Producing AI Innovation and Its Value Implications

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

Abstract

We document that artificial intelligence accounts for a significant and growing share of aggregate innovation produced during the past three decades, and is now diffuse across industries and technology fields. We then study publicly traded firms, finding that firms direct their production of innovation toward AI, motivated by their own, and their customers', labor's exposure to AI technology. We interact exogenously measured innovation capacity and AI exposure to instrument actual AI production. Our central findings are that producing AI increases a firm's future stock returns, supported by both higher profitability and lower risk. The results suggest that AI production increases firm value.

November 17, 2023

JEL classification: G12, G31, G32, J21, J24, O31, O32, O33
Keywords: Artificial intelligence; Innovation; Technology diffusion; Human capital; Productivity; Firm value; Asset pricing; Corporate investment

---

* 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, François Derrien, Charles Martineau, and Mikhail Simutin, and seminar participants at 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 during recent years. In this paper, we use newly available USPTO data on AI patents to study firms that *produce* AI innovation.[1] We document that, in the aggregate, AI is increasingly a prominent subset of all innovation activity during the past three decades. Even in 1990, AI accounted for 5% of all innovation activity, and has risen to 15%-35% of innovation today.[2] Additionally, AI innovation is consistently more valuable than non-AI innovation over time, in terms of both scientific and commercial value. Moreover, AI innovation has already diffused widely over time, across industries and technology fields, as expected from a general purpose technology. For instance, AI today accounts for at least half of all innovation in about 20% of industries. Finally, U.S. publicly traded firms dominate AI, producing close to half of all AI innovation.

Focusing on publicly traded firms during the past three decades, we then study the motivations for and value implications of producing AI innovation. We find that firms are motivated to direct their production of innovation toward AI based on how their own labor, as well as that of their customers, can be substituted or complemented by AI technology. This result is also the basis of our strategy for identifying the effect of AI innovation produced by the firm on the firm's future value, value drivers, and corporate operational outcomes, rather than studying the effect of endogenous actual AI innovation.

We continue to our central findings on value implication. We find that producing AI innovation increases a firm's future stock returns, incrementally to non-AI innovation. These

[1] By contrast, prior literature studies the *adoption* of previously developed AI technology, e.g., Alekseeva, Giné, Samila, and Taska (2020) and Babina, Fedyk, He, and Hodson (2022, 2023).
[2] The share of AI innovation is lowest when measured based on patent counts and highest based on market value of patents.

higher returns are supported by increasing future profitability, and decreasing future risk, with both being measured after, and as a result of, AI innovation successfully produced by the firm. Our final set of findings shows that the firm's AI innovation increases future labor productivity and decreases future physical capital intensity, among other mechanisms that we explore.

Turning to a more detailed exposition of our empirical analysis at the firm level, we begin with our strategy for identifying the causal effect of AI innovation. In so doing, we also examine the motivations of firms to produce AI innovation. We instrument actual AI innovation with two mutually reinforcing incentives for a firm to produce AI innovation: its plausibly exogenous innovation capacity (as induced by tax credits for R&D spending), and its exposure to AI technology (and/or that of its customers). By AI exposure, we are referring to the potential of AI technology to substitute or complement labor for the firm that uses it as a production input. The greater is the firm's innovation capacity, the more its exposure to AI (and/or its customers' AI exposure) will incentivize it to produce AI innovation; and conversely, the greater is the firm's exposure to AI (and/or its customers' AI exposure), the more its innovation capacity will incentivize it to produce AI innovation. Advancing from the relevance condition to the exclusion restriction, it is plausible that the firm's innovation capacity and its AI exposure (and/or its customers') *jointly* only affect the firm's outcomes (e.g., profitability, risk, etc.) through their effects on the firm's AI innovation.

