

Does Corporate Production of AI Innovation Create Value?

ALI AHMADI, AMBRUS KECSKÉS,
RONI MICHAELY, and PHUONG-ANH NGUYEN*

Abstract

Yes, by decreasing firm risk, not by increasing profitability, and with investors taking years to recognize the value created. We start, using novel AI patent data, by documenting significant corporate production of AI innovation as early as 1990. Then, we show that a significant motivation for a firm's AI production is the mutually reinforcing effects of the firm's innovation capacity (exogenous R&D stock) and its labor inputs' AI exposure (both the firm's own and its customers'). We use the interaction of these two effects to instrument for AI production. We find that producing AI creates firm value through a large, permanent *decrease* in risk (cash flow and stock return, systematic and idiosyncratic). Further evidence suggests that AI lowers physical capital intensity and increases bargaining power for producing firms. The initial market reaction to AI patent announcements is economically small, but abnormal stock returns thereafter are significantly *positive* (about 5% per year) for (only) roughly three years, suggesting initial undervaluation followed by gradual correction. We find no evidence of investor learning, except during the past five years. We empirically distinguish producing AI innovation versus AI adoption, automation, general technology, and other potential confounds.

First version: September 28, 2023

This version: November 13, 2025

JEL classification: G12, G31, G32, J21, J24, O31, O32, O33

Keywords: Artificial intelligence; Innovation; Technology; Labor; Firm value; Corporate finance; Asset pricing; Behavioral finance

* Ahmadi (aliahmdi@yorku.ca) and Kecskés (akecskes@yorku.ca) are at the Schulich School of Business, York University; Michaely (ronim@hku.hk) is at the University of Hong Kong Faculty of Business and Economics, and ECGI; and Nguyen (anhvn@yorku.ca) is at the School of Administrative Studies, York University. We greatly appreciate the comments of Pat Akey, Yun Joo An, Kostas Bakalopoulos, David Chambers, Lauren Cohen, Peter Cziraki, Silvia Dalla Fontana, Yao Deng, François Derrien, Laurent Frésard, Koresh Galil, Wei Jiang, Chris James, Kose John, Marcin Kacperczyk, Alex Kim, Charles Martineau, Niranjana Sapkota, Abhijit Sharma, Elvira Sojli, Jungwon Suh, Mikhail Simutin, Noah Stoffman, Lucy Wang, Katarina Warg, Changyu Yang, Hayong Yun, and Linghang Zeng, and seminar participants at the 2025 Academy of Finance Conference, the 2024 Afrimed Finance Society Conference, the 2024 CREDIT International Conference (GRETA), the 2025 Eurasia Business and Economics Society Conference, the 2025 CCA-ESCP Workshop on Financial Institutions and Corporate Finance, the 2024 European Finance Association, the 2024 Finance Symposium, the 2024 FMA, the 2025 FMA Applied Finance Conference, the 2024 FMA Asia/Pacific Conference, the 2025 FMA Consortium on Asset Management, the 2024 International Risk Management Conference, the 2024 Research in Behavioral Finance Conference, the 2025 Vietnam International Conference in Finance, and York University. This research was supported by the Social Sciences and Humanities Research Council of Canada and the Canadian Securities Institute Research Foundation.

1. Introduction

Artificial intelligence, the technology of machine cognition, has grown explosively in recent years. In this paper, we study the *production* of AI innovation, and its value implications for producing firms, using newly available USPTO data on AI patents.^{1,2} We focus on patented innovations because patent protection allows the firm to control the technology it produces. We document that, in the aggregate, AI is increasingly a prominent subset of all innovation activity during the past three decades. In 1990, AI already accounted for 5% of all innovation activity, and has risen to 15%-35% of innovation today.³ Additionally, AI innovations are consistently more valuable over time than non-AI innovations, in terms of both scientific and commercial value, and encompass far more technological breakthroughs. Moreover, as expected from a general purpose technology, AI innovation has diffused widely over time across industries. For instance, AI accounts for at least a majority of all innovation in about 20% of industries today. Finally, U.S. publicly traded firms dominate AI, producing about 70% of AI innovation by U.S. firms.

Focusing on publicly traded firms over the past three decades, we document causally how AI innovation creates value for producing firms, and that the stock market appears to consistently undervalue AI innovation for roughly three years after it is produced. AI innovation still only has a small, brief positive impact on profitability to date. Instead, producing AI innovation creates value principally by substantially and permanently decreasing risk (cash flow

¹ Production of new AI technology is distinct from the *adoption* of previously developed AI technology as studied by prior literature (e.g., Alekseeva, Giné, Samila, and Taska (2020) and Babina, Fedyk, He, and Hodson (2022, 2024a, 2024b)). The distinction between production and adoption is important in both conceptual predictions, notably for firm risk, and in the corresponding empirical effects, as our paper shows.

² As we summarize in Section 2.1, the USPTO data, from Giczy, Pairolo, and Toole (2022), classify patents as "AI" and "non-AI" using stratified machine learning that assigns patents a predicted probability of being AI. AI encompasses the eight AI component technologies identified by the National Institute of Standards and Technology. The machine learning patent classification is validated by AI specialist USPTO patent examiners, and it is shown to outperform alternatives classifications in minimizing false positives and negatives.

³ The share of AI innovation is highest when measured based on the scientific or commercial value of patents.

and stock return, systematic and idiosyncratic). We find that firms direct their innovation capacity toward producing AI technology based on the extent to which AI substitutes for or complements labor, both the firm's and its customers'. Our evidence on mechanisms suggests that AI technology lowers the producing firm's physical capital intensity and increases its bargaining power. We find no evidence over three decades that investors learn more quickly about the value of AI innovation, except for the past five years during which the initial market reaction remains economically small and there is no longer a significant drift after AI patent announcements.

We start our analysis with a brief conceptual framework (with illustrations) of how producing AI innovation affects firm value. AI, as machine cognition technology, can sort through large and multi-dimensional flows of data (text, vision, speech, etc.), learn from its own evolving history, and dynamically update itself. As such, it can substitute for or complement many of the main functions of a firm's employees: information acquisition and forecasting (e.g., estimating customer demand using machine learning); monitoring (e.g., supervising employee-customer interactions with natural language processing and speech recognition); and decision making (e.g., autonomous inventory management or computer vision assisted quality control). AI innovations corresponding to the foregoing applications were being produced (and patented) at least as far back as the early 1990s.⁴ Moreover, there is significant production of (patented) AI innovation even in non-technology industries (e.g., aircrafts, motor vehicles, drugs, petroleum, and industrial equipment).⁵

Turning to firm value implications, producing AI innovation can naturally increase profitability through all of the above applications of the technology. However, AI can also decrease risk, by improving the reliability of the firm's execution and its responsiveness to

⁴ We provide examples of a wide variety of patented machine cognition tasks in Section 2.2.

⁵ This is evident from ranking industries based on patented AI innovation, which we examine in Section 3.2.

changing business conditions. Importantly, when the firm can control its technology (e.g., AI innovations that are patent protected), it can create more value, by, for instance, excluding others from using its technology (e.g., its product market competitors), which will decrease the firm's correlation with the rest of the economy and decrease its risk.

Turning to a detailed exposition of our empirical analysis at the firm level, we begin by examining a potentially significant motivation for firms to produce AI innovation. Our intuition is that firms that benefit more from AI technology will devote more innovation resources to producing it. Firms that benefit more are those whose own labor, or their customers', is more easily substituted for or complemented by AI technology. We refer to such firms as having high AI exposure.

This underlying intuition forms the basis of our empirical strategy for identifying the causal effect of AI innovation. If a firm with exogenously high AI exposure experiences an exogenous increase in innovation capacity, it will produce more AI innovation compared to an otherwise identical firm with exogenously low AI exposure. We instrument actual AI innovation using the interaction of two mutually reinforcing incentives for a firm to produce AI innovation: the firm's (industry level) exposure to AI technology, and the firm's (R&D tax credit induced) innovation capacity. In what follows, we describe the construction of these two components and explain their plausible exogeneity.

We construct the two components of our interaction instrument (a kind of shift-share) as follows. For the first component, we calculate industry-level AI exposure scores using occupation-level (SOC) AI exposure scores⁶ weighted by employment in each occupation in

⁶ From Felten, Raj, and Seamans (2021). These scores capture the extent to which labor in an occupation can be substituted for or complemented by AI technology.

each industry (SIC3) fixed before the start of our sample period.⁷ An important advantage of using time-invariant, pre-sample period, industry-level AI exposure is that it minimizes the possibility that firms choose their (firm-level) AI exposure endogenously with firm outcomes. For the second component, we construct the firm's plausibly exogenous R&D capital stock following the literature.⁸ Specifically, we predict firm-year R&D spending with federal and state R&D tax credits, which vary over time across the firm's R&D hubs located in different cities. We then cumulate predicted R&D spending to R&D capital stock over the firm's prior history. We address potential threats to our identification below, after presenting our main results on valuation implications.⁹

In the first stage of our instrumental variables analysis, we predict the production of AI innovation. We find that for a typical increase in our interaction instrument, AI patent counts increase by about 13% relative to the mean, roughly similarly for both the firm's own AI exposure and its customers' (each exposure varying across industries).¹⁰ We use this natural measure of AI innovation activity, AI patent counts, in our baseline analyses throughout the paper, but the results are similar using alternative scaling variables such as total assets and total patent counts.

We then empirically examine the value created by corporate production of AI innovation, starting with AI patent grant announcements. The initial market reaction is economically small

⁷ We measure the AI exposure of a firm's customers analogously, using our industry AI exposures combined with inter-industry product purchase weights from the BEA's input-output tables.

⁸ For instance, see Wilson (2009); Bloom, Schankerman, and Van Reenen (2013); Hombert and Matray (2018); and Babina and Howell (2024). The latter provide comprehensive evidence for the exogeneity of R&D capital stock thus constructed, including their own evidence as well as that from the prior literature.

⁹ As detailed in Section 4.1., the IV results throughout the paper are incremental to controlling for the direct effects of (tax credit induced) R&D capital stock and AI exposure. We also control for whether the firm produces patents as well as the firm's non-AI patent count, size, and age. Finally, we sweep out persistent differences across firms, three-digit SIC industries (which, indirectly, largely sweep out AI exposure), two-digit SIC industries each year, and the firm's headquarters state each year (which, indirectly, largely sweeps out R&D capital stock).

¹⁰ The results throughout the paper are robust to using either instrument (firm or customers) alone. We use both instruments in our baseline analysis because we have no theoretical basis for favoring either exposure, whereas we can increase the precision of our IV estimates by using both exposures together.

during the first event week (6.4 basis points for the average firm-day). Rather than extending the event window, we focus instead on using annual AI patent counts to predict future annual stock returns because this lower frequency analysis has better statistical and economic properties.¹¹

We start our returns regression analyses using actual AI patent grants, which are readily observable to investors, and form portfolios each year. We find that a high minus low AI stock portfolio earns risk-adjusted returns of roughly 50 basis points per month. This could be consistent with either compensation for risk (if producing AI increases risk), or with investors only gradually impounding into stock prices the positive value effects of producing AI innovation (if producing AI increases profitability and/or decreases risk). By contrast, comparably constructed non-AI patent portfolios do not reliably spread returns.

We then exploit our IV approach to more credibly identify the effect of producing AI innovation on returns in a reduced form setting. We double sort stocks (independently) into portfolios based on firms' (tax credit induced) R&D capital stock and their AI exposure. We again find that portfolios, at the intersection of these double sorts, outperform by about 50 bps per month (i.e., the high minus low AI exposure spread netting out the high minus low R&D capital stock spread). Our portfolio returns results are similar using a wide range of factor models.

Additionally, we implement our IV analysis in monthly Fama-MacBeth cross-sectional regressions using all sample firms over three decades. This more demanding approach allows us to include a battery of control variables like in our first stage IV regressions (e.g., R&D capital stock, AI exposure, and non-AI patent counts) as well as established determinants of stock returns. We again find a significant return spread, comparable in magnitude to the preceding

¹¹ As explained in Section 6.1, the annual frequency improves measurement precision of the firm's AI innovation activities: capturing their scale and synergies while mitigating their persistence, all within the firm; and accounting for their undervaluation by investors. It also makes possible our IV analysis, for identification.

portfolio analysis. The results of our stock returns analyses, taken together, consistently indicate that there is a small initial market reaction to producing AI innovation followed by a significant return spread for roughly three years.

We next examine how producing AI innovation creates firm value using the canonical financial drivers of value: cash flows and risk. Using detailed firm-year panel data spanning three decades, we focus on our IV estimates to identify the causal effect of producing AI innovation. Since the full productivity potential of AI (e.g., as measured by profits) may not be realized by the end of our sample period,¹² our estimates may be understated.

In our first IV results, we find that after producing AI innovation (i.e., during the year after we measure AI patent grants), producer firms become transitorily more profitable. For a typical 10% increase in instrumented AI patent counts, net income increases by roughly 0.7 percentage points relative to total assets. However, this first year effect becomes insignificant by the second year. Recognizing that experience with AI production in the past may facilitate AI production in the future, we explore here the moderating role of experience. The evidence suggests that the effect of AI production on profitability is driven by firms with more experience producing AI innovation (as measured by the firm's past AI patent stock).

More significantly, AI producer firms become permanently less risky. A 10% increase in instrumented AI patent counts decreases the volatility of net return on assets by about 7%. This lower cash flow volatility is also reflected in lower stock return volatility, which decreases by roughly 2%. Both effects persist for at least five years and are also magnified by experience producing AI innovation. Decomposing total stock return volatility into systematic and idiosyncratic components, we find that both decrease permanently. Our finding that AI

¹² E.g., Brynjolfsson, Rock, and Syverson (2019) argue that AI technology is still in the early stages of diffusing across sectors and into complementary technologies.

production decreases risk is consistent with the firm using its (patented) AI technology to improve the reliability of its execution and its responsiveness to changing business conditions. Furthermore, in the second half of our sample period, the decrease in risk is slightly larger (but there is no difference in profitability or future stock returns). At the same time, all of our results (profitability, risk, and returns) are stronger for higher quality AI patents (as captured by higher scientific and commercial value).

Taken as a whole, the results are consistent with investors only gradually impounding into stock prices over several years the positive value effects (mainly, lower risk) of producing AI innovation. Undervaluation taking several years to be corrected is consistent with limits to arbitrage: few people may really understand both AI technology itself and how to value it, during most of our sample period; and valuing long-term risk may be particularly difficult (as opposed to short-term profitability).

We now turn to potential threats to our identification. We do have plausibly exogenous variation in the firm's innovation capacity, as captured by its (tax credit induced) stock of R&D capital (incrementally to our demanding set of fixed effects). However, since our measure of AI exposure only varies at the industry level (not at the firm-year level), the main threat to identification is any omitted factor at the industry level that is correlated with AI exposure and reinforces the effect of the firm's innovation capacity on profitability, risk, or other firm outcomes. We address the possibility that firms respond to random shocks to their innovation capacity differently in different industries, but specifically not because of different AI exposures, using a variety of aggressive specifications (Section 7.3). For instance, we include industry fixed effects interacted with R&D stock or industry-state fixed effects, or we include lagged dependent

variables. We are also able to empirically distinguish the production of AI innovation from potential confounds such as AI adoption, automation, and general technology.¹³

We then provide suggestive evidence on various mechanisms through which producing AI innovation can increase firm value, particularly when the firm controls its technology and exclude others from using it. First, AI can enable the firm to use its physical assets more efficiently or even require less physical capital to produce the same or more output. Indeed, we find a decrease in physical capital (e.g., PP&E) as well as investment (e.g., capex). Second, AI can increase the firm's bargaining power vis-à-vis its customers, employees, and other business counterparties. Some benefits of AI technology can accrue to the producer firm's counterparties downstream and upstream in its supply chain (e.g., AI embedded products), which the producer firm can use to negotiate business deals that are more lucrative or more stable for itself.. We find evidence of more stable sales and costs, and of greater product differentiation. Third, we find that AI innovation increases labor productivity (e.g., profit per employee), but only transitorily so (much like for profitability, as measured by net income to total assets). Moreover, AI innovation does not affect employment or the overall scale of the firm, which suggests that AI complements, rather than substitutes, labor.

In our final analysis, we examine the financial policy implications of AI innovation. Both an increase in expected future profitability and the decrease in risk that we document would enable the firm to be more aggressive with its financial structure. Our results, which include higher leverage and lower cash holdings, indicate that this is the case.

We contribute to the literature on the economics of artificial intelligence. The existing literature focuses on AI adoption and finds that it increases firm growth and product innovation

¹³ We accomplish this by testing the effect of corresponding control variables from the literature, respectively, Babina, Fedyk, He, and Hodson (2024a)), Zhang (2019), and Loughran and Ritter (2004).

(Aleksseeva, Giné, Samila, and Taska (2020) and Babina, Fedyk, He, and Hodson (2024a)) while flattening organizational hierarchies (Babina, Fedyk, He, and Hodson (2022)) and reducing hiring in non-AI positions (Acemoglu, Autor, Hazell, and Restrepo (2022)). Other studies find that, in high skill occupations, potential applications of AI change traditional work procedures (Grennan and Michaely (2021, 2020)). AI tools are also used to estimate latent corporate characteristics such as culture and climate exposure (Li, Mai, Shen, and Yan (2021), Li, Mai, Shen, Yang, and Zhang (2025), and Sautner, van Lent, Vilkov, and Zhang (2023)).

We, instead, focus on the production of AI innovation. We show that it increases a firm's future stock returns, reflecting a large, permanent *decrease* in risk (cash flow and stock return, systematic and idiosyncratic) and a small, transitory increase in profitability. Our results are robust to controlling for AI adoption as captured by the resume or job posting measures used in prior literature. In contrast to our AI production, AI adoption in the literature *increases* systematic risk (market beta increases) and decreases idiosyncratic risk, leaving total risk unchanged (Babina, Fedyk, He, and Hodson (2024b)). These complementary findings are consistent with opposite control strategies for various AI innovations: the producer firm keeping its technology for its own use, as opposed to the producer firm giving it away to other firms for adoption. We summarize the contrasting risk mechanisms in Section 7.3.3. There does not appear to be evidence on profitability in the prior literature.

Moreover, a recent literature on the productivity implications of AI innovation argues that AI may not enhance productivity as much as commonly expected, or its productivity enhancements will take much longer to materialize (Brynjolfsson, Rock, and Syverson (2019)). Our examination of AI production over three decades suggests that, at least for AI producer

firms, both forward looking stock prices and realized corporate operational outcomes reflect some productivity gains from AI innovation.

Our paper uniquely provides causal evidence, from all U.S. publicly traded firms over the past three decades, showing that, and how, producing AI innovation, incrementally to non-AI innovation, increases firm value. Other recent studies examine stock returns around particular events to uncover a moderating role on firm value of the firm's labor's AI exposure: Google's public launch of TensorFlow (Rock (2021)) and OpenAI's ChatGPT (Eisfeldt, Schubert, and Zhang (2023)). Earlier studies in the literature on corporate innovation and stock returns examine the predictive role of R&D intensity (Chan, Lakonishok, and Sougiannis (2001)), innovation efficiency and originality (Hirshleifer, Hsu, and Li (2013, 2018)), commercialization of R&D (Cohen, Diether, and Malloy (2013)), and firm size (Stoffman, Woepffel, and Yavuz (2022)).

Finally, we contribute a novel identification methodology for the production of AI innovation. We use a two variable interaction instrument for actual AI production: R&D capital stock induced by state-level R&D tax credits and industry-level AI exposure. Prior studies in this emerging literature use interaction instruments for local housing prices (Chaney, Sraer, and Thesmar (2012) and Adelino, Schoar, and Severino (2015)) and Chinese development finance (Dreher, Fuchs, Parks, Strange, and Tierney (2021)). The existing AI literature focuses on identifying the causal effect of AI adoption. To this end, Babina, Fedyk, He, and Hodson (2024a) exploit the AI research embedded in university alumni networks at the firms. Grennan and Michaely (2020) use news headline length to predict the usefulness of AI in stock analysis. Rock (2021) and Eisfeldt, Schubert, and Zhang (2023) use product launch events.

