Branstetter Stanford Conference "White Paper"

<u>Quantifying the Impact of AI Innovation and Diffusion</u>

Lee Branstetter, CMU & NBER
Eduard Hovy, CMU & University of Melbourne

Take-aways: (1) We need to use AI to find AI-related patents. (2) AI invention significantly boosts the productivity and employment of AI-inventing firms. (3) The limited numbers of AI experts trained up to the scientific frontier who seek industrial employment in the U.S. may impose a significant human resource constraint on the ability of most firms to turn AI breakthroughs into new products and services.

Our Research in the Context of AI Impact Measurement: What is the impact of AI innovation on productivity? Business leaders have proclaimed a fourth industrial revolution, centered around AI and related advances in information technology. However, earlier industrial revolutions were characterized by 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 stuck at slow rates that limit the growth of American incomes, prosperity, and global power (Benzell et al., 2022). Will AI fail to live up to the enthusiasm of its advocates or are we merely in the early stages of an innovation and adoption process that will take years or decades to unfold? Our research (see Alderucci et al., 2022) seeks to address this question by examining the vanguard of firms that are already introducing AI-related innovations into the marketplace. If these early movers, innovators, and adopters are already reaping significant productivity gains, then this augurs well for the ultimate positive impact of AI on the U.S. economy. In ongoing work, we are also examining the role played by Ph.D.-level academic experts in the creation of AI-related new goods and services and expect to find evidence consistent with the idea that a shortage of advanced human capital profoundly shapes where and by which firms AI-related innovation is advancing. By linking these experts to the firms that employ them, we may obtain empirical leverage around the difficulty of measuring AI invention that does not result in patents and the application of frontier or near-frontier AI ideas to re-engineer existing products and services. This latter phenomenon may lie somewhere between AI "invention" and AI "adoption."

Our work complements many recent streams of research. One stream seeks to measure AI adoption and use through direct surveys of large, representative samples of U.S. firms (Zolas et al., 2020). These valuable efforts have not yet demonstrated a convincing causal relationship between AI adoption and faster productivity growth, and it will take time before these survey data acquire a sufficiently long time series dimension such that researchers can apply the usual econometric techniques for discerning plausibly causal effects from nonexperimental data. A second stream applies randomized controlled trials or quasi-experimental methods to measure the impact of AI on productivity in a particular work context (see Brynjolfsson, Li, and Raymond, 2023; Korinek, 2023; and Noy and Zhang, 2023). Some of these papers have found convincing evidence of a causal impact of AI adoption on productivity, but the results may not generalize from the particular work contexts in which these experiments or quasi-experiments were conducted. A third stream uses data on the recruitment of specialized labor to measure AI use and AI-related innovation (see Babina et al., 2022, forthcoming). Prominent papers in this stream have found positive evidence of this investment on output and product innovation, but have failed to find robust evidence that investments in AI use led to increases in productivity growth. The strong productivity growth effects we document are potentially broader in scope than those found in the experimental literature and point to the optimistic possibility that AI could eventually lead to a significant and persistent acceleration in productivity growth across a broad range of industries.

Methods, Data, and Current Results: We rely on two sources of data related to AI-driven innovation in the U.S. economy – patent data and publication data. We use an ensemble of machine learning algorithms to parse the text of patents granted by the U.S. Patent and Trademark Office (USPTO) and identify AI-related inventions. We then match data on the AI patents to the rich data bases maintained by the U.S. Census Bureau on the AI patent-inventing firms. Using firm fixed effects models and event study approaches, we find striking evidence that AI invention leads to significant growth in output, employment, and productivity. Many other studies have failed to find a significant impact of AI adoption or innovation on productivity – but our approach finds this, and the evidence is quite robust. We also find evidence that AI invention creates greater inequality within the wage distribution of AI-inventing firms.

Figure 1 summarizes the results obtained from firm fixed effects models. The impacts estimated from the initial transition into AI patenting, which capture the extensive margin effects, are denoted "extensive." The impacts estimated from an increase in the stock of AI patents, which capture the intensive margin effects, are denoted "intensive."