We operationalize our interaction instrument as follows. We measure the firm's innovation capacity using its (tax credit induced) stock of R&D capital accumulated from annual R&D spending during the preceding 10 years. More precisely, we use annual R&D spending predicted from a firm-year model that exploits time-varying federal and state R&D tax credits to identify the plausibly exogenous component of R&D spending, which is then cumulated to construct the firm's plausibly exogenous R&D capital stock.[3] We lag R&D stock by two years to reflect the time it typically takes for patents to be granted.

To measure a firm's AI exposure, we use the AI exposure of labor in the firm's industry. Specifically, we use occupation-level (SOC) AI exposure scores, weighted by employment in each occupation in each industry (SIC3), to calculate industry-level AI exposure scores. Our data on AI exposure scores capture the extent to which labor in an occupation can be substituted for or complemented by AI technology. 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. We fix industry-level employment weights before the start of our sample period to maximize the exogeneity of AI exposure. Accounting for data limitations on R&D tax credits, we are able to instrument AI innovation produced from 1990 to 2017.

We use our interaction instrument to predict actual AI innovation. In this first stage of our instrumental variables analytical framework, we find that the mutually reinforcing incentives of a firm's innovation capacity and its AI exposure significantly predict the firm's AI innovation output, economically and statistically. For a typical increase in our interaction instrument, AI patent counts increase by about 13% relative to its mean, roughly similarly for both the firm's own AI exposure and its customers'.[4] We have no theoretical basis for favoring either exposure, whereas we can increase the precision of our IV estimates by using both exposures together.

[3] For instance, see Wilson (2009); Bloom, Schankerman, and Van Reenen (2013); and Hombert and Matray (2018).

[4] These first stage results and all IV results in 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 and 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).

Therefore, we use both instruments in our baseline analysis (and our results are robust to using either instrument alone).

We then examine the value implications of producing AI innovation, starting with stock returns during the year after production. In doing so, we exploit the general tendency of stock prices to be the best available (albeit imperfect) estimate of the discounted present value of future cash flows. Starting with actual AI patent grants, which are readily observable to investors, we find that a high minus low AI stock portfolio earns risk-adjusted returns of roughly 50 basis points per month. We also exploit our IV approach to more credibly identify the effect of AI production on firm value 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). These results are similar using a wide range of factor models.

We also implement our reduced form IV analysis in monthly Fama-MacBeth cross-sectional regressions using all sample firms. This more demanding approach allows us to include a battery of control variables like in our first stage IV regressions (e.g., non-AI patent counts) as well as established determinants of stock returns. By our construction, these Fama-MacBeth regression estimates are incremental to R&D stock and AI exposure. We again find a significant return spread when we double sort on the two components of our interaction instrument, a spread of comparable magnitude to the preceding portfolio analysis. The results of our stock returns analyses, taken together, are incremental to the economic factors captured by R&D stock or non-AI innovation.

We next examine the canonical financial drivers of value: cash flows and risk. From this point forward, we use detailed firm-level data over three decades, and we focus on our IV estimates to identify the causal effect of the firm's production of AI innovation. However, for all baseline analyses, we also report OLS estimates, which are generally smaller in economic magnitude (or negligible) and less statistically significant (often insignificant). Additionally, we are mindful that the full productivity potential of AI (e.g., as measured by profits) may not be realized by the time of writing, let alone throughout our three decade sample period.[5] Thus our estimates may understate the total potential of AI.

In our first IV results, we find that after successfully producing AI innovation (i.e., during the year after we measure AI patent grants), producer firms are more profitable. Specifically, for a typical 10% increase in instrumented AI patent counts, net income increases by roughly 0.8 percentage points relative to total assets, and similarly for net profit margin. Moreover, AI producer firms are less risky. A 10% increase in instrumented AI patent counts decreases the volatility of net return on assets, and similarly net profit margin, by about 5-6%. This lower cash flow volatility is also reflected in lower stock return volatility, which decreases by roughly 2%. Our findings of higher abnormal stock returns and lower risk (even systematic risk) are mutually consistent because we are examining corporate outcomes after a firm successfully produces innovation (i.e., after the firm is granted a patent, rather than when the firm begins spending on R&D).[6]

[5] 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.

