New Approaches to Characterize Industries

AI AS A FRAMEWORK AND A USE CASE

Edited by Brent Orrell

SEPTEMBER 2025

AMERICAN ENTERPRISE INSTITUTE

New Approaches to Characterize Industries

AI as a Framework and a Use Case

Edited by Brent Orrell

Contents

Introduction

Will Rinehart, Brent Orrell, Julia Lane, and Erik Brynjolfsson

The newest artificial intelligence technologies, especially generative AI systems, could fundamentally transform how firms do business and how Americans work. Still, there is little data and evidence to understand how AI is reshaping the economy.

This is coming just as the politics of work are changing. The automation and trade shocks of the early 2000s and the resulting "deaths of despair" have brought political attention to how technological and economic shifts can transform communities and livelihoods. Business leaders, workers, educational and training institutions, and governments need local, timely, and actionable data to help the workforce respond to shocks that are likely to be even greater than those of two decades ago.

Unfortunately, the visibility of the workforce's transformation has also made it clear that traditional data sources are inadequate to inform that response. Even current scientific and industrial classification systems are not fit for this purpose. AI, like many other new and emerging technologies, is neither a well-defined scientific field nor a distinct industry.[1] Moreover, national governments' increasing focus on industrial policy, which is fundamentally reshaping the economy and society, suggests that entirely new data and frameworks may need to be developed to understand how workers and firms interact.

On March 18, 2024, the American Enterprise Institute, Stanford University's Digital Economy Lab, and New York University convened a daylong seminar titled "New Approaches to Characterize Industries: AI as a Framework and a Use Case." Its goal was to begin exploring elements of a theoretical framework that could help meet this data challenge, and seminar participants were tasked with identifying what data, definitions, models, and tools exist or could be developed.

Participants were invited to the workshop for their expertise in data, measurement, and analysis. But more importantly, each presenter had a history of designing, deploying, and using new data systems—such as the Longitudinal Employer-Household Dynamics program at the US Census Bureau, the Institute for Research on Innovation & Science (IRIS) at the University of Michigan, the Texas Workforce Commission, and the New Jersey Department of Labor & Workforce Development—and private job market data. As intended, workshop participants drew on their experiences to identify empirically implementable, dynamic, and flexible approaches for understanding this critical and emerging space.

The following three key takeaways emerged from the formal and informal discussions:

1. An innovative institution should be established to implement a new vision and framework. This independent, nonpartisan institution could be dedicated to producing bottom-up, demand-driven tools and insights for businesses, workers, and governments by connecting advances in AI and other critical technologies to changes in the nature of new and existing jobs, skills, and economic opportunities.

2. Policymakers and practitioners should support the institution in establishing partnerships that directly serve the needs of businesses, workers, and researchers. These partnerships would build an understanding of how AI and other emerging technologies affect local and regional economies and labor markets. An initial focus might include

    a. Providing data and insights to firms as businesses' AI capabilities move from experimental use to enterprise use at scale,

    b. Prototyping and producing customizable tools so current and future workers can acquire skills in response to changing demands, and

    c. Producing tools and analyses that federal, state, and local government agencies would use to allocate programmatic education and training resources.

3. Finally, the new framework and resulting classifications should be designed to inform and be complemented by the federal statistical system's operations; federal, state, and local program and service providers; and scientific research. This can be achieved by bringing together the best minds from each key sector through focused fellowship, training, and competitions.

In this introduction, we analyze and integrate the papers produced for the workshop and the day-of exchanges among the invited experts. The chapter is divided into six key sections that move from the central challenges and opportunities in AI measurement through perspectives informed by research and federal and state agencies, concluding with an analysis of how to better represent AI data through a "follow-the-people" approach to talent flows—rather than using firm surveys to measure AI's implementation. Workshop participants outlined a number of opportunities to strengthen data collection.

The Problem and Promise of Measuring AI

Erik Brynjolfsson, who directs the Stanford Digital Economy Lab, and Susan Athey, a senior fellow at the Stanford Institute for Economic Policy Research, delivered keynotes to begin the workshop. Brynjolfsson started by retelling how Dutch microbiologist Antonie van Leeuwenhoek became the first person to see a microorganism through a microscope. His microscope helped drive other discoveries, igniting a revolution in science that lasted decades. Similarly, Brynjolfsson thought generative AI and other advanced AI models are just at the beginning of their development and will take decades to be fully realized.

New digital tools will require enormous investments and complementary intangible capital to be productive, but with this transition also comes the possibility of better understanding how firms implement technologies. With the right kinds of data collection in place, economists and researchers could better understand how many software engineers are at each company and what skills they have. Starting at this level of analysis would then allow for skills to be connected to firm output and productivity. In turn, this information could be used to better understand job-posting data and create skill and task taxonomies that would be more granular than the current O*NET system. Moreover, data of this type could help locate nascent superstar firms and illuminate markups for AI companies—a critical question for antitrust authorities.

Following Brynjolfsson, Athey offered a complementary perspective by focusing on the practical hurdles all collection systems face. It is often assumed that companies have reliable information about their own processes and customers. But in Athey's experience as the chief economist at the Federal Trade Commission, even the most advanced companies in Silicon Valley face challenges with quantification.

Economists are especially attuned to measurement, estimation, and inference issues. Revealed preference methods, such as surveys, are not always a reliable measure of true preferences or well-being. For example, impulse purchases and the dip in happiness that comes with having children both undermine the connection between respondents' revealed preference and well-being.

Likewise, measuring digital skills and the use of digital tools presents its own challenges. For one, there is a chicken-and-egg problem when it comes to measuring how interventions influence people's digital skills. In one study Athey conducted, an AI system gave book recommendations to students that seemingly helped students learn to read. But Athey and her coauthors discovered that a contest of who can read the most actually produced the biggest lift in reading. In other words, it wasn't the recommendations that mattered but the incentives.

Athey's keynote offered a warning: Even the most careful researchers and the best-paid tech workers are uncertain of the best and most effective ways to engage users. To solve the problem of mismeasurement, researchers and scholars should be looking beyond current survey tools.

Research Perspectives on AI Measurement, Session 1

In the workshop's first session, Jason Owen-Smith of IRIS emphasized the high stakes involved, especially with new legislative requirements like the Creating Helpful Incentives to Produce Semiconductors (CHIPS) and Science Act, which mandates measuring science investments' effects on job creation. Simply put, we lack methodologies, measures, and systems to achieve this goal. Owen-Smith described AI as the "poster child" for these challenges, given the complexities of defining and tracking its impact.

He focused on the measurement challenges AI poses, contrasting traditional classification methods—such as clustering AI-related documents by topic, which he described as largely arbitrary—with the "industries of ideas" approach, which starts with people. This method—pioneered by Julia Lane, a professor in New York University's Wagner Graduate School of Public Service who helped organize the event and presented later in the day—involves identifying industries by tracking domain experts' activities and movements in the workforce.

Owen-Smith suggested that state administrative data could be leveraged to follow AI researchers, thereby defining AI industries based on where these researchers are employed. This method could provide insights into the flow of talent and the translation of government research investments into economic and workforce development. He emphasized the importance of following AI researchers and using state administrative data to build a more accurate picture of AI's impact on regional economies.

Lee Branstetter, the James M. Walton Professor of Economics and Public Policy at Carnegie Mellon University, emphasized the importance of firm-level data in understanding AI's impacts at a granular level. He discussed how machine learning algorithms can be employed to analyze patent texts, enabling AI-related innovations to be identified and quantified. Branstetter highlighted the proliferation of AI patents as a promising metric for measuring AI-driven innovation, noting that firms filing these patents often significantly increase their productivity, employment, and output.

Branstetter also discussed the global landscape of AI innovation, noting that Japan is the closest peer to the United States. Contrary to popular belief, China lags behind in terms of innovation. Reiterating the power of the industries-of-ideas model, he proposed that tracking leading AI researchers and their movements—particularly as they transition from academia to industry—and the patents they produce could be instrumental in understanding AI's evolving landscape. By combining patent analysis with public data, such as wage information, it could become possible to identify top AI scientists and follow their career trajectories.

Branstetter also pointed to potentially using LinkedIn data to track the career paths of students and postdoctoral researchers, who play critical roles in AI development. This approach could further elucidate the economic benefits that arise when a critical mass of AI researchers converges within a startup or other firm.

After Branstetter's presentation, Diane Coyle, who codirects the Bennett Institute, provided critical commentary for Owen-Smith and Branstetter, noting that relying on the Elsevier Corpus, which is a large open-access database of Elsevier's journals, may introduce biases. She suggested that alternative measures, such as the knowledge complexity index, could provide different insights into conceptual distances in the AI field. Coyle also raised concerns about the scalability of both Owen-Smith's and Branstetter's resource-intensive methods, stressing the need for not only accurate but also timely and widespread measurements to fully capture AI's impact. In other words, researchers should be mindful that timely data can often trump more precise data that come months later.

Coyle ended by asking how these measurements could inform our understanding of macroeconomic and long-term effects. She also touched on the current trends in computing power, ending on a positive note. While Moore's law may be decelerating, the rapid decline in computing costs could have significant positive implications for the continued advancement of AI technologies, potentially countering the pessimism surrounding AI's future.

Research Perspectives on AI Measurement, Session 2

Prasanna "Sonny" Tambe, an associate professor of operations, information, and decisions at the Wharton School at the University of Pennsylvania, opened the second session by discussing the significant gaps in measuring AI's impacts, particularly at the firm level. He highlighted the need for better data to understand why some firms succeed in implementing AI while others struggle. AI investment, he noted, remains highly concentrated in a few firms, more so than investments in other digital technologies at so-called frontier firms. But much like other frontier firms, AI-focused firms are investing in infrastructure, workforce training, and high-skill talent acquisition, all of which are unevenly distributed across the economy.

Tambe pointed out that AI infrastructure is often proprietary and rapidly evolving, making it difficult to identify and measure both tangible and intangible AI assets. He called for a consensus on how to track these assets effectively. Additionally, tracking individual AI professionals is a challenge due to a lack of reliable, up-to-date data. Tambe underscored the difficulties of measuring AI's impacts due to the vast amount of data and their rapid depreciation over time.

He also raised questions about AI's effects on the labor market, particularly for low-skilled workers. There is still significant uncertainty around which skills will remain relevant or become obsolete as a result of AI. He highlighted that workforce training was a key factor in helping workers adapt, but the best approaches to prepare workers in such a dynamic environment remain unclear.

Nestor Maslej, research manager at Stanford's Institute for Human-Centered Artificial Intelligence (Stanford HAI), added to the discussion by focusing on the challenges of finding reliable metrics for AI innovation for Stanford HAI's AI Index. He pointed out that traditional bibliometric data, which track the number of AI-related publications and are central to Stanford HAI's AI Index, can be misleading. Paradoxically, while the number of AI publications has decreased recently, innovation in the field appears to be increasing, underscoring the difficulty of using bibliometric data to track AI advancements.

Maslej also highlighted geographic challenges in measuring AI innovation, as bibliometric data do not account for differences in publications' quality or impact. He stressed the importance of identifying the most cutting-edge research rather than treating all publications equally. Despite these limitations, bibliometric data can still play a role in understanding global AI innovation. Like Branstetter, Maslej offered evidence that China significantly lags behind the United States in AI innovation.

Tracking AI legislation is one of the newest developments for Stanford HAI's AI Index, and all evidence points to a tidal wave of new regulations. The Organisation for Economic Co-operation and Development has an AI policy tracker of its own, but it lacks a consistent standard for what constitutes regulation or legislation, and it misses state and local regulations. Like the other presenters, Maslej called for better measurement of government spending on AI research and development (R&D) to gauge public-sector involvement and underscored the need for accurate measures.

Will Rinehart provided commentary on the two presentations, arguing that tracking AI legislation is vital, as regulation can create uncertainty and affect innovation. He also pointed to the need for more granular measurements, including tracking state and local AI regulations such as the California Consumer Privacy Act, which significantly affects AI development.

Perspectives from the National Artificial Intelligence Research Resource

Following the panels, Lane presented on the work of the National Artificial Intelligence Research Resource (NAIRR) and the importance of measuring AI's impact on innovation, diversity, ethics, productivity, and employment. Lane emphasized that Congress is eager to evaluate the returns on its investments in AI, particularly through the CHIPS and Science Act. Effective industrial policy requires timely and accurate data. As much as there is a need for data at the national level, policymakers and researchers need to assess how AI influences labor markets at the state and local levels.

Lane critiqued the use of new PhD graduates as a measure of AI investment, suggesting that while many new AI PhDs are produced each year, their contributions to the field's innovation and advancement are uncertain. To better track the most impactful researchers, she recommended monitoring AI experts who attend conferences, as they are likely to be driving the most significant innovations. This approach supports a shift from traditional sector-based classifications of industries to an industries-of-ideas model, which emphasizes following researchers as they move through academia and into the labor market.

Following Lane's presentation, Manish Parashar—who directs the Scientific Computing and Imaging Institute—focused on the barriers that prevent some companies from adopting AI, despite its potential to enhance their operations. Parashar highlighted that identifying companies that have yet to implement AI and understanding the obstacles they face are key to democratizing AI technology. He emphasized the need for greater access to AI infrastructure, computing power, data, and software—especially for firms that lack the resources to integrate these technologies.

Parashar also stressed the importance of educating firms about AI's capabilities, noting that many companies are unaware of how they can benefit from adopting AI. He advocated for creating a flexible AI R&D ecosystem that fosters collaboration while addressing concerns related to privacy, civil liberties, and national security. By democratizing access to AI R&D, NAIRR could help promote competition, cooperation, and innovation.

He concluded by identifying three critical components of AI development: infrastructure, policy, and measurement. He underscored the importance of providing policymakers with accurate and timely data to ensure informed decision-making. Exploring initiatives such as the national data platform pilot could be a step toward creating a collaborative environment that fosters AI innovation.

Suzette Kent, the former federal chief information officer of the Office of Management and Budget, followed Parashar's presentation. She focused on the complexity of measuring AI and the lack of clear applicability of these measurements to key stakeholders, including state and local governments. Kent stressed the need for accurate data to inform policy decisions and guide AI investments, highlighting the current gaps in understanding AI supply and demand at the national level.

She expressed concern about the uncertainty surrounding whether AI investments are translating into real-world innovation, suggesting that without clear metrics, government and public investment in AI may be reduced. Kent called for the development of a more cohesive and comprehensive national view of AI development and deployment to address these uncertainties.

Perspectives from Data Collection Agencies and Organizations

Adam Leonard, who directs the Texas Workforce Commission (TWC), then spoke on the underutilization of wage records, highlighting how these datasets can offer deeper insights than can individual data points by identifying broader trends. He provided an example of testing media narratives, such as whether Texas truly faces a teacher shortage, using wage and employment data from the TWC.

The TWC has advanced the use of administrative and wage data to better understand labor market dynamics, tracking employment patterns and even helping monitor pandemic-related insurance claims. Leonard argued that wage records could be pivotal in tracing individuals' movement from academia into the private sector, thus illustrating how ideas and innovations disseminate through the economy. Moreover, these records can help identify firm-level investments in AI, particularly through hiring trends.

Lesley Hirsch, the assistant commissioner of research and information at the New Jersey Department of Labor and Workforce Development, discussed AI's potential impacts on the workforce, particularly for low- and mid-skilled workers, who were significantly affected by previous rounds of automation. While AI excels at cognitive nonroutine tasks, Hirsch suggested that jobs requiring employees to work on-site may be less vulnerable to displacement.

Nonetheless, there is great uncertainty surrounding AI's overall effect on the labor market. It could displace jobs but also create new ones, just as computers have led to the creation of entire industries. Hirsch pointed to initiatives like New Jersey's career navigator—which uses machine learning to guide workers toward optimal career choices based on their education, skills, and interests—as an example of AI's potential to assist workers.

Nela Richardson—ADP's chief economist and environmental, social, and governance officer—explored two divergent strategies for measuring AI's impact: following the money or following the people. ADP, which works with 25 million US workers, has focused on matching workers with employers based on relevant skills, with a particular emphasis on AI-related skills. Richardson shared that 85 percent of workers globally expect AI to affect their jobs, but she noted that AI's greatest impact is likely to be at the task level rather than the job level. She advocated for a national task report that breaks down specific tasks AI could replace or enhance. Such a report could provide a granular understanding of how AI interacts with jobs and tasks, complementing other data sources such as job postings.

Karin Kimbrough, LinkedIn's chief economist, shared data on the growing excitement and demand for AI talent. According to LinkedIn's data, discussions about AI have increased by 70 percent in recent months, with millennials and women particularly engaged in the topic. Kimbrough highlighted the rapid 22 percent growth in AI adoption over 2024, emphasizing that AI's impact extends far beyond the tech sector. She divided AI's impact into three categories: augmented, disrupted, and insulated tasks. While many knowledge industries, including economics, sit on the edge of disruption, other industries are poised for augmentation. Kimbrough's data also suggested that women and Gen Z workers are among the most exposed to AI-related changes.

Josh Hawley of Ohio State University presented a critical analysis of LinkedIn data, cautioning that it is not always reliable for tracking workforce trends. He underscored the need for more comprehensive data that go beyond individual states and are pooled across regions. Hawley advocated for developing broader datasets that can provide a fuller picture of AI's labor market impact.

Vipin Arora, the director of the US Department of Commerce's Bureau of Economic Analysis (BEA), argued for a recalibration of how AI is measured, comparing it to knowing the temperature outside without needing exact degrees to decide how to dress a child. While accurate AI measurement is resource intensive, Arora questioned how precise the measurements need to be to make them actionable. He emphasized the importance of communicating AI measurement findings clearly to policymakers and the public, calling this an ongoing challenge that requires significant attention.

BEA Associate Director David Wasshausen focused on how AI is reflected in economic accounts, comparing AI with the internet in terms of its need for vast infrastructure and wide-reaching impacts. Identifying what qualifies as AI and measuring its effects remains a major challenge, especially given the variance in AI's application across different cases. Wasshausen noted that AI could generate enormous economic benefits for firms by enabling custom software development and other innovations, which in turn could boost profits and reduce costs.

In their coauthored chapter, Arora and Wasshausen argued for adding further details to current accounting methods. The "supply and use" tables framework that calculates gross domestic product–style measures can already record information about AI investments, whether these are denominated as intermediate purchases of software and consulting services or final expenditures. Arora and Wasshausen note that "a prerequisite for identifying new approaches to characterize hard-to-measure industries is a comprehensive and accurate understanding of how those industries are currently measured."

Emilda Rivers of the National Center for Science and Engineering Statistics addressed the complexities of measuring AI and its innovation potential, stressing that accurate measurement is critical because it shapes the narrative around AI and influences policy and cultural responses. Rivers advocated for more granular state and local data, highlighting the importance of collaboration and data sharing to capture the full scope of AI's impact.

Nancy A. Potok, who served as the chief statistician of the United States, concluded the session with a discussion on the evolving role of federal statistics in measuring AI's workforce impact. She described the decentralized nature of data collection within federal agencies, which creates challenges for cross-agency analysis. Potok emphasized the need for more granular "microdata" at the local level to enable companies, private individuals, and policymakers to make decisions more effectively. AI's complexity and the lack of coordination across agencies pose significant barriers to understanding and acting on the data currently being collected.

In the chapter she presented on, Potok examined incorporating intellectual property as a category into classification systems, like the North American Industry Classification System and the North American Product Classification System, and more recent attempts to define "factory-less goods" as potential case studies. She pointed out that the BEA uses "satellite accounts" as models for linked sub-accounting of expenditures, which could be applied to AI as well. Potok also suggested using America's Datahub Consortium for a pilot project linking the federal statistical system datasets together. Such an AI-consortium arrangement would require the Office of Management and Budget to establish terms and standards.

Improving AI Representation in Data

Throughout the sessions, participants stressed the need to move away from current survey methods and move toward a "follow the people" approach—what Lane has been calling an industries-of-ideas approach. It involves tracing the movement of grant-funded engineers from academic labs to industry as a way of understanding how ideas shape business. By linking administrative and wage records collected from companies and industries, we can better understand the flow of technological innovation and talent into the economy of goods and services.

The IRIS Universities Measuring the Effects of Research on Innovation, Competitiveness, and Science dataset—consisting of grant administration records for research awards processed by more than 100 US universities—puts the industries-of-ideas model into practice. Securely linking these grant records with individual- and company-level workforce data produces a map of job demand and tech development. There is already a corresponding National Science Foundation pilot project underway in Ohio, which will be crucial to analyze.[2] And to top this off, Leonard details in his chapter how he is testing this idea using data from the University of Texas at Austin.

The follow-the-people approach also aggregates up to the level of companies and industries. Owen-Smith has developed techniques for using "seed sets" of talented AI researchers to identify and map their moves into the economy, and in his chapter, he begins to estimate how employers are using their AI skills. He proposes a threshold for a company's "AI-ness" that could be generalized to companies with similar AI profiles.

"Following the people" means following their companies' records. This raises the question of how to better align terms and timelines with the federal statistical agencies responsible for collecting these records. It also raises the question: Given that the agencies have established frameworks, how can these be remodeled to house new windows on AI development?

Conclusion

Reviewing the contributors' chapters leads to an overwhelming conclusion that an operational definition, or definitions, of AI must be determined from a place of authority. To designate AI engineers, companies, products and services, or related skills, a classification system is required. The United States already has several. But without a definitive statement of terms set via collaboration between government and industry and implemented through statistical agencies and researchers, we risk data and analytical chaos that will leave all segments of the economy and society without sufficient information to make meaningful choices about education, training, and workforce development investments.

Notes

1. See Julia Lane, *Democratizing Our Data: A Manifesto* (MIT Press, 2021).
2. US National Science Foundation, "NSF Launches Pilot to Assess the Impact of Strategic Investments on Regional Jobs," October 26, 2023, https://www.nsf.gov/tip/updates/nsf-pilot-assess-impact-strategic-investments-regional-jobs.

Part I

Research Perspectives on AI Measurement

Will the Real AI Researcher Please Stand Up?

Fields, Networks, and Systems to Measure the Impact of Research Investments

Jason Owen-Smith

Recent legislation that makes federal research investments in critical and emerging technologies a key lever for economic policy on a massive scale has created pressing new measurement challenges. The CHIPS and Science Act (CSA) of 2022 authorized \$81 billion for the National Science Foundation (NSF), allocating about \$1 in every \$4 to the new Directorate for Technology, Innovation and Partnerships (TIP). TIP could become nearly 2.4 times bigger than today's largest NSF directorate (Mathematical and Physical Sciences), controlling nearly 60 percent of the entire NSF research budget.

TIP's authorizing language foregrounds regional innovative capabilities, jobs, workforce and educational capacity, and broad concerns about equity and access. It articulates an ambitious goal—to "advance research and development, technology development, and related solutions to address United States societal, national, and geostrategic challenges, for the benefit of all Americans"[1]—and includes substantial measurement mandates. By 2027, the CSA will require a National Academies review that, among other things, assesses:

- Whether the directorate's "key technology focus areas" offer solutions to "challenges with social, economic, health, scientific, and national security implications";
- "Whether Federal investment in the key technology focus areas have resulted in new domestic manufacturing capacity and job creation"; and
- "Federal investments in education and workforce development to support the key technology focus areas."[2]

Figure 1. Example of a TIP RIE Finalist Region

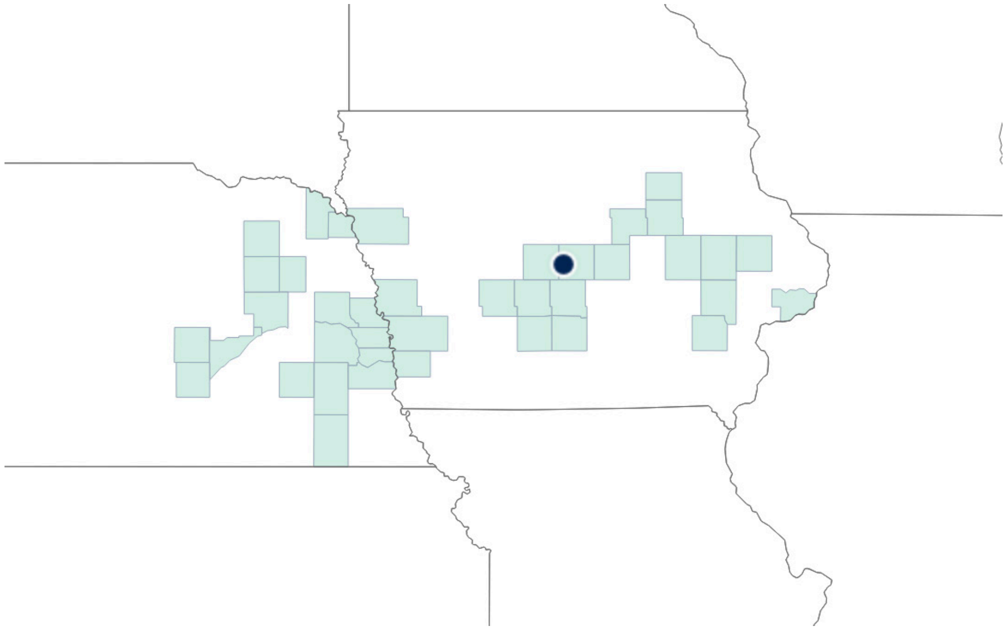

Source: National Science Foundation, "Full Proposal Invitations by Lead Organization Location," January 30, 2025, https://tableau.external.nsf.gov/views/NSFEnginesFullProposalInvitations/MapOverview.

Yet the CSA's funding language for TIP was an authorization, not an appropriation, leaving open what Congress ultimately spends. If Congress antes up, TIP will represent an enormous bet on the economic and social impact of research investments. If Congress does not, the expectations that the CSA created may hollow out traditional NSF research areas in service to TIP's goals, with potentially devastating consequences for the very societal, national, and geostrategic interests the act seeks to bolster.

Either way, the new measurement and reporting requirements are here to stay. The ability to clearly, accurately, and reliably document what these investments do may determine the act's lasting consequences. Today, to the best of my knowledge, nobody can provide that documentation.

Key Measurement Challenges

No one can do this yet because there are real challenges. The first is *classification*. The key technology focus areas are neither scientific research fields nor industries as they are conceptualized in any current classification system. Reliably measuring inputs and outputs even for one field, AI, is difficult. We need to do this at scale, robustly, and for all of them. The current classifications, therefore, need serious rethinking.

The second challenge involves data requirements for *granularity* and *aggregation*. TIP's signature program is Regional Innovation Engines (RIE). In considering applications for funding, RIE currently defines regions in terms of counties and specifies a list of technologies that a candidate's work may involve. Figure 1 presents the "region of service" for one RIE finalist's proposal from TIP's website. RuralSTAMINA—an AI-enabled biomanufacturing project, anchored at Iowa State—proposed to serve a region defined by five separate clusters of counties that span the boundaries of two states.

Meeting CSA mandates for the RIE program requires the capability to identify (1) the fields in which relevant investments are made and innovative solutions produced and (2) the industries in which jobs, manufacturing, education, and workforce capabilities might result. Moreover, that capacity must encompass any possible pairing of one or more CSA technologies with any combination of counties. Granular technology by county data must be aggregable to the national level to examine program effects. Finally, starting to assess impact necessitates, at a bare minimum, the ability to describe change over time.

Ideally, we would want causal estimates based on various identification strategies. So a mature classification system should include mechanisms for research access, complete with the means to ingest and link new data, appropriate privacy protections, essential data security, and all the other technical and legal infrastructure needed to allow responsible restricted data use. This points to the third challenge.

This work will require *new institutional and governance mechanisms* that engage many autonomous, coequal data owners—including states, corporations, federal agencies, and universities—in collaboratives with clear value propositions for each participant.[3] Mechanisms to streamline formation and work by many different partnership configurations to meet different needs will be necessary.

These requirements mirror the regional multisector collaboration for national goals and the logic of the CSA. They depart from traditional organizational models for national measurement, which typically center on the federal statistical agencies.[4] One emerging, complementary alternative is a variety of data federalism characterized by regional networks of public-private partnerships that are anchored in relationships between universities and state agencies.[5] Such networks can productively include multiple types of federal partners.

A Conceptual Framework for Classification

The Industries of Ideas project is building the first end-to-end prototype system that integrates all the core components for measuring the outcomes of federal technology and research investments, piloting measures to meet CSA mandates, and doing broad outreach to develop plans for scaling. Our measurement approach anchors technology and industry classification, measurement, and system design in the people whom research investments support. We propose a new strategy using social science concepts, network science measures, and deep-learning models.

All the essential pieces of this system exist and have been individually proved through at least a decade of use. But a strong and clear classification strategy is essential. That strategy must (1) account for real differences among technologies and (2) work across current CSA technologies and be expandable to new priority areas that arise. Three conceptual steps are necessary to meet that challenge.

Treat Research Fields as a Form of Social Organization

One of the most visible critical and emerging technologies is artificial intelligence, which faces all the challenges outlined above. Technologies like AI are cases of a broader phenomenon long studied by social scientists. Such research areas are institutional fields: recognizable arenas for collaborative and competitive work by diverse people and organizations in and through evolving networks.[6] Understood as forms of social organization, all contemporary research fields have more commonalities than differences. Indeed, one of the most influential early descriptions of such fields famously asked why we consistently observe so much similarity among social and economic competitors.[7]

Research fields share inputs—including funding sources and talent pools—and outputs, such as publications, patents, and trained people. The work that turns inputs into outputs happens under similar rules of the game. Some rules, like those governing federal grant review and conflicts of interest, are more formal. Others, like expectations about what constitutes a strong dissertation or an important finding, are more informal. Research fields sometimes share broad logics of action,[8] like the sensibility—"a rough sense of direction and an imperative to 'get on with it'"—that some attribute to the field of AI.[9]

Different research fields also involve many of the same players. It would come as little surprise if cybersecurity and AI were both defined by work done and people trained at the same 30–50 universities, by the employees of a similar number of firms, and by the investments of a much smaller group of funders. Perhaps more surprisingly, many of the same organizations are likely to be central to synthetic biology and advanced materials research. The organizations that shape different fields are certainly not identical, but substantial overlap is common and leads to similarity.

Careers work similarly. The most successful people in each field move across them throughout their careers, collaborate with one another, hire each other's students and trainees, and build on, critique, or formally review each other's work. They serve together on the program, prize, and study committees that set research agendas and define the formal and informal rules of the game. The best students often gravitate to the most central institutions and well-connected researchers, as do many of the most attractive employers. These dynamics pose significant challenges for efforts to expand equity and broaden participation in research, but they drive similarities in how research fields operate that we can leverage for classification.

Conceptualizing technology areas as research fields and emphasizing similarities provides a solid, general basis for designing, training, validating, and explaining a technology-agnostic classificatory model. Identified divergences, in contrast, provide leverage to capture salient differences in technology-specific model inputs. Anchoring technology-specific inputs and a technology-agnostic model in a common, fairly intuitive conceptual framework offers a convenient balance of flexibility, generalizability, and explicability that serves the needs of this use case.

Shift from Documents to People to Characterize Fields

Speaking generically, most current research classifications

- Begin with a document corpus,
- Extract representation (e.g., topics) of the work they report from some portion of their text,
- Cluster those representations in some abstract knowledge space, and
- Use some algorithmically identified subset of representations to characterize a field.

This conception of a field is fundamentally different from our social and organizational definition. It makes social organization the outcome of relationships among representations of ideas extracted from documents. It also obscures many sources of commonality and stability that enable generalizable tools and clear measures of change. Moreover, it increases sensitivity to rapid alterations in technology, substance, or terminology. These problems are amplified at the frontiers of fast-moving, multidisciplinary research areas such as AI.

In the Industries of Ideas model, we reverse the traditional logic and begin with social organization to identify people working in a field. Rather than saying AI researchers are those who work on specified topics, we determine the topics that constitute AI at a particular time from the work portfolio that AI researchers produce. Our people-centric classification follows the "operational definition" proposed by the One Hundred Year Study on Artificial Intelligence panel: "AI can also be defined by what AI researchers do."[10]

Starting with people has four appealing features. It

- Is resilient to rapid terminological, methodological, and content shifts;
- Allows clear measurements of change in a field's content, because people shift much more slowly than the topics they study;
- Leverages network theory and measures to operationalize key dimensions of group leadership, which supports a clear, explicable approach to data and model design; and

- Supports a framework and data architecture for industry classification that directly connects specific research investments to jobs and employers.

Shifting from documents to people is a fundamental move that distinguishes the Industries of Ideas classification from what has come before.

Follow the People to Identify the Industries

Papers, patents, and other documents are important, but putting people first is useful. It aids with the practical necessities of data linkage and system building and addresses the key role of tacit knowledge in work on research frontiers.[11] In so doing, it provides empirical and conceptual routes to connect investments in specific research fields to outcomes in clearly related industries that aren't identifiable by standard means. Put simply, people are the primary output of research investments.

Empirically, the Industries of Ideas model uses people's careers to integrate a linked data architecture that reaches from grants through university HR records to state workforce and employer information. Conceptually, we treat employers' decisions to hire people trained through those research investments as concrete, and often costly, commitments to the continued development or application of technologies related to the initial research investments.[12] If AI is what AI researchers do, AI industries are those in which employers seek and hire AI researchers.

As Paul Romer has noted: "Universities produce both papers and people. People with specialized problem-solving skills are the essential input into the discovery process, most of which takes place in the private sector. People with these skills are fuel that fires the innovation engine."[13]

I have made a similar argument about how research universities make distinctive contributions to the public good.[14] The framework outlined in this chapter is precisely aligned with the conceptual logic that governed the design of the Institute for Research on Innovation and Science; the Universities: Measuring the Impacts of Research on Innovation, Competitiveness, and Science (UMETRICS) data;[15] and ongoing links between UMETRICS and administrative workforce data maintained by states and shared with federal agencies. Figure 2 presents an overview of the entire prototype system that the Industries of Ideas is building for two research fields, AI and electric vehicles, in a single state, Ohio.

The research field classification, which is the primary input to this system, is our focus in this chapter. However, it is worth briefly discussing the other data this system will generate. The key challenge of any research classification that draws primarily on bibliometric data is that it will miss people who developed relevant skills through research but have not authored papers. That is quite common, either because their roles do not generally lead to authorship (for instance, in the case of staff) or because they are excluded from author lists.[16]

We rectify that problem by linking data on authors and their grants to transaction-level UMETRICS data, which lets us see everyone those grants pay, whether they are faculty, staff, undergraduate students, doctoral students, or postdocs. This move dramatically expands our definition of the impact of research investments and captures a much more extensive "tracer" set of individuals whose later career mobility can allow us to describe job, employer, and workforce implications for newly identified sets of industries.

We classify the industries relevant to research fields using existing data infrastructures that are maintained by states and contain detailed wage records for all employees who are eligible for unemployment insurance. Universities are among the employers, so the people—both those who are identified as authors and those paid on grants who are not—represented in university administrative data can easily be linked to state unemployment insurance wage records. Their post-university employment describes the employers that have bid in to a relevant research field by hiring research-trained people from that field. The people they hire can be connected back to concrete research investments associated with specific technologies.

Figure 2. The Industries of Ideas Data System

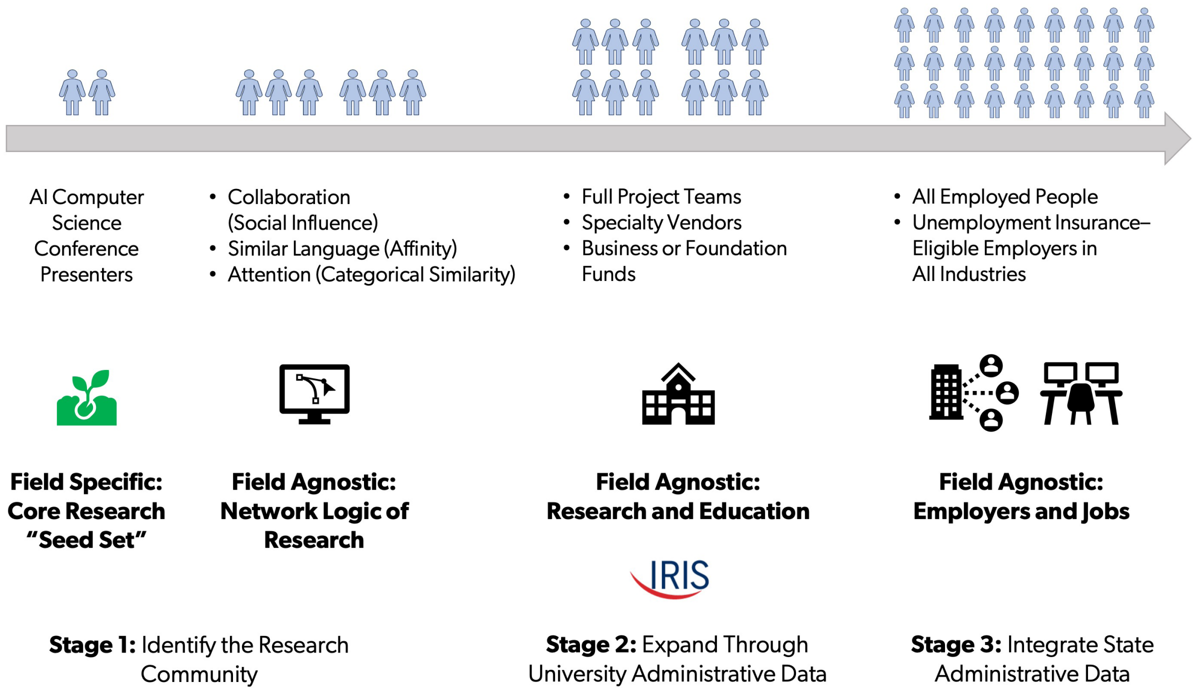

Source: Author.

Once we characterize the employers associated with a research field, we can richly describe all the jobs they support and all the other employers and jobs in their traditionally categorized industries. People and their careers provide the conceptual and empirical through line that makes both the system design and a measurement approach to address CSA mandates possible. With appropriate institutional arrangements and partnerships, links to other types of state data such as K–12, higher education, or social service information may become possible.

A strong classifier is both an essential part of the larger data system and a tool we can use with data from that system to expand our analytic horizons. The results could simultaneously advance knowledge in several fields *and* generate high-value products for many stakeholders, creating a virtuous cycle and accelerating growth to scale.

Implementing a Classification Strategy

A good field-classification strategy should apply to multiple technology areas and be broadly legible to nonspecialists. Specifically, a strong classification approach should do four things:

- **Identify Core Technical Contributors.** In AI, it should confidently classify researchers developing models and techniques at the frontiers of the field's current technical core.

- **Identify Researchers Applying Tools from the Field in Diverse Domains.** AI tools and models are used for many purposes across biomedicine, genetics, physical science, materials, information and social science, engineering, and other areas. Fully measuring the impact of research investments and effectively classifying industries require reliable capture of sophisticated applications in many substantive areas. For some CSA technology areas, including AI, it may also be useful to identify researchers whose work addresses relevant ethical, legal, and social issues.

- **Classify Individuals.** One core task of a classification system is, well, classification. However, the first two goals highlighted above suggest the benefits of a categorical approach rather than a binary approach.

