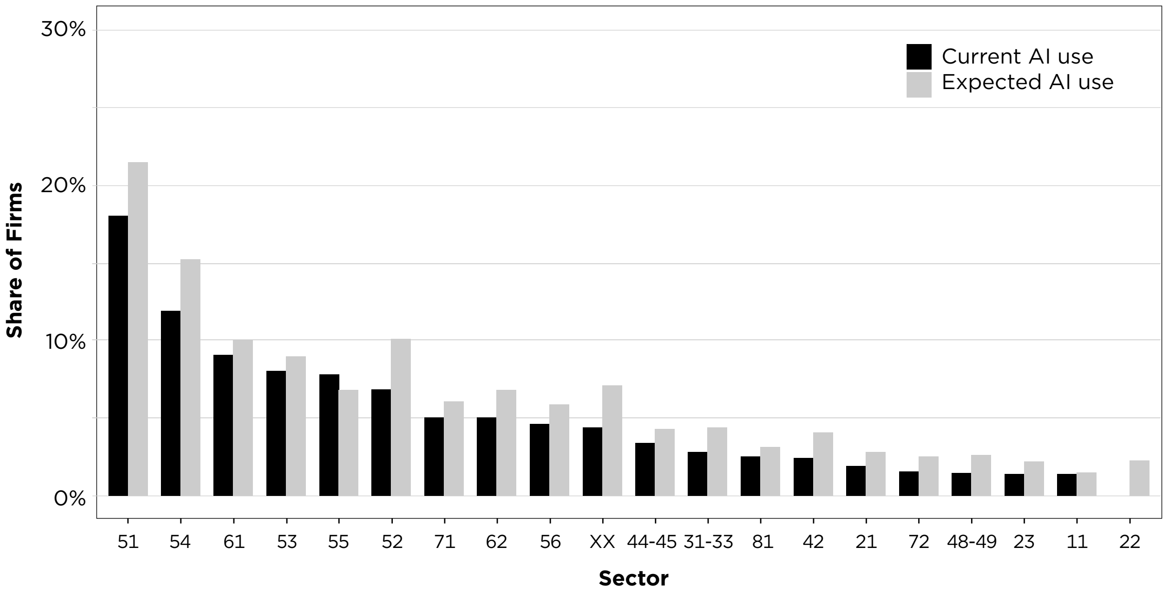

Note: Respondents pooled across 6 two-week panels, with the first and last reference periods ending on December 3, 2023, and February 11, 2024, respectively. "Current AI Use" refers to a business responding "yes" to using AI in producing goods or services during the two-week period prior to data collection. "Expected AI Use" refers to a business responding "yes" to thinking it will be using AI during the 6 months following data collection. Firm-weighted shares should be interpreted as, within each sector, the share of firms using or expecting to use AI. Sector codes: XX (firms with establishments in multiple sectors), 11 (Agriculture, Forestry, Fishing and Hunting), 21 (Mining, Quarrying, and Oil and Gas Extraction), 22 (Utilities), 23 (Construction), 31–33 (Manufacturing), 42 (Wholesale Trade), 44–45 (Retail Trade), 48–49 (Transportation and Warehousing), 51 (Information), 52 (Finance and Insurance), 53 (Real Estate and Rental and Leasing), 54 (Professional, Scientific, and Technical Services), 55 (Management of Companies and Enterprises), 56 (Administrative and Support and Waste Management and Remediation Services), 61 (Educational Services), 62 (Health Care and Social Assistance), 71 (Arts, Entertainment, and Recreation), 72 (Accommodation and Food Services), and 81 (Other Services—except Public Administration). Current AI use for Utilities (NAICS 22) has been suppressed for confidentiality reasons.

Source: Business Trends and Outlook Survey (BTOS) and Bonney et al. (2024a), Figure 2a.

The supplement results allow Bonney et al. (2024a) to dig deeper into the types of AI technologies used.[10] Conditioning on AI use, they find the top three technologies and applications used are for marketing automation, virtual agents, and natural language processing. Only about one-quarter of AI using businesses do so to replace worker tasks (and by far the majority is for a small number of tasks) and about one-fifth do so to replace work by existing capital or software. Turning to employment, most businesses using AI do not have a net change in employment due to AI use. Only 5 percent of firms experience employment change due to AI use and, notably, an increase in employment is more common than a decrease in employment.