[6] At the start of an innovation project, when a firm begins investing in innovation, uncertainty does increase. However, as the project progresses, and if and when the project is successful in producing innovation (e.g., the firm is granted a patent), then uncertainty decreases (and continues to do so if and as the firm successfully commercializes the innovation in subsequent years and decades).

We then provide suggestive evidence on various mechanisms through which AI innovation can increase firm value. First, since AI is a labor enhancing technology, successful AI innovation can increase the producer firm's labor productivity directly, or indirectly by helping the firm better satisfy customer demand. We find that AI innovation increases labor productivity, but we do not find any effect of AI innovation on employment or the overall scale of the firm, suggesting that AI complements, rather than substitutes, labor. Second, 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).

Third, AI can increase the firm's innovation capacity, with the firm using its past AI innovations to improve the efficiency of its future R&D investments. Future innovation output, produced at a lower future cost, can be higher for both AI and non-AI innovation.[7] This is indeed what we find. Fourth, AI can increase the firm's bargaining power vis-à-vis its customers, employees, and other business counterparties. The benefits of successful AI innovation can accrue to the producer firm's counterparties downstream and upstream in its supply chain (e.g., AI embedded products), improving the AI producer's bargaining power and thus increasing its own stability. The evidence, of more stable sales and costs, and greater product differentiation, is consistent with increased bargaining power.

In our final analysis, we examine the financial policy implications of successful AI innovation. Both an increase in expected future profitability and a decrease in risk 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.

[7] See Cockburn, Henderson, and Stern (2019) and Agrawal, McHale, and Oettl (2023) for a discussion.

We contribute to the literature on the economics of artificial intelligence. Focusing specifically on AI production, we show how successful AI innovation affects the producer firm's value. The existing literature focuses instead on AI adoption. Using corporate demand for labor with AI skills as proxied by job postings or resumes, existing studies document recent AI hiring trends in the labor market (Alekseeva, Azar, Giné, Samila, and Taska (2021)). They also find that AI adoption results in higher firm growth and product innovation (Alekseeva, Giné, Samila, and Taska (2020) and Babina, Fedyk, He, and Hodson (2023)) 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)).

Moreover, a recent literature on the productivity implications of AI innovation argues that AI may not enhance productivity as much as widely assumed, or its productivity enhancements will take much longer than expected to materialize (Brynjolfsson, Rock, and Syverson (2019)). Our examination of AI production to date provides evidence that, at least for AI producer firms, both forward looking stock prices and realized corporate operational outcomes already appear to be reflecting 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 producing AI innovation, incrementally to non-AI innovation, increases firm value and how this occurs. Other recent studies examine stock returns around specific 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)), and firm size (Stoffman, Woeppel, and Yavuz (2022)).

Finally, we contribute a novel identification methodology for the production of AI innovation. The existing literature focuses on identifying the causal effect of AI adoption. Babina, Fedyk, He, and Hodson (2023) 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) uses the launch of TensorFlow in an event study.

The rest of this paper is organized as follows. Section 2 characterizes AI innovation. Section 3 presents the methodology. Section 4 examines various motivations for AI production. Section 5 and Section 6 examine, respectively, the value implications and key value drivers of AI production. Section 7 and Section 8 examine the mechanisms underlying AI production and its financing implications, respectively. Section 9 concludes.

2. Characteristics of AI Innovation

We begin our empirical analysis with a simple characterization of AI innovation in the aggregate. 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. Suffice it to say in the present section that we have newly available USPTO data on AI and non-AI patents. We describe these data in Section 3.1. Here, we can use these data to characterize AI innovation during the period 1990-2020.

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]

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 2, both across industries (as captured by three-digit SIC codes) reflecting product markets, and across technology fields (as captured by the section and class of CPC codes). Starting with the share of AI patent grants across industries, already by 2020, in almost 20% of industries, AI accounts for at least half of all innovation (Panel A). Using the lower threshold of AI accounting for at least 10% of all patents, AI is present in roughly half of all industries.