- **Provide Quantitative Estimates.** Given the cross-cutting nature of many CSA technology areas, a regression approach that estimates the intensity of a given researcher's engagement with AI or a continuous measure of proximity to the field's technical core may prove more useful than a classification. For instance, researchers exclusively focused on technical improvements to deep-learning models should be closely associated with AI. Researchers who improve such models and work on non-AI substantive problems or those who address substantive questions about AI using non-AI methods might be less strongly associated with the field. Continuous estimates could be used to construct measures for inferential analysis (e.g., weighted averages) or to allow analysts to vary thresholds for different use cases. If a reasonable measure of direction could be incorporated into a proximity estimate (indicating, for example, that a researcher is moderately distant from the AI core in the direction of social science) to create a vector, even more nuanced analyses might be possible.

Step 1: Defining a "Seed Set" of Researchers

The general approach we follow has two steps. First, consultations with domain experts help identify a plausible "seed set" of researchers. That seed group should offer an easily definable, face-valid input or starting point for a more general, model-based second step, but it need not be comprehensive. It should consider the research field's distinctive social organizational characteristics.

This technology-specific seed definition will typically emphasize the field's core technical areas, but it should also include relevant domain researchers who are especially near that core. We do not intend this to be an exhaustive and exclusive gold standard. Rather, we seek a plausible initial definition of the field's core participants. By including people who might be "network bridges" to far-flung domains where technologies developed in or near the core are commonly applied, we ensure that models can identify researchers applying AI in disparate substantive areas.

For AI, Industries of Ideas researchers defined an initial seed set based on presentations at prominent AI and machine learning conferences. Table 1 presents a list of those conferences. In collaboration with analysts at Elsevier, we identified people who had authored at least one paper presented at any of these meetings since 2010. That highly international group included 94,235 authors.

Step 2: Model-Based Expansion

Seed-set authors and associated data about their publication and grant histories, coauthors, and affiliations provided the input for a technology-agnostic model in the second stage of the Industries of Ideas classification approach. We labeled AI seed researchers as a subset of all authors in the Elsevier publication corpus. Our collaborators at Johns Hopkins University then implemented a shallow Bayesian label-propagation model[17] on the unweighted coauthorship graph to classify unlabeled authors as potential AI researchers.

The model identified an additional 154,096 unique AI authors. It also successfully labeled 21,803 authors who were in the seed set but appeared in full coauthorship networks as unlabeled nodes before their first relevant conference presentation. In other words, more than 23 percent of seed-set authors were publishing actively before their first AI conference paper, were captured by the label propagation model, and were later observed as an author on an AI conference paper.

Table 1. AI Conferences That Define an Initial Researcher "Seed Set"

Conference Title
Annual Meeting of the Association for Computational Linguistics
Association for Advancement of Artificial Intelligence Conference on Artificial Intelligence
Association for Computing Machinery Conference on Recommender Systems
Association for Computing Machinery International Conference on Information and Knowledge Management
Association for Computing Machinery Special Interest Group on Knowledge Discovery and Data Mining Conference on Knowledge Discovery and Data Mining
Conference on Empirical Methods in Natural Language Processing
Conference on Neural Information Processing Systems
Conference on Uncertainty in Artificial Intelligence
Institute of Electrical and Electronics Engineers and Computer Vision Foundation Computer Vision and Pattern Recognition Conference
Institute of Electrical and Electronics Engineers and Computer Vision Foundation Winter Conference on Applications of Computer Vision
Institute of Electrical and Electronics Engineers International Conference on Data Mining
International Association for Computing Machinery Special Interest Group on Information Retrieval Conference on Research and Development in Information Retrieval
International Conference on Computer Vision
International Conference on Learning Representations
International Conference on Machine Learning
International Conference on Principles of Knowledge Representation and Reasoning
International Joint Conference on Artificial Intelligence
International Joint Conference on Natural Language Processing
Joint Conference on Lexical and Computational Semantics
Joint International Conference on Computational Linguistics
Robotics: Science and Systems

Source: Industries of Ideas project.

Seed-set and model-identified authors included 248,331 researchers. About 39 percent (97,379) listed US affiliations in any of their collective 1.96 million research publications. This represents an initial attempt at a person-centric, social, and organizational approach to AI classification that could meet the first three requirements I described for a classification strategy.

These nearly 250,000 people offer a reasonable first approximation of the researchers likely to possess AI-relevant skills globally. Such skills could make them and their students, postdocs, technical staff, and close collaborators attractive candidates for jobs in developing or applying AI tools and technologies to existing or new products, services, or business processes. Due to our definition of the seed set and our model design, this group includes core technical AI scientists and domain application researchers from many fields.

Using these people as a starting point for an industry classification that depends on hiring as a marker of employer engagement with AI will therefore identify firms active in a wide range of traditionally defined industries. We thus address a core challenge: identifying industries that aren't visible in traditional classifications. Nevertheless, this group features a key limitation in that it includes only people who authored papers published in venues indexed by Elsevier. This is why linking the results of a field classifier to UMETRICS data before proceeding to industry classification and associated measurement adds important substantive and technical components to the larger measurement system.

Improving a Classification Strategy

Validation and improvement of the Industries of Ideas approach are ongoing, along with a test of the entire process on a second technology field, electric vehicles. Early results suggest two important needs for refinement. First, this initial model may have difficulty identifying technical AI researchers from outside computer science—and computer scientists developing machine learning as a method to address other substantive issues—who do not attend AI conferences. Second and more importantly, the model is less effective at identifying domain application researchers. Accuracy seems to decline most steeply in biomedical areas.

Both challenges may result partially from model inputs. We may need to expand our definition of the seed set to include more researchers developing and applying AI methods to tackle substantive problems, especially in biomedicine. In this section, I address the possibility that the challenges result from our initial label propagation model's relative algorithmic and empirical simplicity. I suspect that a shallow model relying solely on the coauthorship graph imposes unnecessary limitations. As an alternative, I sketch a more thorough network conceptualization of group membership and outline a more sophisticated semi-supervised deep-learning model.

A Conceptual Approach to Improving the Classification Model

Beginning with people in a socially defined research field allows us to treat classification as, fundamentally, a problem of establishing individual membership in a social group. The basic question we want to answer is: Given some set of information about person X, can we say X belongs (or is more or less closely related) to group Y, for which we have comparable information about a set of known members? Here, X is an author of a scientific paper whose membership status relative to Y is unknown, and Y is the group of AI seed-set authors. The comparable information is a feature set drawn from bibliometric data.

Questions about groups and their members have been fundamental to sociology for more than a century.[18] More recently, network science has dramatically expanded relevant methods and measures.[19] This classification effort aims to establish group membership using networks relevant to socially defined research fields. Viewing it that way offers well-established routes to justify and explain data and model design choices that align closely with our overall conceptual approach. Such theoretical and methodological integration contributes to the clarity, generality, and legibility of the field classification and the larger measurement system of which it is an essential part.

Broadly, group membership is a function of three general mechanisms.

- **Social Proximity.** Traditionally, this is a literal question about physical proximity or kinship. But proximity can be framed socially in terms of nonfamilial relationships regardless of physical copresence or consanguinity. When it is applied to graphs of social relationships among people, social proximity animates most current community-finding algorithms.[20]

- **Affinity.** Group membership can be based on shared likes and dislikes. People who like, discuss, and attend to similar things in similar ways will often identify as members of a group and act according to that identity.[21]

- **Signaling.** Sometimes group membership is as much about others' perceptions as about an individual's actions or beliefs. People's actions can send signals to others (intentionally or not) about their membership in groups. Other people can ascribe membership status (erroneously or not) to individuals based on inferences drawn from their observations.[22]

Typically, group membership means some combination of the following: X is socially close to members of Y; X feels and demonstrates a connection to members of Y based on shared interests, language, and activities; and X makes (often symbolic) claims to membership in Y that others recognize and can accept or reject. This yields a few simple propositions. A researcher is an AI scientist if

- They are professionally connected to known AI scientists (indicating social proximity),

- They work on topics and use language similar to that of known AI scientists (indicating affinity), or

- Their professional activities allow others to infer they are AI scientists (indicating signaling).

Operationalizing Dimensions of Group Membership

Our current model focuses exclusively on one membership dimension, social proximity, measured via coauthorship. This singular focus might limit the model's ability to classify domain application researchers whose social distance from core technical AI could limit their likelihood of coauthoring with seed-set scientists. Social proximity is arguably the most direct measure of membership but may be too limiting for this use case.

Expanding model input data with features designed to operationalize affinity and signaling could address this limitation. Consider two concrete examples.

Measurement Example 1: Affinity

Affinity, framed in terms of language and interests, could be measured with vector embeddings from a pretrained language model. SciBERT,[23] a transformer model trained on a large scientific corpus, could be fine-tuned and used to position labeled and unlabeled authors in a "conceptual space" based on the text of their published abstracts. Distances among them could be calculated in many ways, either in cross-section or longitudinally.[24] Smaller distances would indicate greater affinity.

Distances among pairs of researchers could be represented as a valued network and treated as another route for label propagation. Affinity connections are likely to bridge gaps or collapse distances in coauthorship networks, as many people who are unlikely to ever work together directly study similar things in similar ways. So adding a second network dimension focused on affinity could significantly improve the model's performance.

Moreover, recent research demonstrates the value of fine-tuning multiple versions of transformer models with text from different periods to capture changes in the represented spaces.[25] This measurement approach could capture shifts in individual activities and larger alterations to the field as a whole. People's interests might change over time, but the landscape on which they pursue their interests is also dynamic.[26]

Measurement Example 2: Signaling

Signaling can be operationalized in numerous ways. One example that aligns with our social definition of fields relies on publication venues as markers of participation in particular research areas. Scientists' choices about where to publish their work can be distinct from that work's substance and methodology.

Figure 3. Multidimensional Network Example

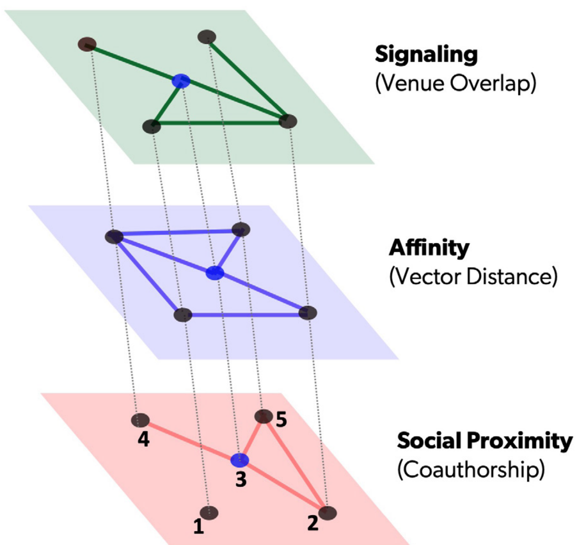

Source: Author.
Note: Node 3 (blue) is a labeled AI scientist identified in the seed set.

These decisions involve what intellectual communities to participate in and what professional identities to build and maintain.

If I have a paper that could fit in a policy, sociology, or management journal (as I often do), I decide where to send it based on which field I want to engage with and which audience I want to reach. My decision depends on the signal I want to send about how to understand the paper. Researchers often read curricula vitae with this kind of signaling in mind, because where one chooses to submit papers and where one succeeds in placing them tell insiders much about the kind of researcher one is or seeks to be.

Since scientific publication is a two-sided affair in which editors and reviewers select manuscripts partially based on fit, publication venues offer a clear example of a signaling mechanism. As a result, we might define group membership by the degree to which two scientists' scholarly works appear in overlapping venues. Greater similarity, by virtue of choices about submission and acceptance, equates to stronger signals of membership in the same group. We could use any of several specific measures calculated on a weighted or unweighted author, by journal matrix, in cross section or longitudinally. All of them could yield another valued network linking researchers.

A Multidimensional Network (with Light Formalization)

To illustrate, consider a simple cross-sectional data structure that could be implemented for any research field. The overall data structure would be a multidimensional graph eminently suited to tensor representation.[27]

Let the multidimensional graph $G = (V, D)$, where V is a set of nodes $(v_i, \ldots, v_N)$ and multiple sets of edges $(E_1, \ldots, E_D)$, connect those nodes in D dimensions. Edges in each dimension can be represented by a set of adjacency matrices $(A_1, \ldots, A_D)$, where $A_d(i, j)$ represents the connection in dimension d between node i and node j, which can be binary or valued. Nodes can also be represented by feature vector X, which, for simplicity, I limit for now to a single binary label that indicates whether a given scientist was identified as a member of the seed set of AI researchers.

Assume that nodes represent individual researchers and edges represent connections defined on three dimensions: (1) social proximity, expressed as $A_1(i, j) = 1$ when scientist i has coauthored with scientist j, and otherwise 0; (2) affinity, expressed as $A_2(i, j)$ = some real number value representing the conceptual similarity between scientist i and scientist j, defined via a vector embedding; and (3) signaling, expressed as $A_3(i, j)$ = some real number value representing the degree of commonality in publication venues between scientist i and scientist j. Figure 3 presents a schematic, hypothetical representation of this kind of graph.

Consider scientist 1, who is isolated in $A1$ by virtue of having no coauthoring relationships. In our initial model, which relied solely on social proximity, there was no path by which the AI label associated with node 3, an identified AI scientist, could propagate. In this multidimensional data structure, however, there are multiple direct and indirect paths for propagation through the signaling [$A_3(1, 3) > 0$] and affinity dimensions.

Likewise, nuance is added even when coauthorship ties are present. Consider the multiplex relationship connecting scientist 3 and scientist 5 [$A_1(3, 5) = 1$, $A_2(3, 5) > 0$]. The pattern of relationships between these scientists indicates that they coauthor together and that their abstracts use similar language but that (besides their coauthored work, which we would likely exclude in defining signaling ties) they publish in nonoverlapping journals.

It seems highly likely that a sophisticated model would classify scientist 5 as an AI researcher. However, this qualitative pattern of ties might intuitively suggest that scientist 5 is a domain researcher whose area of research serves as a case for AI tool development rather than a core AI researcher. Similarly, scientists 4 and 2, who have coauthored with scientist 3 and share affinity and signaling relationships with scientist 3, seem highly likely to be pursuing research closer to the core of the field, even though they are not observed in the seed set.

Modeling Considerations

This data structure is tailor-made for a deep-learning model that operates on graphs. Graph neural networks (GNNs) offer many possibilities for model designs that can consider multistep network neighborhoods, message passing, attention mechanisms, and node and edge features using a growing panoply of operators, including some tailored specifically for multidimensional data.[28]

While a first-pass model should focus on simple classification, the base architecture could readily yield a quantitative estimate of individual scientists' proximity to a field's core via a regression task. This kind of model applied to these kinds of data seems well equipped to satisfy all four criteria I suggested for a strong field classification. More importantly, a unified social scientific conceptual framework offers a plausible basis for expecting that it could be generalized across CSA technology areas.

Deep-learning models are basically opaque to explanation.[29] Nevertheless, the intuitions suggested by Figure 3 provide useful "hooks" for helping lay audiences understand our logic and, ideally, the results. The bedrock empirical insights we derive from well-established theory and findings can help concretize the workings of complex models. In this section, I draw heavily on (and borrow entire formalizations from) a recent review by Michael Bronstein et al.,[30] who partition spatial GNNs into three basic "flavors." I consider two, the convolutional and attentional, in detail to (1) lay the foundation for considering model design in this case and (2) suggest a general strategy for model design that makes this work more legible to nonexpert audiences.

All spatial GNNs use local network structure to update the state of a focal node, u, by appealing to the characteristics of the other nodes, v, to which it connects. The key idea is the network neighborhood $N_u = [v \mid (u, v) \in E]$. Essentially, a given GNN layer iterates across the neighborhoods of every node in a network, updating their states based on their neighbors' characteristics and then passing the new state information to a subsequent model layer for further processing.

For a network with both a node feature vector X and an adjacency matrix A, the GNN layer includes "*permutation equivariant* functions F(X, A) constructed by applying shared *permutation invariant* functions $\varphi(\mathrm{x}_u, \mathrm{X}_{Nu})$ over local neighbourhoods."[31] (Emphasis in original.) Thus, the GNN layer preserves some of the structural features of the larger network embedded in the adjacency matrix. It uses those features to weight the effects of nodes on each other, exploiting the complex interdependencies that are the hallmark of network data. The function φ is often called the "updating function."

I explore common approaches to φ to consider how a GNN classifier anchored in our conceptual and empirical framework might be rendered legible to the audiences that need to understand and use its results. What φ does, at the micro level of a single node and its immediate neighbors, is somewhat akin to label propagation, which ultimately answers our general question about group membership.

Considering this section's focus, that question could be restated: Given some mix of connections to nodes known to be AI researchers, how might differently configured GNN layers update information about whether a given scientist is also an AI researcher? What we are interested in, then, is h_u, the state of node u. In our case, that state might be called "AI-ness." This brings us back to some flavors of spatial GNNs defined by their general approach to φ.

Relying on Graph Structure: Convolutional Approaches

Convolutional GNN layers would address the question of a given scientist's AI-ness, h_u, in the following generic fashion:

$$(1)\ h_u = \varphi[x_u, \oplus_{v \in N_u} C_{vu} \psi(x_u)].$$

Ignore φ and ψ for now. They are generally affine, model-appropriate transformations that are learned through training and often modified by activation functions such as the rectified linear unit.

What is important to understand to get an empirical taste of the convolutional approach is the aggregation function ($\oplus_{v \in N_u}$) and C_{vu}. This function is what determines how you get from many neighbors, each with some state (in our simple case, a binary one) of AI-ness, to an updated state for u. The function $\oplus$ must be permutation invariant, which means its value will not change if you switch the order of the things it operates on.

Sums, averages, and maximums are all common aggregation functions. We will treat $\oplus$ as a sum. The other term, C_{vu}, is what determines the importance of each neighbor for updating u's state. In convolutional approaches, that term is a constant, which is generally defined by the adjacency matrix A, sometimes after transformation.

Stripped of all the gnarly but useful mechanics of machine learning, what these amount to in our simple case is a basic empirical intuition. A scientist's AI-ness is a function of the sum of the AI-ness of the other scientists to whom they are connected. Where our seed-set labels are binary, this basically indicates that a scientist has greater AI-ness when they are connected to more researchers who are themselves AI scientists. The more AI researchers with whom you coauthor, share affinities, and publish in similar venues, the more likely you are to be an AI scientist yourself.

There are complications. We have conceptual reasons to believe the different dimensions have different implications for establishing AI-ness. None of that matters to any deep-learning model we might eventually train. Right now, we cannot know how such a model will learn whatever it learns. Smarter people than me are working on the problem. But there are theory-based changes we could make to address some of our intuitions.

For instance, we could modify C_{vu} to deal with features of edges such as their values. If we make all three dimensions binary weights, we might simply scale the effect of a neighbor's AI-ness on u by the number of dimensions along which the pair are connected. Recall Figure 3 and the difference between scientist 1 (who had one signaling connection to scientist 3, the known AI researcher) and scientists 2 and 4, who each connected to scientist 3 in all three dimensions. More complicated approaches could allow continuous values for all ties and apply further transformations to those.

Regardless, convolutional approaches work primarily from network structure to update node states. In our case, that means whom you are connected to and whether they are part of the seed set is what matters.

Dynamic Weights Based on Attention

A somewhat more complicated flavor of GNN is *attentional*:[39]

$$(2)\ h_u = \varphi[x_u, \oplus_{v \in N_u} a(x_u, x_v)\ \psi\ (x_v)].$$

The only change here is replacing C_{vu} with $a(x_u, x_v)$, a self-attention mechanism. Instead of treating v's influence on u as a constant, driven by observed ties, this approach introduces a learnable parameter that calculates importance coefficients $a_{uv} = a(x_u, x_v)$, which are weights influenced by the features of neighbors. The underlying intuition might be summarized, in our network's social-proximity dimension, by the phrase "Not all coauthors are created equal." Some partners have characteristics that make them more important or salient to *u*'s state than others. The attention parameter creates a variable means to accommodate such differences in an aggregation function.

Reducing this general description to a specific empirical intuition in the highly simplified case (with a single, binary node feature) described above is trivial. So, imagine two possible complications. The first adds features to the node vector *X*. The second builds on the logic of the signaling dimension, A_3. In a research field such as AI, the core empirical intuition is that *u* might be connected to an AI scientist via one or more of our dimensions, but either member of the pair could be unaware of or uninterested in that connection.

Absent a reason for *u* to pay attention to an AI coauthor, the fact of coauthorship alone may have little effect on their state. This might be especially true in the fields where teams tend to be large or where people tend to publish many papers with a wide array of coauthors.

Consider high-energy physics or population genomics. Both are fields in which AI tools are increasingly broadly used. Both are also fields in which papers routinely have hundreds to thousands of coauthors. The fact of coauthorship with an AI scientist on such a paper may have little bearing on the AI-ness of a given physicist or genetics researcher. For what it is worth, things like norms about team size are exactly the research field characteristics that the social and organizational approach we propose could capture and that seed-set definitions might need to accommodate.

In the network that interests us, a node's state change is driven by the intertwined mechanisms of social influence (people become more like those to whom they are connected) and homophily (people are more likely to connect to partners to whom they are already similar). The idea that being connected to an AI scientist indicates one is more likely to be an AI scientist oneself presumes that (1) when non-AI researchers collaborate with AI researchers, their AI-ness increases because of the interaction (social influence) or (2) AI scientists are more likely to collaborate with one another, so observing a tie reveals an otherwise difficult-to-observe categorical similarity (homophily).

These two mechanisms are extremely difficult to tease apart,[32] but both depend on social forms of attention. Whether an existing similarity draws one to connect to a collaborator or an existing connection increases one's similarity to them, people must be aware of the characteristics they share, and those characteristics must be salient enough to factor in decision-making.[33] Two questions follow: (1) What might lead researchers to be more aware of their network neighbors' AI-ness? (2) What characteristics are likely to make some network neighbors more salient than others?

The answers may vary across our three network dimensions. The bar for awareness and salience seems likely to be higher for the affinity dimension, in which the mere fact that two scientists write in fashions that position them near each other in a complex vector space offers no assurance that they will know of each other or read each other's work. Regardless, additional node features or a refined conception of signaling can offer examples of potential answers.

Consider just one class of node features that might have traction. Higher-profile neighbors are noticeable and salient. Though it offers no information about actual GNN outputs, an illustrative attention mechanism based on scientific visibility suggests that neighbors who are more frequently cited, have won high-status awards, or are affiliated with high-visibility institutions or programs would exert greater influence on *u*'s state than others would. As I started writing this chapter, I used all these markers to guide my attention as I immersed myself in a large, complex, and wholly unfamiliar body of literature.

Once again, the benefits of a broadly social conception of research fields are apparent. For instance, we do not need to determine exactly which institutions have the highest status in a field to recognize that institutional affiliations are likely to be important attention features. We don't even need to observe how the ranks of institutions change across fields. We simply need to know that all fields have status hierarchies and that they shape researchers' attention. Knowing that, we construct input data that will increase our confidence and the legibility of a model design that includes attention mechanisms.

Alternatively, we might return to our network's signaling dimension, which fundamentally relies on attention. Recall the intuition behind our proposed measure. Researchers try to place work in venues that reach intellectual communities to which they belong or hope to belong. Publishing in those venues typically leads a researcher to conduct peer reviews for them, which focuses that researcher's limited attention even more tightly on a small set of journals and increases their awareness of others publishing or trying to publish in those venues. As careers progress, ad hoc reviewing turns into program committee or editorial board memberships and sometimes program chair or editorial roles. All these transitions further focus attention and deepen the visibility and the salience of others who publish in those outlets.

These generic social dynamics are common in contemporary research fields. They suggest that $A3$, the signaling adjacency matrix, measures signaling precisely because it represents a general structure of attention. This is entirely consistent with social science research that examines networks, categories, and signals in a wide range of fields.[34] More concretely, this requires that we attend to edge features in the signaling dimension as we think through the empirical logic of attention layers.[35]

There are a few takeaways. First, examining attention mechanisms requires us to consider (1) the conceptual relationship between the social mechanisms (homophily and influence) that underpin our intuitions about state updating at the level of nodes and their neighborhoods, (2) the signaling dimension of our multidimensional graph and its implications for attention in the social organization of research fields, and (3) the way both should shape choices about what node and edge features are necessary to maintain some conceptual, if not actual, explicability. Second, the discussion so far, which has focused solely on making layer-by-layer sense of baseline empirical intuitions for different GNN methods, already suggests implications for model design that, by and large, follow from the general conception of research fields that we propose.

Musing About Model Design

One of the many tricks of designing and training a working model appears to lie in stringing together an appropriate set of layers tuned to accomplish a particular task.[36] The discussion above suggests a model design based on a "block" composed of three layers. Our multidimensional network inputs would include three valued adjacency matrices ($A1$, $A2$, and $A3$) corresponding to the network dimensions (proximity, affinity, and signaling) and a vector of node features (X) composed of scientist-level measures defined by appeal to the social and organizational field definition and the task we want the model to perform.

We want an effective, technology-agnostic model and training data that accord with a conceptual framework that is legible to nonspecialists. These data would provide the initial input states for a convolutional layer, which would update those states based on the structure of connection in each node's one-hop neighborhood and pass the resulting, updated states to an attentional layer emphasizing node features, which in turn would pass the updated states to a second attentional layer focused on edge features.

This is where, as I understand it, the magic of deep-learning models kicks in. The learnable transformations we have been resolutely ignoring (φ, ψ) are updated across layers through the training process, as are the weights that result from each layer. As the input data pass through each layer of the model, the output of a prior layer serves as the input to a subsequent layer. Through training, the sequential application of different flavors of GNNs results in increasingly refined predictions.

The three-layer block we suggest could thus be understood to yield a progressively granular aggregation of neighboring node states, starting with the coarsest measure (connection) and proceeding through two levels of attentional weighting. Depending on specific decisions about input data (particularly the values of edges and the features in node vectors), this proposed architecture might emphasize different features of the broad social conception of group membership that underpins it. Equally importantly, the application of multiple blocks could allow learned weights to reflect broader neighborhoods, so a second block might be understood to encompass weights aggregated at network degree 2 (collaborators of collaborators), a third at network degree 3 (collaborators of collaborators of collaborators), and so forth. Of course, many details need to be worked out, but a general model architecture like that sketched in Figure 4 offers a starting point.

Whether the final details of the technology-agnostic model end up following these suggestions is not important. What is important is that any model architecture in the larger Industries of Ideas data system align with the general principles that helped us articulate this one:

- Emerging technology areas are research fields defined in social and organizational terms.

- People define the boundaries and contents of research fields.

- Field classification is a two-stage process in which

    - Domain experts help define a technology-specific researcher seed set that includes core technical contributors and people pursuing domain applications in a range of areas, and

    - Social theory and network measures inform technology-agnostic data structures and a semi-supervised GNN model to expand the seed set via classification and regression tasks.

- The results of that field classification are essential inputs to an integrated, people-centric data system that

    - Uses UMETRICS data to expand from authors to all grant-employed research personnel working on a particular technology;

    - Links authors and nonauthor university grant employees to state administrative unemployment insurance wage data;

    - Uses employee mobility to identify specific employers that have hired people trained on grants relevant to that technology;

    - Classifies industries by treating hiring as evidence that employers have bid in to pursue work relevant to applying or developing the technology in their products, services, or processes; and

    - Uses "hiring bids" as indications that descriptions and estimations of the impacts of research investments should include employers' traditionally defined industries, the other employers in them, and the jobs and the people who hold them.

Following these general principles to classify fields and industries across CSA technology areas supports a measurement system with all the components needed to rigorously address the challenges of the act.

Brief Thoughts on AI-Relevant Extensions

Integrating a social and organizational classification strategy with the larger Industries of Ideas frame-work and data system creates immediate opportunities to address additional AI-specific measurement concerns.

Measuring the Impact of Platform Investments

One major concern expressed in the National Artificial Intelligence Research Resource Task Force's report is the availability and accessibility of compute and test bed resources.[37] But the impact of such platform investments can be difficult to evaluate. The users of shared platforms typically hail from many fields and institutions. The platforms themselves generally collect information only on users. This gap makes identifying a comparison group impossible, creating a roadblock for assessment. The prototype data system described here, in partnership with compute and test bed operators, is uniquely suited to address this challenge.

UMETRICS data on the complete set of grant-funded people working at universities, coupled with employment information derived from state unemployment insurance wage data, offer many possibilities for developing counterfactual scenarios. Indicators and quantitative measures of AI-ness drawn from a strong field classification add more. What is needed once such a system is constructed is simply the ability to link user information for compute and test bed resources to the larger data infrastructure at the individual level. That integration enables many research designs to assess scientific, technological, and workforce impacts and could guide more effective development and use of shared AI resources.

Software and Model Development as an Impact of Research Investment

One impressive analysis in Stanford University's AI Index tracks technical improvements in model development.[38] This is an essential component of the field that I do not address. It could be systematically examined by integrating another type of data into the Industries of Ideas ecosystem.

The Institute for Research on Innovation and Science, which I cofounded with Julia Lane and Bruce Weinberg, is working on a pilot project to describe the relationship between research funding and the development and use of research software. We extract software mentions[39] from the text of scientific papers and then link named packages, the papers, and the authors who use them to software repository metadata and UMETRICS data. These linked data connect research investments, research publications, and software tools. As with the Industries of Ideas framework, those links follow the people. By virtue of their connection to repository metadata, they also provide many new data points about the code and its use, development, maintenance, updating, and other features.

Figure 4. Sketch of a GNN Architecture

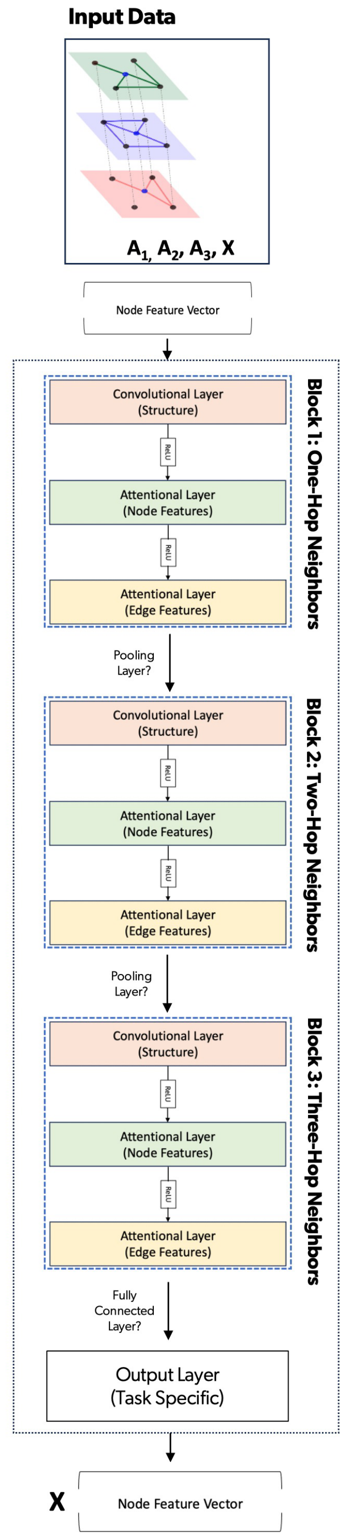

Source: Author.
Note: "ReLu" stands for rectified linear unit.

Research software that implements new AI models, like other intermediate research products such as datasets,[40] is an important research output that is not often systematically studied. In the field of AI—where most papers are openly accessible in full-text formats; two major Python packages (PyTorch and Tensorflow) are the primary basis for developing, expanding, and applying new machine learning models; standard benchmark datasets are commonly used and referenced; and code is generally shared upon the release of working or conference papers—linking code repository metadata into a unified data architecture could offer immense value.

First, such data could expand work like that of Stanford's Institute for Human-Centered Artificial Intelligence, which tracks the characteristics and use of new models. Second, it might allow us to add entirely new classes of inputs, outputs, and, potentially, people to the field and industry classifications we describe here. Third, we might expand affinity and signaling dimensions for field classification input data by using the same tools or datasets.

Finally, to the extent that code developers and maintainers do not completely overlap with authors on papers, we might be able to identify yet another group of people touched by AI research investments whose careers could help classify industries. While such expansions might not be generalizable across all CSA fields, they seem likely to be particularly valuable for those in which code and data are both core outcomes and enabling tools, such as AI, cybersecurity, and distributed ledger technologies.

Modeling the Dynamics of Regional Innovation Ecosystems

The CSA is one of several substantial place-based investments made under the Biden administration. Some, including the CSA, are larger than similar Great Society and even New Deal programs. Associated measurement challenges reach beyond job and employer effects to a program's holistic impact on the dynamics and outcomes of regional innovation ecosystems.

Two features that distinguish successful regional innovation ecosystems from unsuccessful ones are (1) the founding of "second-generation" startups[41] and (2) the formation of dense cross-employer mobility networks. Early research on the success of Silicon Valley highlighted both, citing the "Fairchildren" firms that grew from Fairchild Semiconductor and the ability to change jobs without changing carpools, which explains the resilience of Silicon Valley relative to that of Boston's Route 128.[42] Learning by hiring remains a key source of competitive advantage in technology-intensive regional ecosystems.[43]

A high-profile example from contemporary Bay Area AI clearly illustrates the point. Consider "Attention Is All You Need,"[44] a 2017 paper that helped spark the explosion of large language models and that has been cited almost 184,000 times. It was written by eight authors at Google. By 2021, all eight had departed.

Searching for these authors today reveals that they have collectively worked for or founded nine AI companies since leaving Google. Six of the original eight authors remain in the Bay Area. Seven of the nine companies, most of which were founded after the paper's publication, are also located in that region. The mobility networks among employers created by such moves and the founding of such "later-generation" firms are largely what make regions like the Bay Area such successful and resilient technology ecosystems.

But the data and tools needed to systematically (1) assess when and how various investments might help regions become self-sustaining, (2) support second-generation entrepreneurship, and (3) predict how such networks develop have never existed at scale. The data systems sketched here, which match rich, though restricted, workforce data with detailed bibliometric and university information, could change that dramatically. Long-time series data for established and emerging regions with and without investments from programs like the RIE or the US Economic Development Administration's Technology and Innovation Hubs allow pre- and post-investment analyses. Graph-based deep-learning models trained on those data could offer a wide range of new insights at the regional level.

Models that treat regions as networks and use graph classification to identify structural features that are leading indicators of regional success offer one interesting possibility. Link-prediction tools that focus on the evolution of mobility ties within a region (or into it) offer another. In the latter case, we might envision exciting possibilities for understanding regional ecosystem development around anchor tenants such as universities[45] and the role of key "on-ramp" institutions such as community colleges in ensuring broad access to jobs created by programs like RIE. Combining deep-learning models with granular administrative data could go a long way toward answering questions about educational and workforce capacity and their relationship to regional dynamism. Including state-maintained higher education or K–12 data in the Industries of Ideas ecosystem would dramatically increase the ability to address such questions.

If education data could be included in the mix, it might become possible, for the first time, to systematically assess the essential role of skilled technical workers in creating and sustaining regional success. This analysis would be particularly important for answering questions about manufacturing capacity and manufacturing-oriented CSA technology areas. In the context of reciprocal partnerships and effective governance, the need for strong privacy protections on restricted data could make famously complex and inexplicable models an important selling point.

Notes

1. CHIPS and Science Act, Pub. L. No. 117-167, § 10381.
2. CHIPS and Science Act, Pub. L. No. 117-167, § 10387.
3. Jason Owen-Smith, *AI and Work: How to Build New Data Foundations*, American Enterprise Institute, October 21, 2024, https://www.aei.org/research-products/report/ai-and-work-how-to-build-new-data-foundations/.
4. Julia Lane, *Democratizing Our Data: A Manifesto* (MIT Press, 2020).
5. Wan-Ying Chang et al., "A Linked Data Mosaic for Policy-Relevant Research on Science and Innovation: Value, Transparency, Rigor, and Community," *Harvard Data Science Review* 4, no. 2 (2022), https://pubmed.ncbi.nlm.nih.gov/36313611/; Frauke Kreuter et al., "Change Through Data: A Public Extension Program for Government Employees," *Harvard Data Science Review* 1, no. 2 (2019), https://hdsr.mitpress.mit.edu/pub/0mb0zzlc/release/7?readingCollection=af83430a; and Julia Lane, *Reimagining Labor Market Information: A National Collaborative for Local Workforce Information*, American Enterprise Institute, March 9, 2023, https://www.aei.org/research-products/report/reimagining-labor-market-information-a-national-collaborative-for-local-workforce-information/.
6. Walter W. Powell et al., "Network Dynamics and Field Evolution: The Growth of Interorganizational Collaboration in the Life Sciences," *American Journal of Sociology* 110, no. 4 (2005): 1132–205, https://www.journals.uchicago.edu/doi/10.1086/421508.
7. Paul J. DiMaggio and Walter W. Powell, "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields," *American Sociological Review* 48, no. 2 (1983): 147–60, https://www.jstor.org/stable/2095101.
8. Roger Friedland and Robert R. Alford, "Bringing Society Back In: Symbols, Practices, and Institutional Contradictions," in *The New Institutionalism in Organizational Analysis*, ed. Walter W. Powell and Paul J. DiMaggio (University of Chicago Press, 1991).
9. Barbara J. Grosz et al., *Artificial Intelligence and Life in 2030: One Hundred Year Study on Artificial Intelligence*, Stanford University, September 2016, https://ai100020201023.sites.stanford.edu/sites/g/files/sbiybj18871/files/media/file/ai100report10032016fnl_singles.pdf.
10. Grosz et al., *Artificial Intelligence and Life in 2030*, 13.
11. H. M. Collins, "Tacit Knowledge, Trust and the Q of Sapphire," *Social Studies of Science* 31, no. 1 (2001): 71–85, https://www.jstor.org/stable/285818; and Harry Collins, *Tacit and Explicit Knowledge* (University of Chicago Press, 2013).
12. Julia Lane, *The Industry of Ideas: Measuring How Artificial Intelligence Changes Labor Markets*, American Enterprise Institute, June 9, 2023, https://www.aei.org/research-products/report/the-industry-of-ideas-measuring-how-artificial-intelligence-changes-labor-markets/.

13. Paul Romer, “The Deep Structure of Economic Growth,” Paul Romer (blog), February 5, 2019, https://paulromer.net/deep_structure_growth/.

14. Jason Owen-Smith, *Research Universities and the Public Good: Discovery for an Uncertain Future* (Stanford University Press, 2018).

15. Natsuko Nicholls et al., *2022 IRIS UMETRICS Data Release Summary Documentation*, University of Michigan, Institute for Social Research, Institute for Research on Innovation and Science, 2022, https://iris.isr.umich.edu/wp-content/uploads/2022/04/Release22_Summary_Documentation.pdf.

16. Matthew B. Ross et al., “Women Are Credited Less in Science Than Men,” *Nature* 608 (2022): 135–45, https://www.nature.com/articles/s41586-022-04966-w.

17. Yuto Yamaguchi et al., “SocNL: Bayesian Label Propagation with Confidence,” in *Lecture Notes in Computer Science*, vol. 9077, *Advances in Knowledge Discovery and Data Mining*, ed. Tru Cao et al. (Springer, 2015), https://doi.org/10.1007/978-3-319-18038-0_49.