The rest of this paper is organized as follows. Section 2 describes the measurement of AI innovation, and Section 3 characterizes AI innovation. Section 4 presents the methodology.

Section 5 examines a significant motivation for AI production. Section 6 and Section 7 examine, respectively, the value implications and key value drivers of AI production. Section 8 and Section 9 examine the mechanisms underlying AI production and its financing implications, respectively. Section 10 concludes.

2. Measurement of AI Production

2.1. Measuring AI Production

To measure the production of AI innovation so that we can study its effects, we use patents that are classified as "AI patents" and to which we refer as such throughout the paper. As measures of innovation output, AI patent grants capture the capability of the firm to take commercial advantage of the AI technology that it produces. It can do so by implementing AI in its own operations, or supplying AI to its business counterparties, especially its customers, either directly (e.g., through patent transfers) or indirectly (e.g., embedded in product and services).

To classify patents in the USPTO database as AI and non-AI, we use the recently released classification of Giczy, Pairolero, and Toole (2022). Traditional methods of identifying specific technologies in patent documents are not well suited to identifying AI technology in patent documents. Perhaps the greatest difficulty with AI is that it is a general purpose technology and hence necessarily overlaps technology fields. Consequently, AI cannot simply be captured by a limited, predetermined set of widely used technology classes (e.g., CPCs) or keywords. While previous approaches like these (e.g., see Cockburn, Henderson, and Stern (2019)) tend to be correct about the patents that they identify as "AI", they also tend to miss a large number of patents that are in fact "AI".

As an improvement, Giczy et al. (2022) take a stratified machine learning approach. We provide a summary of their approach here, and refer the reader to Appendix 1 for a description of

the key details. First, AI is broken down into eight component technologies (e.g., knowledge processing and speech recognition). Next, a set of "surely AI" patents is identified as those that are at the intersection of four technology classification systems (CPC, IPC, USPC, and DWPI). Then, a set of "surely non-AI" patents is identified, after excluding patents that are even remotely related to the "surely AI" patents (e.g., through patent family links or citations) and technology classes with abnormally high share of "surely AI" patents.

A machine learning model is trained using the "surely AI" and "surely non-AI" patents, in several passes designed to minimize both false positives and false negatives in the subsequent application to the universe of patents. After training, the model subsequently evaluates all patent documents for their AI content, and assigns them a predicted probability of the patent containing a particular AI component technology. Finally, if the patent is predicted to be AI based on any AI component technology, it is classified as an AI patent.

For a classification of AI and non-AI patents to be accurate, it must naturally minimize both false positive (minimal patents classified as "AI" that are not AI) and false negatives (minimal patents classified as "non-AI" that are in fact AI). With both of these objectives in mind, Giczy et al. (2022) carefully test their patent classification and show that it outperforms the existing alternatives.¹⁴

2.2. Understanding AI Patents

To better understand the patents that we use to measure AI innovation, we provide diverse examples of AI technologies, firms, and industries that use such patents in processes and products. These examples illustrate that (patented) AI technologies include a wide variety of

¹⁴ The authors use four patent examiners at the USPTO, who are specialists in AI, to classify patents as AI or non-AI from 800 randomly selected patent documents. Each patent is reviewed by at least two examiners. If the first two examiners disagree, a third examiner adjudicates. Finally, the patent examiners' annotations are used to evaluate the validity of the authors' prediction model for false positives, false negatives, and a composite measure of the two. The authors' model is compared (and found superior to) existing alternative models.

machine cognition tasks, they can be used by both the producer firm and its customers, and they can be used across many industries. The examples, shown in Appendix Table 1, are inspired by the characterization of firm-level AI innovation in Section 3.2 (especially the ranking of industries based on AI patents in Table 1). We credit USPTO (2020) for several of our illustrative examples.

Unsurprisingly, AI is ubiquitous in the various technology industries. A highly visible example is the virtual assistant systems in consumer electronics produced by all the big technology firms. This includes Siri from Apple (e.g., USPTO patent number 10043516), which relies on speech recognition. AI hardware is also prominently produced and used in the industry itself. An example is IBM's device to improve computational efficiency modeled on the information processing structure of the biological brain (patent 8892487).

Many non-technology industries are also prolific producers and users of AI technology. In transportation, the automated driving systems used in motor vehicles (both personal and commercial) by Toyota and other manufacturers (patent 8384776) are powered by "knowledge processing" AI technology. Similarly, Boeing and other major aircraft manufacturers have been equipping airplanes many decades ago with automated flying systems that are dynamically updated based on historical experience (patent 3987279). In communications, Sprint uses "planning and control" AI technology to assess and improve signal quality for its network (patent 8140069). In drugs, Pfizer uses computer vision to both develop cancer treatments and to evaluate their performance on the body (patent 7231074).

Less traditionally innovative industries are also active producers and users of AI technology. In oil and gas exploration and development, Chevron uses "evolutionary computation" AI technology to estimate reserves (patent 7657494). Delta and other airlines use

AI to forecast unscheduled component orders and labor for repairs (patent 7370001). AI production occurs in even less obvious industries. In consumer products, Coca-Cola uses machine learning to dynamically optimize the operation of its dispensing machines (patent 4827426). Starbucks, among other mass market restaurants, uses AI to identify customers, predict their orders, and begin preparing them (patent 9218633). Overall, there is great diversity in the technologies captured by our sample of AI patents.

3. Characterization of AI Innovation

3.1. Aggregate AI Innovation

We begin our empirical analysis with a simple characterization of AI innovation in the aggregate during the period 1990-2020. Our findings below demonstrate the significance of AI as a unique type of innovation as well as the importance of understanding AI production and its implication for firm value.

First, we examine AI innovation activity, as captured by patent grants. Specifically, we measure innovation activity variously as: patent counts; the scientific value of patents, captured by the number of forward citations made to patents; and the commercial value of patents, captured by the estimates of the market value of patents made available by Kogan, Papanikolaou, Seru, and Stoffman (2017).

[Insert Figure 1 about here]

Figure 1 shows that AI is a prominent subset of all innovation activity. AI constitutes, very roughly, 5% of innovation activity in 1990. However, AI's share grows rapidly during the next three decades, accounting, by 2020, for over 15% of patents by number, 25% by scientific value, and 35% by commercial value.

Additionally, AI patents are also more valuable than non-AI patents, both scientifically and commercially. Even considering the rapid growth of patent counts, the value of the average AI patent is about 50% higher in 2020, both in terms of scientific and commercial value. By comparison, in 1990, the value premiums for scientific and commercial value are 200% and parity, respectively (relative value results not tabulated).

[Insert Figure 2 about here]

We also examine the rate of "breakthrough" innovations in AI versus non-AI technology. Specifically, we use the measure of, and data from, Kelly, Papanikolaou, Seru, and Taddy (2021), who use textual analysis to identify patent grants that are distinct from prior patents but related to subsequent patents (i.e., highly novel and also highly useful). The unconditional rate of breakthrough patents in our sample is roughly 12%. Figure 2 shows that, for AI patents, the breakthrough rate gradually decreased from around 60% during the 1990s, during the early years of AI technological development, to around 40% during the 2000s, and 25% during the early 2010s. However, even in the most recent period, AI innovations are roughly four times as likely to be breakthroughs as non-AI innovations.

[Insert Figure 3 about here]

Second, we examine the diffusion of AI innovation throughout the economy. We would expect to see evidence of widespread diffusion over time from a general purpose technology such as AI. This is indeed what we find in Figure 3, whether we examine the production of AI innovation itself or of innovation that uses prior AI innovation. Panel A shows that, by 2020, at least half of all industries had AI patent grants that accounted for at least 10% of all patent grants in the industry. Even if we consider only industries with at least a majority of AI patents in all patent grants, AI innovation dominated almost 20% of all industries.

Similarly, Figure 3 Panel B shows widespread diffusion of innovation using prior AI innovation, as captured by backward citations of patent grants to prior AI patents. By 2020, about 75% of all industries have patents that build on prior AI technology, if we require at least 10% of patent grants in the industry to cite a prior AI patent. Even if we require at least a majority of patent grants to cite prior AI patents, then prior AI technology is built upon by about 30% of all industries.

Finally, as a suggestive validity check, we compare AI and non-AI patents in terms of their "process innovation" content. Since AI is a labor enhancing technology, we would expect firms to produce AI innovations that improve the productivity of their operations or that of their customers. We can shed light on such innovation using data on the process intensity of patents from Bena and Simintzi (2025). A "process claim" represents an innovation in task performance, whereas a non-process claim represents other types of innovations, including but not limited to product innovations. We find that for AI patents, the share of process claims (relative to all claims) is roughly 50% on average, consistently during the past three decades. By contrast, for non-AI patents, the figure is only 30%. This is broadly consistent with AI patents focusing on improving task performance.

3.2. Firm-Level AI Innovation

[Insert Figure 4 about here]

We examine the importance of publicly traded firms in AI innovation as compared to all patenting entities because publicly traded firms account for a large share of aggregate R&D spending and patent grants. Figure 4 shows that publicly traded firms dominate AI innovation. Relative to all patenting entities (not just firms), publicly traded U.S. firms consistently account

for almost half of all AI patent grants (roughly 45% during the past two decades), compared to only one quarter of non-AI patent grants during the past three decades (Panel A).

Moreover, in the U.S., publicly traded firms, relative to all firms (rather than all patenting entities), account for an even greater share, close to two-thirds, of AI patents, compared to only 55% of non-AI patents (Panel B). Furthermore, we can restrict our sample to innovative public firms, i.e., U.S. public firms with at least one patent, to calculate the share of firms that produce AI innovation, i.e., at least one AI patent. Figure 4 Panel C shows that the proportion of innovative public firms that also produce AI innovation has risen from roughly 15% in 1990 to about 45% in the past decade.

In summary, publicly traded firms have historically, and continue today, to dominate the AI innovation. Additionally, innovative publicly traded firms increasingly include AI in their innovation activities. These findings motivate our focus on, and hence restriction of our sample to, U.S. publicly traded firms in the rest of the paper.

We examine the sensibility of our measure of AI production. We rank industries, from greatest to least, based on their total number of AI patents. We capture industries using three-digit SIC codes. We use all publicly traded firms in our baseline sample and all industries with at least 10 firms per year every year during our sample period.

[Insert Table 1 about here]

Table 1 shows a highly intuitive ranking of industries based on AI production. As one might expect, computer programming, electronic components, and computer equipment have the highest AI production. To illustrate the sensibility of our AI production measure, we rank firms based on average annual AI patent counts. The top 20 firms, tabulated in Appendix Table 3, are technology firms and widely known to be leaders in AI production. Meanwhile, Table 1 shows

that the lowest AI production is in operative builders, clothing stores, and equipment rentals. Importantly, however, even in the top quartile of industries, three of ten are not technology industries (aircrafts, motor vehicles, and drugs), and there are no technology industries in the second quartile (instead, e.g., assorted industries in petroleum and industrial equipment).

Our ranking in Table 1 is also broadly similar if, instead of ranking based on the total number of AI patents, we eliminate the industry size effect by ranking based the mean number of patents per firm. Furthermore, we observe that in some industries, AI production is dominated by a few firms with a disproportionately higher level of AI production than their industry peers. For example, 70% of the AI patents in petroleum refining are owned by Exxon Mobil and Chevron, collectively. Therefore, we exclude from each industry the three firms with the highest number of AI patents, and then rank industries with the remaining firms based on the number of patents and the average number of patents per firm, respectively. The rankings are once again similar.

Following the impression in Table 1 that AI producer firms tend to be big and old firms, we examine the size and age of AI producers to both firms that produce only non-AI innovation and also firms that do not produce any innovation. In every year during our sample period, AI producers are larger on average, twice as large or larger, whether we measure size using total assets, sales, or market capitalization (and even if we exclude the largest 10 or 20 firms in each group being compared). However, we do not find that AI producers are consistently and significantly older, on average.¹⁵

¹⁵ Additionally, we verify that our results throughout the paper are robust to excluding big technology firms. We classify firms as being in the technology industry using their SIC codes following Loughran and Ritter (2004). Various sorting firms based on total assets, sales, and market capitalization, we exclude the largest 20 technology firms. These firms, in all sorts, account for about 60% of all AI patents in our sample. We find that our baseline results are similar if we exclude big technology firms from our sample.

4. Methodology

4.1. Instrumentation of AI Production

While AI patent grants are an observable and straightforward measure of AI innovation, using them directly to study their effects raises potential endogeneity concerns. For instance, while an increase in AI patent grants may lead investors to increase their appraisal of firm value, a firm that anticipates an otherwise unrelated increase in its future value may also be better able to finance its R&D spending and may receive more future AI patent grants. In addition to such cases of reverse causality, omitted factors can generate an observed correlation between AI patent grants and various corporate operational outcomes. In short, endogeneity makes OLS estimates unreliable. For this reason, we do not interpret or draw inferences from our OLS results. Nevertheless, we do tabulate all baseline results implemented as OLS regressions, as we discuss in Section 7.

Our approach is to use an instrument that combines two key lagged components which, together, predict future corporate outcomes, resulting from current AI innovation and plausibly only through it. Our instrumental variable is the interaction of two components.¹⁶ Starting with the first component, in order to produce innovation (AI or non-AI), firms need to have sufficient innovation capacity to direct towards some specific technology such as AI. Empirically, firms with a larger stock of R&D capital are good candidates to invest in and successfully produce AI innovation. Second, the firm must have sufficient incentive to direct its innovation capacity towards AI technology. Since AI is a labor enhancing technology, in empirical terms, firms that are measurably more exposed to AI, through their own labor or that of their customers, are good candidates to produce AI innovation. We develop each of these two measures below.

¹⁶ Prior studies that use interaction instruments include: Chaney, Sraer, and Thesmar (2012); Adelino, Schoar, and Severino (2015); and Dreher, Fuchs, Parks, Strange, and Tierney (2021).

To measure the first component of our instrumental variable, we use the user cost of R&D, implied by time-varying federal and state R&D tax credits, to predict the R&D spending of firms from 1988 to 2015. Specifically, using a panel of firm-years, we predict R&D expenditures by regressing R&D expenditures on the firm's annual user cost of R&D along with firm and year fixed effects. We calculate the firm's R&D user cost as the weighted average of R&D user cost across the firm's R&D hubs, i.e., the states in which its inventors are located, during the previous 10 years. If the firm does not have any patents during this period, we calculate the firm's R&D user cost based on its headquarters location.¹⁷ We then capitalize predicted R&D expenditures for each firm during the previous 10 years at a depreciation rate of 15%. This R&D capital stock is our measure of the firm's plausibly exogenous innovation capacity. Data on the user cost of R&D are from Bloom, Schankerman, and Van Reenen (2013), and our methodology is similar to that of Wilson (2009), Bloom et al. (2013), Hombert and Matray (2018), and Babina and Howell (2024).

We then turn to the second component of our instrument, AI exposure. AI exposure refers to the potential of AI to substitute or complement labor. We measure "AI exposure" at the industry level by calculating the weighted average occupation-level AI exposure using as weights the occupational employment shares within the industry. Occupational AI exposures data are from Felten, Raj, and Seamans (2021), and occupational employment shares data are from the Bureau of Labor Statistics. We describe in detail the construction of occupational AI exposure scores in Appendix 2. Felten et al. (2021) validate their measure by studying job postings data (from Burning Glass Technologies). They find that occupational AI exposure predicts higher AI skill requirements in job postings for the corresponding occupation. Providing

¹⁷ Derrien, Kecskés, and Nguyen (2023) document that, for firms with available data on inventor location, roughly half of inventors are located in the same commuting zone as the firm's headquarters, with a predictably higher share located in the same state.

further validation, Acemoglu, Autor, Hazell, and Restrepo (2022) find that AI exposure aggregated to the establishment level predicts higher AI hiring.

We measure a firm's AI exposure as its industry's labor's exposure to AI. Industries are captured using three-digit SIC codes. We fix employment weights in the 1988-1990 period, before the start of our sample period, to minimize the potential endogeneity of time-varying employment shares. These data first become available in 1988, and only one third of all industries (non-overlapping) are populated in each year during the first three years.¹⁸ We illustrate the sensibility of our AI exposure measure using the most dominant industry based on AI production (see Table 1): computer programming (SIC 737). Appendix Table 4 Panel A shows that the top 20 occupations, ranked by employment share, typically have high AI exposures, with an average exposure percentile of 93.

To measure a firm's customers' AI exposure, we use the purchase share weighted average AI exposure of the firm's industry's customer industries. Specifically, for each industry, we obtain all customer industries from the Bureau of Economic Analysis industry input-output tables along with the product purchase share of each customer industry, i.e., how much of a given industry's products are sold to every possible customer industry. We then calculate, for each industry, the purchase share-weighted average of the AI exposures across customer industries. We again fix product purchase shares before our sample period, in 1987. As an illustration of our customer AI exposure measure, consider once again the most dominant industry based on AI production: computer programming (SIC 737). Appendix Table 4 Panel B shows that the top 20 customers of the computer programming industry, ranked by product

¹⁸ To examine the possibility of firms time-varyingly choosing their industry, and hence AI exposure, endogenously with firm outcomes, we use AI exposure based on industry fixed, alternatively, at the time of firms' first or last appearance in our sample. Our results are robust to both alternatives.

purchase share, are a diverse mix of industries. Computer programming itself has high AI exposure, but so do its typical customer industries, with an average exposure percentile of 97.

We interact these two components – innovation capacity and AI exposure – to construct an interaction instrument. When R&D spending increases (because the user cost of R&D decreases as a result of time-varying federal and state R&D tax credits), firms with greater AI exposure (whether of their own labor or their customers') are more likely to produce AI innovation because such firms benefit more from the labor enhancement of AI technology. We lag our instrument by two years relative to AI patent counts to reflect the time it typically takes for patents to be granted.

Since we can measure both the firm's own AI exposure and that of its customers, we construct two corresponding interaction instruments. We use both interaction instruments together in our baseline analyses because we have no theoretical reason to prefer one over the other, and we can increase the precision of our estimates by using both instruments together. However, we verify that our results are similar if we use each of our two interaction instruments separately. In all main analyses, we report the Hansen J-test for whether the estimated effects of AI patent counts are significantly different using each instrument separately. None of the differences is significant. Additionally, we tabulate all baseline results implemented using each instrument separately.¹⁹ As we discuss in Section 7, our estimates are similar in magnitude.

Our identifying assumption is that firms with different AI exposures will only be affected differentially by changes in the (tax credit induced) R&D capital stock through the impact on AI innovation. To ensure that we identify exclusively off our interaction instrument, all regressions

¹⁹ See Atanasov and Black (2016) for a discussion of this approach, and Angrist and Evans (1998) for an applied example.

directly control for the components of the interaction, i.e., (tax credit induced) R&D capital stock as well as AI exposure.

A fortiori, we include fixed effects for state-years based on the location of the firm's headquarters as well as fixed effects for three-digit SIC industries. This we do so that our results, more broadly, cannot be explained by potentially confounding factors. State-year fixed effects largely absorb R&D capital stock (because innovation activities are concentrated at firm headquarters), but they additionally absorb all commonalities across geographically proximate firms. Similarly, industry fixed effects entirely absorb AI exposure, but they also absorb all additional commonalities across firms competing in proximate product markets.

4.2. Model Specification

Our main analysis begins with examining stock returns for AI producer firms at the portfolio-month and firm-month level. Our analysis then proceeds to the firm-year level, where we regress corporate outcomes (such as cash flow levels) on instrumented AI patent counts.