Instead of examining AI content in current innovation itself, we can examine AI content in the prior innovation upon which current innovation builds. Looking at backward citations to patents, in close to 20% of industries by 2020, AI accounts for at least 50% of all citations to prior innovation (Panel B). Requiring only a minimum of at least 10% of all backward citations, AI is present is about three-quarters of all industries. Analysis of technology fields (Panel C and Panel D) leads to similar inferences, bearing in mind that AI is a technology, so it naturally remains concentrated in particular technology classes (Panel C).

Finally, we examine the importance of publicly traded firms in AI innovation as compared to all patenting entities. Motivating our analysis is the large share of aggregate R&D spending and patent grants attributable to publicly traded firms. We restrict our sample here to U.S. publicly traded firms.

[Insert Figure 3 about here]

Figure 3 Panel A shows that while publicly traded firms consistently account for only one quarter of non-AI patent grants during the past three decades, they account for almost half of all AI patent grants (roughly 45% during the past two decades). Furthermore, we can restrict our sample to innovative public firms, i.e., 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 3 Panel B 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, accounting for roughly half of it. Additionally, innovative publicly traded firms increasingly include AI in their innovation activities. These findings motivate our focus on publicly traded firms in the rest of the paper.

3. Methodology

3.1. Measurement of AI Production

To measure the production of AI innovation so that we can study its effects, we use patents that are classified, and to which we refer, as "AI patents" throughout the paper. As measures of successful 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).

We classify patents in the USPTO database as AI and non-AI using 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 diffuse across 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, Pairolero, and Toole (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. 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 classified as an AI patent.

3.2. Validation of AI Production Measure

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, Pairolero, and Toole (2022) carefully test their patent classification and show that it outperforms the existing alternatives.[8]

Additionally, we perform our own analysis of 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]

[8] 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.

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 2 Panel A, 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.

Our ranking in Table 1 is also broadly similar if, instead of 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.

As a final and 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 (2022). 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. The share of process claims, relative to all claims, can be used to estimate the "process intensity" of a patent. Using this empirical approach, we find that for AI patents, process intensity 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.

3.3. 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 6.

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. 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.

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.[9] 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), and Hombert and Matray (2018).

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

[9] 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.

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 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 avoid the endogeneity that could arise from 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. 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 2 Panel B 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 2 Panel C shows that the top 20 customers of the computer programming industry, ranked by product 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-statistic 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.[10] As we discuss in Section 6, 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 directly control for the components of the interaction, i.e., (tax credit induced) R&D capital stock

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

as well as AI exposure. Therefore, to affect our results, any confounding effect across firms with high versus low AI exposure would have to covary with annual changes in the user cost of R&D.

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 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.

3.4. 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:

$$\ln(0.1+AI\_Patent\_Counts_{i,SIC2,SIC3,s,t}) =$$

$$\alpha_1 \cdot R\&D\_Stock_{i,s,t-2} \times Firm's\_AI\_Exposure_{SIC3} +$$

$$\alpha_2 \cdot R\&D\_Stock_{i,s,t-2} \times Customers'\_AI\_Exposure_{SIC3} +$$

$$\alpha_3 \cdot R\&D\_Stock_{i,s,t-2} + \beta \cdot X_{i,t} + \delta_{s,t} + \delta_{SIC3} + \delta_i + \delta_{SIC2,t} \quad (1)$$

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

$$Outcome_{i,SIC2,SIC3,s,t+1} = \alpha \cdot \ln(0.1+AI\_Patent\_Counts_{i,t}) +$$

$$\beta_1 \cdot R\&D\_Stock_{i,s,t-2} + \beta_2 \cdot X_{i,t} + \delta_{s,t} + \delta_{SIC3} + \delta_i + \delta_{SIC2,t} \quad (2)$$

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 they are 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 3.3. 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 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 1. 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.

3.5. 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 1. 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.