[10] Bonney et al. (2024b) focuses on tasks and employment analyses.

RESULTS FROM BTOS RECENT COLLECTION

While this chapter is a retrospective of work on business use of AI, the BTOS biweekly collection of the core AI questions continued beyond the research cited here. The figure below communicates the results for current and expected business use AI as of the writing of this chapter (summer 2025).

The figure below shows that current AI adoption rate increases from about 3.5 percent to nearly 9.5 percent between September 2023 and May 2025. Similarly, the expected AI use rate in the next 6 months rises to about 11 percent during this period, up from about 6.5 percent. The linear regression based predicted values for biweekly adoption rates confirm these trends. Furthermore, the solid curves based on Locally Weighted Scatterplot Smoothing (LOWESS) applied to biweekly adoption rates indicate that adoption has been accelerating especially in the last 6 months compared to the earlier periods. These curves suggest that AI adoption is still at an early stage—the stage of initial rise in adoption rates based on a typical S-shaped adoption curve.

Current and Expected Future AI Adoption Rates Estimated From Responses to Business Trends and Outlook Survey (BTOS) AI Core Questions

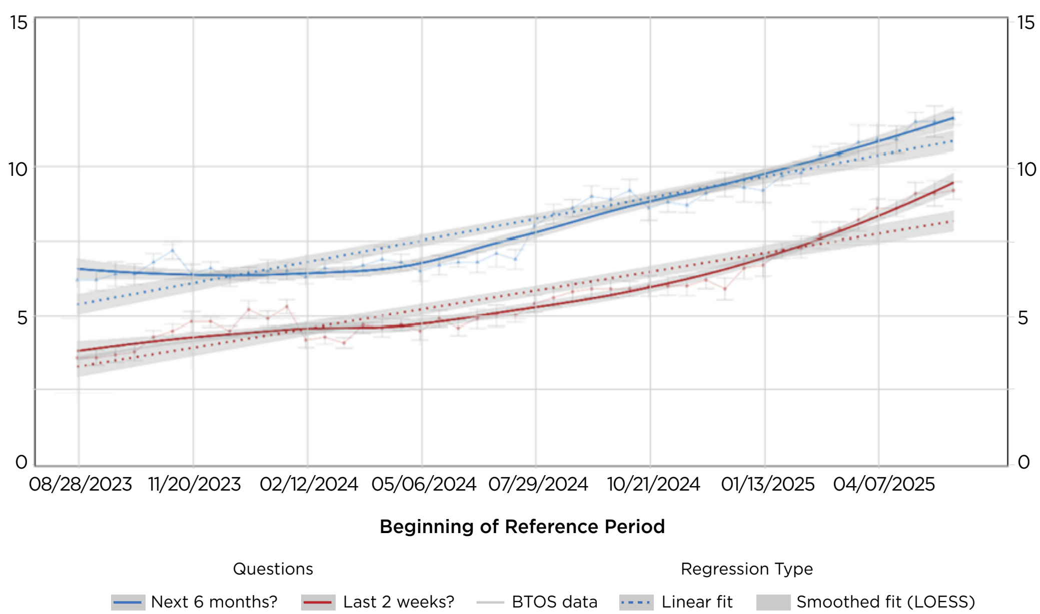

Note: Gray error bars indicate 95 percent confidence intervals around BTOS data given BTOS response standard errors. Gray regions indicate 95 percent confidence intervals of smoothing regressions on BTOS data points.

One-half of AI using businesses did not make organizational changes to accommodate AI. The biggest changes made in the last 6 months were training staff (21 percent) and developing new workflows (20 percent). Interestingly, only 8 percent reported changes in data collection and data management practices. The primary impediment for those businesses that do not expect to use AI in the next 6 months was "AI is not applicable to this business," which was cited by 81 percent of these businesses. As functionality of AI technologies evolves and businesses become more aware of useful applications of AI, the primary impediment to adoption may shift to other factors such as lack of skilled workforce.