The first stage of our instrumental variable regressions is as follows:

$$\begin{aligned} \ln(0.1 + AI_Patent_Counts_{i,SIC2,SIC3,s,t}) = \\ \alpha_1 \cdot R\&D_Stock_{i,s,t-2} \times Firm's_AI_Exposure_{SIC3} + \\ \alpha_2 \cdot R\&D_Stock_{i,s,t-2} \times Customers'_AI_Exposure_{SIC3} + \\ \alpha_3 \cdot R\&D_Stock_{i,s,t-2} + \beta \cdot X_{i,t} + \delta_{s,t} + \delta_{SIC3} + \delta_i + \delta_{SIC2,t} \end{aligned} \quad (1)$$

The AI patent counts predicted from the first stage are then used to explain outcomes in the second stage of our instrumental variable regressions:

$$\begin{aligned} Outcome_{i,SIC2,SIC3,s,t+1} = \alpha \cdot \ln(0.1 + AI_Patent_Counts_{i,t}) + \\ \beta_1 \cdot R\&D_Stock_{i,s,t-2} + \beta_2 \cdot X_{i,t} + \delta_{s,t} + \delta_{SIC3} + \delta_i + \delta_{SIC2,t} \end{aligned} \quad (2)$$

In the equations above, i indexes firms, $SIC2$ and $SIC3$ index two-digit and three-digit SIC industries, respectively, s indexes the state of the firm's headquarters, and t indexes year. $X_{i,t}$ is a vector of firm-level control variables. The parameters δ_i , $\delta_{SIC2,t}$, δ_{SIC3} , and $\delta_{s,t}$ are fixed effects, respectively, for firms, two-digit SIC industry-years, three-digit SIC industries, and state-years. Fixed effects for three-digit SIC industries completely absorb the direct effects of AI exposure (both the firm's and its customers'), so AI exposure is dropped. State-year fixed effects are based on the headquarters location of the firm.

By way of justification, our baseline specification includes a battery of control variables and fixed effects to ensure that we identify exclusively off our interaction instrument and not its components. The components of our interaction instrument, which we use as control variables, we discuss in Section 4.1. A fortiori, we include fixed effects for state-years based on the location of the firm's headquarters as well as fixed effects for three-digit SIC industries. This we do so that our results, more broadly, cannot be explained by potentially confounding factors. State-year fixed effects largely absorb R&D capital stock (because innovation activities are concentrated at firm headquarters), but they additionally absorb all commonalities across geographically proximate firms. Similarly, industry fixed effects entirely absorb AI exposure, but they also absorb all additional commonalities across firms competing in proximate product markets.

To further ensure that generic innovation is not driving our results, we control for the number of non-AI patent grants as well as an innovation dummy variable for whether the firm has at least one patent granted during the previous year. Additionally, we control for total assets and firm age to account for the possibility that larger and older firms are more likely to invest in and adopt advanced technologies. We also include firm fixed effects to rule out the possibility

that time-invariant differences across firms can explain our results. Finally, we include fixed effects for industry-years (using two-digit SIC industry) so that our results cannot be explained by time-varying industrial commonalities.

Finally, in our baseline specification, we cluster standard errors by firm and also by industry-year (using two-digit SIC industry), since firms in similar lines of business tend to behave similarly. Before taking the logarithm of a variable that takes on zero values, we add a constant approximately equal to a small increment of the values of the variable. We indicate these constants in the corresponding results and/or Appendix Table 2. We verify that our results are robust to adding a constant at least one order of magnitude higher or lower. We add a smaller increment of 0.1 to AI patent counts before taking logarithms, rather than 1 as for non-AI patent counts, because firms have roughly one order of magnitude fewer AI patents than non-AI patents. To facilitate comparison across the two AI exposure (the firm's own and its customers'), we standardize them to mean zero and standard deviation one. We winsorize variables whenever appropriate at the 1st and 99th percentiles.

4.3. Sample and Descriptive Statistics

The firms in our sample are publicly traded U.S. operating firms excluding financials and utilities. The data on firms are from CRSP and Compustat. The sample period spans 1990-2017 in terms of year t . We measure AI production using AI patent grants during the 12 months before each fiscal yearend. We start our sample period in 1990 because by then there is a critical mass of AI patent grants each year. We are also limited by the need for 10 years of patent data to construct the tax credit induced R&D stock, which requires inventor locations going back to at least 1978 (i.e., for R&D stock in 1988).

AI patent counts are measured in year t . They are instrumented with R&D stock measured at year $t-2$ and AI exposure fixed before the start of the sample period. Our data on federal and state user cost of R&D end in 2015, which is the last year we are able to calculate the (tax credit induced) R&D stock (and hence AI patent counts in 2017). Outcomes are measured in year $t+1$. Since we need Compustat data from years $t-2$ to $t+1$, we effectively require at least four years of Compustat data for each firm-year. Ultimately, the sample comprises 93,544 firm-year observations from 1990 to 2017 corresponding to 10,362 unique firms.

[Insert Table 2 about here]

Table 2 provides descriptive statistics for the variables used in this paper. Variables are defined in Appendix Table 2. In any given year, on average, 33% of firms have at least one patent grant of any kind, and 10% have at least one AI patent grant (not tabulated). In the average firm-year, AI patent counts are 0.66 compared to 6.5 for non-AI patents, a tenfold multiple.

5. Motivation for Producing AI Innovation and First Stage of IV Regressions

We begin our firm-level analysis by examining a potentially significant motivation for which firms produce AI innovation. Our underlying intuition is as follows. For some firms, their own labor (or their customers') is more substitutable for or complementable by AI technology. If such firms, with higher AI exposure, experience an exogenous increase in their innovation capacity, they will produce more AI innovation. At the same time, this intuition is also the econometric framework in subsequent analyses for our identification of the effect of producing AI innovation. This, the first stage of our instrumental variable regressions, is based on Equation 1.

[Insert Table 3 about here]

Table 3 Panel A presents the results for the regressions of actual AI production, measured by AI patent counts, on our instrumental variable, the interaction of R&D capital stock and AI exposure. Column 1 supports a mutually enforcing effect of the producer firm's innovation capacity, measured by its (tax credit induced) R&D capital stock, and its own AI exposure. Column 2 supports a similar effect, in economic magnitude and statistical significance, when the producer firm's AI exposure is replaced by its customers' AI exposure.

We then combine the two related motivations for AI production by including both instruments in our specification. Panel A Column 3 shows that, overall, each instrument remains economically and statistically significant alongside the other. For a one standard deviation increase in each of the two mutually reinforcing incentives for a firm to produce AI innovation, i.e., (the logarithm of) R&D stock and both AI exposures (mean zero, standard deviation one), AI patent counts increase by 15% ($=2.3 \times (0.024 \times 1 + 0.040 \times 1)$) relative to its mean. Alternatively viewed, this is the estimated magnitude of the reinforcing effect of AI exposure on a given change in R&D stock, and vice versa. We use the specification in Column 3 (both instruments together) in our baseline IV regressions. The results are stronger for the "customers instrument" than the "firm instrument".²⁰ However, our second stage results in this paper do not depend critically on whether we use one instrument, the other, or both together.

We examine the extent to which the results depend on any of the eight component technologies of AI. For firm-years with at least one AI patent, the average patent counts corresponding to each AI component technology is as follows (each AI patent can have multiple

²⁰ This can happen if AI technology that enhances the firm's labor factor of production also enhances its customers' labor factor. As a test, we can mechanically remove the firm's AI exposure that overlaps its customers' AI exposure (since our input-output data indicate positive intra-industry product purchases for most industries), at the expense of a less accurate measure of customers' AI exposure. Specifically, we can exclude the firm's own industry from the list of customer industries before calculating the purchase share-weighted average AI exposure of the firm's customers. In this case, our coefficient estimate decrease in magnitude for the customers instrument, and it increases for the firm instrument (with the firm instrument's t-statistic rising to 2.2) (results not tabulated).

components): knowledge processing 3.43, speech recognition 0.17, AI hardware 2.01, evolutionary computation 0.16, natural language processing 0.35, machine learning 0.78, computer vision 1.21, and planning/control 3.45. For comparison, the average AI patent count for firm-years with non-zero AI patents is 6.5. For each AI component technology individually, we redo Table 3 Panel A (results not tabulated). We find similar results for the individual components compared to our baseline aggregation of all components, especially for the AI component technology with the largest share of AI patents.

To calculate the typical variation in AI production induced by our interaction instrument, we start with our estimates in Table 3 Panel A. Let us fix the logarithm of R&D stock at its mean of roughly 2, and increase AI exposure by one unit (i.e., one standard deviation), for both the firm and its customers. This increases AI patent counts by roughly 13% ($=2 \times (0.024 \times 1 + 0.040 \times 1)$) relative to its mean (Column 3), which we approximate as 10% for ease of interpretation. We use this figure throughout the paper to calculate the estimate effect of a typical change in our instrument on corporate outcomes of interest.

Finally, we emphasize that we are interested in the amount of AI innovation produced by firms, so we use AI patent counts as our baseline measure. However, we also consider alternative variables for scaling AI patent counts. These scaling variables include total assets, total patent stock, total patent counts, and R&D stock (defined in Appendix Table 2). The results using these alternative scaling variables, tabulated in Table 3 Panel B through Panel E, are similar to using unscaled AI patent counts.

6. Value Implications of AI Production for the Producer Firm

We examine the value implications of AI production for the producer firm using realized stock returns. We defer examination of the key drivers of value (i.e., cash flow levels and cash flow risk) to the next section.

6.1. Initial Market Reaction to AI Patent Grants

We begin by examining the initial market reaction to AI patent grants. Patent grants are generally announced by the USPTO once a week, on Tuesdays. We perform perhaps the simplest possible analysis: for every firm and patent grant date pair ("firm-date" hereafter), we calculate the market reaction for the firm-date during the week that starts with the patent grant date. Since there can be multiple patent grants for a given firm-date, we divide the market reaction by the number of patents. Thus far, we have not distinguished between different types of patents, but going forward we only consider the subset of AI patents.

During the first week, the return for the average AI patent is about 1.2 basis points and about 6.4 bps for the average firm-date. During the first four weeks after the average firm-date, cumulative returns increase to roughly 33 bps, then to 73 bps after eight weeks, and to 117 bps after 12 weeks. However, this apparent return drift after patent grants is contaminated by overlapping events. Within one week of a firm-day with at least one AI patent grant, there is a roughly 50% chance of at least one AI patent grant within one week. Within four weeks, this probability rises to about 75%, and then to 85% within 12 weeks. Furthermore, each event-date with at least one patent grant (AI or non-AI) is contaminated by having a different number of AI and non-AI patents.

Our only conclusion thus far is that the initial market reaction to AI patent announcements appears to be economically small compared to the subsequent return drift.

However, the statistical and economic problems with short-run returns after weekly patent grant announcements lead us to focus instead on long-run returns after a period of accumulating AI patent grants. Specifically, we count the number of AI patents for one year, and then we use them to predict returns during the following year. The first and most obvious benefit of this approach is better capturing the scale of the firm's AI innovation activities. With the exception of a few firms with big innovation programs, the presence or absence of a patent in a given week is likely to be a noisy measure of the years-long cycle for a given innovation program, let alone the multiple staggered innovation programs of a given firm. Second, a single patent granted one week is unlikely to capture the synergies across the firm's various innovation programs as well as its assets in place. Third, weekly patent grant announcements are persistent within firms, so multi-week event returns likely inflate the value of early patents by the initial market reaction to later patents, even if investors fully impound the value of a patent within one week of its grant announcement. Finally, a large literature suggests that corporate innovation activities tend to be undervalued by investors (e.g., Chan, Lakonishok, and Sougiannis (2001); Hirshleifer, Hsu, and Li (2013, 2018); and Cohen, Diether, and Malloy (2013)). If this is the case, then the initial market reaction to a patent grant likely underestimates its value to the firm.

By accumulating AI patent grants over a year, we improve the statistical properties of our measure of AI innovation. If anything, we are still likely to underestimate long-run returns due to staleness of our one year accumulation of AI patent grants (at a minimum, we ignore the first 6-18 months of post-announcement drift, depending on whether the patent was granted at the beginning or end of the year). However, our long-run approach allows to address potential confounds with control variables and fixed effects. Our approach also allows us to instrument for

AI production. Neither is possible for short-run returns after patent grant announcement events (because we do not have a daily frequency instrument).

6.2. Returns on Portfolios Sorted by Actual AI Patent Counts

We then examine the returns on portfolios formed based on actual AI patent grants, which are readily observable to investors. For every firm, for every calendar year, we count the number of AI patents granted during the 12 months ending in the month of the most recent fiscal yearend date. At the end of June of the following calendar year, we form portfolios based on AI patent counts. Therefore, we begin using returns with at least a six month lag (for firms with a Dec. fiscal yearend) and up to a 17 month lag (for firms with a Jan. fiscal yearend). These timing differences result from consistently using the same baseline sample construction throughout the paper.²¹ We hold portfolios from July through June of the following calendar year (12 months), at which point we rebalance. Since we observe AI patent grants from 1990 to 2017, we examine returns from July 1991 to June 2019, for a total of 336 monthly observations for each portfolio.

We next describe our portfolio sorts, which effectively create a zero total (AI plus non-AI) innovation group (Q0) and four quartiles based on AI patents (Q1 through Q4). Specifically, we sort firms into four quasi-quartiles (labeled as such because they contain an unequal number of firms): the group of zero AI patents ("zero AI"), and three groups of non-zero AI patent sorted into terciles ("low AI", Q2; "medium AI", Q3; and "high AI", Q4). Terciles (to construct groups Q2 through Q4) are recalculated every year at the time of portfolio formation. We are limited to sorting into these quasi-quartiles because only 10% or so of firm-years have non-zero AI patent counts.

Within the group of firms with zero AI patents, we consider distinguishing between zero innovation and non-zero innovation firms as measured by total (AI plus non-AI) patents. This is

²¹ In our returns analysis, we drop stocks with negative book-to-market ratios and stocks with prices lower than \$1.

because about two-thirds of our sample firms have zero total patents, and of the remainder, a large majority further has zero AI patents. Therefore, we further sort the zero AI patents quartile (defined as having zero AI patents) into two groups: zero total patents (Q0) and non-zero total patents (Q1). We thus have a total of five portfolios to examine. We also form hedge portfolios that long the high AI portfolio (Q4) and, variously, short the low AI portfolio (Q2) or either of the "zero AI" portfolios (Q1 and Q0).

Our portfolios weightings are threefold: equally weighted, value weighted, and size neutral. We use the size neutral approach as our baseline to mitigate the correlation between our sorting variable, AI patent counts, and firm size. As the calculation of size neutral returns demonstrates, this approach balances the equally and value weighted approaches so that returns are neither driven by the smallest or the largest firms.²² We calculate size neutral returns, for any arbitrary portfolio, as follows. We sort stocks in the portfolio into small and large groups, independently, based on the NYSE median size breakpoint. We then value weight stocks within each group within the portfolio, and calculate value-weighted returns for the small group separately from the large group. Finally, we take the simple average of the returns of the small and large groups. This is the size neutral return for the particular portfolio.

[Insert Table 4 about here]

Table 4 presents the results of these time-series portfolio return regressions. We use the Fama and French (2015) five-factor model as our baseline, but the results are robust to using alternative factor models (Section 6.4). The results indicate that our baseline AI patent counts spread returns (Panel A), by about 50 basis points per month in our baseline size neutral specification, for AI patent counts moving from Q1 to Q4. We conservatively choose Q4-Q1

²² For other applications of this approach, see Griffin and Lemmon (2002); Hirshleifer, Hsu, and Li (2013); and Liu, Stambaugh, and Yuan (2019).

(high AI minus zero AI) as our baseline hedge portfolio. Compared to our baseline, spreads are somewhat lower when the short leg of the AI hedge portfolio is firms with low (but non-zero) AI patents (Q4-Q2). If the factor models that we use capture differences in systematic risk between non-zero innovation firms and zero innovation firms, then it is instructive to use zero innovation firms as the short leg of the AI hedge portfolio. In this case (Q4-Q0), spreads are somewhat higher than the baseline hedge portfolio (Q4-Q1). The results are similar if we use alternative scaling variables for AI patent counts (Panel B to Panel E).

Additionally, we compare returns for AI patents to non-AI patents. Table 4 provides a direct comparison using the relative frequency of AI patents to non-AI patents (to be precise, AI patent counts divided by total patent counts). The results in Panel D are similar to the baseline results.

As an alternative but less direct comparison of AI and non-AI, we redo Table 4 sorting firms based on their non-AI patents analogously to our AI patent sorts. The zero innovation group (Q0) is the same as for AI patents, and non-AI patent counts are sorted into four quartiles. Internet Appendix Table 1 presents the results. Comparing the baseline hedge portfolios excluding Panel D,²³ the return spread is higher for AI patents than non-AI patents, with one exception, and is generally higher by 20 basis points per month or more. Comparing high AI (Q4) to high non-AI (Q4) portfolios, the return spread is even bigger and always positive. Moreover, while the AI spreads are reliable in economic and statistical significance, the non-AI spreads are only reliably positive and significant for non-AI patent counts scaled by total assets (Panel B). Taken as a whole, the results suggest that firms with observable higher AI production have higher risk-adjusted returns.

²³ The similarities across the two Panel Ds demonstrate the comparability of the two tables since the non-AI sort is simply the reverse of the AI sort.

We also examine the factor loadings for our hedge portfolios in Table 4. We do not find consistent loadings on any of the factors in our baseline Fama and French (2015) five-factor model (not tabulated). The market factor loading is almost never significant. The size factor loading is occasionally significant, variously on small or big stocks. When the other factor loadings are significant, they tend to characterize the portfolios as growth (rather than value), weak profitability (rather than robust), and aggressive investment (rather than conservative). The only extent to which there is any consistency between Table 4 (sorts based on actual AI patent counts) and Table 5 (sorts based on R&D stock and AI exposure) is in the loadings on value and weak profitability.

6.3. Returns on Portfolios Double Sorted by R&D Stock and AI Exposure

Studying how AI patent grants spread returns has the advantage of using a simple and observable measure of AI production. However, the disadvantage is that AI patent grants are endogenous to corporate outcomes. For instance, while higher future returns may result from innovation output, investor anticipation of innovation output can lower financing costs and thereby further increase the success of innovation efforts.

Motivated by our baseline IV framework, we also take the approach of examining the returns on portfolios formed based on the two components of our interaction instrument. Inspired by the reduced form of our IV regressions, we sort stocks into portfolios based on R&D stock and AI exposure, which allows us to identify the plausibly causal effect of these IV components on returns. Our approach is more complex than spreading returns with AI patent grants, but it can be implemented (information obtained and spreads traded) by sophisticated investors. At the same time, we are mindful of limitations of this quasi-reduced form approach, and we interpret the results suggestively.

Our reduced form approach is analogous to our previous approach. We consistently use the same baseline sample construction throughout the paper. We still use information that is available at the end of a given calendar year (year t), and we form portfolios at the end of June of the following calendar year (year $t+1$). However, instead of using information on actual AI patent counts (from year t) to form portfolios, we use information available on R&D stock and AI exposure. Since R&D stock is lagged by two years relative to AI production, it is measured in calendar year $t-2$. AI exposure is fixed before our sample period.

We sort firms into three quasi-terciles (containing an unequal number of firms) based on (tax credit induced) R&D capital stock, the first component of our interaction instrument. These resulting groups contain zero R&D stock ("low R&D", L), and two halves of non-zero R&D stock ("medium R&D", M; and "high R&D", H) recalculated every year at the time of portfolio formation. Independently, we also sort firms into terciles based on AI exposure (T1 through T3), the second component of our interaction instrument, at the industry level. However, since "AI exposure" comprises the respective AI exposures of the firm and its customers, we need to combine them so that we can sort on a single exposure measure. We do so by taking their first principal component and using it as our measure of AI exposure in our baseline returns analyses. Our portfolios of interest are those at the intersection of the double sorts on R&D stock and AI exposure.