Fig. 1 Impact of AI Patenting on Productivity
Fixed Effects Models (% increase)

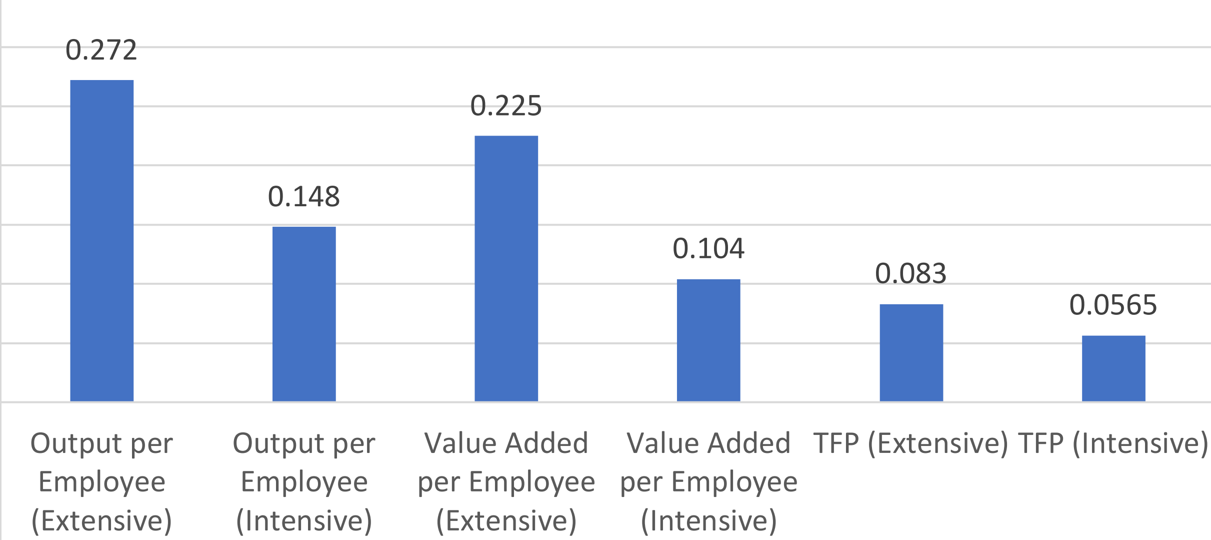

An alternative approach to estimation of the impact of AI patenting on the inventing firms is to associate each AI-inventing firm with its closest same-industry non-AI-inventing peer firm(s) and compare changes in key firm outcome measures between the "treated" (AI-inventing) and "control" (non-AI-inventing) firms after the former firm transitions into AI patenting. Figure 2 summarizes the results obtained from these "event study" models. Here, too, the transition to AI patenting is associated with a statistically significant and economically meaningful increase in employment and in revenue per employee.