18. Georg Simmel, *Conflict and the Web of Group-Affiliations*, trans. Kurt H. Wolff and Reinhard Bendix (Free Press, 1964).

19. M. E. J. Newman, “Modularity and Community Structure in Networks,” *Proceedings of the National Academy of Sciences of the United States of America* 103, no. 23 (2006): 8577–582, https://www.pnas.org/doi/10.1073/pnas.0601602103; M. E. J. Newman and M. Girvan, “Finding and Evaluating Community Structure in Networks,” *Physical Review E* 69, no. 2 (2004), https://journals.aps.org/pre/abstract/10.1103/PhysRevE.69.026113; Vinh Loc Dao et al., “Community Structure: A Comparative Evaluation of Community Detection Methods,” *Network Science* 8, no. 1 (2020): 1–41, https://www.cambridge.org/core/journals/network-science/article/abs/community-structure-a-comparative-evaluation-of-community-detection-methods/3ED39917AAFA103F1AF3A62B0FB91052; and Sara E. Garza and Satu Elisa Schaeffer, “Community Detection with the Label Propagation Algorithm: A Survey,” *Physica A: Statistical Mechanics and Its Applications* 534 (2019): 122058, https://www.sciencedirect.com/science/article/abs/pii/S0378437119312026.

20. Mark Granovetter, “Economic Action and Social Structure: The Problem of Embeddedness,” *American Journal of Sociology* 91, no. 3 (1985): 481–510, https://www.jstor.org/stable/2780199; Mark S. Granovetter, “The Strength of Weak Ties,” *American Journal of Sociology* 78, no. 6 (1973): 1360–80, https://www.jstor.org/stable/2776392; Alejandro Portes, “Social Capital: Its Origins and Applications in Modern Sociology,” *Annual Review of Sociology* 24 (1998): 1–24, https://www.annualreviews.org/content/journals/10.1146/annurev.soc.24.1.1; and Mark T. Rivera et al., “Dynamics of Dyads in Social Networks: Assortative, Relational, and Proximity Mechanisms,” *Annual Review of Sociology* 36 (2010): 91–115, https://www.annualreviews.org/content/journals/10.1146/annurev.soc.34.040507.134743.

21. Jeffrey A. Smith et al., “Social Distance in the United States: Sex, Race, Religion, Age, and Education Homophily Among Confidants, 1985 to 2004,” *American Sociological Review* 79, no. 3 (2014), https://journals.sagepub.com/doi/10.1177/0003122414531776; and Miller McPherson et al., “Birds of a Feather: Homophily in Social Networks,” *Annual Review of Sociology* 27 (2001): 415–44, https://www.jstor.org/stable/2678628.

22. Michael Sauder et al., “Status: Insights from Organizational Sociology,” *Annual Review of Sociology* 38 (2012): 267–83, https://www.annualreviews.org/content/journals/10.1146/annurev-soc-071811-145503; Joel M. Podolny, “Networks as the Pipes and Prisms of the Market,” *American Journal of Sociology* 107, no. 1 (2001): 33–60, https://www.jstor.org/stable/10.1086/323038; and Dick Hebdige, *Subculture: The Meaning of Style* (Routledge, 1979).

23. Iz Beltagy et al., “SciBERT: Pretrained Contextualized Embeddings for Scientific Text” (working paper, ArXiv, March 2019), https://www.researchgate.net/publication/332011375_SciBERT_Pretrained_Contextualized_Embeddings_for_Scientific_Text.

24. Pedro Aceves and James A. Evans, “Mobilizing Conceptual Spaces: How Word Embedding Models Can Inform Measurement and Theory Within Organization Science,” *Organization Science* 35, no. 3 (2024): 788–814, https://pubsonline.informs.org/doi/epdf/10.1287/orsc.2023.1686.

25. Paul Vicinanza et al., “A Deep-Learning Model of Prescient Ideas Demonstrates That They Emerge from the Periphery,” *PNAS Nexus* 2, no. 1 (2023), https://academic.oup.com/pnasnexus/article/2/1/pgac275/6874776.

26. J. Miller McPherson and James R. Ranger-Moore, “Evolution on a Dancing Landscape: Organizations and Networks in Dynamic Blau Space,” *Social Forces* 70, no. 1 (1991): 19–42, https://academic.oup.com/sf/article-abstract/70/1/19/2232723.

27. Manlio De Domenico et al., “Mathematical Formulation of Multilayer Networks,” *Physical Review X* 3, no. 4 (2013), https://

journals.aps.org/prx/pdf/10.1103/PhysRevX.3.041022; Yao Ma et al., "Multi-Dimensional Network Embedding with Hierarchical Structure," in *WSDM '18: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (Association for Computing Machinery, 2018), https://dl.acm.org/doi/10.1145/3159652.3159680; Piotr Bródka and Przemysław Kazienko, "Multilayer Social Networks," in *Encyclopedia of Social Network Analysis and Mining*, 2nd ed., ed. Reda Alahajj and Jon Rokne (Springer, 2018), https://link.springer.com/referenceworkentry/10.1007/978-1-4939-7131-2_239; and Yao Ma et al., "Multi-Dimensional Graph Convolutional Networks," in *Proceedings of the 2019 SIAM International Conference on Data Mining*, ed. Tanya Berger-Wolf and Nitesh Chawla (Society for Industrial and Applied Mathematics, 2019), https://epubs.siam.org/doi/10.1137/1.9781611975673.74.

28. Ma et al., "Multi-Dimensional Graph Convolutional Networks"; Thomas N. Kipf and Max Welling, "Semi-Supervised Classification with Graph Convolutional Networks" (working paper, arXiv, February 22, 2017), https://arxiv.org/abs/1609.02907; Petar Veličković et al., "Graph Attention Networks" (working paper, arXiv, February 4, 2018), https://arxiv.org/abs/1710.10903; Michael M. Bronstein et al., "Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges" (working paper, arXiv, May 4, 2021), https://arxiv.org/abs/2104.13478; Liyu Gong and Qiang Cheng, "Exploiting Edge Features for Graph Neural Networks," paper presented at the IEEE and Computer Vision Foundation Conference on Computer Vision and Pattern Recognition, Long Beach, CA, June 15–20, 2019, https://ieeexplore.ieee.org/document/8954414; Davide Bacciu et al., "A Gentle Introduction to Deep Learning for Graphs," *Neural Networks* 129 (September 2020): 203–21, https://www.sciencedirect.com/science/article/abs/pii/S0893608020302197; and Yu Zhou et al., "Graph Neural Networks: Taxonomy, Advances, and Trends," *ACM Transactions on Intelligent Systems and Technology* 13, no. 1 (2022): 1–54, https://dl.acm.org/doi/10.1145/3495161.

29. Nadia Burkart and Marco F. Huber, "A Survey on the Explainability of Supervised Machine Learning," *Journal of Artificial Intelligence Research* 70 (2021): 245–317, https://www.researchgate.net/publication/345971128_A_Survey_on_the_Explainability_of_Supervised_Machine_Learning; Pantelis Linardatos et al., "Explainable AI: A Review of Machine Learning Interpretability Methods," *Entropy* 23, no. 1 (2021), https://www.mdpi.com/1099-4300/23/1/18; and Alessio Plebe and Giorgio Grasso, "The Unbearable Shallow Understanding of Deep Learning," *Minds and Machines* 29, no. 4 (2019): 515–53, https://philpapers.org/rec/PLETUS-2.

30. Bronstein et al., "Geometric Deep Learning."

31. Bronstein et al., "Geometric Deep Learning," 78.

32. Cosma Rohilla Shalizi and Andrew C. Thomas, "Homophily and Contagion Are Generically Confounded in Observational Social Network Studies," *Sociological Methods & Research* 40, no. 2 (2011): 211–39, https://pmc.ncbi.nlm.nih.gov/articles/PMC3328971/; and Edward McFowland III and Cosma Rohilla Shalizi, "Estimating Causal Peer Influence in Homophilous Social Networks by Inferring Latent Locations," *Journal of the American Statistical Association* 118, no. 541 (2023): 707–18, https://www.tandfonline.com/doi/abs/10.1080/01621459.2021.1953506.

33. Noah E. Friedkin, *A Structural Theory of Social Influence* (Cambridge University Press, 1998); and Noah E. Friedkin et al., "Network Science on Belief System Dynamics Under Logic Constraints," *Science* 354, no. 6310 (2016): 321–26, https://pubmed.ncbi.nlm.nih.gov/27846564/.

34. Ezra W. Zuckerman, "The Categorical Imperative: Securities Analysts and the Illegitimacy Discount," *American Journal of Sociology* 104, no. 5 (1999): 1398–438, https://www.journals.uchicago.edu/doi/abs/10.1086/210178; Helena Buhr et al., "The Authenticity Premium: Balancing Conformity and Innovation in High Technology Industries," *Research Policy* 50, no. 1 (2021), https://www.sciencedirect.com/science/article/abs/pii/S0048733320301633; Elizabeth. G. Pontikes, "Two Sides of the Same Coin: How Ambiguous Classification Affects Multiple Audiences' Evaluations," *Administrative Science Quarterly* 57, no. 1 (2012): 81–118, https://www.sciencedirect.com/science/article/abs/pii/S0048733320301633; G. Hsu et al., "Multiple Category Memberships in Markets: An Integrative Theory and Two Empirical Tests," *American Sociological Review* 74, no. 1 (2009): 150–69, https://www.researchgate.net/publication/237969806_Multiple_Category_Memberships_in_Markets_An_Integrative_Theory_and_Two_Empirical_Tests; Giacomo Negro et al., "Research on Categories in the Sociology of Organizations," *Research in the Sociology of Organizations* 31 (December 2010): 3–35, https://www.researchgate.net/publication/235251408_Research_on_categories_in_the_sociology_of_organizations; and Chad Navis and Mary Ann Glynn, "How New Market Categories Emerge: Temporal Dynamics of Legitimacy, Identity, and Entrepreneurship in Satellite Radio, 1990–2005," *Administrative Science Quarterly* 55, no. 3 (2010): 439–71, https://journals.sagepub.com/doi/10.2189/asqu.2010.55.3.439.

35. Gong and Cheng, "Exploiting Edge Features for Graph Neural Networks"; and Bacciu et al., "A Gentle Introduction to Deep Learning for Graphs."

36. Such a design is beyond the scope of this chapter and, frankly, beyond my skills. But that will not stop me from speculating and asking that anyone who might be interested in collaborating or simply correcting me reach out.

37. Emily Grumbling et al., *Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource*, White House, January 2023, https://www.ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf.

38. Nestor Maslej et al., *Artificial Intelligence Index Report 2024*, Stanford University, Institute for Human-Centered Artificial Intelligence, April 2024, https://hai.stanford.edu/ai-index/2024-ai-index-report.

39. Patrice Lopez et al., "Mining Software Entities in Scientific Literature: Document-Level NER for an Extremely Imbalance and Large-Scale Task," in *CIKM '21: Proceedings of the 30th ACM International Conference on Information & Knowledge Management* (Association for Computing Machinery, 2021), https://dl.acm.org/doi/10.1145/3459637.3481936.

40. Julia Lane et al., "Data Inventories for the Modern Age? Using Data Science to Open Government Data," *Harvard Data Science Review* 4, no. 2 (2022), https://hdsr.mitpress.mit.edu/pub/g6e8noiy/release/2.

41. Elaine Romanelli and Olga M. Khessina, "Regional Industrial Identity: Cluster Configurations and Economic Development," *Organization Science* 16, no. 4 (2005): 344–58, https://pubsonline.informs.org/doi/10.1287/orsc.1050.0131; Maryann Feldman and Elaine Romanelli, "Organizational Legacy and the Internal Dynamics of Clusters: The U.S. Human Biotherapeutics Industry, 1976–2002," in *Knowledge and the Economy*, ed. Peter Meusburger et al. (Springer, 2013); Raphael Suire and Jerome Vicente, "Clusters for Life or Life Cycles of Clusters: In Search of the Critical Factors of Clusters' Resilience," *Entrepreneurship & Regional Development* 26, nos. 1–2 (2014): 142–64, https://www.tandfonline.com/doi/abs/10.1080/08985626.2013.877985; Walter W. Powell et al., "Organizational and Institutional Genesis: The Emergence of High-Tech Clusters in the Life Sciences," in *The Emergence of Organizations and Markets*, ed. John F. Padgett and Walter W. Powell (Princeton University Press, 2012), 434–65; and Michael S. Dahl and Olav Sorenson, "Home Sweet Home: Entrepreneurs' Location Choices and the Performance of Their Ventures," *Management Science* 58, no. 6 (2012): 1059–71, https://www.jstor.org/stable/41499541.

42. AnnaLee Saxenian, *Regional Advantage: Culture and Competition in Silicon Valley and Route 128* (Harvard University Press, 1994); and Timothy Bresnahan et al., "'Old Economy' Inputs for 'New Economy' Outcomes: Cluster Formation in the New Silicon Valleys," *Industrial and Corporate Change* 10, no. 4 (2001): 835–60, https://academic.oup.com/icc/article-abstract/10/4/835/706918.

43. Paul Almeida et al., "Startup Size and the Mechanisms of External Learning: Increasing Opportunity and Decreasing Ability?," *Research Policy* 32, no. 2 (2003): 301–15, https://www.sciencedirect.com/science/article/abs/pii/S0048733302001014; and Raphael A. Corredoira and Lori Rosenkopf, "Should Auld Acquaintance Be Forgot? The Reverse Transfer of Knowledge Through Mobility Ties," *Strategic Management Journal* 31, no. 2 (2010): 159–81, https://sms.onlinelibrary.wiley.com/doi/abs/10.1002/smj.803.

44. Ashish Vaswani et al., "Attention Is All You Need," in *NIPS '17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, ed. Ulrike von Luxburg (Curran Associates, 2017), https://dl.acm.org/doi/10.5555/3295222.3295349.

45. Owen-Smith, *Research Universities and the Public Good*; Ajay Agrawal and Iain Cockburn, "The Anchor Tenant Hypothesis: Exploring the Role of Large, Local, R&D-Intensive Firms in Regional Innovation Systems," *International Journal of Industrial Organization* 21, no. 9 (2003): 1227–53, https://www.sciencedirect.com/science/article/abs/pii/S016771870300081X; and Maryann Feldman, "The Locational Dynamics of the US Biotech Industry: Knowledge Externalities and the Anchor Hypothesis," *Industry and Innovation* 10, no. 3 (2003): 311–29, https://www.tandfonline.com/doi/abs/10.1080/1366271032000141661.

Quantifying the Impact of AI Innovation

Lee Branstetter, Eduard Hovy, and Prasanna "Sonny" Tambe

What is the impact of AI innovation on productivity? Business leaders have proclaimed a fourth industrial revolution centered on AI and related advances in information technology. However, earlier industrial revolutions featured significant and persistent increases in productivity growth that boosted living standards across the income distribution. Despite growing hype and concern over AI applications across the economy, aggregate productivity growth remains slow, which limits the growth of American incomes and prosperity.[1] Will AI fail to live up to its advocates' enthusiasm, or are we merely in the early stages of a process of innovation and adoption that will take years or decades to unfold?

Our research seeks to address this question by examining the vanguard of firms that are already introducing AI innovations to the marketplace.[2] If these early movers, innovators, and adopters are already reaping significant productivity gains, then this augurs well for AI's ultimate impact on the US economy.

In this chapter, we summarize recent research showing that AI invention significantly boosts the productivity and employment of AI-inventing firms. However, the limited number of AI experts trained up to the scientific frontier who seek industrial employment in the US may significantly constrain the human resources that most firms need to turn AI breakthroughs into new products and services.

Recent Research Measuring the Impact of AI

Our work complements several recent streams of research. One stream seeks to measure AI adoption and use through direct surveys of large, representative samples of US firms.[3] These valuable efforts have not yet convincingly demonstrated a causal relationship between AI adoption and faster productivity growth, and it will take time before these survey data acquire a time series dimension long enough that researchers can apply typical econometric techniques for discerning plausibly causal effects from nonexperimental data.

A second stream of research applies randomized controlled trials or quasi-experimental methods to measure AI's impact on productivity in a particular work context.[4] Some of these papers have shown convincing evidence that AI adoption boosts productivity, but researchers may not be able to generalize the results beyond the particular contexts of these experiments or quasi-experiments.

A third stream uses data on specialized labor recruitment to measure AI use and AI-related innovation.[5] Prominent papers in this stream have shown evidence that AI investments boosted output and product innovation but have failed to demonstrate robust evidence that these investments led to increased productivity growth. The strong productivity growth we document is potentially broader than that found in the experimental literature and points to the optimistic possibility that AI could eventually lead to a significant and persistent acceleration in productivity growth across a broad range of industries.

Our Methods, Data, and Current Results

Our work develops two additional streams of research. The first identifies AI-related patents, links these patents to the firms that invent them, and leverages the detailed firm-level data compiled by the US Census Bureau to examine the ways these firms evolve as they generate AI-related innovations.[6] The second identifies AI experts educated up to the technological frontier by the leading scientific experts and then tracks their movement across organizations, geography, and time.

In our research, we use an ensemble of machine learning algorithms to parse the text of US Patent and Trademark Office patents and identify AI-related inventions. We then match data on the AI patents to the rich databases that the US Census Bureau maintains on the production and inputs of US-based firms, including those that are not publicly traded. Using firm fixed effects models and event studies, we find striking evidence that AI invention leads to significant growth in employment and productivity. Many other studies have failed to find a significant impact of AI adoption or innovation on productivity—but we find one, and the evidence is robust.

We also examine the role of PhD-level academic experts in creating new AI-related goods and services. We hypothesize that a shortage of advanced human capital profoundly shapes where, and via which firms, AI innovation is advancing. By linking these experts and their trained students to the firms that employ them, we may obtain empirical leverage around the difficulty of measuring AI inventions that do not result in patents and the application of frontier or near-frontier AI ideas to reengineer existing products and services.[7]

Figure 1 summarizes the results from firm fixed effects models applied solely to US manufacturing firms, because the US Census only calculates firm-level total factor productivity and added value for manufacturing firms. The estimated impacts from the initial transition into AI patenting, which capture the extensive-margin effects, are labeled "Extensive Margin." The estimated impacts from an increase in the number of AI patents, which capture the intensive-margin effects, are labeled "Intensive Margin."

An alternative approach to estimating the impact of AI patenting on the inventing firms is to compare each AI-inventing firm with its closest same-industry, non-AI-inventing peer firms and determine the differences in key outcomes between the treated (AI-inventing) firm and the control (non-AI-inventing) firms after the former transitions into AI patenting. Figure 2 summarizes the results from these event studies. Here, too, the transition to AI patenting is associated with a statistically significant and economically meaningful increase in employment and revenue per employee.

Ongoing Research

Not all innovations are patented, and some industries that invest heavily in AI to generate new products and services hardly patent their inventions at all.[8] To identify investments in AI-related innovation in these sectors, we use publication data from Elsevier to identify the top academic scientists working in domains related to AI. We then use faculty websites, laboratory websites, and the ProQuest database to identify the graduate students whom they supervise.

Figure 1. The Impact of AI Patenting on Productivity: Percentage Increase as Measured by Fixed Effects Models

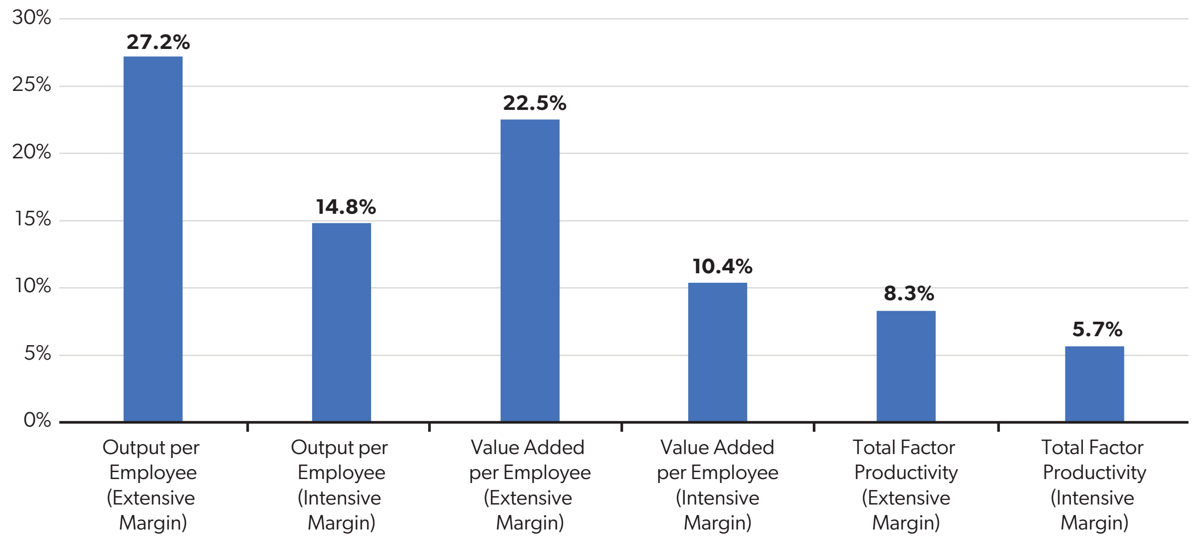

Source: Dean Alderucci et al., "Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Microdata" (working paper). Available from the authors on request.

Once these students are identified, we use data from Revelio Labs to track their movement across geographic areas, organizational boundaries, and time.

We can also use publication data to track the direct interaction between top academic scientists and the companies they work with when that interaction results in a publication. Once we link the star scientists and their students to the firms with which they have worked and trace these links over time, we can leverage US Census micro-data, obtained through our ongoing collaboration with Census microeconomists, to investigate whether these links have given the receiving firms a statistically discernible advantage in output, employment, or, most important, productivity over their same-industry peers that lack these connections.

This line of inquiry is related to the work of Tania Babina and coauthors,[9] but while that research attempts to measure the recruitment of AI workers at all skill levels, our work focuses on identifying elite AI scientists, who may play a disproportionate role in defining the technology frontier, and their doctoral students, who may play a disproportionate role in bringing frontier technology into industrial practice.[10] We then seek to track the direct collaboration between top AI scientists and firms and the movement of their graduated students into those firms. We conjecture that any firm seeking to apply frontier AI to substantively reengineer its products and services or create new ones must create an internal pyramid of AI talent, depicted in the left portion of Figure 3, that includes these high-level experts.[11]

At the lower ranks of the talent pyramid, the firm could productively employ programmers with self-taught AI skills who use standard AI tools and techniques. At the middle levels of the pyramid, the firm might need professionals with bachelor's or master's degrees that include specialized AI training, but these professionals need not have trained at elite universities. However, at the pyramid's apex, a firm seeking to out-engineer its rivals may seek software architects whom elite academic scientists at top universities have trained up to the technology frontier.

As we have noted, Babina and coauthors[12] use data from the entire pyramid, but we focus on star scientists and their students, who could constitute

Figure 2. The Impact of AI Patenting on Employment and Revenue per Employee: Percentage Increase as Measured by Event Study Models

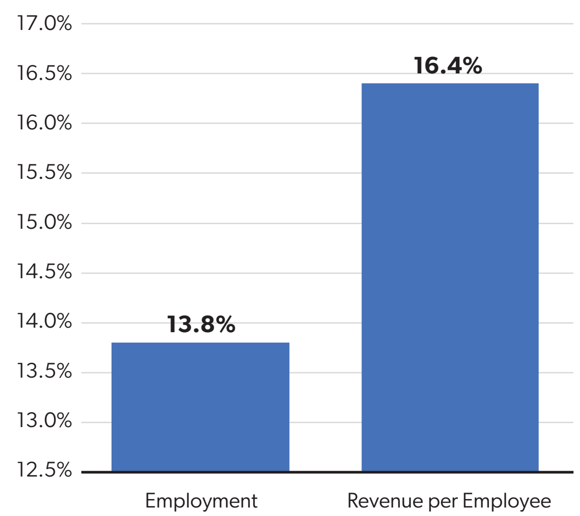

Source: Dean Alderucci et al., "Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Microdata" (working paper). Available from the authors on request.

a disproportionate part of the pyramid's apex. Their role is related to that of the architects described in the theoretical work of Seth Benzell and coauthors.[13] As in that paper, we consider that the limited supply of these software architects could substantially constrain firms' ability to fully leverage frontier AI technologies.

To illustrate, imagine that our data sources identify a Carnegie Mellon University doctorate recipient as an elite academic scientist's PhD advisee. This student's career moves are illustrated in the right side of Figure 3. Their move to Google DeepMind could augment that impressive corporate research operation's intellectual resources. The student begins to specialize at Google DeepMind in applying advanced AI algorithms to medical imaging. Then, they carry this skill to Siemens Healthineers and then to diagnostic imaging startup Arterys.

By following star scientists' students like this one among firms, we can trace their differential impacts, if any, on enhancing host firms' output, employment, and productivity. We can test the hypothesis that these movements predict success using census data drawn from the hiring firms and their same-industry peers that have hired fewer or no advanced AI experts. We can also track the changing impact of these experts as they accumulate experience and transition to higher positions under new employers.

Policy Implications and Trade-Offs

Providing a Strong, Data-Driven Rationale for Investing in AI

The US government possesses limited resources to invest in basic science. If AI boosts corporate profits but fails to boost employment or productivity, that could call into question the desirability of long-term, large-scale investment in AI research. Our results provide strong and robust evidence that the firms investing in AI-related innovation are seeing significant gains in employment and productivity. As the pool of AI innovators expands, these effects will likely appear in aggregate economic data.

Creating a Methodology to Identify AI Innovation Abroad and Compare Its Quality and Impact to Those in the US

Many in the US government worry that adversary nations are developing AI innovation capabilities that rival those of the US and its allies. The data and methods that we have pioneered could be directly applied to foreign (e.g., Chinese) patent data, shedding crucial light on the strengths and weaknesses of AI innovation in China while going far beyond simply counting large numbers of patents of low or average quality.

Tracking the Global Flow of AI Expertise

Preliminary evidence suggests what many industry insiders believe—that there is a global shortage of experts trained up to the scientific frontier who can help companies apply fundamental breakthroughs in the science of AI to develop new goods and services or reengineer existing ones. Our ongoing research will

Figure 3. Tracing the Impact of AI Software Thought Leaders and Architects

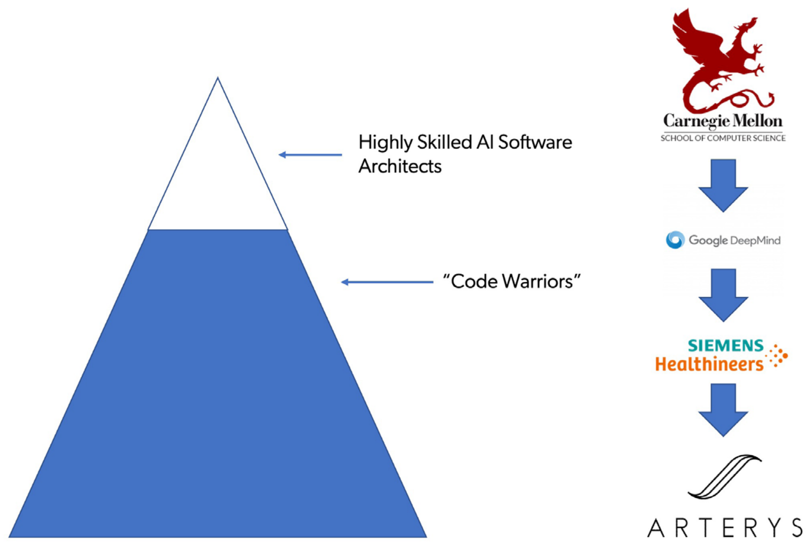

Source: Dean Alderucci et al., “Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Micro-data” (working paper). Available from the authors on request.

create a much more comprehensive database of these AI experts that tracks their movement from leading centers of scientific AI research to innovating companies and organizations worldwide. It will also assess how much the accumulation of this scarce human resource correlates with AI innovation and its impact.

Applications to Other Critical and Emerging Industries

In principle, the basic techniques we apply to AI could be applied to other critical and emerging industries. Our approach of using AI to identify AI-related patents could be adapted to other technological fields. Once identified, these patents could be linked to the inventing firms that generate them, and census micro-data could be leveraged to identify impacts on output, employment, and productivity. Similarly, publication data could be used to identify star scientists in other important technological fields, and the same mix of data from publications, websites, and online resume sites could quantify the interactions of star scientists and their students with firms. However, for this approach to be meaningful, there must be a strong connection between new scientific breakthroughs and their application in a wide range of industrial contexts.

Notes

1. Seth G. Benzell et al., “Digital Abundance Meets Scarce Architects: Implications for Wages, Interest Rates, and Growth” (working paper, Stanford Digital Economy Lab, January 25, 2022), https://digitaleconomy.stanford.edu/publications/digital-abundance-meets-bottlenecks-implications-for-wages-interest-rates-and-growth/.

2. Dean Alderucci et al., “Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Micro-data” (working paper). Available from the authors on request.

3. Kristina McElheran et al., “AI Adoption in America: Who, What, and Where,” *Journal of Economics & Management Strategy* 33,

no. 2 (2024): 375–415, https://onlinelibrary.wiley.com/doi/10.1111/jems.12576.

4. Erik Brynjolfsson et al., "Generative AI at Work," *The Quarterly Journal of Economics* 140, no. 2 (2025): 889–942, https://academic.oup.com/qje/article/140/2/889/7990658; Anton Korinek, "Language Models and Cognitive Automation for Economic Research" (working paper, National Bureau of Economic Research, February 2023), https://www.nber.org/papers/w30957; and Shakked Noy and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science* 381, no. 6654 (2023): 187–92, https://www.science.org/doi/10.1126/science.adh2586.

5. Tania Babina et al., "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition," in Susanto Basu et al., eds., *Technology, Productivity, and Economic Growth* (University of Chicago Press, 2025); and Tania Babina et al., "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics* 151 (January 2024): 103745, https://www.sciencedirect.com/science/article/pii/S0304405X2300185X.

6. This stream builds on Iain M. Cockburn et al., "The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis," in Ajay Agarwal et al., eds., *The Economics of Artificial Intelligence: An Agenda* (University of Chicago Press, 2019); Michael Webb et al., "Some Facts of High-Tech Patenting" (working paper, National Bureau of Economic Research, July 2018), https://www.nber.org/papers/w24793; Alexander. V. Giczy et al., "Identifying Artificial Intelligence (AI) Invention: A Novel AI Patent Dataset," *The Journal of Technology Transfer* 47 (November 2021): 476–505, https://link.springer.com/article/10.1007/s10961-021-09900-2; and Milan Miric et al., "Using Supervised Machine Learning for Large-Scale Classification in Management Research: The Case for Identifying Artificial Intelligence Patents," *Strategic Management Journal* 44, no. 2 (2023): 491–519, https://sms.onlinelibrary.wiley.com/doi/10.1002/smj.3441. A key difference in our work is that it links information on AI-related inventions to the inventing firms and examines how the firms change as they generate these inventions.

7. For related research exploring the impact of the movement of top AI professors into the industry and subsequent entrepreneurship by students at their former institutions, see Michael Gofman and Zhao Jin, "Artificial Intelligence, Education, and Entrepreneurship," *Journal of Finance* 79, no. 1 (2024): 631–67, https://onlinelibrary.wiley.com/doi/10.1111/jofi.13302. For research focusing on the role of co-ethnic academic advisers in linking immigrant talent to US firms, see Caroline Fry and Britta Glennon, "In Good Company: Coethnic Advisors and Career Paths of Immigrant Ph.D. Students" (working paper, National Bureau of Economic Research, May 2025), https://www.nber.org/papers/w33782.

8. Of course, many firms do not invent AI at all but simply adopt other firms' AI inventions. Our methods do not allow us to identify pure AI adoption; thus, we capture only part of AI's impact on the whole economy's productivity.

9. Babina et al., "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition"; and Babina et al., "Artificial Intelligence, Firm Growth, and Product Innovation."

10. Ajay Agrawal and Rebecca Henderson, "Putting Patents in Context: Exploring Knowledge Transfer from MIT," *Management Science* 48, no. 1 (2002): 44–60, https://pubsonline.informs.org/doi/10.1287/mnsc.48.1.44.14279; and Lynne G. Zucker et al., "Intellectual Capital and the Birth of the U.S. Biotechnology Enterprises," *The American Economic Review* 88 (1998): 290–306, https://www.jstor.org/stable/116831.

11. Ashish Arora et al., "Going Soft: How the Rise of Software-Based Innovation Led to the Decline of Japan's IT Industry and the Resurgence of Silicon Valley," *The Review of Economics and Statistics* 95, no. 3 (2013): 757–75; and Lee Branstetter et al., "Get with the Program: Software-Driven Innovation in Traditional Manufacturing," *Management Science* 65, no. 2 (2019): 541–58, https://pubsonline.informs.org/doi/epdf/10.1287/mnsc.2017.2960.

12. Babina et al., "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition"; and Babina et al., "Artificial Intelligence, Firm Growth, and Product Innovation."

13. Benzell et al., "Digital Abundance Meets Scarce Architects."

New Approaches to Characterize Industries Measuring AI

Diane Coyle

The purpose of economic measurement is to guide decision-makers in both the government and private sector. Yet even after some decades of technological transformation, modern economies' structure and trends in consumption and production are all but invisible in available statistics. About four-fifths of advanced economies can be characterized as "hard to measure."[1] This indicates the need for a different approach, which AI's current rapid development and deployment have made more acute, as have other tactics (including innovations such as quantum and biomedical technologies or processes such as additive or bio-manufacturing) in the short to medium term.

The challenges for statisticians are twofold. One is data collection. Existing classifications of sectors or occupations map poorly onto current production and consumption patterns, and there is resistance to updating classifications in ways that are not backward compatible or suitable for low- and middle-income economies. Many researchers have started using a range of innovative data-collection techniques in arrangements with private data companies or methods such as web scraping or processing open satellite data.[2]

Statistical agencies are endeavoring to foster better data access with tech companies, with limited success. For example, there is no official index of the price of cloud computing services because companies do not provide data, so availability indexes are constructed from web-scraped prices.[3] The use of cloud services, not capitalized in company accounts, has shifted business expenditure from capital investment to intermediate consumption, but the scale is unknown.

Recent research taking an engineering-based approach to computation costs—using two different methods to incorporate developments in AI—shows that the pace of decline has been substantially greater than that shown in any official price index.[4] In any case, official statistics in these technology domains are generally not timely enough for decision-making purposes. New measurement systems are essential to fill the empirical gaps.

This chapter is concerned, though, with the second challenge, which is conceptual. I suggest a framework for measuring the economic value of new technologies such as AI, given that general-purpose technologies cause structural shifts in the economy. I also highlight some key economic research questions.

The Economic Value of AI

At the heart of understanding economic progress is understanding the process of turning resources into valued goods and services. One analytical building block is the production function. The capital, labor, energy, materials, and services (KLEMS) approach has become a standard growth accounting tool representing this.[5]

It is incomplete for two reasons. First, it provides a year-by-year snapshot that does not capture the dynamic process—especially as it requires the assumption of constant returns to scale, whereas process innovations are important sources of productivity growth, involving inherently increasing returns (as shown in Table 1). Second, it omits structural change in consumption.

Nevertheless, KLEMS offers a useful starting point. With this approach, the variables are indexed by sector or firm and by time *t*. Rates of growth in aggregate inputs and gross output are weighted averages of their individual components, with the weights given by relative shares of each component in the total. If we assume technology is Hicks neutral (that is, increasing the marginal productivity of all inputs equally), then differentiating the production function with respect to time and using log rates of change gives the familiar decomposition equation:

$$d \ln A/dt = d \ln Y/dt - {}_{sK}\, d \ln K/dt - {}_{sAI}\, d \ln KAI/dt - {}_{sL}\, d \ln L/dt - {}_{sM}\, d \ln M/dt.$$ [6]

Here, AI capital has been distinguished from other capital. Existing produced capital measures in principle account for AI's physical infrastructure, such as servers, chips, and data centers, although there needs to be more focus on collecting and developing these capital stock and services data. Here, the separate category highlights the distinctive new intangible capitals: models and data.

There is growing research literature and a statistical effort on measuring intangible capitals, of which measuring AI is a natural extension.[7] Measuring models' intangible capital services is a new research area, requiring joint work with AI experts, but there is considerable collection underway of useful benchmarks. There is a little more progress on measuring the value of data.[8]

Again, this is ripe for further development. Luca Gamberi and I are working on piloting an approach using the Shannon entropy measure, which is when noise is added to a dataset.[9]

However, measuring models' intangible capitals gives us a first snapshot. The next step is to consider the information enabled by using AI as an input to a knowledge-based production function in an endogenous growth framework. A standard formulation is

$$\Delta A_t/A_t = \theta.H_{At}.At,$$

where *H* is the stock of human capital and $\theta > 0$ is generally interpreted as a research productivity parameter. Empirical applications have focused on codified knowledge, or conventionally measured skills, but tacit know-how and socially embedded capabilities ("organizational capital") would equally drive knowledge production and growth, either through *H* or the parameter θ. In the context of AI and other new technologies, Julia Lane has underlined the need to measure exactly these aspects and provided a framework for statistical implementation.[10]

A related question that labor market and skills data do not address is firms' organizational capital: What distinguishes AI users from those who don't use it? Several researchers have linked growing productivity dispersion at the firm level to different uses of various digital tools, so the question about the barriers to adoption is unresolved.[11] This is an open part of the research agenda, where data collection depends on better understanding the organizational or tacit knowledge barriers to AI use.[12]

Table 1. Examples of Productivity Growth from Process Innovation

Process	Date	Key Technology
American system of manufacturing	Early 19th century	Machine tools
Factory system	Mid- to late 19th century	Steam and rail
Assembly line	Early 20th century	Electricity
Lean manufacturing	Late 20th century	Telecommunications and early digital
Production networks	Late 20th–early 21st century	Information and communications technologies
Digital platforms (production and consumption)	Early 21st century	Information and communications technologies and AI
Novel manufacturing processes	Mid-21st century	AI and additive or bio-manufacturing

Source: Diane Coyle, *The Measure of Progress: Counting What Really Matters* (Princeton University Press, 2025).

The final step in this economic value framework poses the biggest conceptual challenge, and I believe this is a wide-open research question. It concerns the link from production to consumption and how to value AI use in the economy at the final demand level.

The issue is how the revenues counted inside the production boundary are converted into estimates of the "real" economy—in other words, the deflators. There is extensive literature on the well-known challenges of constructing price indexes, particularly when there are large changes in output and consumption patterns or there is an increased number and variety of new goods and services, as there are now.[13] Similarly, issues arising from (incorrectly) assuming homothetic demand when constructing the indexes that statistical agencies generally use are well-known.[14]

The new challenge stems from changes in consumption technologies. Mo Abdirahman and his coauthors noted that the constructed price index for telecommunications services in the UK varied enormously depending on whether revenue or volume weights were selected for combining specific service prices into the sectoral index.[15] Telecom companies charge a higher price per byte of data for traditional services such as fixed-line calls and SMS. The explosion of data use means that a volume-weighted sector price index plunges, while a revenue-weighted index declines more modestly.