[Insert Table 5 about here]

Table 5 presents the results of portfolio return regressions implemented in a quasi-reduced form setting. The Fama and French (2015) five-factor model is again the baseline, but the results are robust to alternatives (Section 6.4). Our interest is in the spread of the spread, and we interpret the results for our baseline size neutral portfolios (Panel C) as follows. We consider

R&D stock moving from low to high together with AI exposure moving from T1 to T3. The results for our baseline size neutral portfolios (Panel C) show that these changes result in higher returns of, very roughly, 50 basis points per month.

We can also infer the AI patent counts corresponding to this return spread by using the results of Table 3. The coefficient estimate on the interaction instrument is approximately 0.06 (Table 3). Let us consider the same increases in R&D stock (low to high, equal to roughly 5 units of $\ln(1+\text{R\&D stock})$) and AI exposure (T1 to T3, equal to about 2 standard deviations) as above. Therefore, the increase in AI patent counts corresponding to a 50 bps/month increase in returns (Table 5 Panel C) is roughly 0.6 units of $\ln(0.1+\text{AI patent counts})$, or a 60% increase relative to the mean ($=0.06 \times 5 \times 2$).²⁴ We are careful to interpret these results suggestively, and we are not comparing them directly to the returns results for AI patent counts (Table 4). However, these results do suggest that innovation capacity and AI exposure, both of which positively affect AI production (Table 3), result in higher risk-adjusted returns.

Once again, we also examine the factor loadings for our hedge portfolios. In Table 5, the only consistent factor loading is growth (rather than value), and somewhat consistently weak profitability (rather than robust). This is, again, the only extent of consistency with Table 4. In Table 5, when significant, the market factor loading tends to be negative, and the size factor positive. The investment factor is never significant.

6.4. Robustness Tests for Portfolio Returns Analyses

We directly eliminate a potentially confounding correlation between AI and non-AI innovation in Fama-MacBeth regressions of monthly stock returns as well as panel regressions throughout the paper. We do so by controlling for various measures of innovation outputs (e.g.,

²⁴ Simply as a reference point, moving from T1 to T3 in Table 4 equals about 5 units of $\ln(0.1+\text{AI patent counts})$.

non-AI patent counts) and inputs (e.g., R&D spending). It is not possible to be as rigorous in portfolio regressions.

We also examine the robustness of our results with respect to alternative factor models proposed in the literature. As tabulated in Panel A through Panel E of both Internet Appendix Table 2 (c.f. Table 4 Panel A) and Internet Appendix Table 3 (c.f. Table 5), we find that the AI return spread remains economically and statistically significant in more demanding factor models, such as the Fama and French (2015) five-factor model with momentum and the Hou, Xue, and Zhang (2015) Q-factor model. In less demanding models, with fewer factors, the results are weaker, which suggests that AI portfolios have less systematic risk as captured by canonical risk factors. It would also be consistent with a firm's total risk being lower as a result of AI innovation (which we also document, in Section 7.2).

Finally, using separately each of our two interaction instruments (i.e., based on the firm's AI exposure versus that of its customers), we redo Table 5 and Internet Appendix Table 3 (both double sorted by R&D stock and AI exposure). We find that our inferences are similar.

6.5. Fama-MacBeth Cross-Sectional Return Regressions

6.5.1. Fama-MacBeth OLS and IV in Year $t+1$

We examine the effect of potentially confounding variables on our estimates of risk-adjusted returns following AI production. We run Fama-MacBeth cross-section returns regressions using the same sample of firm-months that we use in our time-series returns analyses. We implement Fama-MacBeth regressions corresponding to both our portfolio regressions: sorted by actual AI patent counts (Table 4), and double sorted by (tax credit induced) R&D capital stock and AI exposure (Table 5) but implemented here using instrumented AI patent counts. Our explanatory variables of interest are actual and instrumented AI patent

counts in OLS and IV regressions, respectively. We use the respective AI exposures of the firm and its customers together as our baseline measure.

Our OLS and IV implementations of Fama-MacBeth differ as follows. In OLS Fama-MacBeth, we run cross-sectional regressions every month of returns on actual AI patent counts. In IV Fama-MacBeth, we first run cross-sectional regressions every month of actual AI patent counts on our instrumental variables, the interaction of R&D capital stock and each AI exposure measure. We also control for the components of our interaction instrument. We then run the second stage regression for the corresponding month, regressing returns on instrumented AI patent counts. The rest of the Fama-MacBeth procedure is the same for the OLS and IV implementations.

Our battery of control variables includes non-AI patent counts and R&D spending. We also include our innovation dummy variable. We control for variables commonly used in the literature as well as our IV regressions: market capitalization, market-to-book of equity, momentum, short-term reversal, return on assets, capex-to-total assets, stock price, and firm age. As an alternative to unscaled non-AI patent counts and R&D spending, we also include these variables scaled by total assets. Finally, we include fixed effects for industries using the Fama and French 48 industry classification.

[Insert Table 6 about here]

Table 6 shows that, in our panel regressions, actual AI patent counts (OLS) are not consistently significant (Panel A). By contrast, instrumented AI patent counts (IV) are economically and statistically significant (Panel B).²⁵ For a typical 19% increase in AI patent

²⁵ In the first stage of the IV Fama-MacBeth regressions (not tabulated), for a one standard deviation increase in each of R&D stock and AI exposure, AI patent counts increase by 22% ($=2.3 \times (0.029 \times 1 + 0.068 \times 1)$) relative to its mean). This compares to a 15% increase in Table 3 Panel A Column 3.

counts relative to its mean, returns increase by roughly 8 basis points per month.²⁶ As before, this estimated magnitude can be viewed alternatively as the reinforcing effect of AI exposure on a given change in R&D stock, and vice versa. The results are similar if we use each of our two interaction instruments separately (Internet Appendix Table 4).

Recognizing that the persistence of AI production within firms may be correlated with future returns, we additionally control for the mean monthly return during the previous 60 months. This control variable serves as an estimate of the firm's expected future return during the following month. Once again, our results are similar.

6.5.2. Fama-MacBeth IV in Year $t+2$ to Year $t+5$

We also examine the duration of the effect of instrumented AI patent counts on returns. We use the same IV implementation of Fama-MacBeth as before. However, to ensure that we do not attribute future returns to future AI patent grants, we additionally control for the potentially confounding future AI patent grants (unscaled or variously scaled), e.g., during year $t+1$ for returns in year $t+2$, ... , during year $t+1$ to year $t+4$ for returns in year $t+5$. We also control for AI patent stock during year $t-10$ to year $t-1$.

Table 6 Panel C shows that returns are statistically significant until and including year $t+3$. For a typical increase in AI patent counts relative to its mean, returns increase by roughly 5-6 basis points per month in year $t+2$, and then by 7-11 bps in year $t+3$. For AI patent counts, our baseline measure, returns are, respectively, 8, 6, and 7 basis points per month in year $t+1$ to year $t+3$, and thus about 7 bps/month over three years. Beyond that, however, the results become unreliable in terms of both economic and statistical significance.

²⁶ For the IV Fama-MacBeth regressions, we calculate the typical variation in AI production induced by our interaction instrument analogously to our calculations for Table 3 Panel A Column 3 in Section 5. We fix the logarithm of R&D stock at its mean of roughly 2, and we increase AI exposure by its standard deviation, for both the firm and its customers. This increases AI patent counts by roughly 19% ($=2 \times (0.029 \times 1 + 0.068 \times 1)$) relative to its mean.

6.6. Comparison of Magnitudes of Panel Returns and Portfolio Returns

Finally, we compare the magnitudes of the returns estimated in the panel regressions in Table 6 and the corresponding portfolio regressions in Table 5. The calculations above for Table 6, which use a one standard deviation increase in each of R&D stock and AI exposure, are unlike those in Table 5. In the latter, R&D stock increases from low to high, and AI exposure increases from T1 to T3. These changes in Table 5, converted to their equivalent magnitudes in Table 6, are roughly equal to 5 and 2 units of our R&D stock and AI exposure variables, respectively. Their effect is roughly a doubling of AI patent counts ($=5 \times (0.029 \times 2 + 0.068 \times 2)$) in the first stage IV. Therefore, in the second stage IV (Table 6 Panel B Column 1), returns increase by roughly 42 basis points per month ($=0.433 \times 0.97$). This is similar to the risk-adjusted returns in Table 5, even without remarking on the battery of control variables included in Table 6.

7. The Effect of AI Production on the Key Drivers of Firm Value

We argue that AI technology provides the firm with better information acquisition, forecasting, monitoring, decision making, execution, and responsiveness to changing business conditions. Consequently, producing AI innovation can increase firm value through both higher cash flow levels and lower cash flow risk. In the following empirical analysis, we use various measures of profitability and risk, and we regress them on instrumented AI patent counts. The second stage of our instrumental variable regressions is based on Equation 2. We address potential threats to identification, for both our profitability and risk analyses, in Section 7.3, and tabulate the results in Internet Appendix Table 5 and Internet Appendix Table 6, respectively. We summarize our collective inferences in Section 7.4. We conclude with analyses of the time-varying effects of AI production and other clarifications.

7.1. The Effect of AI Production on Cash Flow Levels

[Insert Table 7 about here]

We measure profitability using return on assets. We find some evidence that AI production increases cash flow levels but only transitorily. Table 7 Panel A shows that a 10% increase in AI patent counts relative to its mean (i.e., a typical increase) increases return on assets by 0.7 percentage points, which corresponds to roughly 2.5% of the dependent variable's standard deviation. However, this effect in year $t+1$ decreases each successive year (roughly halving annually), and becomes insignificant by year $t+2$.

By contrast, using endogenous (uninstrumented) AI patent counts, we find no significant effect of AI production on the cash flow levels or cash flow risk. Indeed, we redo all IV regressions implemented as OLS regressions, and tabulate the results in Internet Appendix Table 7 through Internet Appendix Table 12 (corresponding to Table 7 through Table 13, respectively). In contrast to our IV estimates, our OLS estimates are generally much less significant, economically and statistically.

Additionally, the Hansen J-test in Table 7 Panel A indicates the estimated effects of AI patent counts are not significantly different using each instrument separately except in year $t+3$. We also redo all IV regressions implemented using each instrument separately, and again tabulate the results in Internet Appendix Table 7 through Internet Appendix Table 12. Internet Appendix Table 7 shows that our IV estimates in year $t+3$ are not driven by the firm but rather its customers.

For a variety of natural reasons, experience with AI production in the past may facilitate AI production in the future. For instance, large fixed production costs may lower the cost to firms of subsequent innovation activities. Alternatively, corporate knowledge gained from

already having commercialized existing innovations may increase the efficiency of bringing new innovations to market. We therefore explore the moderating role of experience on AI production. We measure AI experience using the firm's past AI patent stock. Mindful of the statistical limitations of cross-sectional contrast analyses, we use the predicted values from the regression of Table 3 Panel A Column 3 for instrumented AI patent counts. Its interaction with past AI patent stock is our variable of interest. We interpret contrast results as providing suggestive evidence.

Table 7 Panel B shows that experience by itself does not have a significant effect. We further divide our experienced sample into low and high halves based on the median of AI patent stock counts as well as the patent stock ratio of AI patents to total patents. Table 7 shows, in Panel C and Panel D, that high experience does significantly further increase the effect of AI production on profitability, particularly when the firm's experience is measured not by AI patent grants by themselves but rather relative to non-AI patent grants (Panel D). The evidence suggests that some of the productivity potential of AI technology is being realized during our sample period, but only for firms with more experience producing AI innovation.

7.2. The Effect of AI Production on Cash Flow Risk

[Insert Table 8 about here]

We first measure risk using the volatility of quarterly return on assets. Our findings suggest that AI production decreases cash flow risk. Table 8 Panel A shows that a 10% increase in AI patent counts relative to its mean decreases the volatility of return on assets by 7% relative to its mean. We also measure risk using the volatility of daily stock returns, and we find a confirmatory reduction of 2% relative to its mean. Similarly across both measures, the increase in AI patent counts corresponds to roughly 3%-5% of the respective dependent variable's

standard deviation. This effect in year $t+1$ decreases in year $t+2$, more so for return on assets than for stock volatility, but it remains highly economically and statistically significant. Indeed, we verify that the results are similar beyond year $t+3$ (not tabulated).

We again also explore the moderating role of past experience with AI production facilitating future AI production (following the same procedure as in Section 7.1). We find that experience by itself does significantly further decrease the effect of AI production on risk, similarly from year $t+1$ to year $t+3$. However, high experience does not have a significantly different effect from low experience.

We further examine the decrease in the total volatility of stock returns from Table 8, decomposing it into its systematic and idiosyncratic components. As we previously argued in Section 1, AI technology can decrease firm risk, for instance, by increasing internal control and external responsiveness. Such risk dampening effects of AI may materialize in a firm's response to the idiosyncratic or systematic shocks that it experiences. At the same time, as a general purpose technology, AI can create commonalities across firms in responding to business challenges, potentially increasing a firm's systematic risk. We explore which effect of AI dominates for systematic risk, and also the effect of AI on idiosyncratic risk, using the fitted and residual values from time-series returns regressions. We again use the Fama and French (2015) five-factor model as our baseline, as in Table 4 and Table 5.

[Insert Table 9 about here]

Table 9 shows that the permanent decrease in total stock return volatility is a consequence of both systematic and idiosyncratic risk decreasing permanently (Panel A and Panel B, respectively). A 10% increase in AI patent counts relative to its mean decreases systematic and idiosyncratic return volatilities by roughly 3% and 1.5%, respectively. The results are very

similar if we replace our baseline model with the CAPM or our alternative factor models, as in Internet Appendix Table 2 and Internet Appendix Table 3.

Since the CAPM is characterized by a single source of systematic risk, we can parsimoniously quantify the decrease in market risk caused by AI production. We find that a 10% increase in AI patent counts relative to its mean decreases the CAPM beta by about 0.035 units (Panel C), or by roughly 4% of the mean of beta and 5% of its standard deviation. Overall, AI production appears to decrease risk, for cash flow as well as stock returns, and in the case of the latter, for both the systematic and idiosyncratic components of risk.

7.3. Potential Threats to Identification

Our identifying assumption, for the production of AI innovation, is that firms respond differently to random shocks to their innovation capacity because they have different AI exposures measured at the industry level. An important advantage of using time-invariant industry-level AI exposure, as we do, is that it minimizes the possibility of firms choosing their (firm-level) AI exposure endogenously with firm outcomes. However, it is possible that firms respond differently because of some differences across industries other than AI exposure, and these differences are correlated with AI exposure. We now address this possibility with specifications aimed at eliminating these omitted factors. However, the aggressiveness of these specifications also tends to eliminate much of the variation in our instrument.

7.3.1. Arbitrary Confounds

We begin by redoing Table 7 (profitability) and Table 8 (risk) adding industry-state fixed effects to capture the differential effect of R&D stock across industries that is specific to particular states within each industry. This eliminates, for instance, variation in R&D tax credits that apply differently to firms in different industries in different states – along with different

government policies, economic conditions, etc. that are specific to each industry-state pair. The Panel A results, in Internet Appendix Table 5 and Internet Appendix Table 6, are similar to our baseline results. This is the case even if we calculate R&D stock only using the firm's headquarters location and ignore all other R&D hubs of the firm (not tabulated).

Next, we directly address the possibility that industry AI exposure, and hence the differential effect of R&D stock across industries, is potentially confounded with other differences across industries, by interacting industry fixed effects with R&D stock. Fixed effects for each value of R&D stock within each SIC3 industry would completely absorb all of the variation in our instrument. Instead, we add SIC3 industry times R&D stock quartile fixed effects, which absorb (incrementally to our baseline fixed effects) much of the differential effect of R&D stock across different industries. Alternatively, we add SIC2 industry times R&D stock decile fixed effects, which mechanically absorb less of the variation in our instrument across industries but more of its variation across R&D stock. In Panel B and Panel C of the two appendix tables, the profitability results become insignificant, but the risk results are similar.

7.3.2. Automation and General Technology Confounds

We then examine two specific omitted factors that may be potentially confounded with our AI exposure measure at the industry level: automation and general technology. Rather than only capturing AI exposure, our measure may also be capturing a firm's *exposure* to automation, or the general characteristic that the firm's operations are in the technology industry. It is noteworthy that AI technology is indeed a type of automation technology (i.e., of non-routine tasks). However, our AI exposure measure could capture exposure to both "AI automation" and "non-AI automation" (e.g., if all the corresponding non-routine and routine tasks, respectively, are not mutually exclusive).

We measure exposure to automation of routine tasks, using data from Zhang (2019), as the share of wages paid to workers in routine task occupations (following Autor and Dorn (2013)) averaged across the firms in each SIC3 industry in 1990 (i.e., at the same time as our AI exposure measure). We classify firms as being in the technology industry using SIC3 industries based on Loughran and Ritter (2004). We add to our baseline specifications each of these measures interacted with R&D stock. Both measures also capture part of the nature of AI technology (i.e., automation of labor produced by technology firms) with the same granularity as AI exposure (i.e., SIC3 industry). Nevertheless, the results in Panel D and Panel E are similar to our baseline results.

As an alternative robustness test for the automation confound, we consider the possibility that AI innovation proxies for innovation in automation technologies. Specifically, we control for automation patents using data from Mann and Püttmann (2023), which are available from the beginning of our sample period until 2014. The results are similar to those in our baseline specifications (not tabulated).

7.3.3. AI Adoption Confound

We also consider the possibility that our results reflect the overall effect of the firm both producing and adopting AI technology. Specifically, our interaction instrument could not only affect AI production but also AI adoption, thereby violating the exclusion restriction. Conceptually, a firm's AI exposure should incentivize it to direct its innovation capacity to produce its own AI technology internally, but the firm could also be thus incentivized to acquire others' AI technology externally. AI scientists are an example of such a confounding factor: they are more likely to be employed by the firm because the firm has greater innovation capacity and

AI exposure (our interaction instrument), and they can potentially both produce more AI innovation and facilitate the firm's adoption of AI technology.

We examine whether the production of AI innovation proxies for AI adoption by controlling for AI adoption using data from Babina, Fedyk, He, and Hodson (2024a). Their AI adoption measures are arguably the most comprehensive available. These measures comprise a primary, resumes-based measure (from Cognism), and a secondary, job postings-based measure (from Burning Glass). Data for both AI adoption measures are available until the end of our sample period, but they are available beginning in 2005 for the primary measure and beginning in 2007 for the secondary measure (but missing for 2008-2009).

We merge our sample with Babina et al.'s and keep only firms common to both samples: about 6,000 and 4,000 firms, respectively, for their primary and secondary measures. For each firm, we linearly interpolate between a value of zero in 1990 and Babina et al.'s actual value for the firm's first year in their sample (and we also linearly interpolate between 2007 and 2010 for their secondary measure). We choose an initial value of zero because, expectedly, Babina et al.'s initial values are overwhelming zero (roughly 90% for their primary measure, and 60% for their secondary measure, rising to 70% if we include minuscule values (less than 0.01%) of the AI share of employees). Indeed, these figures are very similar to the values in the first year of Babina et al.'s sample (2005 and 2007 for their primary and secondary measures, respectively).

We find that Babina et al.'s measures have correlations of approximately 0.3 with our AI production measure (for both measures), which suggests that there is a meaningful distinction between producers and adopters. As for our main corporate outcomes, Panel F shows that the results are similar to our baseline results. We tabulate the results for the primary, resumes-based measure of adoption, but we find similar results for the secondary, job postings-based measure.