The tracking of AI use continues in the BTOS through the present (2025) with (funding contingent) plans to continue this tracking for the future. The latest AI use statistics can be found at <www.census.gov/programs-surveys/btos.html>. As an example of the continuation of the BTOS collection, we show national results for the latest period available as of the writing of this chapter in the text box titled "Results From BTOS Recent Collection."

BUSINESS FORMATION STATISTICS

The survey approach provides the ability to precisely control the content of targeted questions allowing us to measure business use of AI, but the survey approach can mean being locked into obsolete concepts if the technology rapidly evolves. As a higher frequency survey, the BTOS has more flexibility to quickly adapt to the changing technological environment than the MOPS or ABS, but the BTOS also depends critically upon the alignment of respondent understanding and question intent. Dinlersoz et al. (2024) take a complementary approach by leveraging open-ended, self-reported business information from administrative data to capture AI use and/or provision at the time of business entry. In addition, using administrative data enables them to develop a longer time series at annual or higher (e.g., monthly) frequency not typically possible with the survey data.

The Business Formation Statistics (BFS) combines weekly information provided on employer identification number (EIN) applications and the LBD to create measures of business applications and formations (Chapter 2 of "2021 CES Annual Report" at <www.census.gov/library/publications/2022/adrm/2021-ces-annual-report.html> and Bayard et al. [2018] further describe the BFS). Since these applications are submitted on a flow basis with the Census Bureau receiving information every week, they provide an almost real-time glimpse into startup activity. Dinlersoz et al. (2024) use text analysis on applicants' main business line (principal line of merchandise), business name, and trade name on the EIN application to identify business ideas for developing or providing AI technologies or planning to use them. Their sample covers 2004–2023 and includes 60 million applications and 8 million employer startups.

Dinlersoz et al. (2024) describe how they build a comprehensive library of AI-related keywords from multiple sources and use this library to identify business applications related to AI. They check the resulting database for false positives by cross-checking the three business-provided text responses and filtering out cases that do not align with AI use or production or are not related to core business functions (such as a holding company or investment vehicle). However, it is not currently possible for them to precisely distinguish between AI users and producers. Thus, their category is broader than that collected in the surveys noted in the text box titled "Businesses That Produce AI." They also caution about selection effects in this method since businesses whose use of AI is less significant may not include AI text in their application; and, on the other hand, entrepreneurs may use AI text to signal sophistication of their planned business. However, similar selection issues also apply to the surveys discussed above, e.g., nonresponse that is not random, incidental or insignificant use of AI that gets reported as no use, and unknowing use of AI that is not reported.

Figure 3-5.
AI Business Applications Over Time

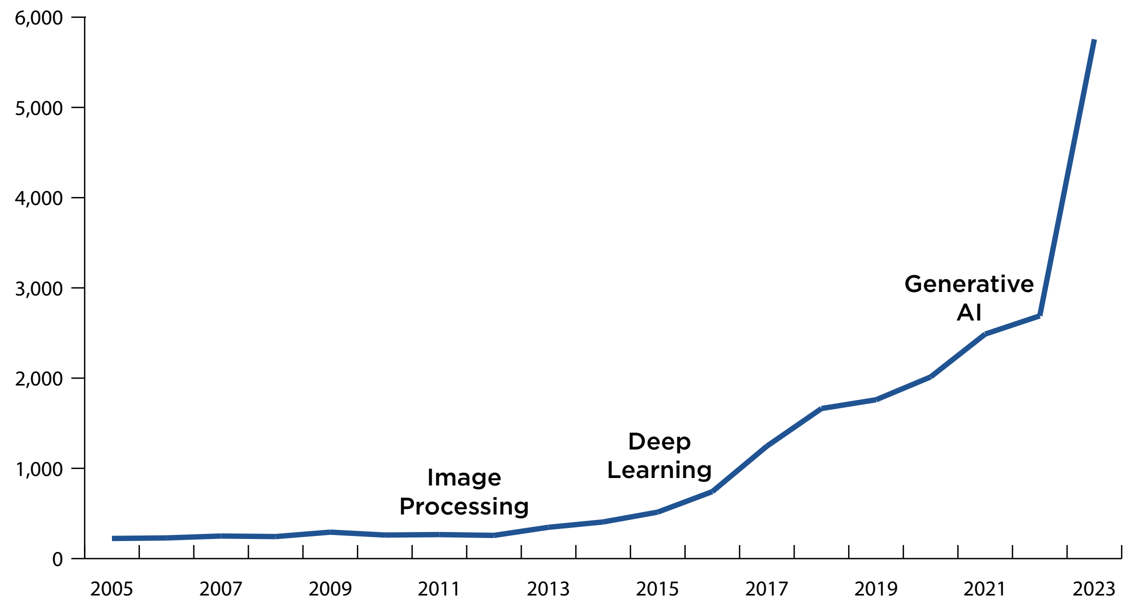

Source: Dinlersoz, et al. (2024).