Which is correct? Neither. We would not want to attribute real economic value to the fact that an operator's price differentiates among similar services (e.g., SMS versus WhatsApp), which is the case with revenue weights. Nor would we want to attribute equal value to every byte because what consumers care about is the byte's content. The conceptual challenge is that the product demanded is a bundle of telecom services, data center services, device services, and (sometimes free) content services.

While some work has looked at the separate elements, how to conceptualize and categorize the relevant economic activities has remained an open question.[16] One promising suggestion is to use Kelvin J. Lancaster's framework to conceptualize a "consumption technology" that shifts economic output.[17] In any case, a starting point is to identify the underlying data needs to explore the relevant "services" that consumers value.

Summary

In this chapter, I have proposed using standard economic models to establish an input-to-output measurement framework for AI. It would be applicable, with modification, to other general-purpose technologies. Table 2 summarizes the stages and data needs for measurement based on these models.

However, some open, deep economic questions remain to understand why and how new general-purpose technologies mean the majority of advanced economies are "hard to measure."

Notes

1. Zvi Griliches, "Productivity, R&D, and the Data Constraint," *The American Economic Review* 84, no. 1 (1994): 1–23, https://www.jstor.org/stable/2117968; and Diane Coyle, "Productivity Measurement: New Goods, Variety, and Quality Change," in *The Measure of Economies: Measuring Productivity in an Age of Technological Change*, ed. Marshall B. Reinsdorf and Louise Sheiner (University of Chicago Press, 2024).

2. For example, see Dave Donaldson and Adam Storeygard, "The View from Above: Applications of Satellite Data in Economics," *Journal of Economic Perspectives* 30, no. 4 (2016): 171–98, https://www.aeaweb.org/articles?id=10.1257/jep.30.4.171; and Liran Einav and Jonathan Levin, "Economics in the Age of Big Data," *Science* 346, no. 6210 (2014), https://www.doi.org/10.1126/science.1243089.

3. David Bryne et al., "The Rise of Cloud Computing: Minding Your P's, Q's and K's," Working Paper No. 25188 (National Bureau of Economic Research, October 2018), https://www.nber.org/papers/w25188; David M. Byrne et al., "Transistors All the Way Down: Viability of Direct Volume Measurement (and Price Indexes) for Semiconductors," paper presented at Summer Institute Conference on Research in Income and Wealth, July 2023, https://www.nber.org/conferences/si-2023-conference-research-income-and-wealth; Diane Coyle and David Nguyen, "Cloud Computing, Cross-Border Data Flows and New Challenges for Measurement in Economics," *National Institute Economic Review* 249 (June 2019): R30–38, https://doi.org/10.1177/002795011924900112; and Diane Coyle and Lucy Hampton, "21st Century Progress in Computing," *Telecommunications Policy* 48, no. 1 (2024), https://doi.org/10.1016/j.telpol.2023.102649.

4. Byrne et al., "Transistors All the Way Down"; and Coyle and Hampton, "21st Century Progress in Computing."

5. Dale W. Jorgenson and Zvi Griliches, "The Explanation of Productivity Change," *Review of Economic Studies* 34, no. 3 (1967): 249–83, https://doi.org/10.2307/2296675; Dale W. Jorgenson et al., *Productivity and U.S. Economic Growth* (Harvard University Press, 1987); and Paul Schreyer and Dirk Pilat, "Measuring Productivity," *OECD Economic Studies* 33, no. 2 (2001): 127–70, https://www.oecd.org/content/dam/oecd/en/publications/reports/2002/01/oecd-economic-studies-volume-2001-issue-2_g1ghg86d/eco_studies-v2001-2-en.pdf.

6. If value-added measures are used rather than gross output, the weights *s* are the factor shares in value added, but this will overstate the rate of total factor productivity growth by a factor of the inverse of the share of value added in gross output. As this has been declining, the degree of overstatement will have increased over time. Total factor productivity accounts for all inputs, not just labor.

7. Carol Corrado et al., "Intangible Capital and Modern Economies," *Journal of Economic Perspectives* 36, no. 3 (2022): 3–28, https://www.aeaweb.org/articles?id=10.1257/jep.36.3.3.

8. For a recent survey, see Diane Coyle and Annabel Manley, "What Is the Value of Data? A Review of Empirical Methods," *Journal of Economic Surveys* 38, no. 4 (2024): 1317–37, https://doi.org/10.1111/joes.12585.

9. Diane Coyle and Luca Gamberi (forthcoming).

10. Julia Lane, "The Industry of Ideas: Measuring How Artificial Intelligence Changes Labor Markets," American Enterprise Institute, June 9, 2023, https://www.aei.org/research-products/report/the-industry-of-ideas-measuring-how-artificial-intelligence-changes-labor-markets/.

11. For example, see Peter Gal et al., "Digitalisation and Productivity: In Search of the Holy Grail—Firm-Level Empirical Evidence from EU Countries," Working Paper No. 1533 (Organisation for Economic Co-operation and Development, February 12, 2019), https://

doi.org/10.1787/5080f4b6-en; Alison Cathles et al., "Digital Technologies and Firm Performance: Evidence from Europe" (working paper, European Investment Bank, April 2020), https://library.oapen.org/handle/20.500.12657/43443; Erik Brynjolfsson et al., "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics* 13, no. 1 (2021): 333–72, https://doi.org/10.1257/mac.20180386; and Diane Coyle et al., "Are Digital-Using UK Firms More Productive?" (working paper, Economic Statistics Centre of Excellence, March 10, 2022), https://www.escoe.ac.uk/are-digital-using-uk-firms-more-productive/.

12. James Bessen, *The New Goliaths: How Corporations Use Software to Dominate Industries, Kill Innovation, and Undermine Regulation* (Yale University Press, 2022).

13. W. Erwin Diewert et al., eds., *Price Index Concepts and Measurements* (University of Chicago Press, 2009); and Diane Coyle, "Productivity Measurement: New Goods, Variety and Quality Change," in *The Measure of Economies: Measuring Productivity in an Age of Technical Change,* eds. Marshall B. Reinsdorf and Louise Sheiner (Chicago University Press, 2024).

14. Thomas A. Stapleford, *The Cost of Living in America: A Political History of Economic Statistics, 1880–2000* (Cambridge University Press, 2009).

15. Mo Abdirahman et al., "Telecoms Deflators: A Story of Volume and Revenue Weights," *Economie et Statistique*, no. 530–31 (2022): 43–59, https://www.doi.org./10.24187/ecostat.2022.530.2063.

16. For consumer device services, see David M. Byrne and Carol Corrado, "Accounting for Innovations in Consumer Digital Services: IT Still Matters," Working Paper No. 2019-049 (Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, June 2019), https://doi.org/10.17016/FEDS.2019.049. For examples of work on free online services, see Erik Brynjolfsson et al., "GDP-B: Accounting for the Value of New and Free Goods in the Digital Economy," Working Paper No. 25695 (National Bureau of Economic Research, March 2019), https://www.nber.org/papers/w25695; and Diane Coyle and David Nguyen, "Free Digital Products and Aggregate Economic Measurement," *Economie et Statistique*, no. 539 (2023): 27–50, https://www.doi.org/10.24187/ecostat.2023.539.2096.

17. Charles Hulten and Leonard I. Nakamura, "Expanded GDP for Welfare Measurement in the 21st Century," in *Measuring and Accounting for Innovation in the Twenty-First Century*, ed. Carol Corrado et al. (University of Chicago Press, 2020); and Kelvin J. Lancaster, "A New Approach to Consumer Theory," *Journal of Political Economy* 74, no. 2 (1966): 132–57, https://www.journals.uchicago.edu/doi/10.1086/259131.

Measuring AI Value: Puzzles, Challenges, and Opportunities

Prasanna "Sonny" Tambe

Why are some firms better than others at producing and using AI capabilities?

Differences in tangible and *intangible* factors may generate different patterns of AI investment and use among US firms. Researchers have begun to explore how firms' accumulation of these underlying, foundational assets drives competitive differences in their abilities to produce AI capital. For instance, although some firms, such as Google, have clearly been particularly successful when using AI, they may have some secret sauce, or, alternatively and more likely, they may simply have developed the right assets to support development of downstream tools for algorithmic decision-making.

The Current State of Play in AI Measurement

In this chapter, I focus on the key issues with measuring firms' investment in AI to drive productivity in three of the most critical areas: (1) infrastructure (data, software, and computing), (2) workforce training, and (3) high-skill talent. A proper accounting of all these factors can improve our understanding of why firms' prospective AI capabilities differ.

Infrastructure: Data, Software, and Computing

Access to key factors for effective AI use—data, software, and processors—is unevenly distributed. Certainly, the volume and variety of data matter. Increasingly, with large AI projects, hardware (chips) has become a notable supply-chain constraint. Much software at the AI frontier is available on open-source platforms, but firms differ in their ability to adapt and deploy these frameworks in their work, due to differences in access to talent, the costs of legacy systems, uncertain regulatory environments in some industries, and other factors.

Workforce Training

A particularly important category of AI investment is the workforce training and reskilling that will likely be needed for employees to use AI. Much of the attention on AI and labor has focused on potential labor displacement. Such research forecasts which occupations are most vulnerable to job displacement based on their tasks. Then, given the occupational

Table 2. Summary of Framework and Data Needs

Framework	Modeling Approach	Data Needs
Production	KLEMS	• Intangible AI capital (e.g., value of models and data) • Enhance collection of physical AI capital data (e.g., data centers and fiber networks) • Price indexes
Knowledge production	Endogenous growth process	• Research and skills (e.g., research and development spending and employment and wages in AI) • Organizational capabilities (e.g., by management survey)
Demand and consumption	Consumption technology	• Consumer device purchase and data usage by category • Price indexes

Source: Author.

distribution in the US and other countries, analysts can estimate AI technologies' impact on the demand for various occupations.

But in the short run, the largest workforce change that employers must make may be reskilling workers in existing occupations to enable human-AI collaboration. As AI diffuses into more jobs, workers in occupations ranging from sales to accounting and customer service will need new skills to be effective. The entry of large language models (LLMs) into the public consciousness has underscored the idea that human-AI collaboration will be the dominant mode of AI use in the medium term. Future call centers will still need customer service agents, but the centers will look different from how they look today.

Talent

Finally, the talent pool for AI development is a critical driver of AI progress. Maintaining the talent supply includes not only hiring highly paid software developers, on whom the media likes to focus, but also developing a robust semiconductor workforce that can support investments in physical capital, such as those related to the 2022 CHIPS and Science Act. Universities like Ohio State and Purdue have initiated training programs for this work, but whether the US can successfully supply well-trained workers to reshored semiconductor initiatives in a way that allows them to replicate foreign manufacturing capabilities is an open question.

Conceptual and Empirical Gaps and Opportunities

Three areas are critical for advancing AI capabilities, and each poses unique measurement challenges. The following section explores these challenges in detail.

Infrastructure: Data, Software, and Computing

Of the various factors distinguishing firms trying to deploy AI technologies, few are measurable and fewer still are measured. How should we account for the value of firms' databases, the recency of their data, and the data's variety? How quickly do data depreciate, and how does this rate vary across industries? For instance, data on consumer preferences may depreciate quickly, but databases of human X-ray images might retain their value longer.

Another complication is that AI software, even on the frontier, is increasingly available from open sources. For example, state-of-the-art LLM capabilities can be accessed through Meta's Llama project, which is free to use even for commercial purposes. So a firm's expenditures on AI software provide a less useful signal of its AI capabilities than that firm's spending on enterprise resource planning provides for its supply-chain capabilities. When firms contribute to, download, or use AI software obtained from open sources, they produce some digital trails, such as GitHub forks and clones, but it is not easy to convert these indicators of digital interest to measures of meaningful software use. Any exercise that attempts to equate these indicators to real outcomes will likely suffer from severe measurement problems.

On the computing front, the growing role of cloud computing in AI projects complicates researchers' ability to measure and track firms' access to computation. In contrast to expenses for a firm's internal servers, which can be tracked, expenses for a cloud are often elusive, and it is not even clear how to convert figures on cloud computing expenditures to meaningful quantities. Depreciation matters here greatly as well. Early figures for AI-specific chips suggest that firms must use the hardware particularly often and rearchitect it frequently to take advantage of AI software innovations, so this hardware may depreciate quickly.

Talent

In other domains of national importance, such as the research and development needed for scientific progress, scholars have focused on the role of highly skilled individuals such as scientists in driving innovation, productivity, and growth. This is often tracked through databases on patenting or federal funding that provide information about scientists' career activities. Just as with basic science, certain individuals and organizations have apparently played an outsized role in pushing the AI frontier, but there may be no ready equivalent in AI for the databases that have been so useful for understanding key players in the scientific ecosystem.

This is because software developers, including star engineers, rarely leave reliable patent trails. Moreover, funding for AI development is often provided internally by corporations, rather than through public institutions like universities, making it harder to track what money is allocated to research and development for AI projects. As a result, researchers are in a different data environment when they seek to understand the relationship between the supply of skilled talent and AI progress—including the roles of AI scientists, skilled software developers, semiconductor engineers, cybersecurity experts, and other tech workers who will eventually be critical to the AI ecosystem.

Workforce Training

Finally, reskilling workers and redesigning jobs are significant investments for employers and critical for enabling firms to increase their productivity with AI technologies. A key gap in measuring these investments is the extent of workforce-level, sub-occupational change in firms. The question here is: In a model in which AI works with humans to improve productivity, how should workers be reskilled to realize these gains? For instance, should workers be taught to determine when LLMs' output is trustworthy? Or will effective salespeople need to be trained in prompt engineering?

As employers learn the answers to these questions and begin to adjust the workforce, large gaps will remain in our ability to measure these changes, because statistical agencies focus on data collection at the occupational level. Short-run opportunities for AI augmentation, on the other hand, are more likely to require changing workers' training than to require changing their job titles.

For instance, software "copilots" are expected to significantly change how developers spend their time, whether or not they change the number of developers needed to do the work. This change is difficult to track, even though upskilling and retraining costs have always been a significant component—often the largest component—of the financial cost of IT transformation. Further, because AI is widely considered a general-purpose technology and therefore relevant to most jobs, the costs to employers of workforce adjustment may be even higher than the costs of the IT-enabled work reengineering that prior waves of business technology have required.

Gaps in measuring workforce needs were already large before 2022, but rapid growth in corporations' use of generative AI tools and LLMs has likely widened them. Generative AI tools require workers to appropriately structure prompts to get the types of output they want and critically assess whether the LLMs' output is usable. There are still many unknowns about how workers must be retrained to work effectively with LLMs, but changes will clearly be needed.

Why These Measurement Gaps Matter

Why are these measurement gaps—which are related to databases, workers, and other factors necessary for building AI systems—important?

Understanding how these assets are distributed among producers, how firms accumulate them, and how quickly they depreciate can be useful for explaining why AI has a much larger impact on some sectors than on others and on some firms than on others. This knowledge is also important for predicting how long any productive gains from AI investments will last as these assets depreciate and require ongoing investment.[1] Another implication of these measurement gaps is that they limit our ability to forecast whether AI technologies will drive democratization or inequality. Particularly given the interest in sources of competitive advantage for Big Tech and "superstar" firms generally, determining whether firm size matters for accessing these critical resources would be useful.

Short-, Medium-, and Long-Term Next Steps

How can we address these measurement gaps?

Workforce Training and Talent

The easiest place to start may be workforce and talent measurement gaps.

Fortunately, corporate data sources can already partially meet the challenges of tracking how workers are being retrained to effectively use AI technologies and measuring the supply of skilled talent. Datasets, already widely used, from providers like Lightcast, LinkedIn, and Revelio Labs provide a window into this type of technical change.[2]

However, academic measurements based on these data sources tend to be ad hoc, with no standard approach, making it difficult to compare measures across papers or time periods. For example, which skills and tools involve AI? One study on the effects of AI technology might incorporate regression analysis techniques, a long-standing staple of statistics, while another may not. The lack of a clear taxonomy mapping technologies to skills has been a long-standing obstacle for all types of measurement in this area.

A useful objective would be to develop a standardized process (and taxonomy) for measuring how AI skills spread across the US workforce. This would allow comparable and reproducible measures to be used across studies, industries, and periods.

Infrastructure

Measuring AI inputs that are not embodied may pose a larger challenge. The data, software, and computing stack required to implement AI technologies has already changed rapidly over the past few years, and this trend will continue and perhaps even accelerate. New tools automate machine learning and reduce the importance of some factors relative to others. Macro factors that affect the supply of these infrastructure resources, such as geopolitical conflict and new export controls, have also become important.

The first short-term step for measuring AI infrastructural inputs is to identify the tangible and intangible assets of interest. If analysts agree that AI productivity reflects an accumulation of a common set of inputs, then developing a consensus on what these critical inputs are in successful AI initiatives is a good starting point.

A longer-term goal is to develop a framework for measuring these assets, both tangible and intangible. This is a considerable challenge. Challenges in accounting for intangible assets, of course, have a long history. Nonetheless, this is a worthy project, and one might use proxy measures for certain inputs—such as using the number of Oracle database administrators and developers to measure investment in Oracle databases—even though numerous measurement problems are likely with this approach.

A better, and longer-term, goal is to improve reporting requirements for firms (e.g., for software investment) and methods to value other inputs that allow consistent measurement of a firm's digital assets.

Important Considerations and Uncertainties

Note that any measurement attempts, whether related to the workforce or infrastructure, must be reproducible across time, because the tools' capabilities are a moving target. LLMs are significantly more capable now than they were even a few months ago, so any prognostication about organizational needs must first accurately forecast these technologies' evolution.

Another source of uncertainty is regulation. We are entering a period of sustained uncertainty when AI technology's uses must be adjusted to meet changing regulations, so these technologies' diffusion path and uses in production are uncertain. For instance, despite LLMs' quick start, concerns about liability or intellectual property violation may significantly slow their use in some industries.

We also need to acquire a much better understanding of how workers will use AI tools and what AI and human collaboration will look like. This, in turn, could provide information about how workers need to be reskilled for such collaboration.

For these reasons, establishing a benchmark now, when the adoption of most AI tools is still limited, and developing a repository of data sources, taxonomies, and approaches to measure this type of workforce change can drive a much better understanding of how the US workforce is changing to meet the need for human-AI collaboration.

Implications for Measuring Other Critical and Emerging Technologies

AI is one of many technologies likely to drive social change in the coming decades. It is probably the most important. However, other emerging technologies will have a significant economic impact, like virtual reality systems for work environments. Virtual reality systems have already significantly affected some domains (e.g., the da Vinci system in robotic surgery).

The approaches outlined above are not specific to AI. Any attempt to measure the impact of emerging technologies on the workforce can benefit from a well-supported infrastructure. Moreover, emerging technologies are likely to be closely interwoven with existing systems, including those that house data and drive AI prediction. For instance, virtual reality systems will integrate AI into their core functions. A proper accounting of AI assets is a foundation for better measurement of all technologies that build on AI.

Notes

1. Prasanna Tambe et al., "Digital Capital and Superstar Firms," Working Paper No. 28285 (National Bureau of Economic Research, December 2020), https://www.nber.org/system/files/working_papers/w28285/w28285.pdf.

2. John J. Horton and Prasanna Tambe, "Labor Economists Get Their Microscope: Big Data and Labor Market Analysis," *Big Data* 3, no. 3 (2015): 130–37, https://pubmed.ncbi.nlm.nih.gov/27442956/.

Part II

Perspectives from the National Artificial Intelligence Research Resource Pilot

Industries of Ideas

Tracing the Links Between Investments in Science, Innovation, and Jobs

Julia Lane

We are witnessing a sea change in government spending on science and technology. Modern industrial policy represents a bet on a new approach—investing in ideas—to transform the economy and create high-wage jobs. Investments in science and technology on the largest scale since the Cold War have anecdotally had massive effects on jobs and earnings. The artificial intelligence revolution alone has set off a gold rush: At the November 2023 Supercomputing Conference in Denver, I was told that the starting salary of new Stanford PhDs in AI was $750,000—plus stock options. Yet firms and workers looking for data on which to base their hiring and career decisions are out of luck. Suzette Kent, a former federal chief information officer and adviser to many tech companies, says in her chapter of this compilation, "The rapid application and adoption of AI in the workplace may have created one of the fastest workforce shifts in history, but reliable data about such a shift are hard to find."

Without relevant data on the links between science and innovation, how are governments, science funders, and scientists to decide where to place their bets? What are the theory, data, and evidence?

The theory got a Nobel Prize in 2018. Governments are investing in people who create ideas—new technologies—that *can* be reused, which is why "the discovery of new ideas lie at the center of economic growth," as Charles Jones said when describing Paul Romer's conceptual framework, for which Romer received the Nobel Memorial Prize in Economic Sciences.

The data seeds were sown almost two decades ago. President George W. Bush's science adviser, John Marburger, sensibly unconvinced of the scientific and practical value of relying primarily on document-based, bibliometric approaches to studying science in understanding its practical effects, called for a "Science of Science Policy."

The evidence is being built by the University of Michigan's Institute for Research on Innovation and Science (IRIS) in a people-centered data infrastructure that draws on the original government-led STAR METRICS (Science and Technology for America's Reinvestment: Measuring the Effect of Research on Innovation, Competitiveness, and Science) program, now called UMETRICS to reflect its university leadership. The infrastructure changes the way data are being *structured*, so that the relevant processes are studied; *classified*, so that levels and trends in funding inputs and subsequent activities can be measured and tracked; *collected*, so that actionable information is available for multiple units of analysis; and *analyzed*, so that governments, science funders, and workers can make informed decisions.

The people-centered IRIS data approach, because it is characterized by a collaborative, bottom-up, and scientifically grounded governance model, is purposefully designed to both respond to the interests of the relevant communities and be used by them. It stands in direct contrast to the current document-centered data approach, which lacks a clear governance structure and scientific framework.

The Framework

The new data *structure* is grounded in understanding the processes by which growth through investments in science is generated. These processes are fundamentally different from investments in capital and labor that produce physical goods and services that, once used, cannot be reused. In practical terms, this is why Stanford PhDs are paid so much, since the firms that hire them expect their ideas to transfer to others in the firm and thus generate more revenue and growth. An operational data structure based on the Romer concept requires joining the dynamic flows of all people funded on research grants at universities—the knowledge workers—with the jobs they get when they move to the private sector. These flows can then be used to trace their effects on both the firms at which they work and the other workers at those firms—as J. Robert Oppenheimer said, "The best way to transmit knowledge is to wrap it up in a human being."

The *classification* system is similarly designed to reflect current scientific and economic activities. Technologies such as artificial intelligence represent clever new ideas of how to better combine existing inputs, and the clustering and measurement of activities represent the clustering of people and the ideas embodied in them. The Industries of Ideas approach deployed by IRIS categorizes firms by the people who created and use the technologies they will adopt. Such a classification framework is a sea change from earlier industrial classifications based on *what* goods are physically produced—like manufacturing and agriculture—or by *how* services and goods are produced—like the delivery of health, financial, and investment services.

Data *collection* is also purposefully designed to be timely, flexible, and useful. People-centric data generated by the administrative processes at universities can capture the organization of people in science at multiple levels (e.g., individuals, teams, projects, and institutions), the multiple sources of university funding (e.g., federal scientific and programmatic agencies, philanthropic foundations, industry, and state and local government), inputs into science from vendors (e.g., computing services, instruments, and biological specimens), and people's career dynamics across time (e.g., individual career earnings and employment trajectories).

Finally, data *analysis* is not centralized but bottom-up, transparent, and collaborative. The IRIS infrastructure has been developed over the past decade. The current production release reflects actual expenditures on more than 535,000 grants, 864,000 employees, and 970,000 vendors paid by more than 80 campuses representing more than 41 percent of total US research and development spending at universities. IRIS, while hosted at the University of Michigan, has a governing board that represents its member institutions. While the core of the data are the administrative records, the infrastructure that provides access to IRIS data are open to all collaborating universities and their approved researchers; over 500 researchers assessed the data for scientific purposes, and hundreds of reports have been generated for science funders, federal and state government agencies, and the participating universities themselves.

What does all this mean in practice? The National AI Research Resources Task Force, on which I served, was charged with developing a roadmap to guide investments in AI compute and data resources to, inter alia, spur innovation. The lack of reliable data identified by Kent was also recognized by this task force; its final report submitted to the president and Congress in January 2023 incorporated the people-centered approach described here as part of its evaluation framework.

Empirical Implementation: AI and EV

One of the use cases is the National Science Foundation's (NSF's) new Technology, Innovation, and Partnership (TIP) Directorate, which has funded a pilot Industries of Ideas project to better understand its regional technology investments. The pilot focuses on two critical and emerging technologies—AI and electric vehicles (EV)—in Ohio and is designed to scale to other technologies and states in subsequent stages. It begins with linking people funded in AI and EV research in Ohio universities with individual- and firm-level state administrative workforce and education data at the Ohio Longitudinal Data Archive.

As Jason Owen-Smith points out, research communities form and can be identified through field-specific activities and collaborations. In the case of AI and EV for the subset of IRIS universities, between 2001 and 2023, NSF invested $8.5 billion in about 13,000 awards to about 3,300 principal investigators. The flow of funding to researchers in that field can be identified through university administrative grant and award data.

The IRIS data then is used to capture all spending on teams: principal investigators (PIs), trainees such as undergraduate and graduate students, postdocs, staff clinicians, and administrative staff. Those grants support 46,385 people—or almost 15 people per PI. Almost half of those (over 21,000) were graduate students, 8,300 were non-PI faculty, 8,000 were staff, almost 3,000 were undergraduates, and 2,000 were postdocs.

Many if not most of these people will never publish a paper or be a PI. But working on new and emerging AI research teaches them about applications to nearly every field NSF supports. It gives them access to specialized professional networks. It makes them competitive and interested in AI jobs. In other words, these hitherto invisible research-funded people are a key "product" of grant-funded research and a way to identify currently unmeasurable workforce effects.

The effects on the private sector are not just the initial earnings of knowledge workers (like the Stanford computer science graduates!). They also include the cumulative knock-on effect on the earnings of all workers in the firms who use the knowledge workers' ideas. Simply put, just as Oppenheimer posited and Romer theorized, the mobility of research-funded staff helps connect the ideas workforce to the firms so that new technologies can be transferred to existing production processes. Their flow to the full economy and the transmission of their ideas are captured when trainees and staff get jobs in the private sector and their earnings and employment are recorded in state administrative data. My back-of-the-envelope estimates from aggregate data suggest that the potential national impact could be up to 36 million workers in 18 sectors; one of the major outcomes of the two-year pilot has been a dashboard presenting firm and worker results for Ohio that is scalable to other technologies.

Finally Making Science Metrics More Scientific

The new approach to understanding the *structure* of data was well captured by Erwin Gianchandani, the TIP assistant director in the pilot's press release, when he said,

> NSF's strategic investments in key technologies warrant innovative tools to accurately assess the impact of these investments across the U.S. The Industries of Ideas project will develop a prototype to better understand the impact of NSF's efforts through the new TIP directorate, providing rich, descriptive

> analyses of the interplay between our investments and people, jobs and regional economies.

There is also a new energy around the *classification* issues raised here: Think tanks, measurement experts, academics, government agencies, and private-sector data providers are considering new approaches to measurement, as evidenced by this workshop.

There is new engagement around data *collection* and *analysis*. One example was the bipartisan H.R. 6655 reauthorization bill, A Stronger Workforce for America Act. That bill specifically called for funding Workforce Data Quality Initiative grants to improve state workforce data capabilities by fostering cross-state collaboration, improving the timeliness and relevance of labor-market data, supporting the adoption of credential navigation tools, and advancing the use of evidence and data to drive decision-making. As noted by Adam Leonard, regional multistate data collaboratives provide a basis for state education and workforce agencies to contribute data and produce new products. The UMETRICS data are increasingly being used for training and research.

In sum, this new approach to constructing data on the links between science and innovation—so urgently needed for governments, science funders, and scientists to make decisions about where to place their bets—is in place. The theory, data, and evidence can now inform our understanding of the impacts of countries' vast investments in AI research and research in many other fields.

Notes

1. Julia Lane, *The Industry of Ideas: Measuring How Artificial Intelligence Changes Labor Markets*, American Enterprise Institute, June 9, 2023, https://www.aei.org/research-products/report/the-industry-of-ideas-measuring-how-artificial-intelligence-changes-labor-markets/.
2. Suzette Kent, *Perspective from NAIRR*, Stanford University, Institute for Human-Centered Artificial Intelligence, Digital Economy Lab, 2024, https://digitaleconomy.stanford.edu/wp-content/uploads/2024/03/Suzette-Kent.pdf.
3. Charles I. Jones, "Paul Romer: Ideas, Nonrivalry, and Endogenous Growth," *The Scandinavian Journal of Economics* 121, no. 3 (2019): 859–83, https://doi.org/10.1111/sjoe.12370;
4. John H. Marburger III, "Wanted: Better Benchmarks," *Science* 3018, no. 5725 (2005): 1087, https://www.science.org/doi/10.1126/science.1114801.
5. Julia Lane, "Let's Make Science Metrics More Scientific," *Nature* 464, no. 7288 (2010): 488–89, https://www.nature.com/articles/464488a.
6. Julia Lane et al., "New Linked Data on Research Investments: Scientific Workforce, Productivity, and Public Value," *Research Policy* 44, no. 9 (2015): 1659–71, https://ideas.repec.org/a/eee/respol/v44y2015i9p1659-1671.html.
7. This is what economists call a production function.
8. Nikolas Zolas et al., "Wrapping It Up in a Person: Examining Employment and Earnings Outcomes for Ph.D. Recipients," *Science* 350, no. 6266 (2015): 1367–71, https://www.science.org/doi/10.1126/science.aac5949.
9. Lane, *The Industry of Ideas.*
10. Robert E. Yuskavage, *Converting Historical Industry Time Series Data from SIC to NAICS*, US Department of Commerce, Bureau of Economic Analysis, 2007, https://www.bea.gov/system/files/papers/P2007-7.pdf.
11. Natsuko Nicholls et al., *Summary Documentation for the IRIS UMETRICS 2022 Data Release*, University of Michigan, Institute for Research on Innovation and Science, April 20, 2022, https://iris.isr.umich.edu/research-data/2022datarelease-summarydoc/.
12. Wan-Ying Chang et al., "A Linked Data Mosaic for Policy-Relevant Research on Science and Innovation: Value, Transparency, Rigor, and Community," *Harvard Data Science Review* 4, no. 2 (2022), https://hdsr.mitpress.mit.edu/pub/u073rjxs/release/3.
13. National Artificial Intelligence Research Resource Task Force, *Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource*, January 2023, https://www.ai.

gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf.

14. US National Science Foundation, "Industries of Ideas: A Prototype System for Measuring the Effects of Research Investments on Regional Firms and Jobs," 2023, https://iris.isr.umich.edu/tip/.

15. Joshua D. Hawley, "Ohio and the Longitudinal Data Archive: Developing Mutually Beneficial Partnerships Between State Government and the Research Community," in *Handbook on Using Administrative Data for Research and Evidence-Based Policy*, ed. Shawn Cole et al. (Abdul Latif Jameel Poverty Action Lab, 2020).

16. Jason Owen-Smith, *Will the Real AI Researcher Please Stand Up? Fields, Networks, and Systems to Measure the Impact of Research Investments*, Stanford University, Institute for Human-Centered Artificial Intelligence, Digital Economy Lab, 2024, https://digitaleconomy.stanford.edu/wp-content/uploads/2024/03/Jason-Owen-Smith.pdf.

17. Julia Lane, *Reimagining Labor Market Information: A National Collaborative for Local Workforce Information*, American Enterprise Institute, March 9, 2023, https://www.aei.org/research-products/report/reimagining-labor-market-information-a-national-collaborative-for-local-workforce-information/.

18. US National Science Foundation, "Industries of Ideas."

19. US National Science Foundation, "NSF Launches Pilot to Assess the Impact of Strategic Investments on Regional Jobs," 2024, https://new.nsf.gov/tip/updates/nsf-pilot-assess-impact-strategic-investments-regional-jobs.

20. A Stronger Workforce for America Act, H.R. 6655, 118th Cong. (2023).

21. Adam Leonard, *Outside of the Box Use of Administrative and Wage Data in Texas*, Stanford University, Institute for Human-Centered Artificial Intelligence, Digital Economy Lab, 2024, https://digitaleconomy.stanford.edu/wp-content/uploads/2024/03/Adam-Leonard.pdf; and National Association of State Workforce Agencies, "Multi-State Data Collaboratives (MSDC)," https://www.naswa.org/partnerships/multi-state-data-collaboratives/about.

22. Tania Babina et al., "Cutting the Innovation Engine: How Federal Funding Shocks Affect University Patenting, Entrepreneurship, and Publications," *The Quarterly Journal of Economics* 138, no. 2 (2023): 895–954, https://academic.oup.com/qje/article-abstract/138/2/895/6979841; Jason Coupet and Yuhao Ba, "Benchmarking University Technology Transfer Performance with External Research Funding: A Stochastic Frontier Analysis," *The Journal of Technology Transfer* 47, no. 2 (2022): 605–20, https://bayuhao.org/research/paper%2010.pdf; Simon Schacter et al., "The Inefficiency of Private Support for Public Health: Comparing Nonprofit Biomedical Research Funding with the NIH," *Social Science Research Network,* no. 4483036 (June 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4483036; and Wei Yang Tham, "Science, Interrupted: Funding Delays Reduce Research Activity but Having More Grants Helps," *PLOS One* 18, no. 4 (2023): e0280576, https://doi.org/10.1371/journal.pone.0280576.

New Approaches to Characterize Industries

Suzette Kent

The rapid application and adoption of AI in the workplace may have created one of the fastest workforce shifts in history, but reliable data about such a shift are hard to find. If the anecdotes about explosive job creation and rapid role-based task disruption are correct, it will be critical to develop models to understand, predict, and support the working world's AI-driven evolution. However, the pace of advancement, variations in adoption of AI capabilities, and nascent methods of measuring actual job impacts related to AI mean that the narrative has been more about trends and headlines than robust model-driven analysis using a comprehensive data framework. Simply put, we lack useful information for making strategic decisions about national workforce matters.

Having been involved in technology-driven business transformation for the past three decades, I have deep experience with the workforce change that occurs when innovative technology becomes operationally pervasive. Relevant examples I have observed include the implementation of image technology in medicine and financial services, robotics in the automotive and transportation sectors, cloud-based services in the technology sector, and the digital and mobile delivery of services across the public and private sectors.

After working with the world's leading technology companies, service companies, institutions of higher education, public company boards, and companies that provide alternative pathways for technical training, I can attest that there is not a single area where the impacts of AI are not at the forefront of strategic discussions. My direct experience in high-level discussions regarding building, deploying, and scaling AI and the ramifications to workforce ecosystems that are relevant to model development and application tells me that regardless of the difference in industries, they share a common challenge—finding knowledgeable people to fuel growth that aligns with our national priorities. There is still a great need to develop a framework to analyze the broader context of workforce evolution and inform the public view of key opportunities and challenges.

Two recent studies illustrate the enormous attention being paid to explosive growth and disruption and demonstrate that current publicly available information is myopic and not based on models with repeatable validity.

In December 2023, the World Economic Forum released another paper in its "Jobs of Tomorrow" series, *Jobs of Tomorrow: Large Language Models and Jobs*, which sought to advance our understanding of large language models' direct impact on specific jobs. The paper offers thoughts on how business leaders could consider large language models in their strategic business planning. Two of its major findings were a forecast that over 40 percent of working hours could be transformed by large language models through automation or augmentation and a prediction about the types of new jobs likely to emerge. Although this information provides a perspective of some entities, the data are self-reported and based on expectations of the respondents, rather than hard data. This same study also notes that

> three out of every four companies across the globe are expected to adopt technologies that include generative AI in the next 3 to 5 years, and 98% of global executives agree AI foundation models will play an important role in their organizations' strategies in that same time period.

While one can find many sources of strong beliefs, shared expectations, fear, and other assortments of media hype, what remains lacking is a forward-looking, data-driven model rooted in comprehensive authoritative jobs and task data. Even if the model development moves in stages, we need a sustainable structure to inform public- and private-sector leaders, practitioners, and the people who develop our workforce about the transformation AI is bringing to the national workforce ecosystem.

In another example, the recent *2023 Europe and US Data, Analytics, and Artificial Intelligence Executive Organization and Compensation Survey* by Heidrick & Struggles surveyed approximately 100 companies in the US and Western Europe. These companies spanned all industries and have had AI-based roles or operational AI capabilities inside their organizations for at least five years. This survey grouped employment into three categories: those who work in data science and analytics, machine learning, and AI; those who work in business analytics and intelligence; and those who work in data engineering, with the third group focused more on the back end or foundational elements of data. Since this survey's report focused on data and analytics, it acknowledged that the operational and oversight roles were also increasing, but those were not directly captured.

I propose a practical approach to address these issues. In the conversation around AI, many aspects of AI have been directly correlated to the generation of jobs. Most visible have been direct hardware and device, software, services, and data activities. The less visible impacts are the secondary and tertiary influences on business-operational infrastructure and governance as AI capabilities move from experimental use to enterprise use at scale.

As AI innovation moves from the idea stage to affecting actual jobs, there are conceptual models that can trace the people and specific work from research, through lab to market evolution, then into business entities. The timely opportunity is to not only trace the practitioners but also understand how changing dynamics of people and capability affect growth in the least visible workforce areas.

In technology industries, advancing products and services using AI has a jobs footprint similar to other transformational technologies. For discussion purposes, this is called the primary impact. In most cases, the workforce involved in the first order of change is technically and functionally skilled for the specific application of AI. These people are software engineers, data scientists, and process specialists and may have specialty skills related to the specific innovation. Identifying these individuals and roles is not complicated.

As AI is scaled within any entity, moving beyond a pilot or experimentation into broader enterprise operational use, secondary jobs are created in business support and business oversight functions. These are risk and compliance roles, model validation roles, and data management and testing roles, all of which require new skills for procurement, legal, privacy, service, and operations tasks. I have directly observed that in many cases, the first step is to assign the AI-related responsibilities to individuals currently holding these business-support-type roles.