As for systematic risk, there is an opposite effect on CAPM beta between the production of AI innovation and the adoption of AI technology. We find that AI production decreases beta (Table 9), consistent with our argument that producing AI improves the reliability of the firm's execution and its responsiveness to changing business conditions. By contrast, Babina, Fedyk, He, and Hodson (2024b) find that AI adoption increases beta, which is consistent with their arguments that adopting AI can increase the firm's fragility during bad economic conditions (because of shared datasets, models, and technological infrastructure) even as it increases the firm's growth opportunities during good economic conditions, increasing beta on both the downside and upside.

7.3.4. Slow Moving Arbitrary Confounds

Finally, potential omitted factors that are correlated with the differential effect of R&D stock across industries may have a slow moving component (e.g., high AI innovation firms may have persistently high risk). To capture this component at the firm-year level, we add to our baseline specifications lagged dependent variables, which will also capture this same slow moving component of our instrument. It is worth noting that such specifications (with firm fixed effects and lagged dependent variables) are not only demanding but are subject to a Nickell (1981) bias that may seriously attenuate our estimates. Panel G in each of the two appendix tables shows that once again the profitability results become insignificant, but the risk results are similar to our baseline.

7.4. Collective Inferences from the Returns, Cash Flow Levels, and Cash Flow Risk Analyses

To summarize our main results, producing AI innovation causes large but transitory positive abnormal future stock returns (for roughly three years), a small and transitory increase in profits (for roughly one year), and a large and permanent decrease in risk. The transitory higher

future returns suggest that the permanent decrease in risk is not immediately and completely reflected in stock prices by investors.

7.5. Time-Varying Effects of AI Production

We consider the evolution, during the past three decades, of the effect of producing AI innovation. The proliferation of AI technology might suggest that producing AI innovation in recent years would be even more profitable and less risky for firms, and more efficiently priced by investors. However, many economic forces can give rise to time trends. For instance, with the significant and growing scale of AI innovation and its diffusion across industries, the returns to AI innovation may decrease as it becomes increasingly competitive to make important technological breakthroughs. At the same time, increasing scale may increase technological agglomeration effects, allowing one firm's innovation activities to benefit from those of other firms, thus increasing returns to AI innovation. Indeed, opposing trends such as these may roughly balance each other out.

We therefore examine our stock returns, profitability, and risk results over time. We do not find significant temporal differences in future stock returns during our sample period. This suggests that investors, in spite of their increasing interest in AI in recent years, have not become better at anticipating the firm value resulting from AI innovation. Similarly, profitability is not significantly different over time. However, we do find a small incremental risk decreasing effect of AI innovation during the second half of our sample period (not tabulated). This would be consistent with technological agglomeration dominating technological competition over time, decreasing firm risk more markedly as a result of AI innovation.

7.6. Patent Quality

Mindful that patents are heterogeneous in terms of quality, we examine our results for higher versus lower quality patents. It is well known that many patents are filed for innovations produced in the normal course of the R&D process but are of low quality because their small expected future value (e.g., from preserving their optionality) exceeds their relatively small incremental filing costs. Filing low quality patents may also keep the inventors working for a firm motivated, by fostering competition among inventors or improving individual career prospects with innovations that are valuable to the individual inventor even if they are low value to the firm. Many low quality patents are also filed because a critical mass of overlapping intellectual property claims can provide a firm with protection against patent litigation.

We capture the quality of patents standardly, using their scientific and commercial value (as in Section 3.1). We classify AI patents as high versus low quality based on their value relative to the median value of all AI patents in the same year. We redo Table 6 (Fama-MacBeth returns) Panel B, Table 7 (profitability) Panel A, and Table 8 (risk) Panel A, replacing (total) AI patents with high quality AI patents and controlling for low quality AI patents. The results (not tabulated) are roughly twice the magnitude for high quality AI patents, compared to our baseline results for total AI patents, using the median to classify patents into high versus low quality, and even larger if we use higher thresholds for classification. The results are also similar using alternative scaling variables for AI patent counts. Only when using commercial value to capture quality and only for stock returns, the magnitudes for high quality AI patents are not as large (but still larger than for total AI patents), but this is consistent with estimating future returns for firms that, by construction, have patents with higher initial returns. Overall, the evidence on higher quality patents supports the firm value hypothesis of AI production.

7.7. Novelty of AI Technology

Since AI technology is in the early stages of diffusion at the beginning of our sample period, we examine whether our results are affected by the novelty of AI technology. Although we find minimal time-varying effects of AI production (Section 7.5), we do find that AI patents are several times more likely to be breakthrough innovations than non-AI patents (Figure 2). We again use the measure and data of Kelly, Papanikolaou, Seru, and Taddy (2021), and we again redo Table 6 Panel B, Table 7 Panel A, and Table 8 Panel A, controlling for breakthrough patents. The results are similar to those in our baseline specifications (not tabulated).

8. Mechanisms

We investigate possible mechanisms underlying the effect of AI production on the value of the producer firm: labor productivity, physical capital intensity, and bargaining power. These mechanisms can directly improve the producer firm's operations (e.g., increase its labor productivity), through its AI production, motivated by its own AI exposure, most naturally decreasing costs (especially labor costs) but also increasing its sales. However, they can also indirectly affect the producer firm, motivated by its customers' AI exposure. Specifically, if the AI innovation motivated by customers' AI exposure allows the producer firm better satisfy demand (e.g., improve measurement, detection, response, etc.), then this lowers the costs of the producer firm's operations (e.g., increase labor output relative to input) and thus increases the firm's profits (separately from any effect of AI on sales). In this way, the producer firm can negotiate more lucrative business deals for itself, but it can also negotiate deals that are more stable, and thus increase profitability or decrease risk. While we frame our exposition of the mechanisms below in terms of the direct effect of producer firm's AI exposure on itself, for

brevity, the abovementioned indirect effect of customers' AI exposure can also result in analogous effects. We therefore consider both exposures here, as in all of our analyses.

8.1. Labor Productivity

As a labor enhancing technology, AI can increase the productivity of the producer firm's operations by improving labor productivity. AI augments earlier automation technologies by automating cognitive tasks that depend on human sensory and decision making abilities. Therefore, even compared to earlier automation technologies, AI can significantly substitute or complement jobs or even entire occupations.

[Insert Table 10 about here]

We first examine labor productivity, which we capture using profit per employee. Table 10 shows that profit per employee increases by roughly \$7,500 as a result of a 10% increase in AI patent counts relative to its mean. This corresponds to about 3% of the standard deviation of profit per employee (which is roughly \$250,000). However, this first year effect for labor productivity, like for profitability (net income to total assets), becomes insignificant by the second year (not tabulated). Therefore, the labor productivity effect is transitory.

We also examine the producer firm's level of employment. However, the effect of AI production here is unclear. If AI is, on balance, a substitute for labor, then employment will decrease. However, if AI complements labor on balance, making existing workers more productive as they work with AI technology, or allowing the firm to hire workers who produce more than they cost thanks to AI technology, then employment will increase.

We find no effect of AI production on the level of employment based on our results in Table 10. Nor do we find any effect on the overall scale of the producer firm, as measured by

total assets. While we have no evidence of AI production hurting employment to date, neither do we find that it helps.

8.2. Capital Intensity

AI technology allows firms to improve the automation and planning of their operations. For instance, it can reduce the need to maintain spare production capacity (not only labor but also capital) and inventory for episodes of surging customer demand. In so doing, AI enables firms to reduce their investment in and maintenance of capital required for development as well as production.

[Insert Table 11 about here]

We therefore examine the capital intensity of AI producer firms along various dimensions. Table 11 shows that firms generally become less capital intensive as a result of AI production. A 10% increase in AI patent counts decreases property, plant, and equipment by roughly 2.5% relative to its mean. Additionally, and consistent with AI improving planning, we find that inventory decreases by about 3.5% relative to its mean. These magnitudes correspond to about 2% of the respective dependent variable's standard deviation.

Additionally, we examine the investment of firms and find that it decreases as a result of AI production. Table 11 shows that capex and R&D spending, decrease by roughly 4% and 7% relative to their means, respectively, corresponding to about 3% of their respective standard deviations. By contrast, acquisitions expenditures increase by about 3.5%, corresponding to roughly 2% of its standard deviation. This suggests that AI technology allows firms to shift some of their investment focus from inside the firm to outside of it. The results for capital intensity are permanent, being generally similar for five years in economic and statistical significance (not tabulated).

8.3. Producer Firm Bargaining Power

In the course of producing AI innovation that can subsequently be commercialized, AI producer firms can also improve their bargaining power vis-à-vis their business counterparties. This can not only increase the spread between the firm's production outputs and inputs (i.e., increase profitability) but also their stability (i.e., decrease risk). Let us elaborate, starting with customers. Products that embed the producer firm's AI technology, or services integrating its AI technology with its customers' operations, can make it costly for customers to shift their business away from the AI producer firm. Similarly, as a safer customer for its suppliers, the AI producer itself may be able to command more reliable, or otherwise better or cheaper, products from its suppliers. Turning to employees, the threat of substitution from AI increases the firm's bargaining power relative to labor, which can allow the firm to lower its labor costs and also to increase its operating flexibility. The latter is particularly valuable in adverse business conditions during which flexibility may be much improved by actually substituting AI for labor. Overall, an AI producer can be more profitable for doing business with, and more costly to switch away from, for its counterparties. At the same time, the greater stability of the AI producer is beneficial both for the firm itself and each of its counterparties.

[Insert Table 12 about here]

In light of the difficulty of measuring bargaining power directly (even more so for business stability than lucrativity), we instead use measures of the stability of the firm's output and input relationships. We find that both increase as a result of AI production. Starting with outputs, Table 12 shows that the volatility of quarterly sales decreases by about 3.5% relative to its mean as a result of a 10% increase in AI patent counts. Also evidencing a more stable relationship between the firm and its customers, product differentiation (vis-à-vis product market

competitors) also increases. Specifically, the Hoberg and Phillips (2016) similarity score, converted to a differentiation, increases by about 5 percentage points.

Proceeding to inputs, Table 12 shows that the volatility of total costs of production decreases by about 3.5% relative to its mean as a result of a 10% increase in AI patent counts. If we break total costs down into their constituent SG&A and COGS, we find that their volatilities decrease by roughly 2.5% and 3%, respectively. Consistently across all of the regressions in Table 12, the estimated magnitudes correspond to about 2-3% of the respective dependent variable's standard deviation. The results for bargaining power are also permanent, with similar statistical significance for five years and somewhat smaller economic magnitude (not tabulated).

9. Financing Implications of AI Production

Having documented that AI producer firms have higher cash flows and lower cash flow risk, we turn to the financing implications of AI production. As a consequence of the effect of AI production on both of these key value drivers, we would expect AI producers to choose more aggressive financial structures. For instance, firms would be incentivized to shield from taxation their higher profits by increasing their leverage. Firms lower financial distress costs resulting from lower risk would similarly motivate them to increase their leverage.

[Insert Table 13 about here]

Table 13 shows that subsequent to a 10% increase in AI patent counts, AI producers increase their leverage by about 3% relative to its mean. Similarly, and also consistent with lower precautionary motives for holding cash, the same variation in AI patent counts lowers cash holdings by about 2.5%. We further investigate the components of the change in leverage to better understand how firms react. We find that net debt issuance increases by roughly 2 percentage points, while equity issuance decreases by roughly 0.8 p.p., even as share repurchases

remain unchanged. These financing results are also similar for several years (not tabulated). Overall, AI production appears to increase financial structure aggressiveness.

10. Conclusion

We document that AI innovation is a prominent form of innovation with widespread applications across different product markets and technology fields. Publicly traded firms dominate the AI production in the economy, and an increasingly high share of innovative publicly traded firms produce AI innovation. We argue that AI production increases firm value for the producer firm, by increasing cash flow levels and decreasing cash flow risk.

In our causal examination of the implications of AI production, we use an instrumental variable that exploits the interaction between the producer firm's plausibly exogenous innovation capacity and AI exposure driven incentives to produce AI innovation. We argue and find that firms produce AI innovation motivated by both their own AI exposure as well as that of their customers. Moreover, AI production causes transitorily higher profitability and permanently lower risk. Furthermore, AI producers have persistently higher future abnormal stock returns for roughly three years, which suggests that investors underestimate the firm value increasing effects of AI innovation. Additionally, we document mechanisms through which AI production affects firm value, including decreased physical capital intensity and increased bargaining power.

Taken together, our findings help inform corporate managers, capital providers, and policy makers who increasingly need to evaluate investment opportunities to develop and deploy AI technology. Producing AI innovation has been value enhancing for producer firms across several operational dimensions, during most of the past three decades.

Appendix 1

Details of the Classification of Patents as AI versus Non-AI

We describe here the key details of Giczy, Pairolero, and Toole (2022)'s machine learning approach for classifying patents as AI versus non-AI. As a starting point, AI is broken down into eight AI component technologies, and the universe of patent documents is evaluated for AI content pertaining to each of the eight components. These components are defined so as to be implementable in patent-level analysis and are motivated by the National Institute of Standards and Technology's definition of AI technology: "software and/or hardware that can learn to solve complex problems, make predictions or undertake tasks that require human-like sensing (such as vision, speech, and touch), perception, cognition, planning, learning, communication, or physical action" (NIST (2019)).

The eight AI component technologies are knowledge processing, speech recognition, AI hardware, evolutionary computation, natural language processing, machine learning, computer vision, and planning/control. These components are not mutually exclusive. For instance, an invention in any one of the components is likely to also exploit machine learning models. The identification algorithm then focuses on each of these eight AI component technologies separately, until, for each component, all patents are assigned a predicted probability of being AI.

To train a machine learning model to identify a patent as AI or non-AI, it is necessary to have one set of patents that are "surely AI" and another set that are "surely non-AI". The set of "surely AI" patents is identified by intersecting four patent classification systems: CPC, IPC, USPC, and Derwent World Patent Index. Each of these systems has its own set of patent classes that allow categorization of every patent as AI or non-AI according to each of the aforementioned eight AI component technologies. Giczy et al. deem a patent to be "surely AI" if

all four patent classification systems agree that the patent belongs to the specific AI component technology under consideration.²⁸

Having thus identified the *training* set of "surely AI" patents, the next step is to identify the set of "surely non-AI" patents. This begins by excluding the set of "surely AI" patents. However, some of the patents that remain may be related to AI. These patents are identified for exclusion in two independent procedures as follows. In the first procedure, patents are excluded if they share a patent family with any patent in the set of "surely AI" patents, and their backward and forward citations are also excluded.²⁹ This step is repeated a second time, but this time the basis of exclusion is sharing a patent family with any patent excluded in the first step (as opposed to the set of "surely AI" patents). In the second procedure, patents are excluded if they belong to a CPC patent class that has an abnormally high share of "surely AI" patents (specifically, if the class' share of "surely AI" patents is more than 50 times the class' share of the universe of patents). The final step in creating the *training* set of "surely non-AI" patents is to randomly select 15,000 of the patents that remain after the foregoing exclusions.

A machine learning model is then trained on the abstract, claims, and citations of the "surely AI" and "surely non-AI" patents. After training, the model subsequently evaluates all patent documents (i.e., not just those of "surely AI" and "surely non-AI" training sets) for their AI content. All patents are thus assigned a predicted probability of containing a particular AI component technology. Finally, if any of the predicted probabilities exceed 0.5 for any of the eight AI component technologies, the patent is classified as an AI patent.

²⁸ For example, to identify the "surely AI" set of patents for computer vision, the following is a list of the patent classes that are intersected from the four patent classification systems. From CPC/IPC: G06K9 (recognition of characters or patterns), G06T3 (image transformation), G06T5 (image enhancement/ restoration), and G06T7 (image analysis). From USPC: 382 (image analysis). From Derwent: T01-J10B (Image Processing), T04-D (Character and signal pattern recognition), and T01-J16 (artificial intelligence).

²⁹ A patent family is a group of patent applications and/or granted patents that share a common applicant/owner and share a similar inventive concept.

Appendix 2

Details of the Construction of Occupational AI Exposure Scores

The AI exposure of an occupation is the extent to which AI can be used to substitute or complement labor in that occupation, and the measure that we use reflects this agnosticism about the effect of AI on labor. Felten, Raj, and Seamans (2021) measure occupational AI exposure starting with estimating the AI exposure of all 52 "workspace abilities" in the Department of Labor O*NET database. These abilities simply describe the skills required to perform the tasks involved in various occupations. O*NET scores each ability, within each occupation, on its relevance and importance (e.g., surgeons receive high scores for arm-hand steadiness and deductive reasoning).

Felten, Raj, and Seamans (2021) conduct a crowd sourced survey via Amazon's mTurk asking respondents if a specified O*NET ability "is related to or can use AI" in 10 "AI applications" defined by the Electronic Frontier Foundation.³⁰ Survey responses (zero-one / no-yes) are averaged within each of 520 ability-application pairs (52×10). Then, within each of 52 O*NET workspace abilities, the survey average AI application scores are summed up, resulting in an AI application score for each workspace ability. Finally, the total AI application scores for O*NET workspace abilities are calculated as a weighted average across each O*NET occupation. The weights used are the initially mentioned O*NET scores for the relevance and importance of each workspace ability specific to the occupation. The final occupational scores are standardized (mean zero, standard deviation one).

³⁰ This focus is chosen for the sake of concreteness and precision of survey responses. The EFF is a digital rights and privacy non-profit that collects statistics about the progress of AI across its applications. The 10 selected AI applications are those for which the EFF has recorded scientific activity since 2010. The applications comprise: abstract strategy games, real-time video games, image recognition, visual question answering, image generation, reading comprehension, language modeling, translation, speech recognition, and instrumental track recognition.

References

- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2020, Artificial intelligence and jobs: Evidence from online vacancies, *Journal of Labor Economics* 40, S293-S340.
- Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2015, House prices, collateral, and self-employment, *Journal of Financial Economics* 117, 288-306.
- Alekseeva, Liudmila, Mireia Giné, Sampsa Samila, and Bledi Taska, 2020, AI adoption and firm performance: Management versus IT, working paper.
- Angrist, Joshua, and William N. Evans, 1998, Children and their parents' labor supply: Evidence from exogenous variation in family size, *American Economic Review* 88, 450-477.
- Atanasov, Vladimir A., and Bernard S. Black, 2016, Shock-based causal inference in corporate finance and accounting research, *Critical Finance Review* 5, 207-304.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2022, Firm investments in artificial intelligence technologies and changes in workforce composition, in Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich Strassner, eds.: *Technology, Productivity, and Economic Growth* (University of Chicago Press)
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024a, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151, 103745.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024b, Artificial intelligence and firms' systematic risk, working paper.
- Babina, Tania, and Sabrina T. Howell, 2024, Entrepreneurial spillovers from corporate R&D, *Journal of Labor Economics* 42, 469-509.

- Bena, Jan, and Elena Simintzi, 2025, Machines could not compete with Chinese labor: Evidence from US firms' innovation, forthcoming *Review of Finance*.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen, 2013, Identifying technology spillovers and product market rivalry, *Econometrica* 81, 1347-1393.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, 2019, Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics, in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds.: *The Economics of Artificial Intelligence: An Agenda* (University of Chicago Press).
- Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock, 2018, What can machines learn and what does it mean for occupations and the economy?, *AEA Papers and Proceedings* 108, 43-47.
- Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57-82.
- Chan, Louis K. C., Josef Lakonishok, and Theodore Sougiannis, 2001, The stock market valuation of research and development expenditures, *Journal of Finance* 56, 2431-2456.
- Chaney, Thomas, David Sraer, and David Thesmar, 2012, The collateral channel: How real estate shocks affect corporate investment, *American Economic Review* 102, 2381-2409.
- Cockburn, Iain M., Rebecca Henderson, and Scott Stern, 2019, The impact of artificial intelligence on innovation: An exploratory analysis, in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds.: *The Economics of Artificial Intelligence: An Agenda* (University of Chicago Press).
- Cohen, Lauren, Karl Diether, and Christopher Malloy, 2013, Misvaluing innovation, *Review of Financial Studies* 26, 635-666.