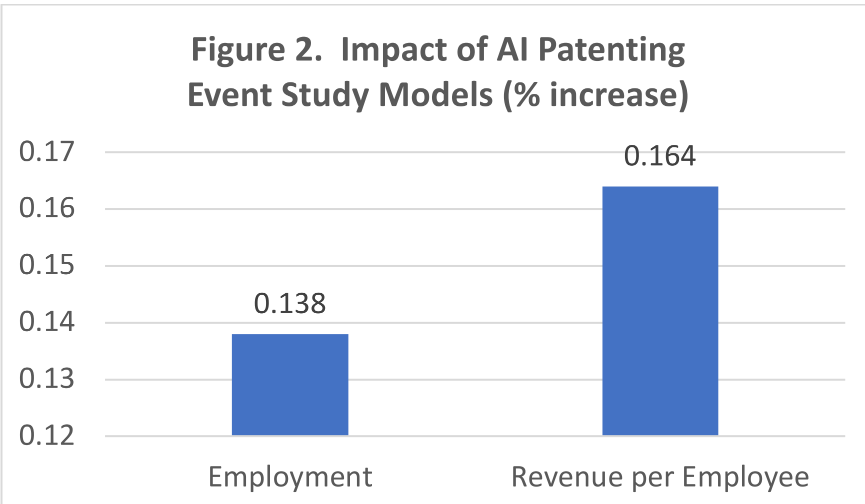

Ongoing Research Efforts: Not all innovations are patented, and some industries investing heavily in AI to generate new products and services hardly patent at all. To identify investments in AI-related innovation in these sectors, we are using publication data from Elsevier to identify the top academic scientists working in domains related to AI and the graduate students whom they supervise and with whom they coauthor. We are then collaborating with Prasana Tambe to use a combination of publication data, website data, and data from professional career profile services and resumes to track the movement of these students of scientific thought leaders across geographic space, 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 can link the "star scientists" and their students to the firms with which they have worked, and trace these linkages over time, we can leverage our access to U.S. Census microdata, obtained through our ongoing collaboration with Census microeconomists, to ask whether these linkages have provided the receiving firms with a statistically discernable advantage over their same industry peers who lack them in terms of output, employment, or, most importantly, productivity.

This line of inquiry is related to the work of Babina et al. (2022, 2024), but focuses on the potentially special role played by elite scientists, who may play a disproportionately important role in the defining the technology frontier, and their doctoral students, who may play a disproportionately important role in bringing this frontier technology into industrial practice (Agarwal and Henderson, 2000; Zucker and Darby, 1998). We can imagine that any firm seeking to apply frontier AI to the substantive reengineering of its current products and services or the creation of new products and services needs to create within itself a "pyramid" of AI talent graphically depicted in the left portion of Figure 3 below (Arora et al., 2013; Branstetter et al., 2019). At the lower ranks of the 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 talent pyramid, the firm might need professionals with bachelors or masters degrees that include specialized AI training, but these professionals need not have trained at elite universities. However, at the very apex of the pyramid, a firm seeking to out-engineer its rivals may seek to acquire "software architects" who have been trained up to the technology frontier by elite academic scientists based at the top universities. The data and approach taken by Babina et al. (2022, 2024) use data on the entire pyramid; our approach focuses on the star scientists and their students who could constitute disproportionately important part of the apex of that pyramid. The role played by these individuals is related to that of the "architects" described in the theoretical work of Benzell et al. (2022). As in that paper, we consider the idea that the limited supply of these software architects could constitute an important constraint on the ability of firms to fully leverage frontier AI technologies.

Figure 3 Tracing the Impact of AI Software "Architects": A Suggestive Illustration

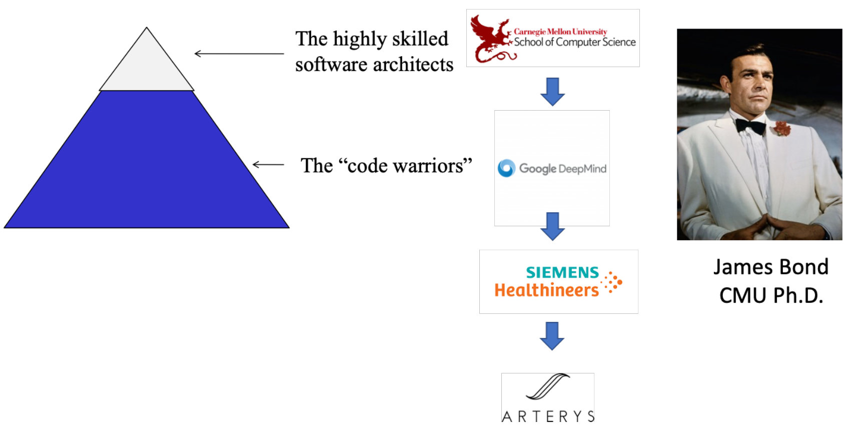

To fix ideas further, imagine that our data sources identify CMU doctoral recipient James Bond as one of the Ph.D. advisees of an elite academic scientist. Dr. Bond's subsequent movement to Google DeepMind could further augment the intellectual resources of this impressive corporate research operation. It is possible that Dr. Bond begins to specialize at Google DeepMind in the application of advanced AI algorithms to medical imaging. Then, he carries this skill to Siemens Healthineers and from there to diagnostic imaging start-up Arterys. By following star scientists' students like Dr. Bond from firm to firm, we could trace out their differential impact, if any, on the enhancement of host firms' output, employment, and productivity. The hypothesis that these movements predict success can be tested using access to Census data on the hiring firms and their same industry peers who have hired fewer or no advanced AI experts.