4. Motivations for Producing AI Innovation and First Stage of IV Regressions

We begin our firm-level analysis by examining the various motivations for which firms produce AI innovation. Firms with both higher innovation capacity and higher AI exposure are more likely to produce AI innovation, whether AI exposure is considered for the producer firm or its customers. At the same time, this economic framework is also the econometric framework for our identification of the effect of AI production in subsequent analyses. This, the first stage of our instrumental variable regressions, is based on Equation 1.

[Insert Table 3 about here]

Table 3 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 motivations for AI production by including both instruments in our specification. 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% (=(0.024×1+0.040×1)×2.3) 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".[11] However, our second stage results in this paper do not depend critically on whether we use one instrument, the other, or both together.

Finally, we use our estimates in Table 3 to calculate the typical variation in AI production induced by our interaction instrument. 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% 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.

[11] 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. 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).

5. 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. We examine realized stock returns to exploit the general tendency of stock prices to be the best available (albeit imperfect) estimate of the discounted present value of future cash flows. This is particular relevant for a technology like AI the full productivity potential of which may only be realized in the future.

5.1. Returns on Portfolios Sorted by Actual AI Patent Counts

We begin by examining the returns on portfolios formed based 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.[12] 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 sort firms into three quasi-terciles (so called because they contain an uneven number of firms): zero AI patents ("zero AI", T1), below the median for non-zero AI patents ("low AI", T2), and above the median ("high AI", T3). Medians are recalculated every year at the time of portfolio formation. We are limited to sorting into these quasi-terciles because only 10% or so of

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

firm-years have non-zero AI patent counts.[13] We further sort the zero AI patents tercile into two groups: zero total patents (T1a) and non-zero total patents (T1b). We therefore have a total of four portfolios to examine. We also form hedge portfolios that longs the high AI portfolio (T3) and shorts either of the "zero AI" portfolios (T1a and T1b).

Our portfolios weightings are threefold: equally weighted, value weight, and size neutral. We use the size neutral approach as our baseline to mitigate the correlation between our sorting variable, AI patents, 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.[14] 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 5.3). The results indicate that AI patent counts spread returns, by about 50 basis points per month in our baseline specification (Panel C), for AI patent counts moving from T1 to T3. Spreads are somewhat higher when the short leg of the AI hedge portfolio is non-innovative firms (no patents) compared to innovative firms with no AI patents.

[13] We nevertheless have a reasonable number of stocks in T2 and T3: an average of about 175 and 150, and a minimum of about 75 and 50, respectively.
[14] For other applications of this approach, see Griffin and Lemmon (2002); Hirshleifer, Hsu, and Li (2013); and Liu, Stambaugh, and Yuan (2019).

The results suggest that firms with observable higher AI production have higher risk-adjusted returns.

5.2. 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 successful innovation, investor anticipation of successful innovation 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 two groups (very roughly, halves) based on (tax credit induced) R&D capital stock, the first component of our interaction instrument: zero R&D stock ("low", L), and non-zero R&D stock ("high", H). Independently, we also sort firms into quintiles based on AI exposure, the second component of our interaction instrument. 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.[15] Our results are similar if we continue to compare the top and bottom groups, but instead of quasi-halves for R&D stock and quintiles for AI exposure, we use more groups, as many as quintiles, for R&D stock, and less groups, as few as terciles, for 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 5.3). While we would generally expect R&D stock to spread returns while holding AI exposure fixed, and vice versa, our greatest interest is in the spread of the spread. 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 Q1 to Q5. The results for our baseline size neutral portfolios (Panel C) show that these changes result in higher returns of about 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

[15] In our baseline portfolio sorts, we again have a reasonable number of stocks: a mean of about 200-400, and a minimum of roughly 100.

(Table 3). Let us consider the same increases in R&D stock (low to high, equal to roughly 3.7 units of ln(1+R&D stock)) and AI exposure (Q1 to Q5, equal to about 2.9 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+AI patent counts), or a 60% increase relative to the mean (=0.06×3.7×2.9).[16] 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 positive affect AI production (Table 3), result in higher risk-adjusted returns.