- Derrien, François, Ambrus Kecskés, and Phuong-Anh Nguyen, 2023, Labor force demographics and corporate innovation, *Review of Financial Studies* 36, 2797-2838.
- Dreher, Axel, Andreas Fuchs, Bradley C. Parks, Austin Strange, and Michael J. Tierney, 2021, Aid, China, and growth: Evidence from a new global development finance dataset, *American Economic Journal: Economic Policy* 13, 135-174.
- Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang, 2023, Generative AI and firm values, working paper.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3-56.
- Fama, Eugene F., and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1-22.
- Felten, Edward, Manav Raj, and Robert Seamans, 2021, Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses, *Strategic Management Journal* 42, 2195-2217.
- Giczy, Alexander V., Nicholas A. Pairolero, and Andrew A. Toole., 2022, Identifying artificial intelligence (AI) invention: A novel AI patent dataset, *Journal of Technology Transfer* 47, 476-505.
- Grennan, Jillian, and Roni Michaely, 2020, Artificial intelligence and high-skilled work: Evidence from analysts, working paper.
- Grennan, Jillian, and Roni Michaely, 2021, Fintechs and the market for financial analysis, *Journal of Financial and Quantitative Analysis* 56, 1877-1907.
- Griffin, John M., and Michael L. Lemmon, 2002, Book-to-market equity, distress risk, and stock returns, *Journal of Finance* 57, 2317-2336.

- Hirshleifer, David, Po-Hsuan Hsu, and Dongmei Li, 2013, Innovative efficiency and stock returns, *Journal of Financial Economics* 107, 632-654.
- Hirshleifer, David, Po-Hsuan Hsu, and Dongmei Li, 2018, Innovative originality, profitability, and stock returns, *Review of Financial Studies* 31, 2553-2605.
- Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, *Journal of Political Economy* 124, 1423-1465.
- Hombert, Johan, and Adrien Matray, 2018, Can innovation help U.S. manufacturing firms escape import competition from China?, *Journal of Finance* 73, 2003-2039.
- Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting anomalies: An investment approach, *Review of Financial Studies* 28, 650-705.
- Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, 2021, Measuring technological innovation over the long run, *American Economic Review: Insights* 3, 303-320.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665-712.
- Li, Kai, Feng Mai, Rui Shen, and Xinyan Yan, 2021, Measuring corporate culture using machine learning, *Review of Financial Studies* 34, 3265-3315.
- Li, Kai, Feng Mai, Rui Shen, Chelsea Yang, and Tengfei Zhang, 2025, Dissecting corporate culture using generative AI, forthcoming *Review of Financial Studies*.
- Liu, Jianan, Robert F. Stambaugh, and Yu Yuan, 2019, Size and value in China, *Journal of Financial Economics* 134, 48-69.

- Loughran, Tim, and Jay Ritter, 2004, Why has IPO underpricing changed over time? *Financial Management* 33, 5-37.
- Mann, Katja, and Lukas Püttmann, 2023, Benign effects of automation: New evidence from patent texts, *Review of Economics and Statistics* 105, 562-579.
- National Institute of Standards and Technology (NIST), 2019, U.S. Leadership in AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools, Gaithersburg, MD.
- Nickell, Stephen, 1981, Biases in dynamic models with fixed effects, *Econometrica* 49, 1417-1426.
- Rock, Daniel, 2021, Engineering value: The returns to technological talent and investments in artificial intelligence, working paper.
- Sautner, Zacharias, Laurence van Lent, Grigory Vilkov, and Ruishen Zhang, 2023, Firm-level climate change exposure, *Journal of Finance* 78, 1449-1498.
- Stoffman, Noah, Michael Woepfel, and M. Deniz Yavuz, 2022, Small innovators: No risk, no return, *Journal of Accounting and Economics* 74, 101492.
- Trajtenberg, Manuel, 2019, Artificial intelligence as the next GPT: A political-economy perspective, in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds.: *The Economics of Artificial Intelligence: An Agenda* (University of Chicago Press).
- United States Patent and Trademark Office (USPTO), 2020, Inventing AI: Tracing the diffusion of artificial intelligence with U.S. patents, Alexandria, VA.
- Webb, Michael, 2020, The impact of artificial intelligence on the labor market, working paper.
- Wilson, Daniel J., 2009, Beggar thy neighbor? The in-state, out-of-state, and aggregate effects of R&D tax credits, *Review of Economics and Statistics* 91, 431-436.

Zhang, Miao Ben, 2019, Labor-technology substitution: Implications for asset pricing, *Journal of Finance* 74, 1793-1839.

Table 1
Industry Ranking Based on AI Production

This table shows the ranking of industries based on their AI production. The sample is all firms in the baseline sample restricted to industries with at least 10 firms per year every year during the sample period. The number of firms in an industry is the annual average number of firms. The number of AI patents is the annual average of the total number of AI patents granted to firms in the industry. The three most AI innovative firms in an industry are the three firms with the highest number of AI patents.

All firms						Excl. three most AI innovative firms	
Rank: Industry total AI patents	SIC3	Industry name	Number of firms	Number of AI patents	Rank: Mean AI patents per firm	Rank: Industry total AI patents	Rank: Mean AI patents per firm
1	737	Computer programming, data processing, and other computer related	312.5	400.14	6	1	3
2	367	Electronic components and accessories	139.9	226.57	3	3	2
3	357	Computer and office equipment	113.9	222.71	2	2	1
4	382	Laboratory apparatus and analytical, optical, measuring, and control	98.4	100.61	8	4	5
5	366	Communications equipment	92.5	87.04	7	5	6
6	384	Surgical, medical, and dental instruments and supplies	139.1	74.39	13	6	8
7	481	Telephone communications	45.8	58.25	4	7	7
8	372	Aircraft and parts	20.5	52.36	1	9	4
9	371	Motor vehicles and motor vehicle equipment	45.5	45.89	9	10	10
10	283	Drugs	271.4	38.89	23	8	14
11	291	Petroleum refining	19.5	24.93	5	11	9
12	353	Construction, mining, and materials handling machinery and equipment	24.5	20.50	10	19	18
13	355	Special industry machinery, except metalworking machinery	36.1	18.89	12	12	11
14	362	Electrical industrial apparatus	18.8	15.36	11	18	17
15	596	Non-store retailers	28.3	11.50	14	20	20
16	138	Oil and gas field services	30.4	11.36	15	22	21
17	873	Research, development, and testing services	39.4	7.21	20	13	13
18	284	Soap, detergents, and cleaning preparations; perfumes, cosmetics, and other toilet preparations	22.6	6.39	16	31	27
19	421	Trucking and courier services, except air	28.4	5.96	17	36	39
20	738	Miscellaneous business services	37.7	5.36	19	15	16

21	369	Miscellaneous electrical machinery, equipment, and supplies	20.8	5.21	18	14	12
22	799	Miscellaneous amusement and recreation services	29.5	4.68	22	24	24
23	483	Radio and television broadcasting stations	25.6	4.57	21	16	15
24	356	General industrial machinery and equipment	39.7	3.93	26	17	19
25	874	Management and public relations services	25.1	2.68	24	23	22
26	504	Professional and commercial equipment and supplies	30.8	1.68	30	21	23
27	208	Beverages	18.9	1.57	28	33	34
28	451	Air transportation, scheduled, and air courier services	20.0	1.57	27	29	29
29	506	Electrical goods	22.6	1.57	25	30	30
30	281	Industrial inorganic chemicals	16.5	1.50	28	27	29
31	131	Crude petroleum and natural gas	98.5	1.39	35	25	32
32	871	Engineering, architectural, and surveying services	18.9	0.75	34	26	31
33	308	Miscellaneous plastics products	29.3	0.64	33	28	25
34	581	Eating and drinking places	69.0	0.64	37	39	33
35	495	Sanitary services	24.9	0.61	31	33	40
36	594	Miscellaneous shopping goods stores	21.9	0.46	32	32	26
37	809	Miscellaneous health and allied services, not elsewhere classified	18.1	0.25	36	34	34
38	736	Personnel supply services	26.8	0.21	38	40	42
39	331	Steel works, blast furnaces, and rolling and finishing mills	27.3	0.14	39	41	38
40	735	Miscellaneous equipment rental and leasing	15.8	0.07	40	38	35
41	153	Operative builders	23.0	0.04	40	37	33
42	565	Family clothing stores	15.5	0.04	41	39	38

Table 2
Descriptive Statistics

This table presents descriptive statistics for the main variables used in this paper. Variables are defined in Appendix Table 2.

	Mean	Standard deviation	25th percentile	Median	75th percentile
Independent variables					
- AI patent counts	0.66	3.46	0.00	0.00	0.00
- Non-AI patent counts	6.5	26.8	0.0	0.0	1.0
- R&D stock [tax credit induced] (\$ M)	127.3	491.0	0.0	1.9	40.2
- Firm's AI exposure	1.23	0.47	0.87	1.21	1.54
- Customers' AI exposure	1.14	0.36	0.85	1.14	1.37
- Total assets (\$ M)	2,313	6,832	62	261	1,208
- Firm age (years)	15	15	4	10	21
- Innovation dummy variable	0.33	0.47	0.00	0.00	1.00
AI patent counts with alternative scaling					
- AI patent counts / Total assets (per \$ B)	0.62	3.12	0.00	0.00	0.00
- AI patent counts / Total patent stock (%)	1.09	5.47	0.00	0.00	0.00
- AI patent counts / Total patent counts (%)	3.60	15.33	0.00	0.00	0.00
- AI patent counts / R&D stock (per \$ B)	1.43	6.64	0.00	0.00	0.00
Variables used in stock returns analysis					
- Monthly stock return (%)	0.66	16.35	-7.81	-0.22	7.73
- Market capitalization (\$ M)	2,319	7,484	46	219	1,082
- Market-to-book of equity	2.7	4.0	0.8	1.4	2.8
- Momentum (%)	9.7	60.6	-26.8	0.6	30.8
- Short-term reversal (%)	0.69	16.29	-7.84	-0.24	7.74
- Stock price (\$)	20.0	22.9	4.3	12.0	27.3
Dependent variables: Profitability					
- Return on assets	-0.057	0.279	-0.084	0.026	0.079
Dependent variables: Risk					
- Volatility of return on assets	0.032	0.054	0.006	0.013	0.033
- Volatility of stock returns (%)	4.18	2.70	2.31	3.43	5.17
Dependent variables: Labor productivity					
- Profit per employee (\$ M per employee)	-0.050	0.247	-0.023	0.005	0.022
- Employment / Total assets (employees per \$ M)	5.4	6.6	1.7	3.4	6.4
- Total assets (\$ M)	2,428	7,102	63	280	1,301
Dependent variables: Capital intensity					
- PP&E / Total assets	0.26	0.22	0.08	0.18	0.37
- Inventory / Total assets	0.14	0.16	0.01	0.09	0.22
- Capex / Total assets	0.061	0.077	0.016	0.036	0.073
- R&D / Total assets	0.067	0.131	0.000	0.004	0.077
- Acquisitions / Total assets	0.030	0.093	0.000	0.000	0.007
Dependent variables: Bargaining power					
- Volatility of (Sales / Total assets)	0.049	0.064	0.012	0.027	0.058
- Product differentiation (%)	96.9	2.2	96.3	97.5	98.3
- Volatility of (Total costs / Total assets)	0.043	0.059	0.010	0.023	0.050
- Volatility of (SG&A / Total assets)	0.012	0.020	0.002	0.005	0.013
- Volatility of (COGS / Total assets)	0.034	0.047	0.007	0.017	0.040

Dependent variables: Financial policies					
- Leverage	0.23	0.23	0.02	0.18	0.35
- Cash holdings / Total assets	0.20	0.23	0.03	0.10	0.29
- Equity issuance / Total assets	0.085	0.270	0.000	0.004	0.022
- Share repurchases / Total assets	0.016	0.040	0.000	0.000	0.008
- Net debt issuance / Total assets	0.026	0.142	-0.023	0.000	0.030

Table 3
First Stage of IV Regressions

This table shows the results of regressions of AI production on the interaction between the producer firm's R&D stock and its own AI exposure or the AI exposure of its customers. Column 3 corresponds to the first stage of the IV regressions. The sample period spans 1990-2017 in terms of year t . AI patent counts are measured in year t . They are instrumented with R&D stock measured at year $t-2$ and AI exposure fixed before the start of the sample period. Scaling variables are measured in, or ending in, year t . Patent counts are measured in year t . Total patent stock and R&D stock are measured during the previous 10 years. The sample and specifications are described in the text. Variables are defined in Appendix Table 2. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Dependent Variable Is AI Patent Counts			
	ln(0.1+AI patent counts)		
	(1)	(2)	(3)
ln(1+R&D stock) [tax credit induced] × Firm's AI exposure	0.053*** (6.53)		0.024* (1.91)
ln(1+R&D stock) [tax credit induced] × Customers' AI exposure		0.058*** (6.69)	0.040*** (3.05)
ln(1+R&D stock) [tax credit induced]	0.037*** (2.61)	0.038*** (2.61)	0.038*** (2.62)
ln(1+Non-AI patent counts)	0.383*** (12.58)	0.384*** (12.60)	0.383*** (12.60)
Innovation dummy variable	0.163** (2.56)	0.163** (2.57)	0.163** (2.56)
ln(Total assets)	0.051*** (6.30)	0.051*** (6.29)	0.051*** (6.32)
ln(Firm age)	-0.036*** (-3.87)	-0.038*** (-4.13)	-0.038*** (-4.12)
Fixed effects			
State × Year?	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes
Firm?	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes
Observations	92,277	92,277	92,277
Adjusted R ²	0.702	0.702	0.702

Panel B: Dependent Variable Is AI Patent Counts Scaled by Total Assets			
	ln(0.01+AI patent counts / Total assets)		
	(1)	(2)	(3)
ln(1+R&D stock) [tax credit induced]	0.080***		0.034*
× Firm's AI exposure	(6.90)		(1.94)
ln(1+R&D stock) [tax credit induced]		0.089***	0.063***
× Customers' AI exposure		(7.04)	(3.32)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	92,119	92,119	92,119
Adjusted R ²	0.574	0.574	0.574
Panel C: Dependent Variable Is AI Patent Counts Scaled by Total Patent Stock			
	ln(0.0001+AI patent counts / Total patent stock)		
	(1)	(2)	(3)
ln(1+R&D stock) [tax credit induced]	0.067***		0.042**
× Firm's AI exposure	(5.41)		(2.35)
ln(1+R&D stock) [tax credit induced]		0.066***	0.034*
× Customers' AI exposure		(4.79)	(1.71)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	92,277	92,277	92,277
Adjusted R ²	0.578	0.578	0.578
Panel D: Dependent Variable Is AI Patent Counts Scaled by Total Patent Counts			
	ln(0.0001+AI patent counts / Total patent counts)		
	(1)	(2)	(3)
ln(1+R&D stock) [tax credit induced]	0.099***		0.049*
× Firm's AI exposure	(5.96)		(1.80)
ln(1+R&D stock) [tax credit induced]		0.106***	0.069**
× Customers' AI exposure		(5.82)	(2.40)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	92,119	92,119	92,119
Adjusted R ²	0.595	0.595	0.595
Panel E: Dependent Variable Is AI Patent Counts Scaled by R&D Stock			
	ln(0.01+AI patent counts / R&D stock)		
	(1)	(2)	(3)
ln(1+R&D stock) [tax credit induced]	0.108***		0.075***
× Firm's AI exposure	(7.51)		(2.95)
ln(1+R&D stock) [tax credit induced]		0.102***	0.045
× Customers' AI exposure		(5.95)	(1.59)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	88,647	88,647	88,647
Adjusted R ²	0.604	0.604	0.604

Table 4
Risk-Adjusted Returns of Portfolios Sorted by Actual AI Patent Counts

This table shows the risk-adjusted returns of portfolios sorted based on actual AI patent counts. The sample and specifications are described in the text. Returns are measured from July 1991 to June 2019 (336 consecutive months). Firms are sorted into four quasi-quartiles: zero, low (Q2), medium (Q3), and high (Q4) AI patents. The zero AI patents quartile is further sorted into two groups: zero innovation (Q0) and non-zero innovation (Q1). Returns are risk-adjusted using the Fama and French (2015) five-factor model. t-statistics are calculated using Newey and West (1987) standard errors with twelve lags. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Alphas from Sorting by AI Patent Counts								
	Innovat'n = 0 & AI = 0 (Q0)	Innovat'n > 0 & AI = 0 (Q1)	Low AI (Q2)	Medium AI (Q3)	High AI (Q4)	Q4 – Q2	Q4 – Q1	Q4 – Q0
Equally weighted	0.07 (0.53)	0.20 (1.53)	0.33** (2.11)	0.52*** (3.14)	0.52*** (3.66)	0.19 (1.46)	0.32** (2.23)	0.45*** (3.35)
Value weighted	-0.19** (-2.41)	-0.09 (-1.18)	-0.04 (-0.41)	0.18 (1.56)	0.27** (2.15)	0.32* (1.90)	0.37** (2.21)	0.46*** (3.42)
Size neutral	-0.17** (-2.06)	-0.02 (-0.33)	0.19* (1.95)	0.25* (1.91)	0.49** (2.39)	0.31* (1.68)	0.52** (2.54)	0.66*** (3.44)
Mean number of stocks	1,945	709	140	78	99			
Panel B: Alphas from Sorting by AI Patent Counts Scaled by Total Assets								
	Innovat'n = 0 & AI = 0 (Q0)	Innovat'n > 0 & AI = 0 (Q1)	Low AI (Q2)	Medium AI (Q3)	High AI (Q4)	Q4 – Q2	Q4 – Q1	Q4 – Q0
Equally weighted	0.07 (0.53)	0.20 (1.53)	0.06 (0.58)	0.50*** (3.19)	0.74*** (2.89)	0.69*** (2.77)	0.54*** (3.06)	0.67*** (3.01)
Value weighted	-0.19** (-2.41)	-0.09 (-1.18)	-0.02 (-0.32)	0.36** (2.16)	0.79** (2.54)	0.82** (2.40)	0.89*** (2.68)	0.98*** (2.96)
Size neutral	-0.17** (-2.06)	-0.02 (-0.34)	-0.14 (-0.91)	0.41*** (2.70)	0.68*** (2.73)	0.82*** (2.84)	0.71*** (3.02)	0.85*** (3.37)
Mean number of stocks	1,945	708	108	106	102			