Figure 1. AI Job Creation "Ripple"

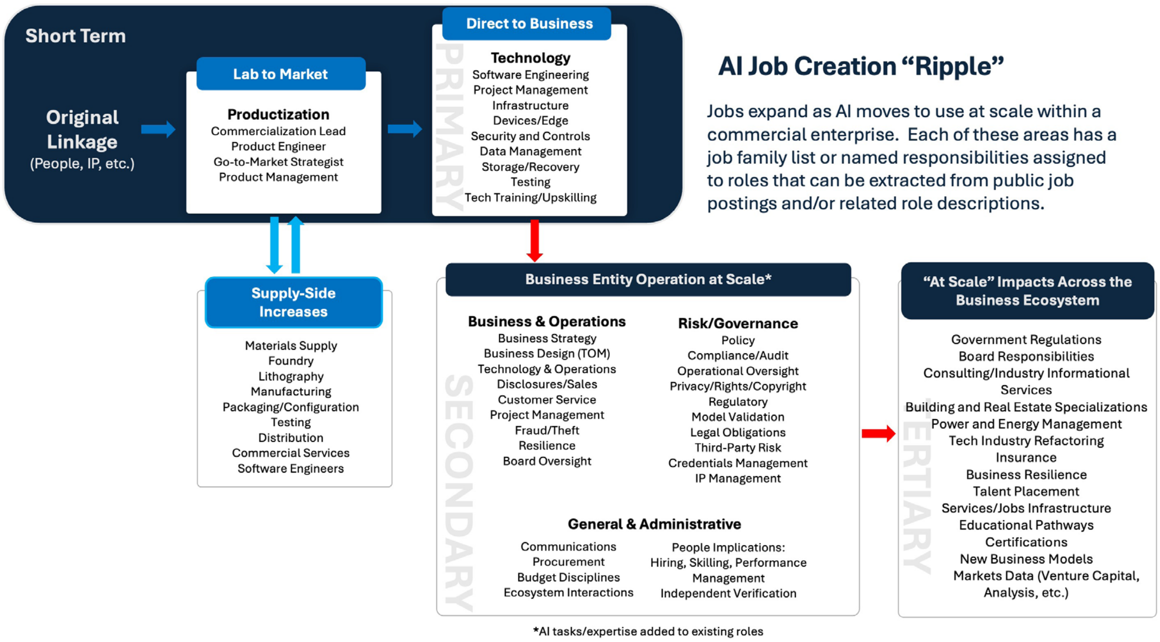

Source: Author.

These individuals usually have little practical or academic experience with applied artificial intelligence. The effect is that either these deployments move slowly due to a steep learning curve or the results of an immature process become evident in business outcomes.

As an entity begins to use AI capabilities more pervasively, one often sees new role criteria developed, with hiring or organizational changes to centralize internal skills. With AI, new trends are emerging, and those trends appear to be different for industry types and data classifications, so even actual experiences are less applicable from use case to use case. For example, the rapid use of AI for language augmentation for customer service functions may seem effective, whereas the same model and data are not effective for legal purposes. Further examples to be examined in the "New Approaches to Characterize Industries: AI as a Framework and a Use Case" workshop include regulated industries, industries with increased restrictions for data use and sharing, and industries and functions with the potential for disruption.

Another ripple or overlay to consider is the growth of responsibilities in oversight, risk, and regulatory areas. These tertiary impacts are in addition to operational-risk-type roles in a standard business-operational paradigm. Examples of these tertiary workforce changes include board responsibilities, government regulatory interactions, third-party risk-assessment activities, and requirements driven by social justice. Tertiary impacts also include the creation of roles reflective of the market dynamics of a transformational technology. These types of roles evolve because of increased merger and acquisition activities, investment market interests, the introduction of new companies, and the development of new business models.

Developing models that are effective at each stage is critical, as they will influence the model's applicability to the following stages. Short-term impact can be measured using data related to people and new lab-to-market capabilities. Trend expectations can be analyzed in comparison with previous technology expansion waves (e.g., digital, cloud, cyber, and mobile).

Medium- and longer-term steps should seek to create models that can relate the transfer of people and capabilities to identifying and quantifying the triggers of scale. Industry and scale can then be correlated to the expansion in other types of roles based on the operational use of AI. For example, in regulated industries, the processes to examine business functions, models, and validation processes are well-defined. As AI is introduced into existing business functions, those oversight processes are also affected.

I have had firsthand involvement with individuals who have regulatory responsibilities, and they have shared that they are challenged to keep up with the pace of change in operational areas. A specific government example might be where drones are used to inspect the safety and soundness of a physical structure versus a human examining the structure. While there might be a high level of confidence about the fidelity of the images gathered, one may question whether there are sufficient data and model capability to interrogate those images and propose a decision. The oversight function must understand the new processes and ensure that the use of AI improves the existing processes without introducing any new risks.

By identifying role types, tasks embedded in job postings, and other public data, we can develop a perspective on the scale of business-operational impact (e.g., people, process type, and regulatory applicability). The scale of impact could then inform forward-looking views about the expansion of secondary and tertiary roles and the prevalence of AI capability as a required skill set.

Consider the correlation of AI skills as an emerging role requirement to the historical pattern of digital skills being required in the workplace. Using public job postings and data contained in role postings, we see a correlation between the primary jobs impact of capability adoption and the secondary and tertiary workforce impacts that follow from operational adoption at scale. In addition, one might use the current O*NET and US Bureau of Labor Statistics taxonomy of over 19,000 tasks, which matches tasks performed to occupations and breaks down the time spent on each task in an occupation. This could become a baseline for calculating year-over-year change as a factor for interpreting disruption and expansion in existing roles. That forecast could then inform future job-demand expectations.

As AI continues to become deeply embedded in how businesses deliver products and services and how governments meet their missions, maturity and scale will directly drive job creation and demand. With detailed data collection aligned to key industry characteristics, we could forecast the types of jobs that will be affected and the scope and scale of the economic impact. A conceptual sample of the waves of job creation was shared in the discussion (Figure 1). These role titles have been validated with industry persons and can be tracked and traced through public job postings.

There are also lessons from the digital and mobile wave, cloud transformation, cybersecurity industry, and technology-as-a-service industry that can help inform the baseline and identify significant areas of deviation. Using that baseline, common milestones can be identified to build demand forecasts for various roles and tasks. These same demand signals can help inform future academic and skills-training planning and approaches for government support.

Notes

1. World Economic Forum and Accenture, *Jobs of Tomorrow: Large Language Models and Jobs—a Business Toolkit*, December 2023, https://www.weforum.org/publications/jobs-of-tomorrow-large-language-models-and-jobs-a-business-toolkit/.

2. World Economic Forum and Accenture, *Jobs of Tomorrow*.

3. Heidrick & Struggles, *2023 Europe and US Data, Analytics, and Artificial Intelligence Executive Organization and Compensation Survey*, December 2023, https://www.heidrick.com/-/media/heidrickcom/publications-and-reports/2023-europe-and-us-data-analytics-and-artificial-intelligence-executive-organization-and-compensatio.pdf.

Job Loss, Displacement, and AI

Anticipating and Preventing Their Costs

Harry J. Holzer

Artificial intelligence will likely have major effects on the job market when it is more fully implemented in the workplace, and its effects will likely be mixed. On the one hand, AI will no doubt augment workers' skills in many cases, enabling them to perform jobs that had not been available to them earlier and making them more productive. On the other hand, millions of workers in many sectors of the economy may be displaced, as AI-enabled automation performs their tasks as well as or better (and more efficiently) than they could before.

If we can better anticipate the occupations and industries in which job loss is likely to occur, we can perhaps prevent some of it by training incumbent workers to perform other tasks. Alternatively, if some will be displaced, we can maybe give them a faster start at retraining for external jobs and ameliorate displacement's worst effects.

But this requires understanding the nature of worker displacements and job losses, as well as where and for whom the displacements are likely to occur.

Worker Displacement

Worker displacement (or worker dislocation) refers to a situation in which workers with at least some minimal level of job tenure (or seniority)—such as that accrued after three years' experience—permanently lose a job due to a workplace reorganization or closure. The cause of such displacements tends to be automation or globalization, both of which can have large effects on the production of goods and services in the US. (Other institutional changes, such as the outsourcing of workers and human resource functions to other companies, can lead to displacement as well.)

In steady state, about one million workers per year are displaced; these numbers tend to rise during recessions. But a general technology like AI could produce higher rates of displacement over time. And displacement is costly to workers and society. Many older displaced workers never return to formal employment; among those who do, earnings losses of 20–30 percent are not uncommon.

Figure 1. Manufacturing Employment in the US Since 1940

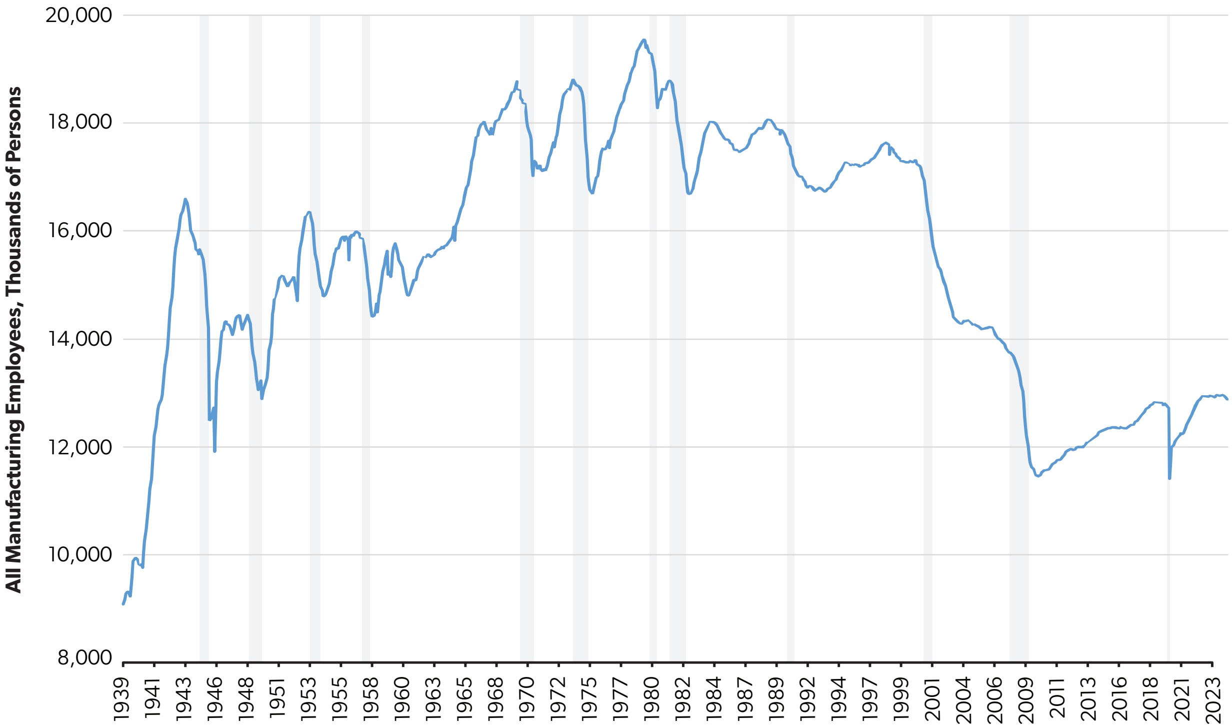

Source: US Bureau of Labor Statistics, "All Employees, Manufacturing (MANEMP)," Federal Reserve Bank of St. Louis, November 1, 2024, https://fred.stlouisfed.org/graph/?g=1dAhk.
Note: Shaded areas indicate US recessions.

When heavily concentrated in particular industries and regions, large displacements and job losses can hurt workers and communities beyond those directly displaced. The digital automation that began in the 1980s reduced employment for production and clerical workers in a number of industries; since these were occupations in which those with only high school diplomas (or less) were reasonably well compensated, these employment reductions reduced earnings for less-educated workers and increased inequality in the labor market between workers with and without college degrees.

Economists have called this phenomenon *skill-biased technical change*, and many view it as the primary culprit in the rising polarization of the US job market between more- and less-educated workers. And when these jobs and industries are highly concentrated in particular regions of the country, the effects can be harmful to these regions' fiscal health and many communities' well-being.

The Case of Manufacturing

One example of such a phenomenon is the large decline in manufacturing employment in the past few decades—especially in durable manufacturing, in which wages for production workers were relatively high. Employment in manufacturing peaked at just under 20 million in 1980; it declined mostly during the 1980s and 1990s before dropping rapidly after 2000. Indeed, manufacturing employment dropped by a remarkable five million between 2000 and 2010, the trough of the Great Recession, and by roughly four million between 2000 and 2019. (See Figure 1.)

Automation in the form of robotics and China's admission to the World Trade Organization (known among economists as the "China shock") were likely responsible.

Whatever its causes, the consequences of the loss of manufacturing jobs have been severe. Not only have earnings and labor force participation declined among less-educated men, but communities and regions in the industrial Midwest have gone into lasting decline, and the working class more broadly has experienced a rise in disability, morbidity, and mortality (which has been dubbed "deaths of despair").

The data and analysis of the effects of manufacturing job loss have all been retrospective, well after the damage has been done and too late to prevent it or help workers in a timely manner. Fortunately, we might be able to do better going forward with job changes associated with AI.

AI and Displacement: Where and Whom

Can we expect future displacements and job losses associated with AI to generate similarly negative impacts? It is not yet clear which occupations, industries, or regions AI will affect over time. It is quite possible that, given the general nature of AI, its overall impacts on employment will be larger than those of digital technologies and the China shock but less concentrated on a particular group or region. AI might well replace particular tasks many workers now do that do not account for most of their job tasks, which raises the likelihood that workers will be retained and retrained by employers rather than displaced.

And at least some evidence suggests that AI's impacts will be less biased against noncollege-educated workers; indeed, it might have larger effects higher up the education scale, given its abilities to do computations, writing, and data analysis, which to date have been performed in jobs by mostly college graduates. On the other hand, while the *incidence* of displacement might be less highly concentrated in the noncollege population, *adjustments* to such potential displacement might occur more rapidly and easily among college graduates, who might be better at anticipating task-performance losses and retooling themselves before such losses occur. Either way, a greater ability for policymakers and practitioners to predict where and for whom such task replacement will occur can likely improve the quality of our responses, helping workers adapt in their current jobs or move more easily to others.

More specifically, a better understanding of who faces task replacement by AI, and in which occupations and industries, will enable a more "rapid response" to potential dislocations, in the lingo of the US Department of Labor. Policymakers and practitioners might help employers retrain workers rather than displace them, overcoming a bias toward replacing workers with machines caused by federal tax incentives for capital investment or a lack of worker "voice" in management decisions. Alternatively, rapid response for those who will be displaced helps them get unemployment insurance more quickly and perhaps enroll sooner in college retraining. Educators could guide students to fields of study and task mastery that make their skills more complementary with AI and less vulnerable to further displacement.

To better understand AI's effects on jobs and workers, we need the right data—and, fortunately, those are becoming more available at the state level. Using quarterly earnings records from the unemployment insurance system, which covers almost all workers, we can identify how AI is affecting specific firms and groups of workers. Once we identify the firms affected—which we can do in a variety of ways—we can monitor which groups of workers they are newly hiring and which are being retained or displaced; then we can measure how each group is faring over time in terms of employment and earnings changes. The punch line: *Available data can help us measure the employment effects of AI in nearly real time* and enable us to respond rapidly to whatever we find.

And if a range of policies are enacted (such as the provision of Lifelong Learning Accounts for workers or more expanded options for training, as well as more voice for employees in the workplace), we could strengthen AI's strong upsides and lessen the negative effects many workers are likely to bear.

Notes

1 Henry S. Farber, "Employment, Hours, and Earnings Consequences of Job Loss: US Evidence from the Displaced Workers Survey," *Journal of Labor Economics* 35, no. S1 (2017): S235–S272, https://doi.org/10.1086/692353.

2. Louis S. Jacobson et al., "Earnings Losses of Displaced Workers," *American Economic Review* 83, no. 4 (1993): 685–709, https://www.jstor.org/stable/2117574.

3. David H. Autor et al., "Trends in U.S. Wage Inequality: Revising the Revisionists" (working paper, National Bureau of Economic Research, March 2007), https://www.researchgate.net/publication/24096173_Trends_in_US_Wage_Inequality_Revising_the_Revisionists.

4. Alan B. Krueger and Lawrence H. Summers, "Reflections on the Inter-Industry Wage Structure" (working paper, National Bureau of Economic Research, June 1986), https://www.nber.org/system/files/working_papers/w1968/w1968.pdf.

5. George J. Borjas and Richard B. Freeman, "From Immigrants to Robots: The Changing Locus of Substitutes for Workers" (working paper, National Bureau of Economic Research, January 2019), https://www.nber.org/papers/w25438; Daron Acemoglu and Pascual Restrepo, "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica* 90, no. 5 (2022): 1973–2016, https://www.econometricsociety.org/publications/econometrica/2022/09/01/Tasks-Automation-and-the-Rise-in-US-Wage-Inequality; and David H. Autor et al., "The China Shock: Learning from Labor Market Adjustment to Large Changes in Trade" (working paper, National Bureau of Economic Research, January 2016), https://www.nber.org/papers/w21906.

6. David Autor and Melanie Wasserman, *Wayward Sons: The Emerging Gender Gap in Labor Markets and Education*, Third Way, March 20, 2013, https://thirdway.imgix.net/downloads/wayward-sons-the-emerging-gender-gap-in-labor-markets-and-education/Third_Way_Report_-_NEXT_Wayward_Sons-The_Emerging_Gender_Gap_in_Labor_Markets_and_Education.pdf.

7. Autor et al., "The China Shock"; and Benjamin Austin et al., "Saving the Heartland: Place-Based Policies in 21st Century America" (working paper, Brookings Institution, March 8, 2018), https://www.brookings.edu/articles/saving-the-heartland-place-based-policies-in-21st-century-america/.

8. Anne Case and Angus Deaton, *Deaths of Despair and the Future of Capitalism* (Princeton University Press, 2020).

9. See, for example, Michael Webb, "The Impact of Artificial Intelligence on the Labor Market" (working paper, Stanford University, January 2020), https://www.michaelwebb.co/webb_ai.pdf. Another prediction that AI will affect highly educated workers in industries like finance and technology can be found in Society for Human Resource Management and Burning Glass Institute, *Generative Artificial Intelligence and the Workforce*, February 1, 2024, https://shrm-res.cloudinary.com/image/upload/v1706729099/AI/CPR-230956_Research_Gen-AI-Workplace_FINAL_1.pdf.

10. US Department of Labor, Employment and Training Administration, "Rapid Response Services," https://www.dol.gov/agencies/eta/layoffs.

11. Daron Acemoglu et al., "Does the US Tax Code Favor Automation?" (working paper, National Bureau of Economic Research, April 2020), https://www.nber.org/papers/w27052; and Henri Haapanala et al., "Robots and Unions: The Moderating Effect of Organized Labour on Technological Unemployment," *Economic and Industrial Democracy* 44, no. 3 (2023): 827–52, https://journals.sagepub.com/doi/abs/10.1177/0143831X221094078.

12. Julia Lane, "Industries of Ideas: Tracing the Links Between Investments in Science, Innovation, and Jobs," in *New Approaches to Characterize Industries: AI as a Framework and a Use Case*, ed. Brent Orrell (American Enterprise Institute, 2025).

13. For more on Lifelong Learning Accounts, see Alastair Fitzpayne and Ethan Pollock, *Lifelong Learning and Training Accounts: Helping Workers Adapt and Succeed in a Changing Economy*, Aspen Institute, May 2018, https://www.aspeninstitute.org/wp-content/uploads/2018/05/Lifelong-Learning-and-Training-Accounts-Issue-Brief.pdf. For evidence on retraining dislocated workers in funding from the US Department of Labor, see Fredrik Andersson et al., "Does Federally Funded Job Training Work? Nonexperimental Estimates of WIA Training Impacts Using Longitudinal Data on Workers and Firms," *Journal of Human Resources* 59, no. 4 (2024): 1244–83, https://jhr.uwpress.org/content/wpjhr/59/4/1244.full.pdf. If the kinds of funding for training and household stipends that Trade Adjustment Assistance provides could also be provided to workers displaced by automation, its positive effects could be broadened. See Benjamin Hyman, "Can Displaced Labor Be Retrained? Evidence from Quasi-Random Assignment to Trade Adjustment Assistance" (working paper, Federal Reserve Bank of New York, January 10, 2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3155386.

Part III

Perspectives from Data Collection Agencies and Organizations

Leveraging Wage Data and Other Administrative Records in Texas to Answer Questions About the Economy, Policy, and Prosperity

Adam Leonard

During the "New Approaches to Characterize Industries: AI as a Framework and a Use Case" workshop, experts considered new empirical methods to grouping firms into industries. One possible approach advocated by Julia Lane is to track investments in critical technologies by tracing where research-funded individuals go in the private sector; the assumption is that these individuals take what they have learned—their ideas—and deploy them as founders or employees. She called this the "industry of ideas," and she has advocated using it to examine the transformative effect artificial intelligence will have on the economy.

One potential way to trace these research-funded individuals is by using the universe of wage records that employers file quarterly to state unemployment insurance agencies, which agencies use to administer their programs. These records are a potential treasure trove of data that can track the transmission of ideas and their spread through the economy.

Starting with federal grant expenditure data (that is, administrative data), it is possible to identify who is funded to do research from a grant. If we were to connect the data identifying those doing research on a grant to university student administrative data or a state longitudinal data system (SLDS), sometimes called a P20W data system, we could track the spread of ideas from faculty to students.

It would then be possible to move longitudinally to trace where these grant workers and students go when they leave the university and become employed (or start a business and employ others) by looking at wage records over time. Further, by starting with grants, we could understand how these technologies spread through industries via grant workers' hiring in ways that are not always observable through surveys conducted in accordance with US Bureau of Labor Statistics standards, such as the Quarterly Census of Employment and Wages (QCEW).

Though QCEW includes detailed industry and "primary product or service" information for a firm, those data are limited in telling how a firm develops or uses new technology. Examining new employers that list AI as their primary business activity provides an implausibly low picture of the AI field in a state like Texas, where only 298 firms between January 2020 and January 2024 so identified, totaling 1,021 workers. Further, the primary product or service data are woefully underreported by employers, making it, at best, slightly useful.

Most of the biggest players in the AI space are established companies that are adding AI to their portfolio—their QCEW profile does not reflect this shift. Other companies that are investing heavily in AI are not even in the traditional tech sector but rather finance, insurance, bioscience and pharmaceuticals, and oil and gas. Through wage records, we can measure more directly the type of investments firms are making in AI—by looking at the number of employees they hire with AI backgrounds (former grant-funded researchers and their students), their overall employee head count, and how those numbers change over time.

In this chapter, I discuss the feasibility of such an empirical approach using Texas workforce data.

The Texas Experience

The Texas Workforce Commission (TWC) has long realized that its greatest untapped resource is its data. Data represent our past, inform our present, and can shape our future. The TWC has invested heavily in trying to tap into that resource to help Texas employers, individuals, families, and communities achieve and sustain economic prosperity—that is, "Data for Prosperity." The remainder of this chapter describes how Texas has used administrative and wage data for alternative purposes and how TWC conducted a limited tests of some of the ideas put forward in this workshop for tracking the spread of ideas to particular firms and clustering them into Lane's "industries of ideas."

Making and Maintaining Employment Connections

The Texas Workforce Commission operates dozens of programs intended to help employers, individuals, families, and communities. Its core principle is ensuring that everyone who wants to work has the skills necessary to fill an employer's needs and that employers can find the qualified workers they need. While there are many ways to measure these outcomes, none of the statutory methods get to the core concept of helping connect workers to employers.

As part of that work, TWC uses wage records to better understand our program results, the experiences of employers and others we have served, and the labor market more broadly. This moves beyond simply measuring performance or being limited to the Bureau of Labor Statistics' standards in our work. TWC has long used wage data to better comprehend industries' employment patterns.

One key concept TWC recently built focuses on identifying new employment connections between workers and employers (presumed to exist when an employer's quarterly report lists an employee who had not been reported in the prior two quarters). We then see what percentage of those new employment connections are maintained for at least two subsequent quarters. We're able to examine these data by industry and can better characterize the industry's nature and identify positive and negative outliers. This approach also creates an objective, independent picture of natural churn in the full labor market and gives programs a way to evaluate their success in helping people retain employment—that is, a way to measure the quality of the employment connections we help create relative to those generally experienced in the industry.

Evaluating Popular Narratives

Beyond that, TWC has found that combining administrative and wage data can be extremely helpful in evaluating narratives.

Figure 1. Teacher Trends in Texas

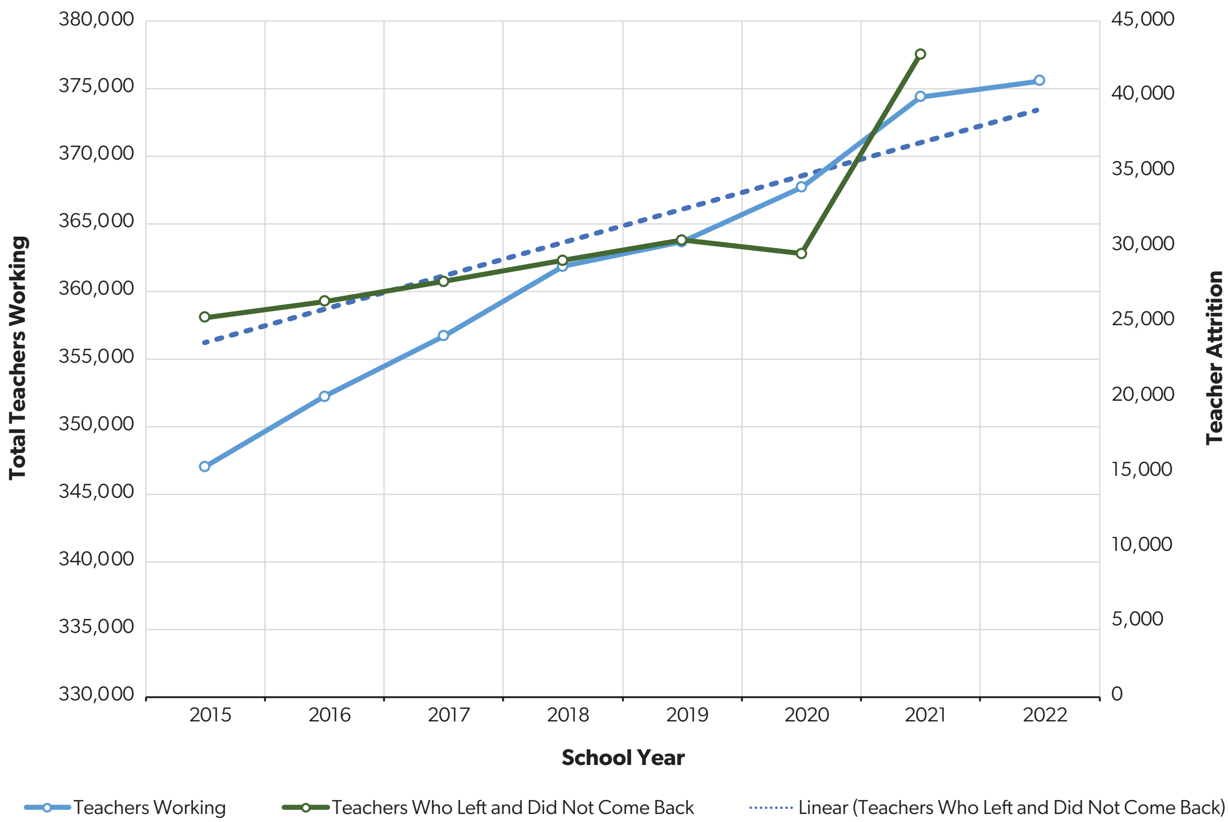

Source: Author.

For example, before the pandemic, the media published numerous stories about Texas schools struggling to find teachers and the large number of teachers leaving the profession. Using teacher data from the Texas Education Agency (TEA) through the Texas Tri-Agency Workforce Initiative, TWC identified Texas teachers who left the profession each year from 2015 through 2022 and then pinpointed where they went through analyzing wage records. This analysis contextualized and tempered the prevailing narrative in several ways.

First, the data showed that the number of teachers working in Texas grew every year over the period examined, and the growth curve was largely unchanged from pre- to post-pandemic. (Indeed, there was a significant bump in 2021, shown in Figure 1.)

Second, the data showed that a high percentage of teachers "leaving" were still employed in education services. This caused us to go back to TEA for more data on the roles these former teachers now held. The vast majority were in more senior education roles, such as assistant or full principal, counselor, department head, or teacher supervisor or facilitator.

Third, a significant portion of the teachers who left the profession did not show up in wage records at all, suggesting that some portion of them retired. Matching these data to those of the Texas Retirement System might provide additional clues in that regard.

Together, these data do not refute the difficulty that many school districts may have with finding enough teachers. However, they do provide a more nuanced picture of the extent of the issue and the factors contributing to it.

Proof of Concept for Industry of Ideas

TWC also has begun experimenting more directly with the industry-of-ideas approach using a quick and simple proof of concept. We started with University of Texas at Austin employees who were employed there in 2021 but not in 2022. From that subset, we then isolated the employees working elsewhere in 2022 who had quarterly earnings of $50,000 or more, assuming that people who previously worked at a university and then immediately started earning $200,000 a year or more are likely to be in high-demand fields such as AI or computer science. This is an imperfect assumption, but it's workable for the exercise. There were 140 such individuals.

We then examined these individuals' new employers and the industries those employers were in. While we can't be sure all 140 former university employees were in AI or other advanced computer work, many were at firms that identified as being in the "custom computer programming" industry. What was more interesting was the other industries we found these workers in. Some, such as "offices or lawyers," were probably not relevant to our research, but many others likely were. For example, many of these former grant-funded university employees were working for some of the biggest e-commerce, internet, AI hardware, biochemical or pharmaceutical, and materials science firms in the world—all of which makes sense given AI's potential to transform those industries.

Now that we have identified those firms, it is possible to examine the educational background of most of those firms' workers. States have long used SLDSs to track students through K–12, postsecondary education, and their careers. These systems combine schools' administrative data about students (where they went to school, what they studied, and what credentials they achieved) with wage data to show where they are working, how much they earn, and how those earnings change over time. But that has been primarily an education evaluation system and a way to help students make informed decisions about which schools to attend and programs to enroll in—not as a mechanism to examine industries and the labor market.

Conclusion

Combining administrative and wage records data provides an avenue for closely examining employment and earnings trends, the spread of ideas, and the impact of federal investments as people move through their careers and firms. If we were to combine SLDS data with grant-tracking data, we could use them to do a more macroeconomic analysis of industries and technologies and how they change over time.

Notes

1. Julia Lane, "The Industry of Ideas: Measuring How Artificial Intelligence Changes Labor Markets," American Enterprise Institute, June 9, 2023, https://www.aei.org/research-products/report/the-industry-of-ideas-measuring-how-artificial-intelligence-changes-labor-markets/.

2. Data for Prosperity was a concept I coined that applied to TWC's data strategy and promoting data use to improve decisions that can lead to prosperity.

3. Caroline Wilburn, "Struggling to Find Teachers Close to Home, Texas Schools Are Looking Overseas for Help," *The Texas Tribune*, December 20, 2023, www.texastribune.org/2023/12/20/texas-teacher-shortage-cultural-exchange-aldine-isd/; Andrea Lucia, "Extended Coverage: Texas Teachers Are Leaving the Profession in Record Numbers. They Told Us Why.," CBS News, September 8, 2023, www.cbsnews.com/texas/news/texas-teachers-are-leaving-the-profession-in-record-numbers-we-got-firsthand-accounts-on-why/; and Sarah Guthery and Lauren P. Bailes, "Texas Tried to Fix Its Teacher Shortage by Lowering Requirements—the Result Was More New Teachers, but at Lower Salaries," *The Conversation*, November 2, 2023, www.theconversation.com/texas-tried-to-fix-its-teacher-shortage-by-lowering-requirements-the-result-was-more-new-teachers-but-at-lower-salaries-213962/.

Using Online Labor Market Data to Measure AI Impact

Layla O'Kane

As artificial intelligence continues to reshape how we live and work, measuring its size and impact on the economy will become increasingly important for policymakers, workers, learners, and employers who are trying to understand how to incorporate AI into their decision-making. This measurement challenge holds true for a wide range of emerging and important technologies, including those used in green jobs, digital jobs, and other key areas of investment for a future-ready workforce. Real-time labor market data, such as online job postings and online social profile data, can help support this measurement challenge and provide insight into high-level growth trends and granular information about which occupations, geographies, and sectors have been most affected by AI. This report outlines key variables available in these data, a working definition of AI using a "basket-of-skills" approach, challenges and opportunities with this methodology, and potential next steps and applications to other technologies.

Approaches to AI Measurement Using Online Labor Market Data

Two primary online labor market data sources are job posting data and social profile data. Job posting data comprise an aggregate of online job advertisements and job postings that are scraped from individual employer websites, government sources, job boards, and other aggregators with high frequency, often daily. Social profile data are sourced from online professional profiles and include information similar to that of a resume or curriculum vitae. More information on the information available in each of these data sources can be found in Table 1.

One of the main challenges for AI measurement is that AI does not reside in only one industry or occupation but rather cuts across many sectors, employers, and types of jobs. However, job posting data and social profile data can use skills and competencies instead of sectors to create a definition of AI. This makes for a flexible and dynamic measurement system that can look across the labor market. Specifically, one can define AI jobs as all jobs that require an AI-related skill (e.g., natural language processing,

Table 1. Information Available from Online Labor Market Data

Information Extracted from Job Posts	Information Extracted from Social Profiles or Resumes
• Employer name and industry • Job title and occupation • Skill requirements, including AI technologies or AI-related skills • Education, certification, and experience requirements • Compensation • Job location • Job type: full-time, part-time, permanent, temporary, internship, remote, or other • Duration of post	• Worker name (generally anonymized during processing and aggregation) • Residential location and job locations • Current and former job titles and occupations • Time spent in each role • Current and former employers and industries • Skills and competencies, including AI-related skills and technologies • Certifications, degree attainment, and education programs

Source: Adapted from Asia-Pacific Economic Cooperation, Policy Support Unit, *Big Data for the Labor Market: Sources, Uses and Opportunities*, December 2021, https://www.apec.org/docs/default-source/publications/2021/12/big-data-for-the-labor-market-sources-uses-and-opportunities/221_psu_big-data-for-the-labor-market.pdf.

robotics, and autonomous driving). This basket-of-skills approach allows for flexibility in the definition to account for the following:

- Distinguishing between developers of artificial intelligence tools and users of those tools. For example, this definition could include a software engineer building a machine learning model and a curriculum writer using ChatGPT to support their copyediting.
- Direct technical skills (e.g., "Python") and subject-matter expertise (e.g., "computer vision").
- Subclusters of AI skills, such as generative AI, machine learning, visual image recognition, robotics, and neural networks.

An example of the skills that could be included in an AI measurement definition can be found in Appendix A.

Using this approach can answer questions about the impact of a new policy change, such as the following:

- Was there an increase in demand for AI jobs (measured in total job postings)?
- Which occupations, locations, industries, and employers saw the biggest increase in demand for AI jobs and AI skills?
- Which AI skills were most demanded before and after the policy change?
- What new AI-focused companies were created after the policy change?
- What job titles have had the highest growth in AI skills, and what other skills frequently co-occur with these AI skills (e.g., leadership, management, and collaboration)?
- What is the salary premium associated with certain AI skills, and has that changed over time?
- Which certifications, majors, and education programs are demanded most by employers and advertised most by workers?

In addition to the granularity and flexibility that this measurement approach offers, there are other advantages to using real-time online labor market data, especially when compared with more traditional methods of measurement such as survey data. These data, by nature, are updated in real time, which means researchers can learn about new skills and technologies quickly. In contrast, survey data are slow to collect and often have lags of at least one month. Real-time online data show emerging trends, especially with new technologies.

Similarly, due to the nature of "big data" sample sizes (upward of 30 million postings in the United States annually), sample size issues become less of a concern for measuring results relative to survey data. If a policy is implemented that targets only one or two new technologies, this effect can still be tracked. The data also make for easy comparison across countries, industries, and other subsectors, especially if there is taxonomy alignment, which is rare in publicly available survey data across countries.

This approach to AI measurement has already gained traction, with variations of this basket-of-skills approach in several recent publications on AI:

- The Stanford Institute for Human-Centered Artificial Intelligence annually publishes the *Artificial Intelligence Index Report*, which uses this approach in its "Economy" chapter to highlight growth in AI job postings over time.[1]

- The Organisation for Economic Co-operation and Development (OECD) and Lightcast October 2023 joint working paper, "Emerging Trends in AI Skill Demand Across 14 OECD Countries," showed an increase in demand for socio-emotional, foundational, and technical skills across AI employers and growth in AI in information and communications technology and professional services across 14 OECD countries.[2] This was followed by the *OECD Skills Outlook* 2023, which highlighted the need for ethical AI skills.[3]

- A compendium of analysis published in 2022 by the European Network on Regional Labour Market Monitoring titled "The Relevance of Artificial Intelligence in the Digital and Green Transformation of Regional and Local Labor Markets Across Europe" showed the growth in AI in the labor market across large EU countries.[4]

- The 2022 paper "Artificial Intelligence and Jobs: Evidence from Online Vacancies" finds rapid growth in AI-related jobs driven by firms whose workers' tasks are compatible with current AI technologies.[5]

- The 2020 paper "Artificial Intelligence in Health Care? Evidence from Online Job Postings" finds that there has been relatively little AI adoption in the health care industry despite substantial media discussion of AI in health care.[6]

- The 2020 paper "Artificial Intelligence, Firm Growth, and Product Innovation" finds that firms that invest in AI experience higher growth in sales, employment, and market valuations.[7]

In addition to these policy and academic publications, this basket-of-skills approach has been used in many shorter-form articles and dashboards, including Lightcast's *2023 Global Talent Playbook*,[8] a heat map of AI in the UK,[9] and articles on ethical AI[10] and generative AI.[11]

Challenges and Gaps

While the approach described above can provide substantial insight into AI measurement, there are some gaps and challenges inherent to these data sources. Online labor market data vary in terms of representativeness by industry and country. Some industries, such as information technology, are likely to post jobs online, whereas others, such as retail, may rely on help wanted signs or other in-person job ads. While it is likely that many positions involving new AI technologies will be posted online, it is possible that some jobs to which AI technologies are supplementary or complementary are less likely to be posted online.

Similarly, there is imperfect representation in online social profile data. Certain professionals such as doctors and academics are much less likely to have online professional profiles relative to others such as CEOs and sales workers. Since a substantial portion of AI technological development is likely accomplished by workers who have some academic affiliation, this may constitute a larger gap relative to other datasets.

Additionally, while this measurement approach allows for insight on the results of investment in technology, it does not allow for direct tracking of funding. This means there is likely to be latent AI change that has not yet become a new skill in the labor market, which will not be captured by this approach. For example, a substantial funding allocation to the use of generative AI in classroom curricula may not be directly captured by job postings data, but it will likely eventually be captured as the knowledge of uses for generative AI grows and employers start to request skills related to it more frequently. Other inputs pivotal to policy conversations are also easy to miss in online labor market data, such as education credentials and especially micro-credentials, since employers tend to request skills but are less likely to request specific ways for workers to demonstrate they have those skills.

As AI becomes increasingly popular, it is also possible that employers or job seekers will add AI-related skills to their postings and profiles because of hype rather than a true demand for or possession of that skill. To hedge against this, researchers have categorized AI skills as being "generic" and "specific" and required at least one specific AI skill in a posting or profile to count as an AI job or profile.[12]

To track these skills, the data need a flexible and robust skills taxonomy, which is often provided in various forms by data collection companies. This taxonomy enables tracking of key AI skills and, with regular updates, new and emerging technologies. Without access to a skills taxonomy, a similar approach could be used relying on keyword searching, though this can be prone to false positives and require substantial quality review.

Last but certainly not least, while these data are widely available, they are not publicly available. Online labor market data are largely collected by the private sector and expensive to aggregate, clean, and maintain. They also require relatively frequent taxonomy updates to keep abreast of current changes and allow for new occupations to be measured and emerge. These updates are best taken on by a dedicated team that works with these data and can understand how best to synthesize and aggregate text data for use by researchers (itself an AI problem).