Results

Dinlersoz et al. (2024) find about 22,000 AI business applications over the period 2004–2023 (which is less than 0.05 percent of all applications in their database). As noted above, one of the strengths of the administrative data is the ability to provide a time series. Figure 3-5 shows the time series of AI business applications with annotations at inflection points on technological breakthroughs (this is a simplified version of Figure 1 found in Dinlersoz et al. [2024a]). As they note, applications are relatively low and steady from 2005–2011 and then they identify three inflection points: 2012, 2016, and 2023. In terms of AI business formations, they find about 1,300 business startups through 2019.

One of the strengths of the BFS is the ability to examine the transition from application to formation. The BFS includes series on high-propensity business applications (HBA) which are applications that are more likely to transition to employer businesses. The HBA transition to employer businesses at a rate of 24 percent (as compared to 5 percent rate for non-HBA applications) over the sample period. Thus, it is of interest to know whether the AI business applications are more likely to be HBA. Through regression analysis, Dinlersoz et al. (2024) find that AI business applications disproportionately possess characteristics associated with being HBA. An open question is whether these HBA applications have in fact transited to actual employer businesses. Dinlersoz et al. (2024) highlight that these AI-oriented applications are concentrated in NAICS sectors 51 (Information) and 54 (Professional, Scientific and Technical Services). Decker and Haltiwanger (2024) also provide evidence of a surge in actual business formation in these sectors in the post-pandemic economy.

Dinlersoz et al. (2024) examine patterns in size and sector (and state) for AI business applications. The measure of size is the anticipated size of the business based on information provided in the EIN application about the "highest number of employees expected in the next 12 months." They find that AI business applications have higher expected employment than non-AI business applications, and this relationship holds in regression analysis where they control for industry, location, and time of application. Recall that the ABS results found a positive relationship between AI use and size.

Dinlersoz et al. (2024) further track the performance of employer businesses that originate from AI-related business applications and find that these businesses outperform other businesses emerging from non-AI-related business applications in terms of employment, revenue, payroll, and average wage. However, they also have higher labor share and higher exit rates. At the same time, they do not exhibit statistically different labor productivity compared to other businesses. These findings are conditional on several application characteristics, detailed industry, business age, and year effects. Overall, the findings indicate an up-or-out dynamic for AI businesses and suggest that AI businesses represent a more dynamic segment of the firm population, potentially contributing to business dynamism if the observed patterns continue to hold into the future.

Finally, in terms of sectors, AI business applications are concentrated—in descending order of application share—in NAICS 54, 52, and 51. Almost 75 percent of AI business applications are in these sectors. This pattern is broadly consistent with the results discussed above for ABS and BTOS. However, within this grouping, the largest share is in Professional, Scientific, and Technical Services (56 percent of AI business applications) while the Information share is much lower (less than 10 percent, as is the Finance and Insurance share). This is in marked contrast to ABS and BTOS where the Information sector dominates.

LESSONS LEARNED

While there is some sensitivity of results to the survey instrument and specific questions about AI use, there are some common lessons that emerge from this body of work. Currently, the fraction of firms that use AI is relatively low but growing. There are sectors and types of firms where AI use is higher, such as the Information sector and larger firms, respectively. Given the size pattern in adoption, the percentage of workers exposed to AI at work is much higher than percentage of firms using AI. There is also an age pattern in adoption: high-growth young firms are more likely to adopt.

In terms of productivity, we are still in the early days. While there is clear evidence that adoption of advanced technologies is associated with higher productivity, this is evident in more mature technologies like cloud computing, robotics, and specialized software. This is not surprising given that AI has not fully diffused across the economy and firms are still experimenting with how to implement AI technologies. Patterns thus far are consistent with the J-curve hypothesis, where if anything there may be an initial drag on productivity during the experimental phase, where intangible and complementary capital investments in implementation are important.