Panel C: Alphas from Sorting by AI Patent Counts Scaled by Total Patent Stock								
	Innovat'n = 0 & AI = 0 (Q0)	Innovat'n > 0 & AI = 0 (Q1)	Low AI (Q2)	Medium AI (Q3)	High AI (Q4)	Q4 – Q2	Q4 – Q1	Q4 – Q0
Equally weighted	0.07 (0.53)	0.20 (1.53)	0.08 (0.87)	0.48** (2.54)	0.73*** (3.52)	0.64*** (3.26)	0.53*** (3.29)	0.66*** (3.92)
Value weighted	-0.19** (-2.41)	-0.09 (-1.18)	-0.05 (-0.72)	0.24 (1.55)	0.53*** (2.84)	0.57*** (2.97)	0.62*** (2.72)	0.71*** (3.61)
Size neutral	-0.17** (-2.06)	-0.02 (-0.33)	0.16 (1.02)	0.30* (1.80)	0.54*** (2.90)	0.38 (1.49)	0.56*** (2.93)	0.71*** (4.27)
Mean number of stocks	1,945	709	108	107	102			
Panel D: Alphas from Sorting by AI Patent Counts Scaled by Total Patent Counts								
	Innovat'n = 0 & AI = 0 (Q0)	Innovat'n > 0 & AI = 0 (Q1)	Low AI (Q2)	Medium AI (Q3)	High AI (Q4)	Q4 – Q2	Q4 – Q1	Q4 – Q0
Equally weighted	0.07 (0.53)	0.20 (1.53)	0.20* (1.92)	0.42** (2.26)	0.67*** (3.31)	0.47*** (2.61)	0.47*** (2.93)	0.60*** (3.76)
Value weighted	-0.19** (-2.41)	-0.09 (-1.18)	-0.01 (-0.18)	0.18 (1.07)	0.53*** (2.71)	0.54*** (2.63)	0.62*** (2.68)	0.72*** (3.49)
Size neutral	-0.17** (-2.06)	-0.02 (-0.33)	0.27* (1.91)	0.23 (1.51)	0.56*** (2.82)	0.29 (1.25)	0.59*** (2.87)	0.73*** (4.22)
Mean number of stocks	1,945	709	109	109	99			
Panel E: Alphas from Sorting by AI Patent Counts Scaled by R&D Stock								
	Innovat'n = 0 & AI = 0 (Q0)	Innovat'n > 0 & AI = 0 (Q1)	Low AI (Q2)	Medium AI (Q3)	High AI (Q4)	Q4 – Q2	Q4 – Q1	Q4 – Q0
Equally weighted	0.07 (0.53)	0.28** (1.97)	0.28** (2.31)	0.45*** (3.11)	0.70*** (3.05)	0.42** (2.29)	0.41*** (2.78)	0.63*** (3.29)
Value weighted	-0.19** (-2.41)	0.05 (0.58)	0.04 (0.57)	0.30* (1.85)	0.32 (1.64)	0.28 (1.39)	0.27 (1.15)	0.51** (2.50)
Size neutral	-0.17** (-2.06)	0.10 (1.26)	0.30** (2.25)	0.36** (2.27)	0.40** (2.50)	0.10 (0.71)	0.30* (1.88)	0.57*** (3.71)
Mean number of stocks	1,945	615	99	97	94			

Table 5
Risk-Adjusted Returns of Portfolios Double Sorted by R&D Stock and AI Exposure

This table shows the risk-adjusted returns of portfolios double sorted independently based on (tax credit induced) R&D capital stock and AI exposure. The sample and specifications are described in the text. Returns are measured from July 1991 through June 2019 (336 consecutive months). Firms are sorted into three quasi-terciles based on R&D capital stock: zero R&D stock ("low R&D", L), and two halves of non-zero R&D stock ("medium R&D", M; and "high R&D", H). Independently, firms are sorted into terciles based on AI exposure at the industry level measured as the first principal component of the respective AI exposures of the firm and its customers. Returns are risk-adjusted using the Fama and French (2015) five-factor model. t-statistics are calculated using Newey and West (1987) standard errors with twelve lags. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Alphas of Equally Weighted Portfolios				
	Exposure T1 (low)	Exposure T2 (medium)	Exposure T3 (high)	Exposure (T3 – T1)
Low R&D stock (L)	-0.23 (-1.57)	-0.45** (-2.57)	-0.07 (-0.62)	0.15 (1.34)
Medium R&D stock (M)	-0.08 (-0.54)	-0.01 (-0.08)	0.55** (2.42)	0.63** (2.53)
High R&D stock (H)	-0.25* (-1.79)	0.00 (0.04)	0.79*** (4.05)	1.04*** (3.77)
R&D stock (H – L)	-0.02 (-0.17)	0.46** (2.37)	0.87*** (3.99)	0.89*** (3.71)
Panel B: Alphas of Value Weighted Portfolios				
	Exposure T1 (low)	Exposure T2 (medium)	Exposure T3 (high)	Exposure (T3 – T1)
Low R&D stock (L)	-0.37*** (-3.07)	-0.26** (-2.38)	-0.29*** (-2.85)	0.09 (0.55)
Medium R&D stock (M)	-0.32 (-1.56)	-0.27 (-1.19)	0.20 (1.57)	0.52** (2.32)
High R&D stock (H)	-0.20 (-1.62)	-0.03 (-0.27)	0.27*** (3.13)	0.47*** (2.86)
R&D stock (H – L)	0.18* (1.69)	0.23 (1.24)	0.56*** (4.04)	0.38** (2.06)
Panel C: Alphas of Size Neutral Portfolios				
	Exposure T1 (low)	Exposure T2 (medium)	Exposure T3 (high)	Exposure (T3 – T1)
Low R&D stock (L)	-0.35*** (-3.44)	-0.39*** (-3.22)	-0.27*** (-2.93)	0.08 (0.71)
Medium R&D stock (M)	-0.35* (-1.83)	-0.30 (-1.36)	0.27** (2.13)	0.62*** (2.93)
High R&D stock (H)	-0.26** (-2.01)	-0.07 (-0.63)	0.44*** (3.71)	0.70*** (3.52)
R&D stock (H – L)	0.09 (1.15)	0.32** (1.98)	0.71*** (4.77)	0.62*** (3.26)

Panel D: Mean Number of Stocks in Each Portfolio				
	Exposure T1 (low)	Exposure T2 (medium)	Exposure T3 (high)	Exposure (T3 – T1)
Low R&D stock (L)	258	422	652	
Medium R&D stock (M)	115	134	537	
High R&D stock (H)	103	149	598	

Table 6
Fama-MacBeth Regressions of Stock Returns on AI Production

This table shows the results of Fama-MacBeth regressions of individual monthly stock returns on AI production. The sample and specifications are described in the text. Returns are measured from July 1991 to June 2019 (336 consecutive months). The controls variables for non-AI patent counts and R&D spending are scaled by total assets, and the small constants added before taking logarithms are suitably adjusted. The other control variables are market capitalization, market-to-book of equity, momentum, short-term reversal, return on assets, capex-to-total assets, stock price, and firm age. Variables are defined in Appendix Table 2. t-statistics are calculated using Newey and West (1987) standard errors with twelve lags. The coefficient of determination, F-statistic for instrument, and p-value of Hansen J-statistic are the time-series averages of the corresponding statistics across the cross-sectional regressions. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Actual AI Patent Counts (OLS)					
Dependent variable is monthly stock return					
Scaling variable for AI patent counts					
	None	Total assets	Total patent stock	Total patent counts	R&D stock
Year t+1					
ln(0.1+AI patent counts)	0.002 (0.13)	0.026** (2.14)	0.018* (1.72)	0.013* (1.66)	0.016 (1.61)
ln(1+Non-AI patent counts)	-0.081** (-2.10)	0.037* (1.81)	0.040 (0.81)	0.041** (2.03)	-0.006 (-0.31)
ln(1+R&D spending)	0.139*** (3.81)	0.111*** (3.35)	0.115*** (3.41)	0.113*** (3.39)	0.117*** (3.38)
Innovation dummy variable	0.120 (1.52)	-0.160** (-2.27)	-0.052 (-0.77)	-0.166** (-2.33)	0.033 (0.28)
Control variables?	Yes	Yes	Yes	Yes	Yes
FF48 industry fixed effects?	Yes	Yes	Yes	Yes	Yes
Observations	997,382	997,382	997,394	997,394	956,844
R ²	0.090	0.090	0.090	0.090	0.090
Panel B: Instrumented AI Patent Counts (IV) and Returns in Year t+1					
Dependent variable is monthly stock return					
Scaling variable for AI patent counts					
	None	Total assets	Total patent stock	Total patent counts	R&D stock
Year t+1					
ln(0.1+AI patent counts) [instrumented]	0.433*** (3.12)	0.308*** (3.00)	0.275*** (3.11)	0.229*** (3.11)	0.245*** (2.95)
Firm's AI exposure	-0.019 (-0.39)	-0.024 (-0.48)	-0.025 (-0.50)	-0.025 (-0.49)	-0.030 (-0.58)
Customers' AI exposure	0.065 (1.36)	0.065 (1.34)	0.062 (1.30)	0.064 (1.33)	0.060 (1.18)
ln(1+R&D stock) [tax credit induced]	0.220*** (5.86)	0.218*** (5.86)	0.221*** (5.90)	0.216*** (5.86)	0.225*** (5.83)
ln(1+Non-AI patent counts)	-0.268*** (-3.21)	-0.160*** (-2.89)	-0.124*** (-2.63)	-0.159*** (-2.89)	-0.191*** (-3.00)
ln(1+R&D spending)	-0.132*** (-3.06)	-0.133*** (-3.06)	-0.143*** (-3.23)	-0.140*** (-3.17)	-0.113*** (-2.67)
Innovation dummy variable	0.149* (1.69)	-0.062 (-0.81)	-0.133 (-1.45)	-0.088 (-1.00)	-0.085 (-0.99)
Control variables?	Yes	Yes	Yes	Yes	Yes
FF48 industry fixed effects?	Yes	Yes	Yes	Yes	Yes
Observations	997,382	997,382	997,394	997,394	956,844
F-statistic for instrument	36.3	48.8	41.2	35.3	44.5
p-value of Hansen J-statistic	0.450	0.440	0.445	0.446	0.452

Panel C: Instrumented AI Patent Counts (IV) and Returns in Year t+2 to Year t+5					
Dependent variable is monthly stock return					
Scaling variable for AI patent counts					
	None	Total assets	Total patent stock	Total patent counts	R&D stock
Year t+2					
ln(0.1+AI patent counts) [instrumented]	0.736** (2.22)	0.319** (2.07)	0.321** (2.48)	0.279** (2.51)	0.234* (1.74)
Observations	910,899	910,830	910,953	910,899	857,221
Year t+3					
ln(0.1+AI patent counts) [instrumented]	1.322** (2.45)	0.644** (2.47)	0.600** (2.38)	0.558** (2.37)	0.705* (1.90)
Observations	835,074	834,948	835,247	835,074	774,358
Year t+4					
ln(0.1+AI patent counts) [instrumented]	0.379 (0.71)	0.509* (1.83)	0.664** (2.25)	0.563** (2.39)	0.950 (1.55)
Observations	757,043	756,889	757,392	757,043	692,923
Year t+5					
ln(0.1+AI patent counts) [instrumented]	0.653 (1.20)	0.212 (0.59)	0.317 (1.21)	0.387* (1.68)	0.587** (2.00)
Observations	676,957	676,797	677,429	676,957	612,509

Table 7
The Effect of AI Production on the Producer Firm's Profitability

This table shows the results of regressions of cash flow levels on AI production. AI patent counts are instrumented with the interaction between the producer firm's R&D stock and its own AI exposure as well as that of its customers. The sample period spans 1990-2017 in terms of year t . AI patent counts are measured in year t . They are instrumented with R&D stock measured at year $t-2$ and AI exposure fixed before the start of the sample period. Outcomes are measured in year $t+1$ and subsequent years. The sample and specifications are described in the text. In Panel B, AI patent stock is measured at $t-3$ during the previous 10 years. In Panel C and Panel D, positive AI patent stock is sorted into below and above its median (respectively, "low" and "high"). Variables are defined in Appendix Table 2. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Current AI Patent Counts Only, Without Conditioning on Past AI Patent Stock			
	Dependent variable is return on assets		
	Year $t+1$	Year $t+2$	Year $t+3$
$\ln(0.1 + \text{AI patent counts})$ [instrumented]	0.074** (2.18)	0.034 (1.33)	0.019 (0.82)
$\ln(1 + \text{R\&D stock})$ [tax credit induced]	0.006** (2.57)	0.007*** (3.18)	0.006*** (2.86)
$\ln(1 + \text{Non-AI patent counts})$	-0.038*** (-2.89)	-0.017* (-1.72)	-0.008 (-0.87)
Control variables?	Yes	Yes	Yes
Fixed effects			
State $\times$ Year?	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes
Firm?	Yes	Yes	Yes
SIC2 industry $\times$ Year?	Yes	Yes	Yes
Observations	92,275	85,721	79,813
F-statistic for instrument	25.7	25.0	25.3
p-value of Hansen J-statistic	0.991	0.312	0.036

Panel B: Conditioning on Past AI Patent Stock Being Positive			
	Dependent variable is return on assets		
	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	0.081*** (2.66)	0.036 (1.59)	0.023 (1.04)
ln(0.1+AI patent counts) [instrumented] × Dummy variable for past AI patent stock is positive	0.008 (1.28)	0.010 (1.62)	0.007 (1.19)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	92,275	85,721	79,813
Panel C: Conditioning on Past AI Patent Stock Counts Being Positive and Separately Above the Median			
	Dependent variable is return on assets		
	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	0.086*** (2.78)	0.039* (1.69)	0.026 (1.14)
ln(0.1+AI patent counts) [instrumented] × Dummy var. for past AI patent stock is positive	0.000 (0.05)	0.003 (0.49)	0.004 (0.57)
ln(0.1+AI patent counts) [instrumented] × Dummy variable for past AI patent stock is high	0.016** (2.27)	0.012* (1.73)	0.005 (0.80)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	92,275	85,721	79,813
Panel D: Conditioning on Past AI Patent Stock Ratio Being Positive and Separately Above the Median			
	Dependent variable is return on assets		
	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	0.085*** (2.75)	0.040* (1.71)	0.023 (1.03)
ln(0.1+AI patent counts) [instrumented] × Dummy var. for past AI patent stock is positive	-0.004 (-0.61)	-0.007 (-0.97)	-0.008 (-1.03)
ln(0.1+AI patent counts) [instrumented] × Dummy variable for past AI patent stock is high	0.024** (2.56)	0.031*** (3.15)	0.030*** (3.07)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	92,275	85,721	79,813

Table 8
The Effect of AI Production on the Producer Firm's Risk

This table shows the results of regressions of cash flow and stock return volatilities on AI production. AI patent counts are instrumented with the interaction between the producer firm's R&D stock and its own AI exposure as well as that of its customers. The sample period spans 1990-2017 in terms of year t . AI patent counts are measured in year t . They are instrumented with R&D stock measured at year $t-2$ and AI exposure fixed before the start of the sample period. Outcomes are measured in year $t+1$ and subsequent years. The sample and specifications are described in the text. In Panel B, AI patent stock is measured at $t-3$ during the previous 10 years. In Panel C and Panel D, positive AI patent stock is sorted into below and above its median (respectively, "low" and "high"). Variables are defined in Appendix Table 2. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Current AI Patent Counts Only, Without Conditioning on Past AI Patent Stock						
	ln(Volatility of return on assets)			ln(Volatility of stock returns)		
	Year $t+1$	Year $t+2$	Year $t+3$	Year $t+1$	Year $t+2$	Year $t+3$
ln(0.1+AI patent counts) [instrumented]	-0.710*** (-4.39)	-0.454*** (-3.38)	-0.470*** (-3.45)	-0.229*** (-4.21)	-0.178*** (-3.34)	-0.187*** (-3.34)
ln(1+R&D stock) [tax credit induced]	-0.013 (-0.84)	-0.019 (-1.41)	-0.011 (-0.77)	-0.002 (-0.43)	-0.001 (-0.19)	-0.002 (-0.39)
ln(1+Non-AI patent counts)	0.278*** (4.42)	0.170*** (3.20)	0.179*** (3.30)	0.082*** (3.94)	0.051** (2.45)	0.053** (2.34)
Control variables?	Yes	Yes	Yes	Yes	Yes	Yes
Fixed effects						
State $\times$ Year?	Yes	Yes	Yes	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes	Yes	Yes	Yes
Firm?	Yes	Yes	Yes	Yes	Yes	Yes
SIC2 industry $\times$ Year?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	91,251	84,909	78,943	91,854	89,416	81,842
F-statistic for instrument	25.7	25.0	25.3	25.9	26.4	25.6
p-value of Hansen J-statistic	0.732	0.247	0.017	0.780	0.455	0.190

Panel B: Conditioning on Past AI Patent Stock Being Positive						
	Dependent variable is					
	ln(Volatility of return on assets)			ln(Volatility of stock returns)		
	Year t+1	Year t+2	Year t+3	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	-0.657*** (-5.17)	-0.397*** (-3.39)	-0.425*** (-3.53)	-0.200*** (-4.11)	-0.149*** (-2.98)	-0.157*** (-3.01)
ln(0.1+AI patent counts) [instrumented] × Dummy var. for past AI patent stock is positive	-0.062*** (-2.62)	-0.064*** (-2.72)	-0.067*** (-2.72)	-0.021*** (-2.70)	-0.029*** (-3.62)	-0.034*** (-3.96)
Control variables and fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	91,251	84,909	78,943	91,854	89,416	81,842

Panel C: Conditioning on Past AI Patent Stock Counts Being Positive and Separately Above the Median						
Dependent variable is						
	ln(Volatility of return on assets)			ln(Volatility of stock returns)		
	Year t+1	Year t+2	Year t+3	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	-0.651*** (-5.07)	-0.403*** (-3.43)	-0.422*** (-3.49)	-0.189*** (-3.84)	-0.135*** (-2.70)	-0.146*** (-2.79)
ln(0.1+AI patent counts) [instrumented] × Dummy var. for past AI patent stock is positive	-0.053* (-1.88)	-0.036 (-1.24)	-0.042 (-1.42)	-0.009 (-0.97)	-0.012 (-1.30)	-0.018* (-1.74)
ln(0.1+AI patent counts) [instrumented] × Dummy variable for past AI patent stock is high	-0.024 (-0.74)	-0.037 (-1.14)	-0.029 (-0.91)	-0.017 (-1.58)	-0.022** (-2.01)	-0.016 (-1.37)
Control variables and fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	91,251	84,909	78,943	91,854	89,416	81,842
Panel D: Conditioning on Past AI Patent Stock Ratio Being Positive and Separately Above the Median						
Dependent variable is						
	ln(Volatility of return on assets)			ln(Volatility of stock returns)		
	Year t+1	Year t+2	Year t+3	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	-0.668*** (-5.20)	-0.415*** (-3.52)	-0.428*** (-3.56)	-0.193*** (-3.89)	-0.144*** (-2.83)	-0.153*** (-2.89)
ln(0.1+AI patent counts) [instrumented] × Dummy var. for past AI patent stock is positive	-0.074** (-2.35)	-0.074** (-2.31)	-0.094*** (-2.69)	-0.014 (-1.21)	-0.028** (-2.29)	-0.033** (-2.46)
ln(0.1+AI patent counts) [instrumented] × Dummy variable for past AI patent stock is high	0.031 (0.79)	0.020 (0.50)	0.048 (1.17)	-0.007 (-0.48)	0.001 (0.05)	-0.001 (-0.07)
Control variables and fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	91,251	84,909	78,943	91,854	89,416	81,842

Table 9
The Effect of AI Production on the Producer Firm's Stock Return Volatility: Decomposition

This table shows the results of the same regressions of stock return volatility as in Table 8 Panel A but with slight modifications, as indicated.