Next Steps

The technological challenges associated with maintaining these big data systems are not negligible, but the good news is that many states and government agencies already have access to these data through third parties.

In the short term, greater access to clean postings and profile data could help workforce agencies, state and local governments, and other key stakeholders make decisions more quickly. While a working definition of an AI sector is built into Lightcast's Analyst tool, not all stakeholders have access, and this definition may not suit every priority. Funding should be expanded to ensure that these agencies can access data on top in-demand AI skills and jobs so that they can know where best to use other available dollars to shore up training to support AI.

In the medium term, measurement of AI impact should extend beyond a purely economic endeavor and move to include other considerations, such as ethics. Recent research published jointly by my team at Lightcast and the OECD looked at job postings for AI workers in 14 OECD countries.[13] Within those job postings, we looked for mentions of ethical AI, including keywords such as "AI ethics," "responsible AI," "ethical AI," and others across languages used in these countries. A tiny fraction (less than 2 percent) of AI job postings in all countries studied requested their AI workers have these skills. In the United States, where new generative AI technologies are developing rapidly, this proportion is only 0.5 percent. Without an understanding of how and with what other skills these tools are being used, many workers may continue to fear their work being supplanted by AI technology.

In the long term, flexibility will be the key to ensuring AI measurement continues to become more precise and insightful. Updating the skills that are relevant to or frequently use AI and continuously assessing if there are new opportunities to measure AI in different ways will require more collaboration between data collection sources. Additionally, adding new AI job titles, occupations, and sectors to government-maintained taxonomies would help measure other variables such as employment. This would require additional collaboration between stakeholders, including data aggregators and official government statistical offices, such as the Bureau of Labor Statistics.

Informing Other Emerging Technologies

In addition to measuring AI, there are several other critical or emerging technologies to which this basket-of-skills approach using online labor market data has already been applied. This section reviews key emerging technologies—digital jobs and green jobs—and provides examples of published works that use this approach to measure the technological impact on the economy.

Relevant research on digital jobs includes the following:

- *APEC Closing the Digital Skills Gap Report: Trends and Insights* looks at digital job growth in Australia, Canada, New Zealand, Singapore, the UK, and the US.[14]

- *Digitalization in the German Labor Market: Analyzing Demand for Digital Skills in Job Vacancies* analyzes digital skill growth in Germany and finds digitalization is correlated with socioeconomic factors such as education level, gender, and income.[15]

- *No Longer Optional: Employer Demand for Digital Skills* shows how pervasive digital skills have become across a wide range of jobs in the UK.[16]

Relevant research on green jobs includes the following:

- "Green Energy Jobs in the US: What Are They, and Where Are They?" uses a skills approach to define green jobs, finding that solar and wind jobs have more than tripled since 2010 and growth of renewable energy leads to the creation of relatively high-paying jobs.[17]

- Working Nation's *Green Jobs Now* report shows substantial growth in green jobs and estimates that 51 million workers could be reskilled into green jobs in the future.[18]

- "Who's Fit for the Low-Carbon Transition? Emerging Skills and Wage Gaps in Job Ad Data" finds that low-carbon job ads have higher skill requirements, particularly for technical skills, and the wage premium for low-carbon job ads has declined over time.[19]

- "Workers and the Green-Energy Transition: Evidence from 300 Million Job Transitions" finds that the rate of transition from carbon-intensive ("dirty") jobs to non-carbon-intensive ("green") jobs has risen rapidly in recent years.[20]

These examples demonstrate that online labor market data can help measure new, emerging technologies and provide a wide range of insight into how economies are changing in real time as a result of these technologies. This basket-of-skills approach is flexible, dynamic, and adaptable, lending itself well to as-yet-undefined or undiscovered technologies.

Conclusion

While online labor market data have their challenges, they can provide insight that other available data sources cannot. Their two main strengths, frequent (even daily) updates and large sample sizes, are particularly relevant for new technologies that may be just beginning to affect small pockets of the economy. These data are also often forward-looking, with employers listing job postings for a future hire, representing a view of the labor market to come rather than the current stock of the labor market. This can help policymakers and other stakeholders forecast effects of new technologies and track rates of adoption and impact in real time. Alongside other datasets, online labor market data analyzed using a basket-of-skills approach should be a key metric for measuring artificial intelligence and other emerging sectors.

Appendix A

A working definition based on AI skills and skill clusters for measuring AI is provided below, based on Lightcast's skills taxonomy.[21] This group of skills can be used to define a set of job postings as "AI jobs" if these skills are requested in the job posting.

Artificial Intelligence. AI/ML Inference, AIOps (Artificial Intelligence for IT Operations), Applications of Artificial Intelligence, Artificial General Intelligence, Artificial Intelligence, Artificial Intelligence Development, Artificial Intelligence Markup Language (AIML), Artificial Intelligence Systems, Azure Cognitive Services, Baidu, Cognitive Automation, Cognitive Computing, Computational Intelligence, Cortana, Ethical AI, Expert Systems, Explainable AI (XAI), IPSoft Amelia, Intelligent Control, Intelligent Systems, Interactive Kiosk, Knowledge Engineering, Knowledge-Based Configuration, Knowledge-Based Systems, Multi-Agent Systems, Open Neural Network Exchange (ONNX), OpenAI Gym, Operationalizing AI, Reasoning Systems, Watson Conversation, Watson Studio, Weka

Autonomous Driving. Advanced Driver Assistance Systems, Autonomous Cruise Control Systems, Autonomous System, Autonomous Vehicles, Guidance Navigation and Control Systems, Light Detection and Ranging (LiDAR), OpenCV, Path Analysis, Path Finding, Remote Sensing, Unmanned Aerial Systems (UAS)

Generative Artificial Intelligence. ChatGPT, Generative Adversarial Networks, Generative Artificial Intelligence, Large Language Modeling, Prompt Engineering, Variational Autoencoders

Natural Language Processing (NLP). AI Copywriting, ANTLR, Amazon Textract, Apache OpenNLP, BERT (NLP Model), Chatbot, Computational Linguistics, Conversational AI, Dialog Systems, fastText, Fuzzy Logic, Handwriting Recognition, Hugging Face (NLP Framework), Hugging Face Transformers, Intelligent Agent, Intelligent Virtual Assistant, Kaldi, Language Model, Latent Dirichlet Allocation, Lexalytics, Machine Translation, Microsoft LUIS, Natural Language Generation, Natural Language Processing, Natural Language Programming, Natural Language Toolkits, Natural Language Understanding, Optical Character Recognition (OCR), Screen Reader, Semantic Analysis, Semantic Parsing, Semantic Search, Sentiment Analysis, Seq2Seq, Speech Recognition, Speech Recognition Software, Speech Synthesis, Statistical Language Acquisition, Text Mining, Text-to-Speech, Tokenization, Voice Assistant Technology, Voice Interaction, Voice User Interface, Word Embedding, Word2Vec Models

Neural Networks. Apache MXNet, Artificial Neural Networks, Autoencoders, Caffe2, Chainer (Deep Learning Framework), Convolutional Neural Networks, Cudnn, Deep Learning, Deep Learning Methods, Deeplearning4j, Evolutionary Acquisition of Neural Topologies, Fast.ai, Keras (Neural Network Library), Long Short-Term Memory (LSTM), OpenVINO, PaddlePaddle, Recurrent Neural Network (RNN), TensorFlow

Machine Learning. AWS SageMaker, AdaBoost (Adaptive Boosting), Adversarial Machine Learning, Apache MADlib, Apache Mahout, Apache SINGA, Apache Spark, Association Rule Learning, Automated Machine Learning, Autonomic Computing, Azure Machine Learning, Boosting, CHi-Squared Automatic Interaction Detection (CHAID), Classification and Regression Tree (CART), Cluster Analysis, Collaborative Filtering, Confusion Matrix, Cyber-Physical Systems, Dask (Software), Data Classification, Dbscan, Decision Models, Decision Tree Learning, Dimensionality Reduction, Dlib (C++ Library), Ensemble Methods, Feature Engineering, Feature Extraction, Feature Learning, Feature Selection, Gaussian Process, Genetic Algorithm, Google AutoML, Gradient Boosting, H2O.ai, Hidden Markov Model, Hyperparameter Optimization, Inference Engine, K-Means

Clustering, Kernel Methods, Kubeflow, Loss Functions, MLOps (Machine Learning Operations), Machine Learning, Machine Learning Algorithms, Machine Learning Methods, Machine Learning Model Monitoring and Evaluation, Machine Learning Model Training, Markov Chain, Matrix Factorization, Meta Learning, Microsoft Cognitive Toolkit (CNTK), MLflow, ModelOps, mlpack (C++ Library), Naive Bayes Classifier, Perceptron, Predictive Modeling, PyTorch (Machine Learning Library), PyTorch Lightning, Random Forest Algorithm, Recommender Systems, Reinforcement Learning, Scikit-Learn (Python Package), Semi-Supervised Learning, Soft Computing, Sorting Algorithm, Supervised Learning, Support Vector Machine, Test Datasets, Theano (Software), Torch (Machine Learning), Training Datasets, Transfer Learning, Transformer (Machine Learning Model), Unsupervised Learning, Vowpal Wabbit, Xgboost

Robotics. Advanced Robotics, Bot Framework, Cognitive Robotics, Motion Planning, Nvidia Jetson, Robot Framework, Robot Operating Systems, Robotic Automation Software, Robotic Liquid Handling Systems, Robotic Programming, Robotic Systems, Servomotor, Simultaneous Localization and Mapping (SLAM) Algorithms

Visual Image Recognition. 3D Reconstruction, Activity Recognition, Computer Vision, Contextual Image Classification, Digital Image Processing, Eye Tracking, Face Detection, Facial Recognition, Gesture Recognition, Image Analysis, Image Matching, Image Recognition, Image Segmentation, Image Sensor, Imagenet, Machine Vision, Motion Analysis, Object Recognition, OmniPage, Pose Estimation

Notes

1. Stanford University, Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2023*, April 2023, https://hai-production.s3.amazonaws.com/files/hai_ai-index-report_2023.pdf.
2. Francesca Borgonovi et al., "Emerging Trends in AI Skill Demand Across 14 OECD Countries" (working paper, Organisation for Economic Co-operation and Development, October 17, 2023), https://doi.org/10.1787/7c691b9a-en.
3. Organisation for Economic Co-Operation and Development, *OECD Skills Outlook 2023*, November 6, 2023, https://www.oecd.org/en/publications/2023/11/oecd-skills-outlook-2023_df859811.html.
4. Christa Larsen et al., "The Relevance of Artificial Intelligence in the Digital and Green Transformation of Regional and Local Labour Markets Across Europe: Perspectives on Employment, Training, Placement, and Social Inclusion," *Forschung Und Deren Anwendung Im Bereich Des Regionalen Und Lokalen Arbeitsmarktmonitorings Im Transnationalen Vergleich* (*Veröffentlichungen Des IWAK*), 2022, https://doi.org/10.5771/9783957104113.
5. Daron Acemoglu et al., "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics* 40, no. S1 (2022), https://doi.org/10.1086/718327.
6. Avi Goldfarb et al., "Artificial Intelligence in Health Care? Evidence from Online Job Postings," *AEA Papers and Proceedings* 110 (May 2020): 400–04, https://doi.org/10.1257/pandp.20201006.
7. Tania Babina et al., "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics* 151 (January 2024): 103745, https://doi.org/10.1016/j.jfineco.2023.103745.
8. Lightcast, *2023 Global Talent Playbook*, https://4906807.fs1.hubspotusercontent-na1.net/hubfs/4906807/Lightcast_Talent-Playbook%20(100%25).pdf.
9. Lightcast, "Heatmap of AI Jobs," https://aiskills.lightcast.io/heatmap.
10. Scott Bittle, "Few Employers Ask for 'Ethical AI' Skills in Job Postings," Lightcast, July 19, 2023, https://lightcast.io/resources/blog/ethical-ai-in-job-postings.
11. Layla O'Kane, "4 Ways Generative AI Will Change the Job Market," Lightcast, October 11, 2023, https://lightcast.io/resources/blog/4-ways-generative-ai-will-change-the-job-market.
12. Borgonovi et al., "Emerging Trends in AI Skill Demand Across 14 OECD Countries."
13. Borgonovi et al., "Emerging Trends in AI Skill Demand Across 14 OECD Countries."
14. Asia-Pacific Economic Cooperation, Human Resources Development Working Group, *APEC Closing the Digital Skills Gap Report: Trends and Insights*, December 2020, https://www.apec.org/docs/default-source/publications/2020/12/apec-closing-the-digital-skills-gap-report/220_hrd_apec-closing-the-digital-skills-gap-report_rev.pdf.
15. Layla O'Kane et al., *Digitalization in the German Labor Market: Analyzing Demand for Digital Skills in Job Vacancies*, Bertelsmann Stiftung, July 2020, https://www.bertelsmann-stiftung.de/fileadmin/files/user_upload/STUDIE_Burning_Glass_EN_FINAL.pdf.
16. Julia Nania et al., *No Longer Optional: Employer Demand for Digital Skills*, June 2019, https://assets.publishing.service.gov.uk/media/5cfe713fed915d097daca4b5/No_Longer_Optional_Employer_Demand_for_Digital_Skills.pdf.
17. E. Mark Curtis and Ioana Marinescu, "Green Energy Jobs in the US: What Are They, and Where Are They?" Working Paper No. 30332 (National Bureau of Economic Research, August 2022), https://doi.org/10.3386/w30332.
18. Rachel Sederberg et al., *Green Jobs Now*, Lightcast and WorkingNation, April 21, 2023, https://workingnation.com/wp-content/uploads/2023/08/WorkingNation_Lightcast_GreenJobsNow_August2023.pdf.
19. Aurélien Saussay et al., "Who's Fit for the Low-Carbon Transition? Emerging Skills and Wage Gaps in Job Ad Data," Working Paper No. 31 (Fondazione Eni Enrico Mattei, October 27, 2022), https://doi.org/10.2139/ssrn.4260227.
20. E. Mark Curtis et al., "Workers and the Green-Energy Transition: Evidence from 300 Million Job Transitions" (working paper, National Bureau of Economic Research, August 2023), https://doi.org/10.3386/w31539.
21. Lightcast, "Lightcast Open Skills Library," https://lightcast.io/open-skills.

Harnessing the Power of AI to Understand Changing Skill Demands

Lesley Hirsch

For at least the past two decades, two trends have left an indelible imprint on the American economy: digitization and globalization. American policymakers responded to these trends' polarizing effects only after they became apparent and, largely, too stubborn to address. Automation replaced jobs that could be done more cheaply and with less labor. Similarly, firms sent jobs offshore to workers earning lower wages to cut production costs and maximize revenues.

The jobs most susceptible to both forces required workers to undertake routine cognitive and manual tasks. Workers with more than a high school diploma but less than a four-year college degree held many of these jobs. The impact on jobs requiring less than a high school education was mixed. While the twin trends were highly likely to affect low-wage jobs requiring manual labor, the demand for low-wage service jobs grew during this period.

Jobs requiring a four-year college degree were relatively unaffected, especially "knowledge jobs" in technology and professional services. These divergent impacts were compounded by the geographical concentration of the most-hard-hit industries, especially manufacturing. The result was a devastating loss of jobs in Midwestern and Northeastern urban centers and a burgeoning service economy on the coasts. Many economists and policymakers agree: Efforts to address these harms were too little, too late.[1]

Policymakers do not want to make the same mistake again. To help them, economists are using novel data sources and methods to predict how the emergence of artificial intelligence will affect the labor market. Which jobs are most likely to be displaced, and which will AI simply augment?

The following section distills the major insights from the existing literature that can inform policymakers' approach to anticipating, measuring, and mitigating any negative impacts of future technological shocks like the rise of AI. Following this overview is a description of a groundbreaking application, the NJ Career Navigator, that applies causal modeling and natural language processing to generate personalized recommendations for jobs, training opportunities, and career transitions for job seekers. This application is a strong proof of concept showing that states can deliver labor market information differently and meet the challenges of future technological shocks like AI head-on with timely, local, and granular intelligence about emerging skill needs.

Major Insights About AI's Predicted Impacts

Like automation, AI leverages the comparative advantages of computers over people, such as in computing large amounts of data quickly. A central observation of the literature on AI, however, is that AI differs from previous technological advancements in a critical respect: While computers are good at routine tasks, AI can infer tacit relationships that the underlying software or human user does not fully specify. As a result, AI can perform cognitive nonroutine tasks well.[2] Thus, AI is becoming increasingly proficient at generating narrative and audiovisual content, as seen in the past year through the emergence of OpenAI's GPT-4 and DALL-E, Anthropic's Claude, Google's MusicLM, and ElevenLabs' AI Voice Generator, to name a few examples.

Much of the research on AI and the workforce breaks work into component tasks and then estimates the degree to which AI can perform them. The researchers vary in how they define AI capabilities, such as by using predefined categories provided by proprietary job-ad-parsing software, the Electronic Frontier Foundation's list of AI applications, subject-verb pairings in patents, and definitions crowdsourced through Amazon Mechanical Turk.[3]

The researchers then cross-walk the tasks with occupations' skill and role requirements, often based on information from O*NET OnLine's production database.[4] Their results vary somewhat but share this primary takeaway: Occupations most exposed to AI require a higher-than-average degree of education, require pay above a median annual wage, and entail cognitive nonroutine tasks. Table 1 lists occupations the literature identified as being more and less susceptible to replacement by AI.

Despite this research's important contributions, it presents policymakers with several conundrums to address before they can predict the nature and extent of AI's impact. The first difficulty is in defining AI capabilities and mapping them to skills and occupations. Since researchers do not yet agree on how to identify these capabilities, the lists of the occupations they predict to be most susceptible to AI's emergence vary.

Moreover, as the rapid advancement of generative AI has recently demonstrated, the list of AI's capabilities is bound to change and grow. Even if AI's capabilities do not evolve, mapping skills to occupations is an inexact science. Occupational requirements vary even within a detailed, six-digit standard occupational code. For example, web developers (Standard Occupational Classification code 15-1254) include front-end developers, who might be required to know JavaScript and responsive design, and back-end developers, who do not need either skill but may have to be proficient with Python and MySQL.

Second, most of the research does not take an explicit stand on which tasks AI will replace versus augment. However, it is reasonable to infer that AI will likely replace those occupations in which a majority of tasks are highly exposed to it, whereas it may simply augment occupations with a lower percentage of highly exposed tasks.[5]

Third, since labor markets are complex social phenomena shaped by interconnected factors, accurately predicting their trajectories poses a considerable challenge. For example, when ATMs became prevalent, many believed the demand for bank tellers would decrease. Exactly the opposite happened. Freed from dispensing cash, bank tellers' tasks changed to encompass more sales and customer service duties. Bank branches proliferated, and there are more bank tellers today than before the ATM.

New industries emerge with technologies too: There were no airplane maintenance technicians before air travel and no movie producers before moving pictures. We are still unsure how other major economic, demographic, and natural trends—for example, immigration, aging, climate change, global conflict, and large-scale investments like those related

Table 1. Occupations with Low and High Susceptibility to Replacement by AI

Low Susceptibility	High Susceptibility
Animal caretakers, except farmers	Actuaries
Animal scientists	Brokerage clerks
Archaeologists	Budget analysts
Construction laborers and helpers	Clinical lab technicians
Dancers	Coders
Dining room and cafeteria attendants	Concierges
Fitness trainers	Credit authorizers
Food preparation workers	Customer service representatives
Heavy truck and bus drivers	Financial advisers
Home health aides	Genetic counselors
Iron workers	Judicial law clerks
Massage therapists	Mathematicians
Plasterers and stucco masons	Mechanical drafters
Plumbers and pipefitters	Morticians
Public address system announcers	Office clerks
Slaughterers and meat-packers	Optometrists
Special education teachers	Paralegals
Vocational education teachers	Retail salespersons

Source: Erik Brynjolfsson et al., "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?," *AEA Papers and Proceedings* 108 (May 2018): 43–47; Kweilin Ellingrud et al., *Generative AI and the Future of the Workforce*, McKinsey & Company, July 26, 2023, https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america; Edward W. Felten et al., "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal* 42, no. 12 (2021): 2195–217: https://onlinelibrary.wiley.com/doi/full/10.1002/smj.3286; Ian Shine, "These Are the Jobs That AI Can't Replace," World Economic Forum, May 17, 2023, https://www.weforum.org/agenda/2023/05/jobs-ai-cant-replace; Michael Webb, "The Impact of Artificial Intelligence on the Labor Market" (working paper, Social Science Research Network, January 11, 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3482150; and Ali Zarifhonarvar, "Economics of ChatGPT: A Labor Market View on the Occupational Impact of Artificial Intelligence," *Journal of Electronic Business & Digital Economics* 3, no. 2 (2024): 100–16, https://www.emerald.com/insight/content/doi/10.1108/jebde-10-2023-0021/full/html.

to the Bipartisan Infrastructure Law, the CHIPS and Science Act, and the Inflation Reduction Act—will affect the labor market, let alone interact with one another and AI's advancement.

Finally, researchers must better identify the differential impacts of AI on populations and geography subgroups. For example, labor-enhancing AI in the hands of a highly paid, highly educated professional might increase that person's value in the labor market, exacerbating income inequality. Conversely, if AI replaces highly paid, highly educated professionals, their value will decrease, and income inequality could shrink.

Also, there is a nonzero amount of occupational segregation by race, ethnicity, and gender in the US, so differential impacts are likely. Similarly, impacts

Figure 1. My Career NJ Home Page

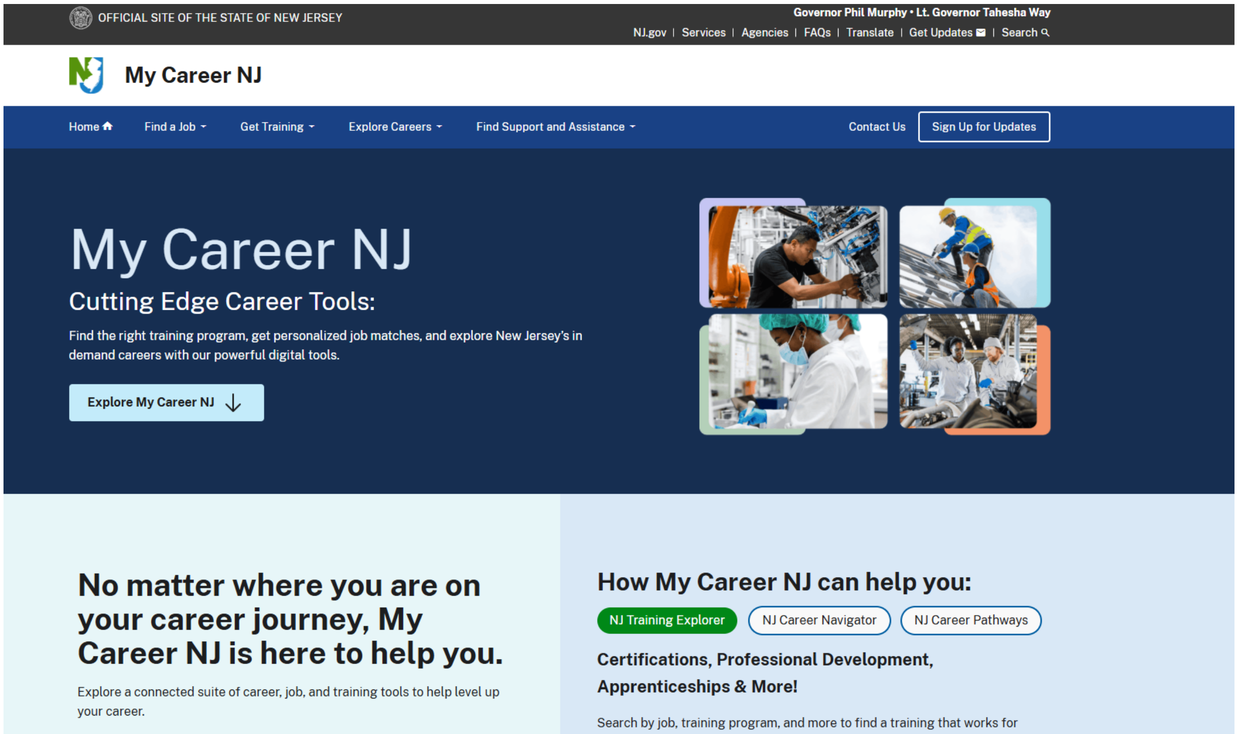

Source: My Career NJ, website, https://mycareer.nj.gov/.

might be more pronounced in areas of the country with higher concentrations of people employed in the affected occupations.

Using AI to Help Job Seekers: A New Jersey Case Study

Ironically, AI can bridge the data gap and provide granular, timely information about skills in close to real time. AI is uniquely capable of quickly revealing the latent structures in multidimensional occupational skills data, despite the vast amounts of data from online job ads and publicly available sources like O*NET. AI can parse job ads into meaningful components including skill requirements and use natural language processing and clustering techniques to signal what skill sets are emerging and declining.

New Jersey provides a suite of tools for workers called My Career NJ. (See Figure 1.) In 2024, the New Jersey Department of Labor and Workforce Development (NJDOL), Governor Phil Murphy's Office of Innovation, and the nonprofit RIPL launched a cloud-based application called the NJ Career Navigator.[6] (See Figure 2.) Along with its sister application, the NJ Training Explorer, the NJ Career Navigator is part of the My Career NJ toolkit. The Career Navigator provides job seekers with AI-generated recommendations on new careers, job postings, and training programs.

The data powering this AI recommendation engine include millions of wage records, providing earnings and industry information on all workers covered by unemployment insurance in New Jersey firms; employment and wage outcomes for hundreds of thousands of graduates of New Jersey occupational skills training programs; several years of NLx Research Hub online job postings; and resumes from 400,000 New Jersey residents, which are used to understand the interplay of education and careers, identify workers' transitions,

Figure 2. NJ Career Navigator Log-In

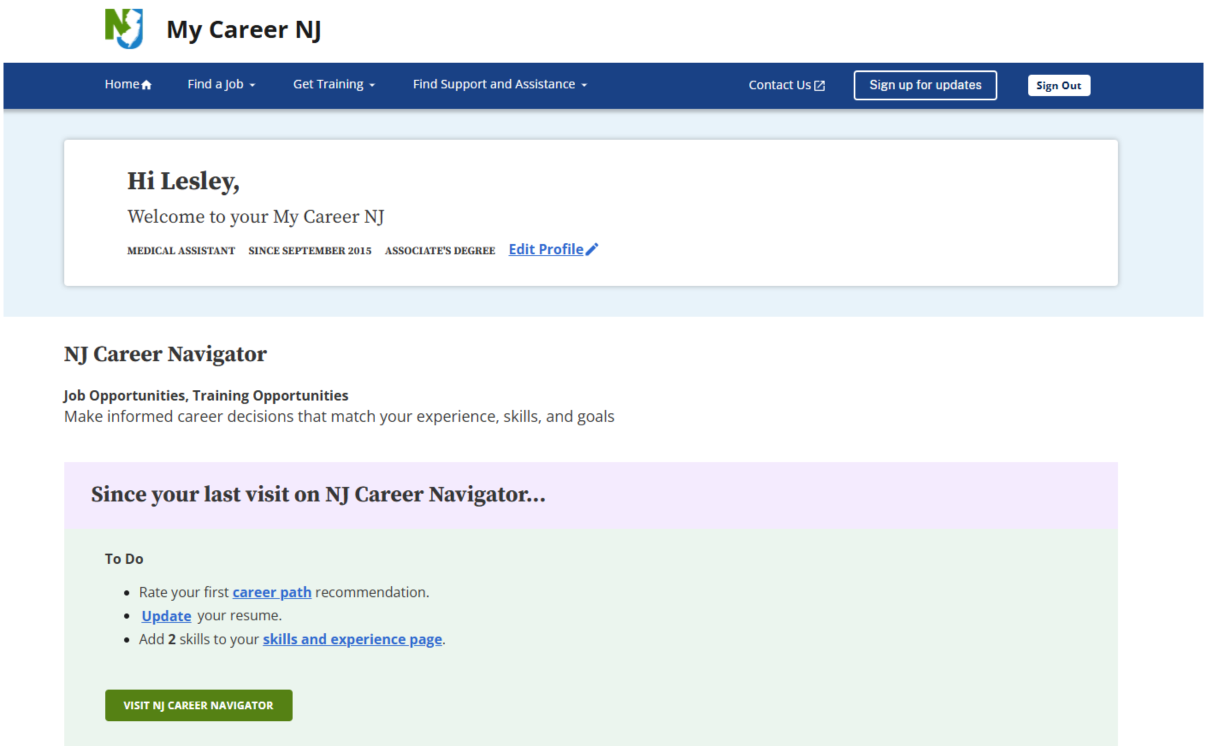

Source: New Jersey, "NJ Career Navigator," https://mycareer.nj.gov/navigator/.

and enrich the information connecting skills to occupations. These data have been modeled to identify successful job changes, career transitions, and training program completions that resulted in increased earnings for previous job seekers.

This information is combined with measures of skill similarity across occupations and occupational demand, as measured by the volume of online job ads in the National Labor Exchange. The underlying science relies on robust, unbiased causal machine learning models to estimate the earnings boost from career transitions, job choices, and training programs. Natural language processing was used to measure the skill similarities among occupations.

Career seekers log in to the NJ Career Navigator and upload their resumes or enter one or more of their latest job titles and level of educational attainment. The NJ Career Navigator automatically generates a list of skills most closely associated with the users' most recent jobs, which the user can curate to include up to a dozen skills selected from thousands. Next, career seekers can select the type of recommendation they want: a job opening right now, a training program that will help them advance their career, or a career transition.

The NJ Career Navigator then generates the appropriate type of recommendations. If the user desires training recommendations, the results will provide links to programs on the NJ Training Explorer, and the user can contact a school directly. If the user desires a career transition, the application offers career ideas that leverage their skills, provide opportunities for increased earnings, and are in high demand. If the user requests job recommendations, the NJ Career Navigator generates links to live job ads. (See Figure 3.) The results can be sorted and filtered by several variables, such as skill similarity, job ad volume, and location.

Other states can replicate the NJ Career Navigator. Along with its sister site, the NJ Training Explorer, it

Figure 3. The NJ Career Navigator's New Career Recommendations for a Medical Assistant

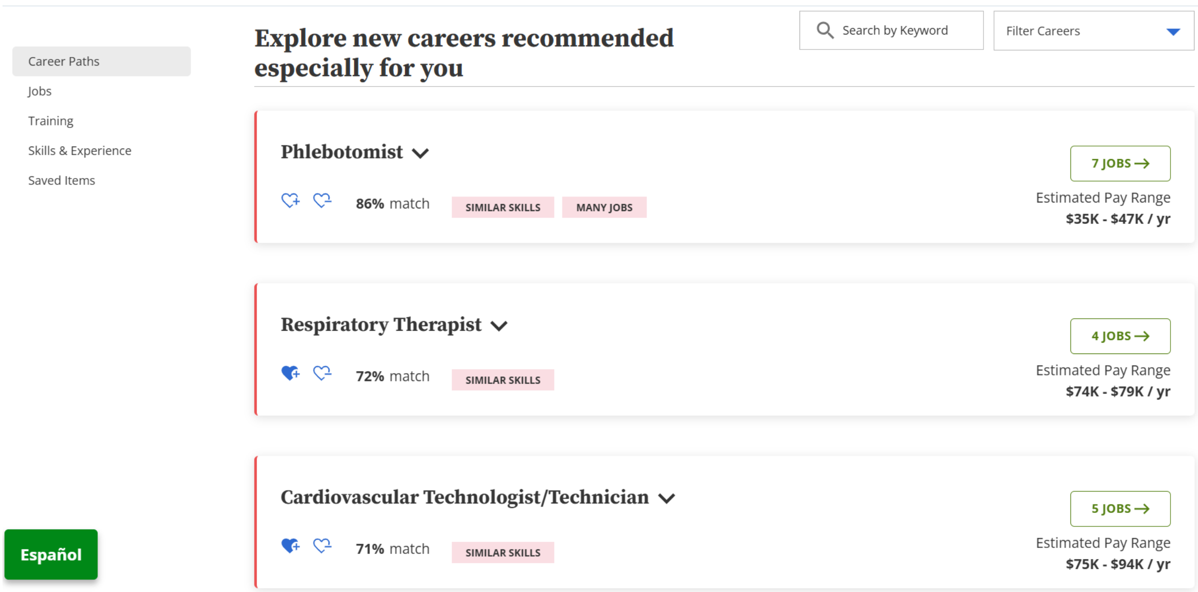

Source: New Jersey, "NJ Career Navigator," https://mycareer.nj.gov/navigator/.

was created using human-centered design with open code in a GitHub repository available to other states on request.[7] This is a powerful approach that governments can leverage to gain insights into evolving skill demands in the labor market.

A New Kind of Labor Market Information

Currently, governments rely on administrative data and large-scale surveys administered by the federal statistical system. While these sources are statistically sound and provide an essential overview of the national economy, they are released on a lag, are not as useful at the state and local levels, and are not designed to surface new technologies and skill requirements. State, local, and tribal governments pay annual licensing fees to access software that enables them to analyze online job ads by using proprietary models and algorithms to generate reports. It is time for governments to harness the power of their own data and use AI to generate a new kind of labor market information without subscription fees.

Marrying government employment data with natural language processing of online job postings is a promising approach for policymakers and workforce agencies to better understand and anticipate skill trends. It's an example of how governments can and should leverage big data and AI techniques to gain valuable labor market intelligence. Online job ads contain a wealth of information about the specific skills and qualifications employers are looking for in new hires. By training large language models on the full text of these job ads, states can reveal the underlying patterns and relationships among job titles, required skills, and other attributes. For example, the model could learn that data analyst jobs are increasingly requiring skills in R but decreasingly requiring proficiency with SQL or proprietary statistical software.

The advantage of using these large language models is that they can capture nuanced, contextual relationships between skills and jobs that rigidly defined taxonomies or keyword searches might miss. The models can be retrained and updated frequently to stay on top of rapid changes. Simply by applying clustering techniques, analysts can identify which skill sets are ascendant and which are in decreasing demand, including AI and non-AI skills.

Conclusion

By leveraging the power of machine learning and natural language processing, we can gain unprecedented visibility into the changing landscape of skill demand. This powerful tool can help reinvent workforce development practice, educational program design, and career guidance for job seekers. As in New Jersey, governments can facilitate career transitions by helping job seekers replace unnecessary skills in the workplace with other skills in adjacent careers requiring minimal additional educational investment.

Using this approach, governments can empower secondary and postsecondary institutions to develop program content that quickly aligns with the evolving demands of the labor market. Governments can shore up the workforce's resiliency by informing workers about which fields are more and less resistant to replacement by AI. Finally, this approach can help governments locate the most negative impacts so they can mitigate them through tax-and-transfer policies and other social safety-net programs.

The time has come for state, local, and tribal governments to harness the power of these transformative analytical capabilities and deliver timely, actionable insights to their constituents. Admittedly, the analytical capacity required to build and implement such sophisticated models remains beyond the reach of most workforce agencies today. However, strategic partnerships can bridge this gap: By leveraging the expertise of university-based researchers and nonprofit innovators, governments can cultivate the scientific foundations needed to revolutionize education and workforce development services.

If policymakers embrace a collaborative, open-source approach, the burden of developing this groundbreaking methodology need not be the responsibility of multiple jurisdictions at once. Instead, through openly sharing code, a handful of pioneering states can lay the groundwork, which their counterparts nationwide can readily adapt and build on. Thus, state, local, and tribal governments can collectively unlock a new era of data-driven, responsive policymaking that empowers workers, strengthens communities, and positions the American work-force for long-term resilience.

Notes

1. David H. Autor and David Dorn, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *The American Economic Review* 103, no. 5 (2013): 1553–97, https://www.aeaweb.org/articles?id=10.1257/aer.103.5.1553; and Maarten Goos et al., "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *The American Economic Review* 104, no. 8 (2014): 2509–26.

2. David Autor, "Polanyi's Paradox and the Shape of Employment Growth" Working Paper No. 20485 (National Bureau of Economic Research, September 2014), https://www.nber.org/papers/w20485.

3. See, for example, the different approaches to identifying AI capabilities in Edward W. Felten et al., "A Method to Link Advances in Artificial Intelligence to Occupational Abilities," *AEA Papers and Proceedings* 108 (May 2018): 54–57, https://www.aeaweb.org/articles?id=10.1257/pandp.20181021; Liudmila Alekseeva et al., "The Demand for AI Skills in the Labor Market," *Labour Economics* 71 (August 2021): 102002, https://www.sciencedirect.com/science/article/abs/pii/S0927537121000373; Erik Brynjolfsson et al., "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?," *AEA Papers and Proceedings* 108 (May 2018): 43–47; and Michael Webb, "The Impact of Artificial Intelligence on the Labor Market" (working paper, Social Science Research Network, January 11, 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3482150.

4. O*NET OnLine is a comprehensive source of information on hundreds of occupations, including required skills and education,

pay, and outlook, supported by the US Department of Labor's Employment and Training Administration.

5. Daron Acemoglu et al., "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics* 40, no. S1 (2022): S293–S340, https://www.journals.uchicago.edu/doi/abs/10.1086/718327; Alekseeva et al., "The Demand for AI Skills in the Labor Market"; Brynjolfsson et al., "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?"; Erik Brynjolfsson et al., "How Will Machine Learning Transform the Labor Market?," Hoover Institution, May 6, 2019, https://www.hoover.org/research/how-will-machine-learning-transform-labor-market; Edward W. Felten et al., "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal* 42, no. 12 (2021): 2195–217, https://onlinelibrary.wiley.com/doi/full/10.1002/smj.3286; Webb, "The Impact of Artificial Intelligence on the Labor Market," 54–56; and Ali Zarifhonarvar, "Economics of ChatGPT: A Labor Market View on the Occupational Impact of Artificial Intelligence," *Journal of Electronic Business & Digital Economics* 3, no. 2 (2024): 100–16, https://www.emerald.com/insight/content/doi/10.1108/jebde-10-2023-0021/full/html. For another approach matching subject-verb pairings in patent to online job postings, see Giacomo Damioli et al., "AI Technologies and Employment: Micro Evidence from the Supply Side," *Applied Economics Letters* 30, no. 6 (2023): 816–21, https://www.tandfonline.com/doi/full/10.1080/13504851.2021.2024129.

6. Office of New Jersey Governor Phil Murphy, "ICYMI: Governor Murphy, Labor Commissioner Asaro-Angelo & Chief Innovation Officer Cole Announce Launch of My Career NJ," press release, March 14, 2024, https://nj.gov/governor/news/news/562024/approved/20240314a.shtml.

7. While New Jersey has a license to the underlying science, other states interested in replicating this application must arrange licensing directly with RIPL.