Implications for Policy Options and Trade-offs: 1. *Providing a strong data-driven rationale for investing in AI.* The U.S. government possesses limited resources with which it can invest in basic science. If AI boosts corporate profits but fails to boost employment or productivity, then the desirability of a long-term, large-scale investment in AI research could be called into question. Our results provide strong and robust evidence that the firms investing in AI-related innovation are seeing significant gains in output, employment, and productivity. As the set of AI innovators expands, these effects are likely to show up in aggregate economic data. 2. *Creating a methodology that can identify AI innovation abroad and compare its quality and impact to that of the U.S*. Many groups within the U.S. government worry that adversary nations are developing AI innovation capabilities that may rival those of the U.S. and its allies. The data and methods pioneered by our team could be applied to foreign (e.g., Chinese) patent data, shedding crucial light on the real strengths and weaknesses of AI innovation in China in a way that goes far beyond simply counting large numbers of patents of low average quality. 3. *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 the development or new goods and services or the re-engineering of existing goods and services. Our project will create a much more comprehensive data base of these AI experts that tracks their movement from leading centers of AI scientific research to innovating companies and organizations across the world. It will also assess the degree to which the accumulation of this scarce human resource is correlated with AI innovation and the impact of AI innovation.

Because human resources appear to be a critical constraint on the ability of firms to turn these scientific breakthroughs into new products and services, we believe our project also provides a data-driven rationale for a large and immediate increase in the allocation of H1-B visas and/or green cards to the foreign graduates of top U.S. AI-related Ph.D. programs and to the similarly skilled graduates of foreign programs.

Applications to Other Critical and Emerging Industries: In principle, the basic techniques we are applying to AI could be applied to other critical/emerging industries. The iterative approach through which we built an ensemble of machine learning algorithms to identify AI-related patents could be adapted to other technological fields. Once identified, these patents could be linked to the inventing firms, and the Census microdata could be leveraged to identify impacts on output, employment, and productivity. Similarly, publication data could identify star scientists associated with this new technology, and the same mix of publication data, website data, and "LinkedIn-like" data could quantify the interaction of these star scientists and their students with these firms. However, for this approach to be meaningful, there would need to be a strong connection between new scientific breakthroughs and their application in a wide range of industrial contexts.

<u>Papers Cited</u>

Agarwal, Ajay and Rebecca Henderson, "Putting Patents in Context: Exploring Knowledge Transfer from MIT," Management Science 48 (1), pp. 44-60.

Alderucci, Dean, Lee Branstetter, Eduard Hovy, Andrew Runge, and Nikolas Zolas, 2022, "Quantifying the Impact of AI on Productivity and Labor Demand: Evidence from U.S. Census Microdata," working paper.

Arora, Ashish, Lee Branstetter, and Matej Drev, 2013 "Going Soft: How the Rise of Software-Based Innovation Led to the Decline of Japan's IT Industry and the Resurgence of Silicon Valley." *Review of Economics and Statistics*, vol. 95, no. 3, pp. 757-775.

Babina, Tania, Anastassia Fedyk, Alex He, and James Dodson, 2022, "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition," working paper, Columbia Business School.

Babina, Tania, Anastassia Fedyk, Alex He, and James Dodson, forthcoming, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economic*s

Benzell, Seth, Erik Brynjolfsson, and Guillame Saint-Jacques, 2022, "Digital Abundance Meets Scarce Architects: Implications for Wages, Interest Rates, and Growth," Working Paper.

Branstetter, Lee, Matej Drev, and Namho Kwon, 2019, "Get with the Program: Software-driven Innovation in Traditional Manufacturing," *Management Science*, vol. 62, no. 2, pp. 541-558.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond, 2023, "Generative AI at Work," NBER Working Paper 31161.

Korinek, Anton, 2023, "Language Models and Cognitive Automation for Economic Research," NBER Working Paper 30957.

Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz, 2020, "Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey," NBER Working Paper 28290.

Zucker, Lynne, Michael Darby, and Michael Brewer, 1998, "Intellectual Capital and the Birth of the U.S. Biotechnology Enterprise," American Economic Review 88 pp. 290-306