Panel A: Systematic Component of Stock Return Volatility from Fama and French (2015) Five-Factor Model			
	Dependent variable is ln(Stock return volatility)		
	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	-0.316*** (-4.54)	-0.292*** (-4.12)	-0.306*** (-4.11)
ln(1+R&D stock) [tax credit induced]	0.007 (0.98)	0.003 (0.34)	-0.002 (-0.21)
ln(1+Non-AI patent counts)	0.089*** (3.13)	0.075** (2.56)	0.086*** (2.74)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	91,870	89,418	81,844
F-statistic for instrument	25.9	26.4	25.6
p-value of Hansen J-statistic	0.531	0.316	0.235
Panel B: Idiosyncratic Component of Stock Return Volatility Fama and French (2015) Five-Factor Model			
	Dependent variable is ln(Stock return volatility)		
	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	-0.186*** (-3.43)	-0.136** (-2.52)	-0.146** (-2.56)
ln(1+R&D stock) [tax credit induced]	-0.007 (-1.26)	-0.005 (-0.86)	-0.005 (-0.85)
ln(1+Non-AI patent counts)	0.067*** (3.23)	0.035* (1.66)	0.036 (1.57)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	91,870	89,418	81,844
F-statistic for instrument	25.9	26.4	25.6
p-value of Hansen J-statistic	0.747	0.905	0.500
Panel C: Beta from CAPM			
	Dependent variable is beta		
	Year t+1	Year t+2	Year t+3
ln(0.1+AI patent counts) [instrumented]	-0.346*** (-4.75)	-0.360*** (-4.66)	-0.339*** (-4.33)
ln(1+R&D stock) [tax credit induced]	0.008 (1.12)	0.000 (0.05)	-0.008 (-0.94)
ln(1+Non-AI patent counts)	0.070** (2.33)	0.089*** (2.75)	0.091*** (2.79)
Control variables and fixed effects?	Yes	Yes	Yes
Observations	91,870	89,418	81,844
F-statistic for instrument	25.9	26.4	25.6
p-value of Hansen J-statistic	0.107	0.080	0.163

Table 10
Mechanisms Underlying the Effect of AI Production: The Producer Firm's Labor Productivity

This table shows the results of regressions of labor productivity measures on AI production. AI patent counts are instrumented with the interaction between the producer firm's R&D stock and its own AI exposure as well as that of its customers. The sample period spans 1990-2017 in terms of year t . AI patent counts are measured in year t . They are instrumented with R&D stock measured at year $t-2$ and AI exposure fixed before the start of the sample period. Outcomes are measured in year $t+1$. The sample and specifications are described in the text. Variables are defined in Appendix Table 2. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is		
	Profit per employee	ln(Employment / Total assets)	ln(Total assets)
ln(0.1+AI patent counts) [instrumented]	0.073* (1.94)	-0.075 (-0.84)	-0.030 (-0.73)
ln(1+R&D stock) [tax credit induced]	0.005** (2.41)	0.019** (2.16)	-0.002 (-0.54)
ln(1+Non-AI patent counts)	-0.027* (-1.91)	0.043 (1.26)	0.021 (1.29)
Control variables?	Yes	Yes	Yes
Fixed effects			
State $\times$ Year?	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes
Firm?	Yes	Yes	Yes
SIC2 industry $\times$ Year?	Yes	Yes	Yes
Observations	90,899	90,741	92,323
F-statistic for instrument	25.2	26.1	25.6
p-value of Hansen J-statistic	0.957	0.760	0.886

Table 11
Mechanisms Underlying the Effect of AI Production: The Producer Firm's Capital Intensity

This table shows the results of regressions of capital intensity and investments on AI production. AI patent counts are instrumented with the interaction between the producer firm's R&D stock and its own AI exposure as well as that of its customers. The sample period spans 1990-2017 in terms of year t . AI patent counts are measured in year t . They are instrumented with R&D stock measured at year $t-2$ and AI exposure fixed before the start of the sample period. Outcomes are measured in year $t+1$. The sample and specifications are described in the text. Variables are defined in Appendix Table 2. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is $\ln(\bullet / \text{Total assets})$				
	Property, plant, and equipment	Inventory	Capex	R&D spending	Acquisitions expenditures
$\ln(0.1 + \text{AI patent counts})$ [instrumented]	-0.236** (-2.41)	-0.335** (-2.09)	-0.377*** (-2.83)	-0.636*** (-3.50)	0.348* (1.67)
$\ln(1 + \text{R\&D stock})$ [tax credit induced]	-0.002 (-0.16)	0.039*** (2.66)	-0.027** (-2.18)	0.167*** (9.00)	0.011 (0.50)
$\ln(1 + \text{Non-AI patent counts})$	0.099** (2.55)	0.145** (2.28)	0.117** (2.21)	0.269*** (3.74)	-0.103 (-1.24)
Control variables?	Yes	Yes	Yes	Yes	Yes
Fixed effects					
State $\times$ Year?	Yes	Yes	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes	Yes	Yes
Firm?	Yes	Yes	Yes	Yes	Yes
SIC2 industry $\times$ Year?	Yes	Yes	Yes	Yes	Yes
Observations	91,860	92,323	92,316	92,332	92,336
F-statistic for instrument	26.0	25.6	25.6	25.6	25.6
p-value of Hansen J-statistic	0.116	0.967	0.367	0.733	0.839

Table 12
Mechanisms Underlying the Effect of AI Production: The Producer Firm's Bargaining Power

This table shows the results of regressions of the volatility of various production inputs and outputs on AI production. AI patent counts are instrumented with the interaction between the producer firm's R&D stock and its own AI exposure as well as that of its customers. The sample period spans 1990-2017 in terms of year t . AI patent counts are measured in year t . They are instrumented with R&D stock measured at year $t-2$ and AI exposure fixed before the start of the sample period. Outcomes are measured in year $t+1$. The sample and specifications are described in the text. Variables are defined in Appendix Table 2. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is				
	ln(Volatility of (Sales / Total assets))	Product differentiation	ln(Volatility of (Total costs / Total assets))	ln(Volatility of (SG&A / Total assets))	ln(Volatility of (COGS / Total assets))
ln(0.1+AI patent counts) [instrumented]	-0.471*** (-3.87)	0.500* (1.90)	-0.509*** (-4.06)	-0.355*** (-3.00)	-0.472*** (-3.79)
ln(1+R&D stock) [tax credit induced]	0.005 (0.42)	-0.023 (-0.81)	0.005 (0.41)	0.003 (0.27)	-0.003 (-0.26)
ln(1+Non-AI patent counts)	0.177*** (3.47)	-0.155 (-1.53)	0.187*** (3.73)	0.137*** (2.76)	0.174*** (3.47)
Control variables?	Yes	Yes	Yes	Yes	Yes
Fixed effects					
State $\times$ Year?	Yes	Yes	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes	Yes	Yes
Firm?	Yes	Yes	Yes	Yes	Yes
SIC2 industry $\times$ Year?	Yes	Yes	Yes	Yes	Yes
Observations	89,671	84,120	88,503	80,867	88,080
F-statistic for instrument	25.6	23.0	25.8	29.7	26.3
Hansen J-statistic	0.555	0.243	0.474	0.811	0.727
p-value of Hansen J-statistic	-0.471***	0.500*	-0.509***	-0.355***	-0.472***

Table 13
The Effect of AI Production on the Producer Firm's Financial Policies

This table shows the results of regressions of various financing variables on AI production. AI patent counts are instrumented with the interaction between the producer firm's R&D stock and its own AI exposure as well as that of its customers. The sample period spans 1990-2017 in terms of year t . AI patent counts are measured in year t . They are instrumented with R&D stock measured at year $t-2$ and AI exposure fixed before the start of the sample period. Outcomes are measured in year $t+1$. The sample and specifications are described in the text. Variables are defined in Appendix Table 2. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is				
	ln(Leverage)	ln(Cash holdings / Total assets)	Net debt issuance / Total assets	Equity issuance / Total assets	Share repurchases / Total assets
ln(0.1+AI patent counts) [instrumented]	0.283* (1.70)	-0.250** (-2.13)	0.024* (1.70)	-0.082** (-2.56)	0.009 (1.60)
ln(1+R&D stock) [tax credit induced]	0.020 (1.21)	0.002 (0.15)	0.002 (1.40)	-0.008*** (-2.88)	0.000 (0.30)
ln(1+Non-AI patent counts)	-0.045 (-0.68)	0.095** (2.06)	-0.004 (-0.69)	0.044*** (3.47)	-0.001 (-0.31)
Control variables?	Yes	Yes	Yes	Yes	Yes
Fixed effects					
State $\times$ Year?	Yes	Yes	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes	Yes	Yes
Firm?	Yes	Yes	Yes	Yes	Yes
SIC2 industry $\times$ Year?	Yes	Yes	Yes	Yes	Yes
Observations	92,323	92,322	92,337	92,337	92,337
F-statistic for instrument	25.6	25.6	25.6	25.6	25.6
Hansen J-statistic	0.402	0.136	0.327	0.746	0.766
p-value of Hansen J-statistic	0.283*	-0.250**	0.024*	-0.082**	0.009




Figure 1. Share of AI innovation in aggregate innovation activity. This figure shows the annual share of AI patent grants in all patent grants (AI and non-AI). Innovation activity is measured variously as patent counts, forward citations to patents, and the market value of patents.





Figure 2. Share of breakthrough innovations in AI vs. non-AI innovations. This figure shows the annual rate of patents that are classified as breakthrough patents. Classification is based on textual analysis identifying patent grants that are distinct from prior patents but related to subsequent patents.

Panel A: Share of Industries that Produce AI Innovation



Panel B: Share of Industries that Produce Innovation using Prior AI Innovation






Figure 3. Diffusion of AI innovation across industries. This figure shows the diffusion of AI innovation across industries (SIC3s). Panel A shows the share of industries with AI patent grants exceeding various thresholds, e.g., industries where AI patents account for at least 10% of all patents. Panel B shows the share of industries with patent grants that cite prior AI patents, where such citing patent grants exceed various thresholds, e.g., accounting for at least 10% of patents in the industry.

Panel A: Share of Publicly Traded U.S. Firms, Relative to All Patenting Entities, of AI vs. Non-AI Patents



Panel B: Share of Publicly Traded U.S. Firms, Relative to All U.S. Firms, of AI vs. Non-AI Patents



Panel C: Share of Innovative Publicly Traded U.S. Firms With At Least One AI Patent




Figure 4. The importance of publicly traded firms in AI innovation. This figure shows the share of publicly traded firms in AI patent grants separately from their share in non-AI patent grants. The share of publicly traded U.S. firms is calculated relative to all patenting entities (Panel A) and also relative to all U.S. firms (Panel B). The figure also shows, within the sample of publicly traded U.S. firms with at least one patent, the share of firms with at least one AI patent (Panel C).

Appendix Table 1
Illustrative Examples of AI Patents

This table shows illustrative examples of AI patents in the baseline sample restricted to industries with at least 10 firms per year every year during the sample period. Patent numbers and titles reference patents from the USPTO. Industries are ranked, from greatest to least, based on their total number of AI patents, as in Table 1.

Patent number	Patent title	Firm	SIC3	Industry AI patent rank
10043516	Intelligent automated assistant	Apple	366	5
8892487	Electronic synapses for reinforcement learning	IBM	737	1
8384776	Detection of topological structure from sensor data with application to autonomous driving in semi-structured environments	Toyota	371	9
3987279	Automatic performance reserve (APR) system	Boeing	372	8
8140069	System and method for determining the audio fidelity of calls made on a cellular network using frame error rate and pilot signal strength	Sprint	481	7
7231074	Method for determining the efficacy of an anti-cancer treatment using image analysis	Pfizer	283	10
7657494	Method for forecasting the production of a petroleum reservoir utilizing genetic programming	Chevron	291	11
7370001	Method and system of forecasting unscheduled component demand	Delta	451	29
4827426	Data acquisition and processing system for post-mix beverage dispensers	Coca-Cola	208	27
9218633	Cooking management	Starbucks	581	34

Appendix Table 2
Variable Definitions

Independent Variables Common to All Regressions	
Name	Definition
- AI patent counts	The number of AI patent grants during the 12 months before the fiscal yearend date
- Non-AI patent counts	The number of non-AI patent grants during the 12 months before the fiscal yearend date
- R&D stock [tax credit induced]	R&D spending predicted using the user cost of R&D implied by federal and state R&D tax credits, capitalized during the previous 10 years at a depreciation rate of 15%. See Section 4.1 for details.
- Firm's AI exposure	The producer firm's industry's labor's exposure to AI. AI exposure scores for each occupation are from Felten, Raj, and Seamans (2021) and aggregated at the industry level using employment shares between 1988 and 1990. Firms are assigned to the industry-level AI exposure of their primary industry. See Section 4.1 for details.
- Customers' AI exposure	The producer firm's customers' industries' labor's exposure to AI. AI exposure scores for each occupation are from Felten, Raj, and Seamans (2021) and aggregated at the industry level using employment shares between 1988 and 1990. Customer industries and their product purchase shares are identified using industry input-output tables. Customers' AI exposure is calculated as the product purchase weighted average of the industry-level AI exposures of customer industries. See Section 4.1 for details.
- Total assets	AT from Compustat
- Firm age	Years since the date the firm began trading publicly according to CRSP
- Innovation dummy variable	Dummy variable for whether the firm has at least one patent granted during the preceding 12 months
AI Patent Counts with Alternative Scaling Variables	
Name	Definition
- AI patent counts / Total assets	AI patent counts as defined above scaled by AT from Compustat (in \$ billions). Small constant added before taking logarithm: 0.01.
- AI patent counts / Total patent stock	AI patent counts as defined above scaled by the cumulative number of total patent (AI and non-AI) grants during the 10 years before the fiscal yearend date. Set to zero when the denominator is zero. Small constant added before taking logarithm: 0.0001.
- AI patent counts / Total patent counts	AI patent counts as defined above scaled by total patent (AI and non-AI) counts defined analogously during the same period. Set to zero when the denominator is zero. Small constant added before taking logarithm: 0.0001.
- AI patent counts / R&D stock	AI patent counts as defined above scaled by R&D stock as defined above but using actual (rather than predicted) R&D spending (in \$ billions). When the denominator is zero, set to zero and missing, respectively, if total (AI plus non-AI) patent counts is, respectively, zero and positive. Small constant added before taking logarithm: 0.01.
Variables Used in Stock Returns Analysis	
Name	Definition
- Monthly stock return	RET from CRSP
- Market capitalization	Stock price multiplied by shares outstanding from CRSP
- Market-to-book of equity	Market capitalization at the end of December from CRSP scaled by the book value of common equity in the same year from Compustat. The latter is constructed as the Compustat book value of stockholders' equity, plus balance-sheet deferred taxes and investment tax credit, minus the book value of preferred stock. See Fama and French (1993) for details.
- Momentum	Cumulative stock return during months [-12,-2]

- Short-term reversal	Stock return during the previous month from CRSP
- Stock price	Stock price from CRSP lagged by two months
Dependent Variables	
Name	Definition
Profitability	
- Return on assets	NI/AT from Compustat
Risk	
- Volatility of return on assets	Standard deviation of quarterly NIQ/AT during the 12 months after the fiscal yearend date. From Compustat.
- Volatility of stock returns	Standard deviation of daily stock returns during the 12 months after the fiscal yearend date. From CRSP.
Labor productivity	
- Profit per employee	NI/EMP from Compustat
- Employment / Total assets	EMP/AT from Compustat
Capital intensity	
- PP&E / Total assets	PPENT/AT from Compustat
- Inventory / Total assets	INVT/AT from Compustat. Small constant added before taking logarithm: 0.001.
- Capex / Total assets	CAPX/AT from Compustat. Small constant added before taking logarithm: 0.001.
- R&D / Total assets	XRD/AT from Compustat. Small constant added before taking logarithm: 0.001.
- Acquisitions / Total assets	AQC/AT from Compustat. Small constant added before taking logarithm: 0.001.
Bargaining power	
- Volatility of (Sales / Total assets)	Standard deviation of quarterly SALEQ/AT during the 12 months after the fiscal yearend date. From Compustat.
- Product differentiation	Hoberg and Phillips (2016) average product similarity score, subtracted from 1, multiplied by 100
- Volatility of (Total costs / Total assets)	Standard deviation of quarterly (COGSQ+XSGAQ)/AT during the 12 months after the fiscal yearend date. From Compustat.
- Volatility of (SG&A / Total assets)	Standard deviation of quarterly XSGAQ/AT during the 12 months after the fiscal yearend date. From Compustat.
- Volatility of (COGS / Total assets)	Standard deviation of quarterly COGSQ/AT during the 12 months after the fiscal yearend date. From Compustat.
- Volatility of stock returns	Standard deviation of daily stock returns during the 12 months before the fiscal yearend date. From CRSP.
Financial policies	
- Leverage	(DLC+DLTT)/AT from Compustat. Small constant added before taking logarithm: 0.01.
- Cash holdings / Total assets	CHE/AT from Compustat. Small constant added before taking logarithm: 0.01.
- Net debt issuance / Total assets	(DLCCH+DLTIS-DLTR)/AT from Compustat
- Equity issuance / Total assets	SSTK/AT from Compustat
- Share repurchase / Total assets	PRSTKC/AT from Compustat

Appendix Table 3
Example: Top Publicly Traded Firms by AI Patent Grants

This table shows the top 20 publicly traded firms by AI patent grants, along with their AI patent counts and industry classification.

AI patent counts (annual mean)	Firm	SIC3	Industry name
1,499	IBM	737	Computer programming, data processing, and other computer related
722	Microsoft	737	Computer programming, data processing, and other computer related
704	Google	737	Computer programming, data processing, and other computer related
297	HP	357	Computer and office equipment
277	GE	351	Engines and turbines
247	Intel	367	Electronic components and accessories
240	Facebook	737	Computer programming, data processing, and other computer related
227	HP	357	Computer and office equipment
205	Amazon	596	Non-store retailers
190	Xerox	357	Computer and office equipment
186	Oracle	737	Computer programming, data processing, and other computer related
170	AT&T	481	Telephone communications
158	Apple	366	Communications equipment
134	Lucent	737	Computer programming, data processing, and other computer related
131	Sun	357	Computer and office equipment
130	Qualcomm	367	Electronic components and accessories
101	Cisco	357	Computer and office equipment
89	Yahoo	737	Computer programming, data processing, and other computer related
77	Adobe	737	Computer programming, data processing, and other computer related
75	Verizon	481	Telephone communications

Appendix Table 4**Example: Top Occupations in, and Top Customer Industries of, the Computer Programming Industry**

This table shows, for the computer programming industry (SIC 737), the top 20 occupations (Panel A) and the top 20 customer industries (Panel B). Industry SIC 737 is chosen because it has the most AI patent grants.

Panel A: Top Occupations in Industry SIC 737			
Employment share (%)	Occupation name		AI exposure (percentile)
15.0	Computer programmers		89
8.8	Systems analysts		77
5.1	Computer engineers		86
4.7	General managers & top executives		79
4.6	Data entry keyers, except composing		67
3.3	Secretaries, except legal & medical		83
3.1	Computer operators, except peripheral equipment		73
2.9	Engineering, mathematical & natural sciences managers		82
2.3	Data processing equipment repairers		53
2.1	General office clerks		80
2.1	First line supervisors, clerical & administrative		82
2.0	Bookkeeping, accounting & auditing clerks		82
2.0	Salespersons, scientific products & services		77
2.0	Marketing/advertising/public relations managers		91
1.9	Sales agents, business services		91
1.9	Electrical & electronic engineers		83
1.9	Electrical/electronic technicians & technologists		61
1.8	Computer programmer aides		89
1.8	Other professional, paraprofessional/technicians		79
1.5	Other computer scientists & related		81
Panel B: Top Customer Industries of Industry SIC 737			
Product purchase share (%)	SIC3	Industry name	AI exposure (percentile)
13.2	737	Computer programming, data processing, and other computer related	93
12.9	602	Commercial banks	92
12.5	872	Accounting, auditing, and bookkeeping services	100
5.8	874	Management and public relations services	89
5.3	735	Miscellaneous equipment rental and leasing	69
3.8	801	Offices and clinics of doctors of medicine	87
3.8	603	Savings institutions	93
2.6	806	Hospitals	66
2.6	481	Telephone communications	81
2.6	871	Engineering, architectural, and surveying services	91
2.1	621	Security brokers, dealers, and flotation companies	99
1.5	802	Offices and clinics of dentists	65
1.4	491	Electric services	62
1.1	451	Air transportation, scheduled, and air courier services	56
1.1	606	Credit unions	94
0.8	541	Grocery stores	48
0.8	272	Periodicals: publishing, or publishing and printing	88
0.7	633	Fire, marine, and casualty insurance	96
0.6	631	Life insurance	96
0.6	822	Colleges, universities, professional schools, and junior colleges	86