In terms of the overall impact on the economy, the papers are consistent with concerns about a growing divide between firms with the resources and capabilities to adopt AI versus those firms that do not adopt AI (for reasons that can include feasibility or applicability). Currently, nimble startups and digitally sophisticated, large incumbents are leading adopters. However, as AI becomes easier to use and requirements for skilled users lessen, these adoption patterns may change. The evidence from the BFS suggests that AI may play a growing role in business formation and business dynamism.

There are also operational lessons learned. This chapter describes how the CES technology team worked with internal and external partners to capture the diffusion of AI across businesses. We show how a proactive approach to collection enabled the Census Bureau to measure the early stages of AI adoption using both survey and administrative data and how collections can be adapted to keep pace with a rapidly changing technology. We discuss how the distinct approaches complement each other. The ABS and MOPS surveys provide a detailed view but are less timely; the BTOS provides less context but is better at capturing dynamics. The BFS also helps us capture dynamics in a timely fashion as well as allowing us a view into emerging technologies. Underpinning some of the empirical analyses using either survey or administrative data is the data infrastructure previously developed by CES, especially the LBD.

FUTURE WORK

We hope future survey collections can continue to enable us to trace the trajectory of AI adoption from its early stages through diffusion and maturity. We continue to explore the possibility of capturing AI use across establishments within firms through surveys (such as the Annual Integrated Economic Survey). We also continue work with cognitive testing teams to refine our definition of AI. The CES tech team continues to work with the programmatic team on edits to questions on a possible, future BTOS AI supplement.

We have discussed how administrative data can be an important complement to this approach. The work of Dinlersoz et al. (2024) shows the promise of BFS micro data in identifying emerging-technology startups in a timely fashion. Currently, this work is being extended to understanding the nature of business applications in robotics and other advanced technologies by using machine learning techniques. CES will continue to build out the infrastructure that enables the use of administrative data to understand the diffusion of this important technology.

One source of data the CES technology team looks forward to exploring in more depth in the future is private data. The team conducted some exploratory work using text analysis of resumes on job posting websites to track the rise of AI adoption through workers (building on Babina et al. [2023]). The idea had been to combine this information with data from LEHD. Unfortunately, we did not have the resources to pursue this idea but may be able to revisit it in the future. Further along is the work of Akcigit, Chikis, and Goldschlag who use Microsoft Academic Graph (MAG) data to identify scientific publications related to AI. They then link the authors in the MAG data to individual data on person characteristics and employment histories in the LEHD. This allows for a comprehensive view of the career lifecycle of AI researchers and the allocation of AI talent across sectors in the U.S. economy.

We could also revisit the idea of getting information on business use of AI from AI producers noted in the text box titled "Businesses That Produce AI" by using information provided directly to the Census Bureau from the producing companies rather than through survey collection.

Finally, future work could focus on productivity, technology, and tasks. A different CES team works collaboratively with researchers at the Bureau of Labor Statistics (BLS) to improve our understanding of productivity dynamics (Chapter 2 of "2014 CES Annual Report" at <www.census.gov/library/publications/2015/adrm/2014-ces-research-report.html> describes this long-running collaboration). To that end, the productivity team has combined Census Bureau micro data on business activity with task and occupation data from BLS to examine more closely drivers of productivity growth (Blackwood et al., 2025). An ambitious project could bring together the work of these two teams to examine more closely the interaction between technology adoption and task assignment and the resulting impact on productivity.

This chapter was written by humans but, given the subject of this chapter, we experimented with our own "business use of AI" by using AI to write a version of this chapter. For a clean comparison, the experiment was run only after the first draft of the chapter was completed. We focused on the main text of this chapter, asking ChatGPT to summarize research results and methodological challenges from the eight main papers cited in the results sections and providing it with an example annual report chapter that had been solo authored by one of us. We ran this a few times, changing the prompts slightly. As expected, AI provided a useful synthesis of research results and overview of methodological challenges. The AI version organized the chapter by results; whereas our version is organized by instrument since we are trying to simultaneously highlight the infrastructure required to measure AI use. ChatGPT focused more attention on broader economic implications than we did in this chapter. With more refined prompts, AI was able to capture the importance of collaborations between the Census Bureau and academic experts. Based on the feedback from these multiple versions, we added a clarifying sentence to the introductory paragraph and a new paragraph about broader economic implications to the Lessons Learned section.