Preparing the Workforce for Generative AI

Insights and Implications

Karin Kimbrough and Mar Carpanelli

Throughout history, technological advancements such as robotic assembly lines, tax preparation software, and—more recently—AI technologies like chatbots and self-driving cars have sparked concerns about the automation of human labor. Despite these concerns, demand for labor has remained steadfast, defying doomsday predictions. In general, the creation of new jobs and the evolution of old jobs to incorporate new skills and technologies overwhelmingly offset the costs of technological shifts. For example, the introduction of ATMs in the late sixties was feared to displace bank tellers; since then, the number of bank tellers in the US more than doubled, with their skills changing tremendously from clerking to sales and customer service.[1] Moreover, it created new related roles—such as banking relationship manager, which grew by 7 percent in the US in the past year alone, according to our data.

More generally, AI and generative AI (GAI) innovations could usher in a new era of labor productivity and economic growth. The gains from productivity could help address the puzzling pre-COVID period of elusive productivity growth.[2] It could also help address other longer-term trends in many advanced economies, such as demographic challenges from an aging workforce or labor shortages.

Regardless of technological innovation, whether employment in a given occupation expands or contracts is ultimately a function of supply and demand. On the supply side, emerging technologies often enable workers to do their jobs more efficiently. Think of language translators, whose jobs are already being redefined by automation. Translators have increasingly relied on software tools to create draft translations, which they refine for nuance and quality. AI-powered tools enhance these drafts, reducing the need for extensive human intervention. On the demand side, productivity (and therefore price) changes have moved demand toward both directions. Following the translator example, while bulk-translation jobs may decrease, high-end projects like literature may increase opportunities for human translators thanks to human-AI collaboration's cost-effectiveness.

Although it is too early to know GAI's net impact on the labor market or precisely measure productivity gains in specific occupations, it is certain that this technological change will lead to a shift in many of the skills that define most jobs. In this chapter we assert that GAI will have a substantial impact on a broad swathe of the workforce, evidenced by an evolution in the occupational skills that workers bring and employers seek. We anticipate that this skill evolution will generally elevate the focus on skills in the workplace and specifically heighten the importance of people skills.

Using LinkedIn's Economic Graph, a unique digital representation of the global labor economy capturing over 950 million professionals, we identify the skills that stand to be affected by and those that will likely complement new technologies. Workers and companies can use insights from this analysis to be more strategic about planning and acquiring new skills, reskill and upskill on AI-powered technologies, and build the people skills necessary to stay ahead. This is an advantage not available to workers and employers in past technological revolutions.

Most Jobs Can Leverage GAI

GAI is a new technology that uses AI models to create new content—resembling some people skills. We identify jobs, companies, and industries where GAI is likely to have the most significant impact using our own skill-based framework. It highlights where companies and workers can benefit most from developing complementary skills to remain productive, agile, and competitive in the face of technological change.

A Skills-Based Framework to Understand GAI's Impact on the Workforce

Despite uncertainty on the full spectrum of GAI capabilities, there is value in building a conceptual framework to explain how these technologies can affect skills—and consequently jobs. In this framework, we identify skills that can likely be carried out by GAI technologies ("GAI-replicable skills") and skills that intrinsically rely on human proficiency and can likely complement these technologies ("GAI-complementary skills"). We then rank every occupation in LinkedIn's taxonomy according to how likely it is to require these types of skills among its core 100 skills.

GAI-Replicable Skills

There are over 500 skills most likely to be replicable by GAI among LinkedIn's taxonomy of approximately 38,000 skills. The skills LinkedIn members most frequently added are the following:

- Communication and media skills including writing, editing, documentation, translation, video, photography, music, and content creation;
- Business and industry skills including financial reporting, email marketing, and data analysis;
- Engineering skills including software development, programming, and data science; and
- People skills such as knowledge of time management tools.

GAI-Complementary Skills

There are over 800 skills that—currently—can only be performed by people and that typically serve as complements to GAI technologies. The skills LinkedIn members most frequently added are the following:

- Communication and media skills including oral presentations and influencing;
- Business and industry skills including entrepreneurship, maintenance and repair, and military strategy;
- Engineering skills including software innovation and product innovation; and
- People skills including leadership, teamwork, negotiation, problem-solving, people

Figure 1. Anticipated Job Skill Changes, 2015–2030

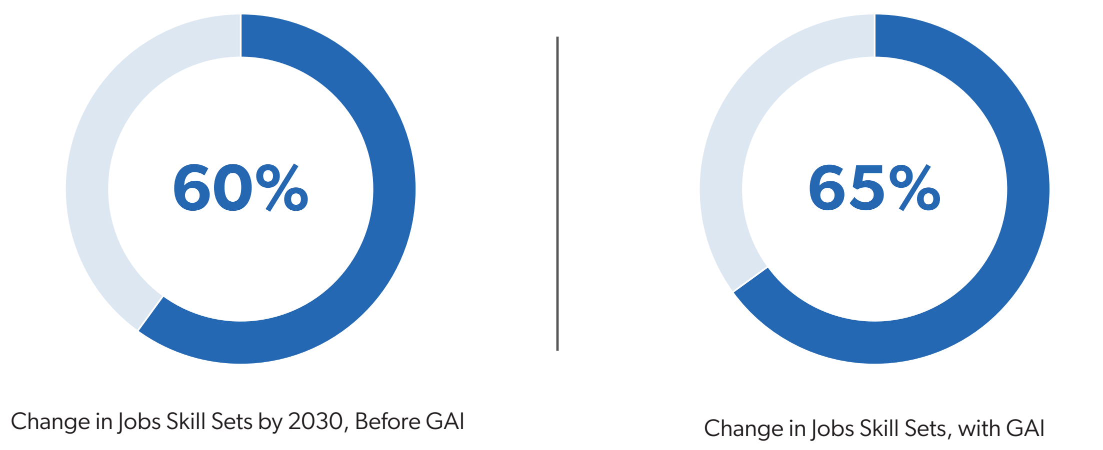

Source: Authors' own data.

management, relationship building, creativity, and emotional intelligence.

This framework was designed to be generalizable and dynamic. We thus expect results to change over time, following the advances of AI and GAI and changes in occupations and their skills beyond AI and GAI.

Findings

GAI Is Accelerating the Shift to a Skills-Based Labor Market

LinkedIn's *Skills-First: Reimagining the Labor Market and Breaking Down Barriers* report previously found that members have seen the skills related to their jobs change by 25 percent since 2015.[3] At that rate, workers could expect their jobs to change by nearly 60 percent between 2015 and 2030. However, with new GAI technologies and tools emerging every day, we now forecast the pace and scale of change to jobs to accelerate by an additional 5 percentage points to reach at least 65 percent by 2030 (Figure 1).

This underscores how GAI is both transforming the core skills required for many jobs and accelerating the pace of transition in the workforce. Implicit in the acceleration is the idea that productivity gains will be realized as workers rotate away from certain tasks and redeploy toward more productive, challenging areas. It also implies some degree of workforce disruption over time, consistent with other technological innovations whereby jobs lost bring associated challenges while new jobs are created. By taking a skills-based approach to understanding tasks and occupations across the economy, our work suggests that fundamental changes to the way we work and allocate our time across tasks will materialize before we see changes in occupations and titles.

Most Jobs Require Skills That Can Potentially Be Performed by GAI Technologies, but Not Every Job Will Be Affected in the Same Way

GAI may have a broad impact on the workforce. By analyzing occupations' core skill compositions, we find that 84 percent of US LinkedIn members are in occupations that could have at least one-quarter of their core skills affected by GAI technologies (Figure 2).

We expect jobs to change by incorporating new technologies that will reduce the time spent on

Figure 2. US Workforce Exposure to GAI

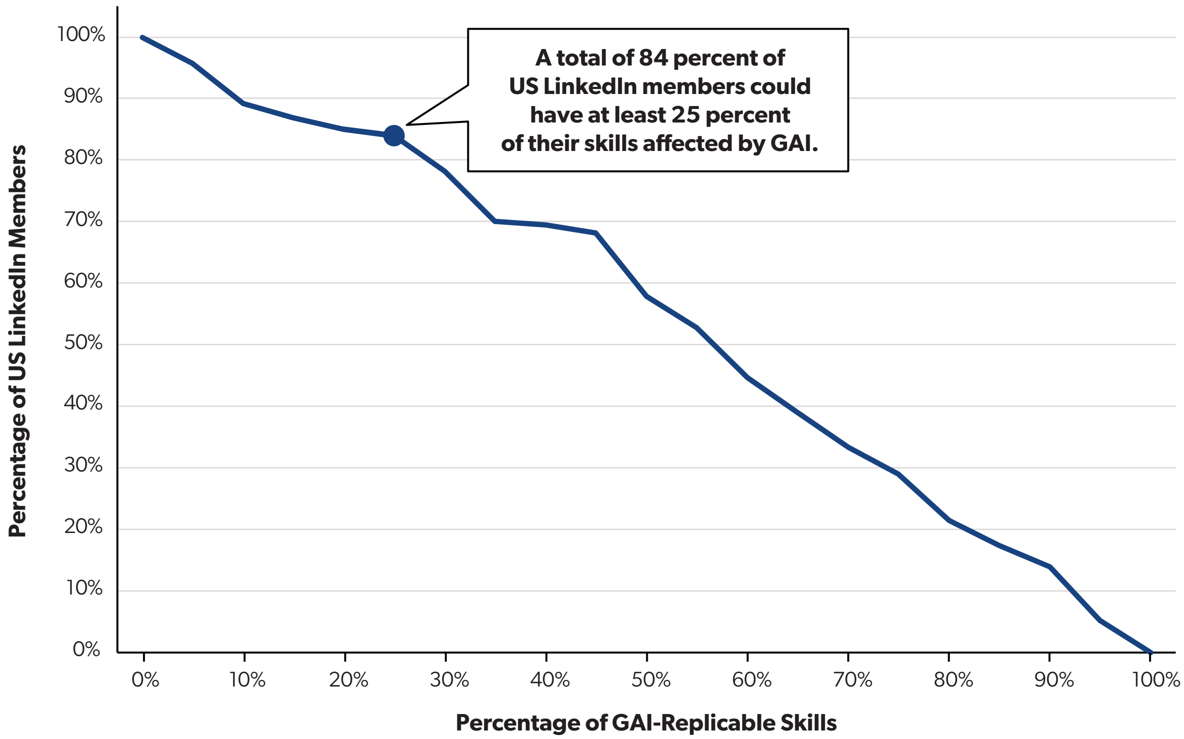

Source: Authors' own data.

applying some skills associated with routine tasks and make other skills much more important. For example, saving time writing or analyzing data unlocks more time for other meaningful work like creative thinking, problem-solving, and team leading. Notably, no occupation is fully composed of skills that stand to be affected by GAI, which is why we do not expect any occupation to disappear in the near term.

LinkedIn's skills-based framework allows us to classify each occupation by the percentage of core skills potentially replicable by GAI and the share of core skills complementary to GAI. This categorization results in three groups of occupations (Table 1).

Based on this framework, Figure 3 categorizes over 600 occupations according to how likely their core skills are to be similar to what GAI can currently perform (replicable by GAI) and aligned with intrinsic human skills that complement GAI (complementary to GAI).

Our analysis suggests that most jobs require skills GAI technologies can perform, but not every job will be affected in the same way. Each dot in Figure 3 represents an occupation. (Some have been labeled for illustrative purposes.) Each occupation has been decomposed into its 100 core skills. Those skills were mapped to our sets of GAI-replicable and GAI-complementary skills, yielding the percentage of core skills in each occupation that can be considered GAI replicable and GAI complementary. We later normalize these values to yield each occupation's relative stance across these dimensions.

This process identifies three relative areas for occupations: (1) jobs likely to be augmented by GAI because their core skills include a large share of both of GAI-replicable and GAI-complementary skills, (2) jobs that could be disrupted by GAI as their core skills include a large share of GAI-replicable but a relatively low share of GAI-complementary skills, and

Table 1. Three Groups of Occupations

Group	Impact on Occupations	Example
Augmented by GAI		
These jobs' core skills include a large share of both GAI-replicable and GAI-complementary skills.	GAI may affect a relatively large portion of the skills in these jobs, leaving more time for higher value-added complementary skills	Data analysts automate the computation and interpretation of metrics with GAI, enabling them to focus their time on GAI-complementary skills, such as cross-functional influencing and stakeholders' engagement.
Disrupted by GAI		
These jobs' skills include a large share of GAI-replicable and a relatively low share of GAI-complementary skills.	As GAI is adopted more broadly, these jobs will undergo reskilling, possibly leading to more innovation.	Language translators' skills shift from doing translations from scratch to reviewing and certifying machine-generated translations or to specializing on specific legal or literary domains.
Insulated from GAI		
These jobs have a relatively small proportion of GAI-replicable skills in their core skills.	As these jobs are relatively protected from the influence of GAI, their core skills are likely to remain unchanged in the near term. Some of these jobs tend to be susceptible to other forms of automation, such as robotics.	Real estate agents might use GAI for writing house descriptions, but core relationship management skills would be insulated from GAI.

Source: Authors' own assessment.

(3) jobs that may be insulated from GAI because they have a relatively small proportion of GAI-replicable skills in their core skills. Since the data are normalized, these classifications are relative. In this analysis, occupations in either the augmented or disrupted sections will be referred to as relatively exposed to GAI.

This framework tells us, for example, that GAI might augment data analysts' work in the near term for tasks such as reporting. Moreover, data analysts may be able to lean more into their complementary skills, which include influencing and stakeholder management.

GAI's Most Immediate Impact Is Concentrated Among Higher-Paid and Better-Educated Workers

The advent of GAI will likely have unequal effects on workers, determined by their skill level. Unlike most previous technological advancements, which primarily affected workers in lower-paid roles or jobs with lower education requirements, the GAI technological wave will likely affect some of the highest-skilled and highest-paid jobs most (Figure 4). According to our research, some of the most affected occupations involve a high degree of skills and expertise. Examples include finance officer, mathematician, and web designer.

Women and Younger Workers Are Likely to Be Disproportionately Affected by GAI

GAI technologies' potential impacts are likely to be experienced to different degrees across the workforce. When we evaluate the distributional implications, we find that women and younger workers are likely to be disproportionately affected. This can be in part attributable to occupational segregation—differences in the likelihood of men and women to work in certain occupations. As a result, jobs vary in terms of their

Figure 3. Occupational Composition by GAI-Replicable and GAI-Complementary Skills

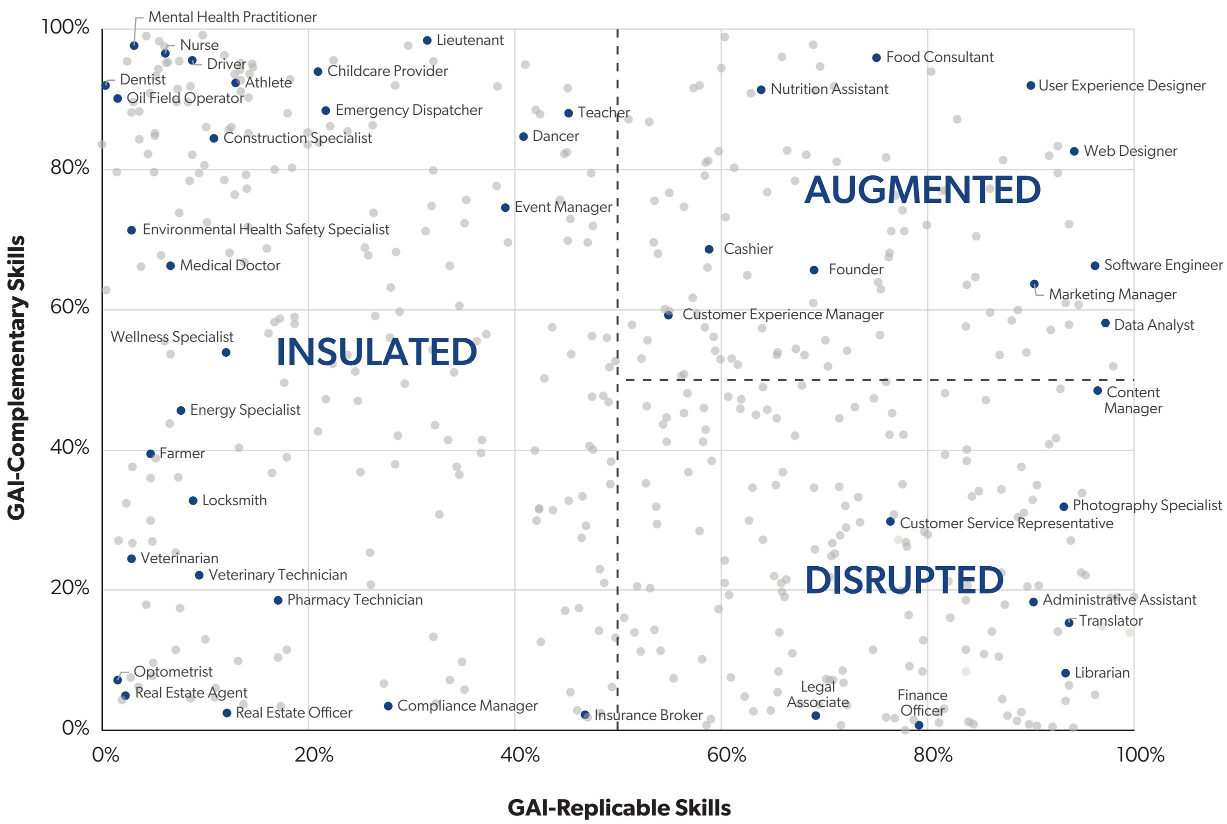

Source: Authors' own data.

gender and age talent composition. These differences are reflected in how these groups in aggregate may be affected by GAI.

Joint research by LinkedIn and the World Economic Forum indicates that multiple gaps remain between men and women in the labor market, such as in holding leadership positions or work participation in STEM fields.[4] Our analysis suggests there may be yet another gender gap related to GAI exposure: Women are underrepresented in jobs that are likely to be insulated or augmented by GAI and overrepresented (57 percent of US members) in jobs that will likely be disrupted by these new technologies (Figure 5).

Examples of occupations where women are overly represented that could be disrupted include medical administrative assistant (91 percent female), office manager (88 percent), and legal assistant (87 percent). On the other hand, examples of occupations where men are overly represented that could be augmented are electrical engineer (94 percent male), mechanical engineer (89 percent), and computer network technician (88 percent).

Our research indicates that all generations face some exposure to GAI. However, younger generations face slightly higher exposure, especially in occupations that can be either augmented or disrupted by GAI (Figure 6). This is not surprising. Those starting careers tend to be relatively better represented in roles requiring GAI-replicable skills, including writing and analytics, while they develop the people skills that can complement technology, such as

Figure 4. GAI's Relative Impact by Income and Education

9%

US LinkedIn members in jobs in the top quartile of GAI-replicable skills live in communities with median incomes 9 percent higher than those in the bottom quartile of GAI-replicable skills.*

2X

US LinkedIn members in jobs in the top quartile of GAI-replicable skills are 2 times more likely to hold a bachelor's degree or higher than those in the bottom quartile of GAI-replicable skills.

Source: Authors' own data.
Note: * Community income is estimated from US Census Bureau, "Selected Characteristics of Health Insurance Coverage in the United States," 2023, https://data.census.gov/table/ACSST5Y2023.S2701?g=860XX00US90804.

leadership and negotiation. Examples of jobs predominantly held by Gen Z that could be augmented by GAI include graphic design assistant (49 percent of US members are Gen Z), academic tutor (46 percent), and marketing assistant (42 percent). Among jobs held by Gen Z that could be disrupted by GAI are clinical research assistant (57 percent of US members are Gen Z), industrial design specialist (47 percent), and library science specialist (46 percent). In their favor, younger workers who are earlier in their careers can see greater return on investment in upskilling to adapt to the demands of technological change, whether it is from GAI or future innovations.

By contrast, baby boomers are relatively overrepresented in occupations that could be insulated from exposure to GAI. With more career experience, these workers are more likely to be in more senior roles that predominantly call on people skills, like leadership and management. Examples of jobs typically held by boomers that could be insulated from GAI include board member (33 percent of US members are boomers), managing partner (21 percent), and supply-chain or environmental consultant (17 percent).

Certain Jobs Are More Likely to Be Relatively Insulated from GAI

Despite GAI's potential to have wide-ranging impacts on the workforce, some industries and occupations are likely to be relatively insulated in the near term. Jobs that rely on license mandates, government regulations, physical skills, ethical considerations, or empathic engagement all demand such a high degree of essential human functions that GAI is not likely to immediately transform these roles.

Compared to the Average Job, Roles That Mandate Licenses or Are Government Regulated May Be 19 Percent Less Exposed to GAI

Jobs like emergency dispatcher, real estate agent, and veterinarian tend to be relatively shielded from GAI due to their reliance on people skills.[5] These jobs have critical requirements of public safety and professional accountability, requiring skills that only humans possess, such as moral judgment, ethical decision-making, and critical thinking. Furthermore, these roles, predominantly in health care and customer-facing industries, involve substantial human interaction and relationship building.

Figure 5. Gender Composition by GAI Segment

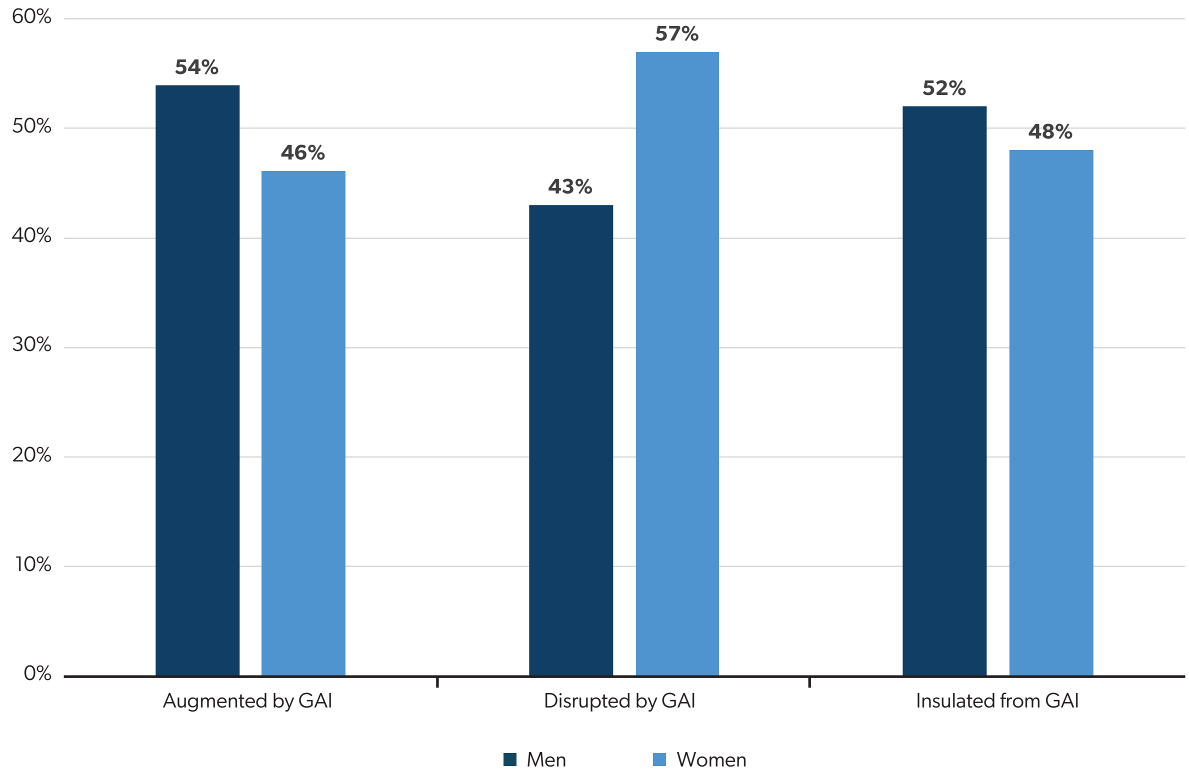

Source: Authors' own data.

Jobs That Rely on Physical Skills to Some Extent May Be 10 Percent Less Exposed to GAI Compared to the Average Job

These roles require extensive human interaction, dexterity, and creativity skills that GAI is not likely to supplant. For example, childcare specialists provide nurture and care at an irreplaceable level of human connection and empathy, landscapers work on-site with a high degree of customization and creativity, and dancers perform tasks that demand artistic expression and physical execution.

Jobs Requiring Green Skills May Be 7 Percent Less Exposed to GAI, Compared to the Average Job

As covered in LinkedIn's *Global Green Skills Report*, occupations such as environmental health safety specialist, construction specialist, and farmer frequently involve specialized knowledge and the on-site application of skills that rely heavily on human judgment, critical thinking, and consideration of complex environmental trade-offs—skills GAI would struggle to replicate.[6] Moreover, the green sector's focus on sustainability, energy efficiency, and ecological impact necessitates a high degree of regulatory compliance, adherence to safety protocols, and a deep understanding of environmental regulations. However, beyond GAI, new AI applications may be capable of helping combat climate change, potentially augmenting many green jobs in the future.[7]

GAI's Potential Impact Varies Across Industries and Goes Beyond the Technology Sector

Throughout history, new technologies have affected sectors and industries differently, and GAI is no exception. Our skills-based framework allows us to

Figure 6. Generation Distribution by GAI Segment

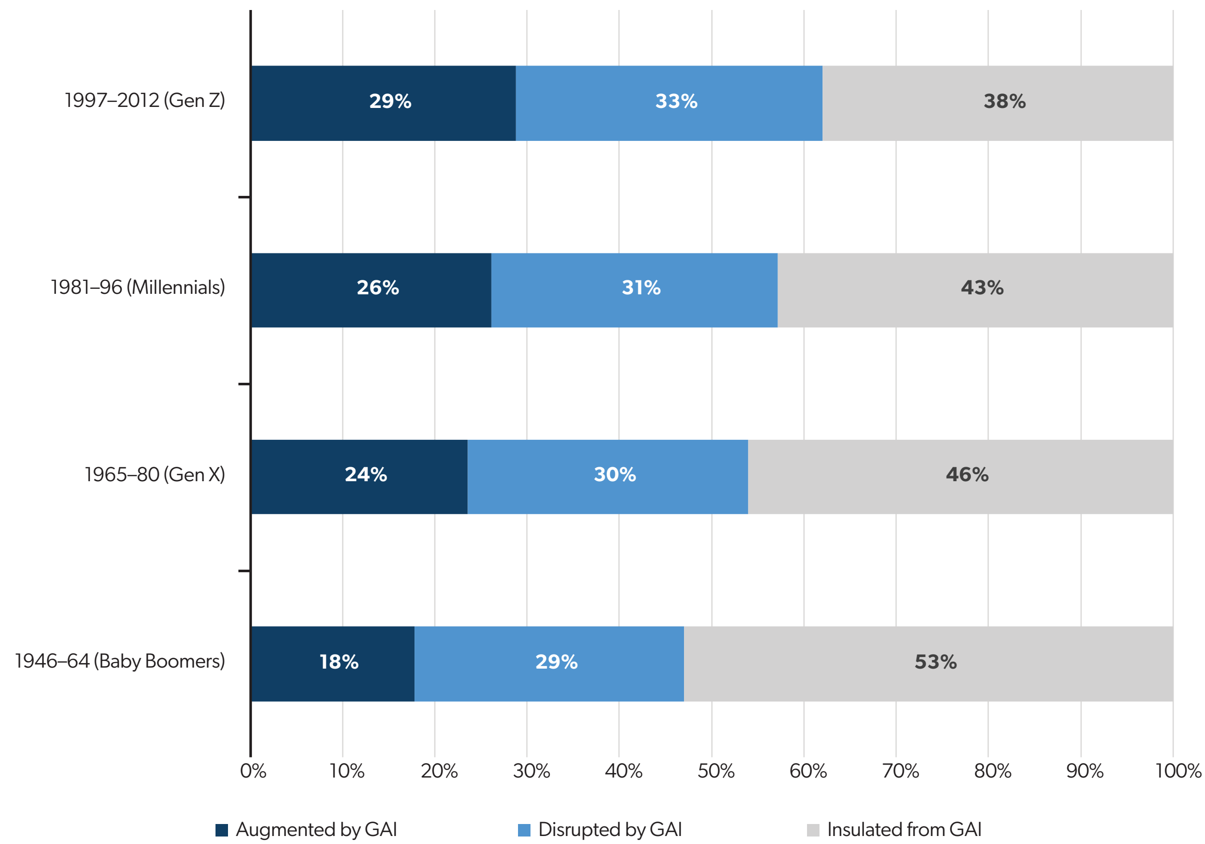

Source: Authors' own data.

describe industry-level aggregate exposure to GAI-replicable and GAI-complementary skills (Figure 7).

We estimate that the industries with the largest concentration of workers anticipated to be augmented or disrupted by GAI—where occupations have a relatively high exposure to GAI-replicable skills—are technology, information, and media (41 percent of US workers are poised to be augmented by GAI, and 36 percent stand to be disrupted by GAI); accommodation and food services (54 percent and 18 percent, respectively); wholesale (26 percent and 46 percent, respectively); financial services (21 percent and 50 percent, respectively); and professional services (31 percent and 38 percent, respectively). Notably, at 33 percent, the education sector ranks third in share of US members that could be augmented by GAI.

Over time, technological inflection points such as the one stemming from GAI often result in widespread adoption and impact across entire economies. While it is not surprising that the technology, information, and media industry leads in terms of exposure to GAI-replicable skills—since most AI and GAI innovations come from this sector—it is remarkable that the impact of GAI on the workforce extends across workers in all industries.[8] Moreover, if the percentages in Figure 7 change over time, it may be informative regarding the potential speed of adoption as GAI transitions from early adopters to mainstream users. New jobs directly related to emerging technologies will also arise, but most of the impact will likely be on transforming existing roles, which will persist but evolve.

Figure 7. Industry Composition by GAI Segment

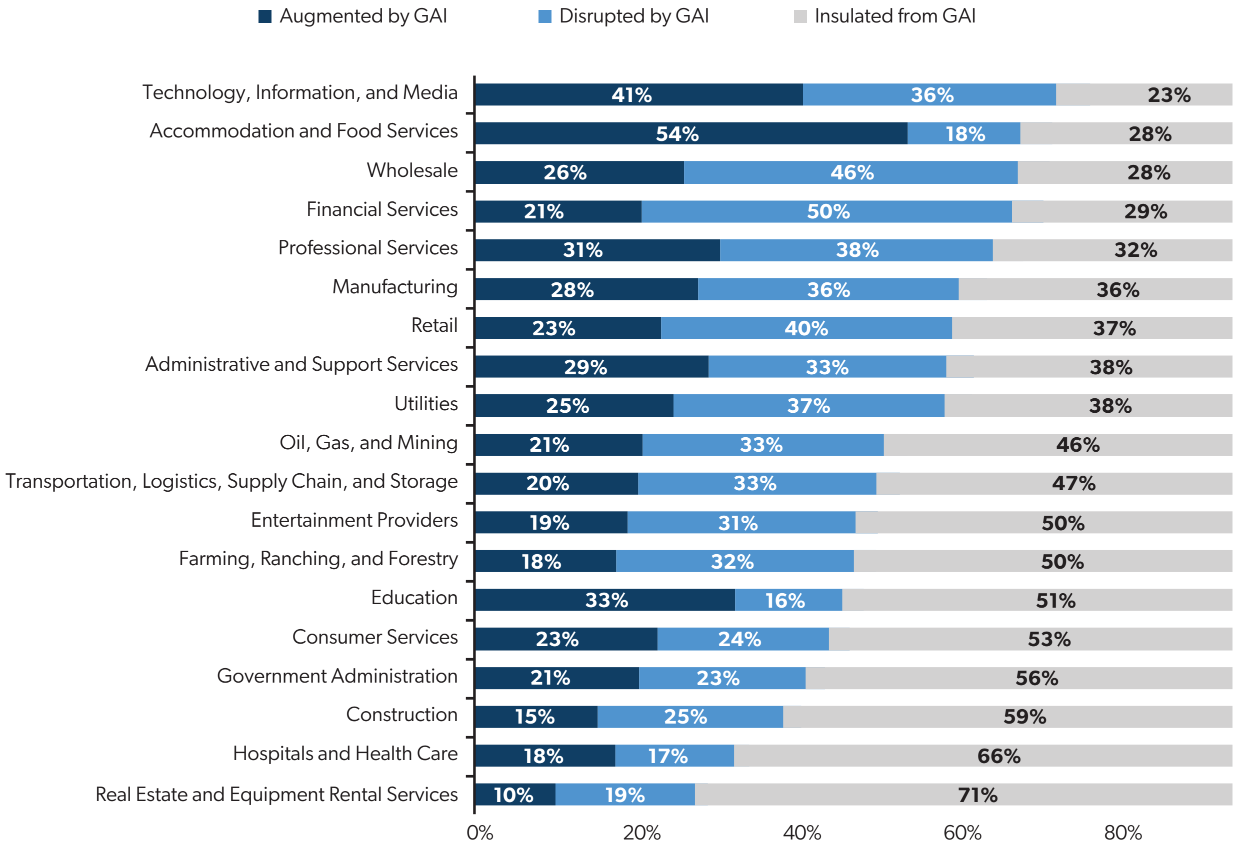

Source: Authors' own data.

In the Future of Work, the Conversation Will Be About Skills

While GAI's long-term impact on the labor market will depend on how much and how quickly the market adopts GAI-augmented technologies, we expect that jobs will fundamentally change. In the short term, we expect jobs to change by incorporating these new technologies, which will reduce the time spent applying some skills while making others much more important.

This is why the conversation about GAI's impact on jobs should focus on building the right skills to adopt and complement GAI technologies and not on job displacement. Leaders' and workers' responses to GAI, not exposure to GAI, is more likely to shape the future of the workforce.

Implications for Leaders

A Skills-Based Approach to Strategic Workforce Planning

Leaders should comprehensively understand GAI's impact on their organizations and the economy. New jobs will emerge as a direct result of this emerging technology, but the impact will be felt mostly in existing roles, which will be transformed by GAI. By acknowledging the skills GAI will most affect and those that complement GAI across roles and functions, leaders can optimize their hiring strategies and make informed decisions on the reskilling and upskilling programs needed to remain competitive in this changing environment.

In this context, it is critical to recognize that using AI in hiring can help tackle labor market inequalities through a skills-based recruitment approach. The labor market's oversight of exceptional candidates presents missed opportunities for companies, the economy, and society. Emphasizing a skills-based hiring approach not only invites more participants into the workforce but also levels the playing field for underrepresented groups. As global markets change, a skills-based approach can direct policymakers to enhance workforce initiatives, allowing businesses to access a broader, more diverse talent spectrum and promote inclusivity in hiring.

Productivity Gains via Collaborative Human-AI Workflows

Leaders should recognize GAI technology's potential to augment human capabilities. Building a culture that promotes collaborative human-AI workflows and adopts responsible guidelines for AI use can unlock new avenues for productivity and innovation. Empowering employees to leverage AI technologies via increased AI literacy may be crucial for success in a digitally driven economy and ensure that AI and GAI enhance human expertise rather than supplant it.

The rise of automation and AI has emphasized the growing importance of people skills, which have always been vital for job success. As hard skills' half-life diminishes, the significance of people skills, such as creativity and leadership, continues to grow. With AI technologies increasingly automating tasks formerly performed by humans, there's a shift toward collaborative intelligence whereby humans and intelligent machines work together to achieve superior outcomes. This evolution requires companies to shift their hiring focus to essential people skills and employees to hone new capabilities, including engaging in effective AI communication, employing data interpretation, and spotting optimization opportunities.

Promotion of an Equitable Distribution of AI Benefits

Leaders face an opportunity to address GAI's disparate impacts on different workforce segments, such as its potentially disproportionate effect on women and the youth. Collaborative efforts are needed to ensure an equitable distribution of AI benefits and the mitigation of socially undesirable effects. This entails prioritizing transparent research, combating systemic biases, and advocating inclusive AI education. As technology, especially AI, becomes more integrated into our daily lives, a new baseline of AI literacy will emerge—distinct from traditional digital literacy, which focuses on using and understanding digital technologies. AI literacy should encompass understanding AI fundamentals, recognizing ethical implications like privacy and bias, becoming familiar with AI-driven tools in the workplace, and grasping AI's broader socioeconomic implications.

With the accelerating deployment of generative AI, businesses must monitor labor market trends to remain competitive. Meanwhile, regulators and governments should support workforce development in digital skills, and education systems must evolve to produce an AI-ready workforce. By fostering global collaboration across research bodies, governments, civil society, and the private sector, leaders can navigate AI's complexities and lay the foundation for a more just and inclusive future.

Appendix A

GAI-Replicable and GAI-Complementary Skills

We identify GAI-replicable and GAI-complementary skills with the following steps:

1. We give ChatGPT 3.5 (in February 2023) the following prompts:

    - "What are the 100 top skills that AI technologies (ChatGPT, Dall-E, LaMDA, etc.) can perform very well?"

    - "What are the 100 top skills that can currently exclusively be performed by humans?"

2. We map these lists to LinkedIn's taxonomy with LinkedIn's taxonomy application programming interface, and we refined matches manually. We expand coverage further by applying skill similarities based on skill embeddings to score skills that are similar to those flagged in each list and by manually reviewing the skills in the popular skill groups containing the skills from the previous steps.

3. For external validation, we ingest and map to our taxonomy three exposure scores from the academic literature.[9] We use these scores to train a model that learns which skills contribute more to these three rankings, and we use this model to score all skills in LinkedIn's taxonomy.

Occupations Exposed to GAI and Complementary Skills

To calculate the percentage of skills that are exposed to GAI by occupation, we use each occupation's skills genome.[10] An occupation's skills genome is the ranking of its top 30 most relevant skills based on a term frequency—inverse document frequency model. In this model, skills are relevant when they tend to be disproportionately added by members in this occupation compared to other occupations.

The thresholds for classifying occupations into high and low exposure to GAI and GAI-complementary skills are based on the metrics' medians.

Segments Exposed to GAI and Complementary Skills

Based on the classification of occupations by GAI-complementary exposure, we compute the share of LinkedIn members in each category as a share of all members in that segment, gender, generation group, and so on. We report these shares and run linear regressions to compare GAI exposure with dimensions of interest, such as skill type, education, and experience.

Notes

1. David Autor, "Will Automation Take Away All Our Jobs?," TED, March 29, 2017, https://ideas.ted.com/will-automation-take-away-all-our-jobs/.
2. Robert J. Gordon and Hassan Sayed, "Productivity Growth Before and During the Pandemic," National Bureau of Economic Research, October 2022, https://www.nber.org/system/files/working_papers/w30267/w30267.pdf.
3. LinkedIn Economic Graph, *Skills-First: Reimagining the Labor Market and Breaking Down Barriers*, 2023, https://economicgraph.linkedin.com/content/dam/me/economicgraph/en-us/PDF/skills-first-report-2023.pdf.
4. World Economic Forum, *Global Gender Gap Report 2023*, June 2023, https://www3.weforum.org/docs/WEF_GGGR_2023.pdf. While we acknowledge that gender is a spectrum, due to data limitations, we restrict our analysis to the binary classification of men and women.
5. Classification based on National Conference of State Legislatures, "National Occupational Licensing Database," 2022, https://app.powerbi.com/view?r=eyJrIjoiNzA0MGM4NGMtN2E5Ny00NDU3LWJiNzMtOWFlMGIyMDczZjg2IiwidCI6IjM4MmZiOGIwLTRkYzMtNDEwNy04MGJkLTM1OTViMjQzMmZhZSIsImMiOjZ9&pageName=ReportSection.
6. LinkedIn, Economic Graph, *Global Green Skills Report 2023*, 2023, https://economicgraph.linkedin.com/research/global-green-skills-report.
7. Megan Mastrola, "How AI Could Help Combat Climate Change," *The Hub*, March 7, 2023, https://hub.jhu.edu/2023/03/07/artificial-intelligence-combat-climate-change/.
8. Stanford University, Human-Centered Artificial Intelligence, *Artificial Intelligence Index Report 2024*, 2023, https://aiindex.stanford.edu/report/.
9. Edward W. Felten et al., *How Will Language Modelers like ChatGPT Affect Occupations and Industries?,* SSRN, March 6, 2023, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4375268; Edward W. Felten et al., "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal* 42, no. 12 (2021): 2195–217, https://onlinelibrary.wiley.com/doi/full/10.1002/smj.3286; and Michael Webb, *The Impact of Artificial Intelligence on the Labor Market*, SSRN, November 15, 2019, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3482150.
10. Zhichun Jenny Ying, "How We Mapped the Skills Genome of Emerging Jobs," *Engineering Blog*, December 10, 2019, https://www.linkedin.com/blog/engineering/skills-graph/how-we-mapped-the-skills-genome-of-emerging-jobs.

Measuring and Accounting for Artificial Intelligence in the US Bureau of Economic Analysis's National Economic Accounts

Vipin Arora and David Wasshausen

As with most challenging economic measurement issues, defining exactly what we mean by artificial intelligence is top of mind. *Encyclopaedia Britannica* defines AI as "the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings," while Stanford University's Institute for Human-Centered AI defines AI as "the science and engineering of making intelligent machines."[1] Both definitions highlight the challenge that national accountants face when trying to answer questions like "How is AI measured in gross domestic product?" or "What is the value of AI in gross domestic product?"

In this chapter, we first share how we think about and define AI. Next, we present an example of an alternative method we use to estimate hard-to-measure activities—specifically for intangible assets. This example aims to provide insights and promote further ideas about alternative ways to measure critical and emerging technologies. We then focus on a handful of case studies that illustrate how both the production and use of AI are reflected in the National Economic Accounts' statistics produced by the US Bureau of Economic Analysis (BEA).

Analyzing the impacts of AI-related transactions through various case studies not only helps convey how economic accountants mechanically view AI's role in the context of measuring gross domestic product (GDP) but also serves to identify how our view of AI may be incomplete and how it differs from others. Presenting the framework in which economic accountants view this challenge while highlighting conceptual and empirical gaps sets the stage for researchers working to improve the measurement of not only AI but other critical and emerging technologies.

What Is AI, and How Do We Identify It?

First, AI is not entirely new, although its production and use seemingly exploded in recent years. Computers could long perform complicated tasks traditionally associated with humans. For example, Deep Blue, the chess-playing IBM supercomputer, defeated world champion Garry Kasparov in a six-game match in 1997. Computer video games have long had a built-in "computer" opponent that challenges gamers. There have been, however, substantial advances in AI that are likely having a profound economic impact across many sectors and industries.

We do not view AI itself as an asset type, although the pursuit and application of AI each spur capital formation. For example, research and development (R&D) is needed to create the complex algorithms that underlie AI. R&D expenses are treated as capital expenditures in BEA's National Economic Accounts. Moreover, continued R&D in semiconductor manufacturing enables the production of high throughput chips that are required to execute modern AI algorithms.

That R&D, along with advanced machinery required to produce those chips, is considered a capital expenditure—as is the production of software applications that embed AI. Finally, AI's production and use are reflected in intermediate inputs and captured directly in value added for the producing industries. We sought to provide examples in the case study analysis section of this chapter to illustrate the many ways that BEA's economic statistics capture AI-related activity.

Identifying new, emerging products in a timely fashion is a long-standing challenge in the federal statistical system. For example, the GDP benchmarking processes have notable lags, which is our primary opportunity for introducing emergent technology products. As another example, we released the results of our 16th comprehensive update in September 2023, with GDP benchmarked to the 2017 Economic Census.[2] Moreover, emerging technologies by nature are rapidly changing, making it difficult to define new product classes and modify existing surveys in a way that allows businesses to report accurate and consistent economic statistics (such as receipts, expenses, capital expenditures, and prices charged).

Alternative Methods Deployed for Difficult-to-Measure Activities

BEA faces a number of measurement challenges and has a proven track record of developing innovative methods that produce comprehensive measures of GDP and related statistics. For example, in the 1999 comprehensive update, BEA introduced software as a capital expenditure. Purchased software was reclassified from an intermediate expense to a capital expenditure.

Moreover, BEA introduced measures of own-account software, which captured the value of in-house production of software for in-house use (i.e., not to be sold). Because no market transactions are associated with this production, alternative methods had to be developed. In this case, BEA used detailed occupational employment data from the Bureau of Labor Statistics to derive measures of own-account software production.[3]

More recently, BEA has developed experimental measures of data as an asset, building on the methodology used to estimate own-account software. Machine learning models are used to identify occupations associated with producing data and specific time-use factors associated with the types of data-producing activities that we consider capital in these experimental statistics.[4]

Case Studies

Artificial intelligence is produced and used by businesses, the government, and nonprofit institutions serving households. In this section, we begin with three examples illustrating how BEA's GDP statistics capture the production of AI, and we conclude with an example of how GDP could reflect AI use. We use "supply and use" tables to illustrate these economic transactions and orient readers who are unfamiliar with these tables.[5] We do not attempt to capture second-order

Table 1. Use Table Illustrating Production of AI That Is Sold

	Intermediate Purchases			Final Expenditures		
Commodities	Computer System Design Industries	All Other Industries	Sum	Private Fixed Investment	Total (GDP)	Commodity Output
Software	0	0	0	100	100	100
All Other Commodities	20	0	20	0	0	20
Total Intermediate	20	0	20	100	100	120
Employee Compensation	50	10	60	0	0	0
Gross Operating Surplus	30	10	40	0	0	0
Value Added	80	20	100	0	0	0
Industry Output	100	20	120	0	0	120

Source: Authors.

effects related to AI adoption, including changes in productivity, shifts in employment, and so on.

Case 1: Production of AI That Is Sold

Supply and use tables enable one to analyze how individual economic transactions are reflected in all three measures of GDP (production, income, and final expenditures) in a fully integrated, consistent framework. In Table 1, the columns under "Intermediate Purchases" show the derivation of GDP by industry using the production approach, specifically purchases of intermediate inputs by the using industry. Total intermediate output is taken from the supply table (not shown), and the value added is derived residually:

$$\text{Value Added} = \text{Gross Output} - \text{Intermediate Inputs.}$$

We can also derive GDP or value added as the sum of the incomes associated with production; those values are presented in the "Employee Compensation" and "Gross Operating Surplus" rows of Table 1. The data in the "Final Expenditures" columns also yield a measure of GDP. By definition, total commodity output (120) equals total industry output (120).

In our first case study, shown in Table 1, a telecommunications company pays a computer systems design company $100 to develop and implement a new AI-enabled customer service software application.

In this first example, the computer system design company purchased $20 of intermediate inputs (rent, electricity, maintenance and repair, etc.) to produce the AI-enabled custom software application that sold for $100. The value added for computer system design companies is $80, the value added from other industries that supplied the $20 of intermediate input is $20, and the sum of the value added is $100. The purchase of the $100 AI-enabled custom software application is recorded as a private fixed investment in software, contributing $100 to GDP. The value of AI is embedded in the software's value and not separately identifiable.

Case 2: Own-Account Production of AI

In this case study, shown in Table 2, the telecommunications company develops and implements a new AI-enabled customer service software application in-house using existing resources—no new intermediate inputs are required.

The telecommunications company will first engage in R&D to develop the AI required for its new AI-enabled customer service software application. The value of that AI-related R&D is $50, which is recorded as a private fixed investment in own-account R&D.[6]

Next, the telecommunications company builds the new customer service software application in-house, implementing the AI that was developed in its R&D. That software is valued at $50 and recorded

Table 2. Use Table Illustrating Own-Account Production of AI

	Intermediate Purchases		Final Expenditures		
Commodities	Telecommunications Industries	Sum	Private Fixed Investment	Total (GDP)	Commodity Output
Software	0	0	50	50	50
R&D	0	0	50	50	50
All Other	0	0	0	0	0
Total Intermediate	0	0	100	100	100
Employee Compensation	0	0	0	0	0
Gross Operating Surplus	100	100	0	0	0
Value Added	100	100	100	100	0
Industry Output	100	100	0	0	100

Source: Authors.

as a private fixed investment in own-account software. The own-account creation of a new asset is recorded as gross operating surplus for the creating industry. In this example, the gross operating surplus (and value added) for the telecommunications company is $100.[7]

Case 3: Production of AI That Is Sold and Own-Account Production of AI

This example, shown in Table 3, reflects a hybrid of case studies 1 and 2, in which the telecommunications company pays a computer system design company a $15 consulting fee to evaluate current customer service systems and develop recommendations for deploying an AI-enabled customer service software application. The computer system design company has previously invested in AI-related R&D and has the expertise to provide this consulting service.

Transactions associated with the telecommunications company paying the computer system design company $15 for AI-related consulting services are recorded in orange: a $15 purchase of intermediate inputs by the telecommunications company and $15 of value added that is attributed to the computer system design company ($10 in compensation paid and $5 of gross operating surplus). The telecommunications company's $15 intermediate input expense is reflected in the final expenditures for telecommunications services (e.g., household or government purchases). At this stage, there is no new capital formation.

Using the AI consulting service's recommendations, the telecommunications company builds and implements a new AI-enabled customer service software application. These transactions are shown in green: The private fixed investment in own-account software is recorded as $65, and the own-account creation of this new asset is recorded as $65 in gross operating surplus for the telecommunications industry.

Case 4: Using AI to Produce Software

In our final example, shown in Table 4, we examine a case in which a computer system design company uses an AI-based code generator to produce custom software that is then sold to a third party.

The computer system design company pays $25 for the rights to use the AI-based code generator. For this example, we assume the $25 licensing fee covers three months and treat it as an intermediate input as opposed to a capital expenditure. Transactions associated with this purchase are shown in orange: a $25 purchase of intermediate inputs by the computer system design company and $25 of value added that is attributed to the AI services provider ($20 in compensation paid and $5 of gross operating surplus). The computer system design company produces custom software applications that are sold

Table 3. Use Table Illustrating Production of AI That Is Sold and Own-Account Production of AI

	Intermediate Purchases			Final Expenditures			
Commodities	Telecommunications Industries	Computer System Design Industries	Sum	Private Fixed Investment	All Other	Total (GDP)	Commodity Output
Software	—	—	—	65	—	65	65
AI Consulting Services	15	—	15	—	—	—	15
Telecommunication Services	—	—	—	—	15	15	15
Total Intermediate	15	—	15	—	—	—	—
Employee Compensation	—	10	10	—	—	—	—
Gross Operating Surplus	65	5	70	—	—	—	—
Value Added	65	15	80	65	—	80	—
Industry Output	80	15	95	—	—	—	95

Source: Authors.

Table 4. Use Table Illustrating the Use of AI to Produce Software

	Intermediate Purchases			Final Expenditures		
Commodities	AI Services Provider Industries	Computer System Design Industries	Sum	Private Fixed Investment	Total (GDP)	Commodity Output
AI-Based Code Generator	—	25	25	—	—	25
Custom Software	—	—	—	100	100	100
Total Intermediate	—	25	25	—	—	—
Employee Compensation	20	55	75	—	—	—
Gross Operating Surplus	5	20	25	—	—	—
Value Added	25	75	100	—	—	—
Industry Output	25	100	125	—	—	125

Source: Authors.

for $100. These transactions are shown in green: $100 for private fixed investment in custom software and $75 in value added ($55 in compensation paid and $20 in gross operating surplus) that is attributed to the computer system designer that produced that $100 software.

In this example, it's tempting to value the generative AI embedded in the custom software at $25. However, note that this ignores spillover gains and likely increased productivity associated with using the generative AI programmer. Those gains should be reflected in the computer system designer's gross operating surplus but are not separately identifiable.

Conclusion

This chapter was originally written for the "New Approaches to Characterize Industries: AI as a Framework and a Use Case" workshop, which aimed to use AI as a use case for considering new approaches to characterize industries. With that goal in mind, we have sought to provide insights and identify challenges associated with measuring AI from a federal statistical agency's perspective.

As with most economic measurement challenges, we started with sharing our definition of AI and how we would approach identifying it. Next, we shared a salient example in which we deployed alternative, innovative techniques to estimate hard-to-measure economic activity. This chapter's focus is four case studies, in which we used supply and use tables to understand how BEA's GDP, gross domestic income, and gross value added statistics would capture each scenario. A prerequisite for identifying new approaches to characterize hard-to-measure industries is a comprehensive and accurate understanding of how those industries are currently measured.

Notes

1. *Encyclopaedia Britannica*, "artificial intelligence," accessed January 2, 2025, https://www.britannica.com/technology/artificial-intelligence; and Christopher Manning, "Artificial Intelligence Definitions," Stanford University Human-Centered Artificial Intelligence, September 2020, https://hai-production.s3.amazonaws.com/files/2020-09/AI-Definitions-HAI.pdf.

2. For more information, see Lisa S. Mataloni et al., "Improved Estimates of the National Economic Accounts: Results of the 2023 Comprehensive Update," *Survey of Current Business*, November 9, 2023, https://apps.bea.gov/scb/issues/2023/11-november/pdf/1123-nea-comprehensive-update.pdf.

3. For more information on the methodology, see Robert Parker and Bruce T. Grimm, *Recognition of Business and Government Expenditures for Software as Investment: Methodology and Quantitative Impacts, 1959–98*, US Department of Commerce, Bureau of Economic Analysis, May 2000, https://www.bea.gov/research/papers/2000/recognition-business-and-government-expenditures-software-investment.

4. For more information on this ongoing research, see José Bayoán Santiago Calderón and Dylan G. Rassier, "Valuing the U.S. Data Economy Using Machine Learning and Online Job Postings," Working Paper No. WP2022-13 (US Bureau of Economic Analysis, October 2022), https://www.bea.gov/research/papers/2022/valuing-us-data-economy-using-machine-learning-and-online-job-postings.

5. Though we implicitly used supply and use tables to create the data, we provide only the use tables in this chapter.

6. BEA's estimates of private fixed investment in R&D are based primarily on R&D expenditures data from three National Science Foundation (NSF) annual surveys: the Business Enterprise Research and Development Survey, the Annual Business Survey, and the Higher Education Research and Development Survey. (The first two surveys are jointly conducted with the US Census Bureau.) Beginning with 2019, NSF separately identifies AI-related R&D.

7. For a discussion on how the income associated with own-account investment is recorded, see Shelly Smith et al., *Preview of the 2013 Comprehensive Revision of the National Income and Product Accounts: Changes in Definitions and Presentations*, US Department of Commerce, Bureau of Economic Analysis, March 2013, 14–18, https://apps.bea.gov/scb/pdf/2013/03%20March/0313_nipa_comprehensive_revision_preview.pdf.

New Approaches to Characterizing Industries

The Perspective from Federal Statistical Agencies

Nancy Potok

The current federal statistical approach to measuring industries, products, jobs, and other aspects of the US economy is evolutionary rather than revolutionary. For several reasons, changes in measurement methodology and classification often take years to become part of official statistics. Yet there is significant pressure to innovate how federal statistics are produced and disseminated to ensure the agencies meet their mandate to produce relevant, timely, accurate, and objective data.

AI Measurement: The Current State of Play

The American Statistical Association and George Mason University have undertaken a research project to measure annually the health of the Federal Statistical System (FSS).[1] Timeliness and relevance of data releases and the ability to innovate are among the key metrics.[2]

Innovation in federal statistics requires investing agency resources, skilled staff, and sustained high-level attention from leadership. The ability to bring these characteristics together differs significantly by agency. The Census Bureau has roughly 7,500 permanent employees (of whom about 2,320 are statisticians and economists) and an annual budget of $1.5 billion. At the other end of the spectrum is the Bureau of Justice Statistics, with 52 employees (of whom about 33 are statisticians) and an annual budget of $42 million.[3] The FSS agencies have varying degrees of control over their resource levels, hiring, and autonomy to set priorities, define their data products, and develop statistical methodology. (Note that since this chapter was originally written in 2024, the circumstances of the federal statistical agencies have rapidly changed, with significant contract cancellations, staff and leadership layoffs, and other major disruptions to ongoing statistical activities beginning in January 2025. These are not accounted for in this chapter.)

In addition, government policymakers often overlook the FSS when their activities involve open data and AI. For example, Executive Order 13859, "Maintaining American Leadership in Artificial Intelligence,"required federal agencies to take major steps in supporting AI research and development (R&D), prioritize efforts to grow an AI-ready workforce, and make federal data assets available to the public to advance the development of AI in the US.[4] Missing from the executive order was any mention of funding the development of methods to measure the outcomes of these new or stepped-up efforts.

This omission was particularly noticeable because the president had signed the Foundations for Evidence-Based Policymaking Act, known as the Evidence Act, just one month before the release of the executive order, in January 2019. The Evidence Act requires agencies to establish learning agendas and build capacity for evaluating programs. However, as of 2023, no funding had been made available to the FSS agencies specifically to establish standard methods for measuring AI's effects on the labor force and economy. Subsequent AI-related executive orders also lack statistical measurement provisions.[5]

Thus, it should be no surprise that the development of statistical measures to assess the effects of AI on the economy are nascent. The largest of the 13 FFS agencies—the US Census Bureau, which has a relatively robust R&D budget—has taken steps to measure automation, technology, and the associated workforce, teaming up with the National Center for Science and Engineering Statistics (NCSES) to add questions on AI, beginning with the 2019 Annual Business Survey.[6] This survey collects information from over 300,000 firms on the use of five technologies: AI, robotics, dedicated equipment, specialized software, and cloud computing. According to 2019 survey data, adoption is concentrated on large and young firms.[7] Another survey conducted jointly by the Census Bureau and NCSES is the Business Enterprise Research and Development Survey,[8] which is the primary source of information on R&D expenditures and R&D employees of for-profit, publicly or privately held nonfarm businesses with 10 or more employees in the United States that have performed or funded R&D either domestically or abroad.

In addition, in 2022, the Census Bureau transformed its Business Pulse Survey, initiated as an experiment during the COVID pandemic, into the (also experimental) Business Trends and Outlook Survey (BTOS).[9] The BTOS collects data biweekly and publishes estimates every other Thursday. It collects information on a wide range of business conditions. Among other topics, sampled companies are asked about their current performance and changes in revenue, employment, hours worked, location, operating status, supply-chain impacts, demand, and prices. The survey asks businesses about the previous two weeks and for a six-month projection.

Beginning in September 2023, the BTOS sample included all employer businesses (single and multilocation) in the US, excluding a few North American Industry Classification System (NAICS) codes. During August 2023, the Census Bureau initiated an AI supplement to the survey with several questions about current and future AI use and its effect on the business's workforce during the past six months and the next six months. However, the BTOS questions are fairly basic, keeping the response time for businesses to about 17 minutes. While this brevity is helpful in garnering high response rates, it cuts down on the level of detail that the survey respondents provide.

The Census Bureau's Center for Economic Studies has been extending its Longitudinal Business Database[10] by matching additional datasets from various sources. The center is developing experimental Business Dynamics Statistics (BDS) products using newly linked data. The basic BDS data measure the net change in employment at the establishment level. The new linked datasets are intended to provide additional public information about how firm characteristics relate to employment flows.

The BDS-Innovative Firms data describe firms engaged in innovation. Currently available is the

Business Dynamics Statistics of US High Tech Industries (BDS-HT), an experimental data product[11] that merges industry-level information on STEM occupation intensity.[12] According to the Census Bureau, the BDS-HT classifies industries as high-tech using the concentration of STEM occupation employment as described in a 2020 paper and a 2005 paper.[13] To create this product, the Census Bureau used 11 years of data from the 2007–17 Economic Censuses and identified firms with at least five times the national average STEM employment in six of the 11 years. The industry list included 11 NAICS codes. Other data products, including experimental data products on patenting and trademarking, are under research using records from the US Patent and Trademark Office.

While the Census Bureau has taken significant steps to create data products that measure the impact of technologies on business and employment, its work is tied to the five-year cycles of the Economic Census and NAICS code updates. This helps the bureau track changes over time and shift the 11-year window after each Economic Census.

Of course, the downsides of this gradual change are that it becomes hard to capture a rapidly changing environment in real time and that the surveys rely on respondents' cooperation. In the trade-off between changing the survey more rapidly and keeping it respondent friendly, the data collected may not be sufficiently detailed for some analyses. Moreover, continuing to create experimental data products makes for slow and limited progress in measuring AI. The rate at which agencies can adapt their ongoing data collections and begin new ones that focus on technological change is highly dependent on each agency's resources, staff, and ability to innovate within the federal bureaucracy.

Conceptual and Empirical Gaps and Opportunities

Some efforts to define the effect of AI on the economy and workforce while maintaining standardization across the FSS recall efforts to measure intellectual property during the early 2000s. There existed a

> widespread perception that IP, rather than "brick and mortar" and other physical assets have been a major force in the rapid growth of GDP [gross domestic product], productivity, and wealth that occurred during the 1990's in the U.S. Among economists, for example, the recent productivity and growth accounting literature has been intensely focused on testing the hypothesis that much of the large unexplained quotient of long-term economic growth (total factor productivity) can be accounted for by better measurement of IP and other intangible assets that are regarded as components of "knowledge capital" inputs to industry production processes.[14]

Knowledge capital was divided into four categories in *Intangibles: Management, Measurement, and Reporting* and three categories in *Measuring Capital in the New Economy*, capturing tangibles and intangibles.[15] This work led to new classifications in the NAICS and North American Product Classification System (NAPCS). Examining this process could be helpful in devising a roadmap in the medium term for including AI in the industry, product, and occupation classification systems if warranted.

Another example of FSS research was the attempt by the Census Bureau and BEA to measure offshoring of manufacturing and define non-factory goods beginning in 2010.[16] "Factoryless" production was defined based on three main attributes: ownership of intellectual property, ownership and control of finished products, and outsourcing transformation activities. Ownership of intellectual property was measured as R&D expenditures, number of patents, and number of trademarks. Ownership and sales of finished goods were measured as revenue. Incidence of borderless production arrangements was measured as imports, and incidence of "headquarter" activity encompassing strategic or organizational planning and decision-making activities was measured as employment in NAICS codes 54 and 55. Of importance was the location of the activities—especially the application of IP in manufacturing overseas. The lessons learned from this effort may help researchers and policymakers define and measure AI and develop a roadmap and use cases for it.

An example of a more rapid way to drive experimental data collection and analysis in the FSS is the Bureau of Economic Analysis (BEA) satellite accounts. The BEA defines satellite accounts as

> supplemental accounts that expand the analytical capacity of the main system of accounts by focusing on a particular aspect of economic activity. Satellite accounts are linked to the main accounts but have greater flexibility in providing more detailed information or in using alternative definitions, concepts, and accounting conventions.[17]

The BEA's digital economy research is also an important avenue to explore, given the bureau's experience and skills in developing new ways to measure the economy and workforce and its access to multiple data sources.[18]

A factor that could slow down researchers' ability to correctly define and measure AI is that many of the Census Bureau's microdata files are accessible to only a small number of them, paralleling the problems of limited access to IRS tax data. Creating linked datasets at the Census Bureau that can be accessed only through the cumbersome process of using a Federal Statistical Research Data Center (FSRDC) will benefit the small number of researchers who are adept at using FSRDCs, but this approach will not further the goal of democratizing data by increasing access. Measuring AI's effects on the economy and workforce will take many people with different perspectives and questions that drive public policy. Relying on a few statistical agencies to create new measures will not result in rapid advancement with broad acceptance of the data. It will also not meet the intent of the Evidence Act.

The Evidence Act seeks to harness the potential of evidence-based policymaking by institutionalizing key principles and practices, and it includes 11 recommendations from the Commission on Evidence-Based Policymaking.[19] The act highlights the significance of public transparency and accessibility, stressing the value of sharing evidence and data with the broader community.[20] It demonstrates a commitment to enhancing data-driven governance and aligning governmental actions with empirically grounded insights.

The Evidence Act mandated the formation of a two-year Advisory Committee on Data for Evidence Building and charged it with giving recommendations to the director of the US Office of Management and Budget (OMB) on implementing the act, with a special focus on establishing a National Secure Data Service. This committee, composed of experts from inside and outside government in data analysis, privacy, and governance, provided valuable guidance on how federal agencies could navigate the intricacies of data sharing and integration while upholding ethical standards and addressing privacy concerns. The advisory committee's *Year 2 Report*, issued in October 2022,[21] focused on expanding access to data for evidence building, facilitating data sharing, enabling data linkage, and developing privacy-preserving techniques. It also provided a vision for how the National Secure Data Service could offer coordination and capacity-building services. The report's recommendations aimed to facilitate the secure exchange of sensitive data among federal agencies, researchers, and policymakers.

Subtitle F, Section 10375 of the CHIPS and Science Act of 2022 established a five-year demonstration project for the National Secure Data Service.[22] Congress charged NCSES with running the project to develop, refine, and test models for full implementation, according to the recommendations of the Commission on Evidence-Based Policy and the Advisory Committee on Data for Evidence Building. NCSES has established America's Data Hub Consortium[23] and has awarded some of the $9 million a year authorized for the National Secure Data Service demonstration project to proposals advancing knowledge that could inform future implementation. America's Data Hub could be a vehicle for a record-linkage project that partners with statistical agencies, other levels of government, and academic institutions to research the economic impact of AI.

Title III of the Evidence Act updated the Confidential Information Protection and Statistical Efficiency

Act of 2002 (CIPSEA). Since 2002, the Census Bureau, BEA, and the Bureau of Labor Statistics have been authorized to significantly improve the nation's economic statistics through sharing tax data that the Census Bureau has acquired from the IRS. However, the improvements have never been realized because the corresponding changes were not made in the IRS statues to enable this sharing. Although the Evidence Act provided an opportunity to change the IRS statute to enable CIPSEA to be implemented, no such changes were included, indicating again that relying solely on the FSS to take the lead in developing AI metrics and creating accessible data products to inform public policy will be a slow, evolutionary process that complements but can't lead needed efforts.

The viability of state-level data sharing has been successfully demonstrated by the Midwest Consortium, which has a strong governance structure and uses a secure platform to share sensitive data on education, the workforce, and training, greatly increasing the data's value to policymakers.[24] Building multistate capacity for data sharing is important to supporting evidence-based policymaking. Paired with collaborative data-sharing partnerships that already exist and have proved valuable, this capacity can be used to address the challenge of measuring AI.

Another area to explore is the need for federal standards on collecting information on AI. The director of OMB has delegated to the chief statistician of the US in the Executive Office of the President the authority to develop OMB standards for federal agencies' data collection. Thus, we have standards for geography through metropolitan and micropolitan statistical areas; race, ethnicity, and gender identification; and NAICS and NAPCS codes.

While the OMB needs to issue these standards, it would be helpful early on to consider conducting collaborative, community-based research in a way that feeds into early standardization of terms. One serious barrier to combining data is a lack of standardization in how terms are defined. Many OMB standards arose after problems in inconsistent data collections already existed. Developing standards informed by research on early use cases could significantly accelerate the efforts to assure high-quality data products.

Finally, research partnerships need to be expanded and encouraged, including in academics and all levels of government.

Short-, Medium-, and Long-Term Next Steps

To advance statistical measurement of AI's effect on the nation's economy and workforce, several actions need to be taken. These short-, medium-, and long-term steps, as laid out below, should be accomplished in a collaborative environment that engages the agencies and the user community.

Short-Term Steps

- Develop two or three use cases to explore AI measurement in the national economy and workforce. Identify vehicles for funding such research, which could use existing infrastructure such as Universities: Measuring the Impacts of Research on Innovation, Competitiveness, and Science; America's Data Hub; and other National Science Foundation and philanthropic funding mechanisms.

- Engage with government policymakers to identify pressing needs for information that may inform the selection of the most valuable use cases.

- Connect with the BEA about developing satellite accounts for AI.

Medium-Term Steps

- Engage in outreach and share results of use cases to attract collaborators and funding to continue the research.

- Raise awareness among policymakers of the data's availability.

- Explore with the Census Bureau and the Bureau of Labor Statistics how collaborative research can inform experimental data products being developed with survey and other data.
- Explore with the chief statistician of the US the need for federal standards on AI statistical data collection and how research could inform such standards, as well as changes to the NAICS and NAPCS and occupational codes.

Long-Term Steps

- Continue to expand research with input from the broader community.
- If feasible, continue to collaborate with the BEA on developing a satellite account for AI, coordinating research results.
- Create a blueprint from lessons learned and best practices for measuring AI that can be applied to other rapid changes in the economy and society.

Notes

1. The FSS is composed of 13 designated statistical agencies and includes roughly 100 units scattered throughout agencies that conduct statistical activities but are not formally recognized. The 13 agencies are the Bureau of Economic Analysis, Bureau of Justice Statistics, Bureau of Labor Statistics, Bureau of Transportation Statistics, Economic Research Service, Energy Information Agency, National Agricultural Statistics Service, National Center for Education Statistics, National Center for Health Statistics, National Center for Science and Engineering Statistics, Office of Research and Evaluation, Statistics of Income, and US Census Bureau.
2. American Statistical Association, "Assessing the Health of the Principal Federal Statistical Agencies," https://www.amstat.org/policy-and-advocacy/assessing-the-health-of-the-principal-federal-statistical-agencies.
3. Executive Office of the President, Office of Management and Budget, *Statistical Programs of the United States Government Fiscal Years 2021/2022*, 2024, https://www.whitehouse.gov/wp-content/uploads/2024/02/statistical-programs-20212022.pdf.
4. Executive Office of the President, "Maintaining American Leadership in Artificial Intelligence," *Federal Register* 84, no. 31 (February 14, 2019): 3967–72, https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence.
5. Executive Office of the President, "Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government," *Federal Register* 85, no. 236 (December 8, 2020): 78939–43, https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government; and Executive Office of the President, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," *Federal Register* 88, no. 210 (November 1, 2023): 75191–226, https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.
6. National Science Foundation, National Center for Science and Engineering Statistics, "Annual Business Survey (ABS) | 2022," https://ncses.nsf.gov/surveys/annual-business-survey/2022#survey-info; and US Census Bureau, ABS Tables, November 13, 2024, https://www.census.gov/programs-surveys/abs/data/tables.html.
7. Daron Acemoglu et al., "Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey," Working Paper No. 30659 (National Bureau of Economic Research, November 2022), https://www.nber.org/papers/w30659.
8. National Science Foundation, National Center for Science and Engineering Statistics, "Business Enterprise Research and Development (BERD) Survey | 2021," https://ncses.nsf.gov/surveys/business-enterprise-research-development/2021#methodology.
9. US Census Bureau, "Business Trends and Outlook Survey (BTOS)," https://www.census.gov/programs-surveys/btos.html.
10. US Census Bureau, "Longitudinal Business Database (LBD)," https://www.census.gov/programs-surveys/ces/data/restricted-use-data/longitudinal-business-database.html.
11. US Census Bureau, "Business Dynamics Statistics of U.S. High Tech Industries (BDS-HT)," https://www.census.gov/programs-surveys/ces/data/public-use-data/experimental-bds/bds-high-tech.html.

12. Nathan Goldschlag and Javier Miranda, *Business Dynamics Statistics of High Tech Industries*, US Census Bureau, December 2016, https://apps.bea.gov/fesac/meetings/2017-12-15/Goldschlag-BDS-High-Tech.pdf.

13. Nathan Goldschlag and Javier Miranda, "Business Dynamics Statistics of High Tech Industries," *Journal of Economics & Management Strategy* 29, no. 1 (2020): 3–30, https://onlinelibrary.wiley.com/doi/abs/10.1111/jems.12334; and Daniel E. Hecker, "High-Technology Employment: A NAICS-Based Update," *Monthly Lab Review* 128 (July 2005): 57, https://www.bls.gov/opub/mlr/2005/07/art6full.pdf.

14. Michael F. Mohr and John B. Murphy, *NAPCS Discussion Paper: An Approach for Identifying and Defining Intellectual Property (IP) and Related Products in Product Classification Systems*, Voorburg Group on Service Statistics, September 2002, https://www.voorburggroup.org/documents/2002%20nantes/papers/2002-064.pdf.

15. Baruch Lev, *Intangibles: Management, Measurement, and Reporting* (Brookings Institution Press, 2001); and Carol Corrado et al., "Measuring Capital and Technology: An Expanded Framework," in *Measuring Capital in the New Economy*, ed. Carol Corrado et al. (University of Chicago Press, 2005), 11–46.

16. Fariha Kamal, "A Portrait of U.S. Factoryless Goods Producers," Working Paper No. 25193 (National Bureau of Economic Research, October 2018), http://www.nber.org/papers/w25193.

17. Bureau of Economic Analysis, "Satellite Accounts," April 26, 2018, https://www.bea.gov/help/glossary/satellite-accounts.

18. Bureau of Economic Analysis, "Digital Economy," October 3, 2024, https://www.bea.gov/data/special-topics/digital-economy.

19. Nancy Potok, "Deep Policy Learning: Opportunities and Challenges from the Evidence Act," *Harvard Data Science Review* 1, no. 2 (2019), https://doi.org/10.1162/99608f92.77e63f8f.

20. Nancy Potok, "Data Usage Information and Connecting with Data Users: U.S. Mandates and Guidance for Government Agency Evidence Building," in "Democratizing Data: Discovering Data Use and Value for Research and Policy," special issue, *Harvard Data Science Review* 4 (April 2024), https://hdsr.mitpress.mit.edu/pub/605i9bgg/release/3.

21. Advisory Committee on Data for Evidence Building, *Advisory Committee on Data for Evidence Building: Year 2 Report*, October 14, 2022, https://www.bea.gov/system/files/2022-10/acdeb-year-2-report.pdf.

22. CHIPS and Science Act of 2022, Pub. L. No. 117-167.

23. America's Data Hub Consortium, website, https://www.americasdatahub.org/.

24. Jessica Cunningham et al., "A Value-Driven Approach to Building Data Infrastructures: The Example of the MidWest Collaborative," *Harvard Data Science Review* 4, no. 1 (2022), https://hdsr.mitpress.mit.edu/pub/mfhpwpxq/release/2.

About the Authors

Vipin Arora is the director of the US Bureau of Economic Analysis. Arora oversees the production of closely watched economic statistics that provide an objective and timely picture of the US economy.

Lee Branstetter is the James Walton Professor of Economics and Public Policy at Carnegie Mellon University's Heinz College. He is also a research associate at the National Bureau of Economic Research, affiliated with the programs on productivity, innovation, and entrepreneurship and international trade and investment. Branstetter holds a PhD in economics from Harvard University.

Erik Brynjolfsson is the Jerry Yang and Akiko Yamazaki Professor and senior fellow at the Stanford Institute for Human-Centered AI and director of the Stanford Digital Economy Lab. He is also the Ralph Landau Senior Fellow at the Stanford Institute for Economic Policy Research, professor by courtesy at the Stanford Graduate School of Business and Stanford Department of Economics, and a research associate at the National Bureau of Economic Research.

Mar Carpanelli is a senior staff data scientist and head of AI and skills research at LinkedIn.

Diane Coyle is the Bennett Professor of Public Policy at the University of Cambridge and codirector of the Bennett Institute for Public Policy. Her research focuses on economic measurement, the digital economy, and competition policy.

Lesley Hirsch is the assistant commissioner of research and information at the New Jersey Department of Labor and Workforce Development.

Harry J. Holzer is the John LaFarge Jr. SJ Professor of Public Policy at Georgetown University, a nonresident fellow in Economic Studies at the Brookings Institution, and a former chief economist of the US Department of Labor.

Eduard Hovy is the executive director of Melbourne Connect and a professor in the School of Computing and Information Systems at the University of Melbourne. He is also an adjunct professor at Carnegie Mellon University. Hovy is a leading scholar in natural language processing with more than 460 published papers. He holds a PhD in artificial intelligence from Yale University.

Suzette Kent is a global business transformation executive. She served as the federal chief information officer for the United States from 2018 until July 2020 and has held leadership roles ranging from partner at Accenture and EY to managing director at JPMorgan. She currently leads her own advisory business through which she continues strategic transformation work with clients around the world, across industries, and in the public and private sectors.

Karin Kimbrough is chief economist at LinkedIn.

Julia Lane is a tenured full professor at the New York University Wagner Graduate School of Public Service. She was a senior adviser in the Office of the Federal Chief Information Officer at the White House, supporting the implementation of the Federal Data Strategy. She recently served on the Advisory Committee on Data for Evidence Building and the National AI Research Resources Task Force and currently serves on the secretary of labor's Workforce Innovation Advisory Committee and the National Science Foundation's Advisory Committee on Cyberinfrastructure. She has authored over 80 scientific publications, edited or coauthored 13 books, and received over $180 million in grants and contracts from national and international agencies and foundations.

Adam Leonard is the chief analytics officer for the Texas Workforce Commission and a self-styled data evangelist, working to leverage data to improve decisions and promote the conditions that lead to attaining and sustaining prosperity.

Layla O'Kane is the senior director of data products at Opportunity@Work and was formerly the director of research at Lightcast.

Brent Orrell is a senior fellow at the American Enterprise Institute specializing in job training and workforce development with a special focus on disconnected and disadvantaged populations, including youth, justice-involved, veterans, and neurodivergent persons. His recent work has focused on the workforce opportunities and challenges resulting from generative AI and automation.

Jason Owen-Smith is the cofounder and executive director of the Institute for Research on Innovation and Science, professor of sociology and public policy, and research professor in the Institute for Social Research at the University of Michigan.

Nancy Potok is the CEO of NAPx Consulting. She served as the chief statistician of the United States from 2017 to 2020. Before that, she served as the deputy director and chief operating officer of the US Census Bureau and the deputy under secretary for economic affairs at the US Department of Commerce. She has also led private and nonprofit social science research and data analytics organizations. She earned her PhD from The George Washington University Trachtenberg School of Public Administration and Public Policy.

Will Rinehart is a senior fellow at the American Enterprise Institute, where he focuses on the political economy of technology and innovation. His research covers policy areas such as diagnostic testing regulation, federal agency regulatory guidance, the Creating Helpful Incentives to Produce Semiconductors and Science Act, and the regulation of artificial intelligence.

Prasanna "Sonny" Tambe is a professor of operations, information, and decisions at the Wharton School at the University of Pennsylvania and faculty codirector of AI at Wharton. His research focuses on the economics of technology and labor. Tambe received a PhD in managerial science and applied economics from the Wharton School.

David Wasshausen is the associate director for National Economic Accounts at the US Bureau of Economic Analysis. Wasshausen oversees the calculation of official national economic statistics, including gross domestic product.

American Enterprise Institute
1789 Massachusetts Avenue, NW
Washington, DC 20036
202.862.5800
www.aei.org