Lastly, we challenged ChatGPT to create a haiku summarizing the eight papers. It let us down by returning three lines that did not adhere to the 5-7-5 syllable format (but then politely acknowledged the mistake). It overemphasized code in most of its attempts, however, we were impressed with how well it captured the tension between the importance of measurement and the difficulty of measuring an elusive concept. With some back and forth, we present to you our hybrid (AI and human creators) haiku:

To measure the world
Surveys could miss what firms do
AI slips the net

Echoing the research summarized in this chapter, we found AI a useful *complement* to our human input on this humorous "work" task, but it is not yet a *substitute* for human input.

REFERENCES

Acemoglu, Daron, Gary Anderson, David Beede, Catherine Buffington, Eric Childress, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas, “Advanced Technology Adoption: Selection or Causal Effects?”, *AEA Papers and Proceedings*, 2023, 113, pp. 210–214.

Acemoglu, Daron, Gary W. Anderson, David N. Beede, Catherine Buffington, Eric E. Childress, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John C. Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas, “Automation and the Workforce: A Firm-Level View From the 2019 Annual Business Survey,” in eds. Susanto Basu, Lucy Elridge, John Haltiwanger, and Erich Strassner, *Technology, Productivity, and Economic Growth*, 2025.

Alderucci, Dean, Lee Branstetter, Eduard Hovy, Andrew Runge, and Nikolas Zolas, “Quantifying the Impact of AI on Productivity and Labor Demand: Evidence From U.S. Census Microdata,” presentation at the Allied Social Science Associations 2020 Annual Meeting, 2019.

Atrostic, B.K. and Sang Nguyen, “Computer Networks and U.S. Manufacturing Plant Productivity: New Evidence From the CNUS Data,” Center for Economic Studies Working Paper 02-01, 2001.

Babina, Tania, Anastassia Fedyk, Alex X. He, and James Hodson, “Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition,” NBER Working Paper No. 28290, 2023.

Basker, Emek, Randy A. Becker, Lucia Foster, T. Kirk White, and Alice Zawacki, “Addressing Data Gaps: Four New Lines of Inquiry in the 2017 Economic Census,” Center for Economic Studies Working Paper 19-28, 2019.

Basker, Emek, Lucia Foster, and Shawn Klimek, “Customer-Employee Substitution: Evidence From Gasoline Stations,” *Journal of Economics and Management Strategy*, 26, 2017, pp. 1–21.

Bayard, Kimberly, Emin Dinlersoz, Timothy Dunne, John Haltiwanger, Javier Miranda, and John Stevens, "Early Stage Business Formation: An Analysis of Applications for Employer Identification Numbers," NBER Working Paper No. 24364, 2018.

Becker, Randy A., John Haltiwanger, Ron S. Jarmin, Shawn D. Klimek, and Daniel J. Wilson, "Micro and Macro Data Integration: The Case of Capital," in eds. Dale W. Jorgenson, J. Steven Landefeld, and William D. Nordhaus, *A New Architecture for the U.S. National Accounts*, 2006.

Blackwood, G. Jacob, Cindy Cunningham, Matthew Dey, Lucia Foster, Cheryl Grim, John Haltiwanger, Rachel Nesbit, Sabrina Wulff Pabilonia, Jay Stewart, Cody Tuttle, and Zoltan Wolf, "Opening the Black Box: Task and Skill Mix and Productivity Dispersion," in eds. Susanto Basu, Lucy Elridge, John Haltiwanger, and Erich Strassner, *Technology, Productivity, and Economic Growth*, 2025.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen, "What Drives Differences in Management?", *American Economic Review*, 109, 2019, pp. 1648–1683.

Bonney, Kathryn, Cory Breaux, Catherine Buffington, Emin Dinlersoz, Lucia Foster, John Haltiwanger, Nathan Goldschlag, Zachary Kroff, and Keith Savage, "Tracking Firm Use of AI in Real Time: A Snapshot From the Business Trends and Outlook Survey," Center for Economic Studies Working Paper No. 24-16, 2024a.

Bonney, Kathryn, Cory Breaux, Catherine Buffington, Emin Dinlersoz, Lucia Foster, John Haltiwanger, Nathan Goldschlag, Zachary Kroff, and Keith Savage, "The Impact of AI on the Workforce: Tasks Versus Jobs?" *Economics Letters*, 244,2024b, 111971.

Brynjolfsson, Erik, Catherine Buffington, Nathan Goldschlag, J. Frank Li, Javier Miranda, and Robert Seamans, "The Characteristics and Geographic Distribution of Robot Hubs in U.S. Manufacturing Establishments," Center for Economic Studies Working Paper 23-14, 2023.

Brynjolfsson, Erik and Andrew McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W.W. Norton Company, 2014.

Brynjolfsson, Erik and Kristina McElheran, "Data in Action: Data-Driven Decision Making and Predictive Analytics in U.S. Manufacturing," Rotman School of Management Working Paper No. 3422397, 2019.

Brynjolfsson, Erik and Kristina McElheran, "The Rapid Adoption of Data-Driven Decision-Making," *AEA Papers & Proceedings*, 106, 2016, pp. 133–139.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 13, 2021, pp. 333–372.

Buffington, Catherine, Lucia Foster, and Colin Shevlin, "Measuring Business Trends and Outlook Through a New Survey," *AEA Papers and Proceedings*, 113, 2023, pp. 140–144.

Buffington, Catherine, Lucia Foster, Ron Jarmin, and Scott Ohlmacher, "The Management and Organizational Practices Survey (MOPS): An Overview," *Journal of Economic and Social Measurement*, 42, 2017, pp. 1–26.

Cockburn, Iain M., Rebecca Henderson, and Scott Stern, "The Impact of Artificial Intelligence of Innovation: An Exploratory Analysis," in Ajay K. Agrawal, Joshua Gans, and Avi Goldfarb (eds.), *The Economics of Artificial Intelligence: An Agenda*, 2019.

Chow, Melissa C., Teresa C. Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White, "Redesigning the Longitudinal Business Database," NBER Working Paper No. 28839, 2021.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda, "Declining Business Dynamism: What We Know and the Way Forward," *American Economic Review*, 106, 2016, pp. 203–207.

Decker, Ryan A. and John Haltiwanger, "Surging Business Formation During the Pandemic: Causes and Consequences," *Brookings Papers on Economic Activity*, 2023, pp. 249–302.

Dinlersoz, Emin, Can Dogan, and Nikolas Zolas, "Starting Up AI," Center for Economic Studies Working Paper 24-09, 2024.

Dinlersoz, Emin and Zoltan Wolf, "Automation, Labor Share, and Productivity: Plant-Level Evidence From U.S. Manufacturing," Center for Economic Studies Working Paper 18-39, 2018.

Dunne, Timothy, "Plant Age and Technology Usage in U.S. Manufacturing Industries," *RAND Journal of Economics*, 25, 1994, pp. 488–499.

Hatzius, Jan, Joseph Briggs, Devesh Kodnani, and Giovanni Pierdomenico, "The Potentially Large Effects of Artificial Intelligence on Economic Growth," *Goldman Sachs Economic Research*, 2023.

McElheran, Kristina, J., Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas, "AI Adoption in America: Who, What, and Where," *Journal of Economics & Management Strategy*, 33, 2024, pp. 375–415.

McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson, "The Rise of Industrial AI in America: Microfoundations of a Productivity J-Curve," Center for Economic Studies Working Paper 25-27, 2025.

McKinsey Global Institute, "The Economic Potential of Generative AI: The Next Productivity Frontier," 2023.

National Science Foundation, "Artificial Intelligence in the Business Sector: R&D, Use, and Impact on Employees," InfoBytes, NSF-25-312, 2024.

Seamans, Robert, "Robot Census: Gathering Data to Improve Policymaking on New Technologies," 2021.

Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David Beede, Catherine Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz, "Advanced Technologies Adoption and Use by U.S. Firms: Evidence From the Annual Business Survey," Center for Economic Studies Working Paper 20-40, 2020.