5.3. Robustness Tests for Portfolio Returns Analyses

We directly eliminate the possibility of a 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., 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 Panels A through E of both Appendix Table 3 (c.f. Table 4) and Appendix Table 4 (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

[16] Simply as a reference point, moving from T1 to T3 in Table 4 equals about 4.5 units of ln(0.1+AI patent counts).

would also be consistent with a firm's total risk being lower as a result of successful AI innovation (which we also document, in Section 6.2).

Our inferences are also similar if we redo the results of Table 5 and Appendix Table 4 (both doubled sorted by R&D stock and AI exposure) using, separately, each of our two interaction instruments (i.e., based on the firm's AI exposure versus that of its customers). Consistent with a decrease in precision from using only one instrument or the other, the results are somewhat less economically and statistically significant (not tabulated).

5.4. Fama-MacBeth Cross-Sectional Return Regressions

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, in Fama-MacBeth regressions, both our portfolio regressions sorted by actual AI patent counts (Table 4) and doubled sorted by R&D stock and AI exposure (Table 5). The erstwhile sorting variables are now our explanatory variables of interest. We again combine the respective AI exposures of the firm and its customers by taking their first principal component and using it as our baseline measure of AI exposure.

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, turnover, and firm age. As an alternative to unscaled non-AI patent counts and R&D spending, we also include these variables scaling 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 are not statistically significant alongside our battery of control variables (Columns 1 and 2). By contrast, the interaction instrument is economically and statistically significant (Columns 3 and 4). For a one standard deviation increase in each of the two mutually reinforcing incentives for a firm to produce AI innovation, i.e., R&D stock and AI exposure, returns increase by roughly 7 basis points per month (=0.030×2.3×1). 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 (Appendix Table 5). It is noteworthy that our estimates on our interaction instrument are, by construction, incremental to our estimates on R&D stock.

5.5. 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 Q1 to Q5. These changes in Table 5, converted to their equivalent magnitudes in Table 6, are equal to 3.7 and 2.9 units of our R&D stock and AI exposure variables, respectively. Therefore, in Table 6, changes in R&D stock and AI exposure that are comparable to the change in Table 5, result in higher returns of roughly 32 basis points per month (=0.030×3.7×2.9). This is similar to the risk-adjusted returns in Table 5, even without remarking on the battery of control variables included in Table 6.

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

We examine how successful production of AI innovation affects the key drivers of firm value: cash flow levels and cash flow risk. To this end, 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.

6.1. The Effect of AI Production on Cash Flow Levels

[Insert Table 7 about here]

We measure profitability using return on assets and profit margin (both with net income in the numerator). These findings suggest that AI production increases cash flow levels. Table 7 shows that a 10% increase in AI patent counts relative to its mean (i.e., a typical increase) increases profitability by 0.8 percentage points, similarly for return on assets and profit margin, which corresponds to roughly 2.5% of each dependent variable's standard deviation. Even though the full productivity potential of AI may not be realized by the time of writing, there is already some evidence of it during the past three decades that comprise our sample period.

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 the first column of Appendix Table 6 through 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-statistic indicates the estimated effects of AI patent counts are not significantly different using each instrument separately. We also redo all IV regressions implemented using each instrument separately, and again tabulate the results in Appendix Table

6 through Appendix Table 12, in the second and third columns. Our estimates are similar in magnitude.

6.2. The Effect of AI Production on Cash Flow Risk

To better understand the effect of AI production on cash flow risk of the producer firm, it is noteworthy that we study firms subsequent to the successful grant of their AI patents. Hence, we distinguish between the innovation effort and innovation outputs. Innovation effort, characterized by R&D investments, involves upfront investments with highly uncertain outcomes. By contrast, innovation output, measured by patent grants, moderates the initial uncertainty associated with innovation efforts.

[Insert Table 8 about here]

We measure risk using the volatilities, respectively, of quarterly return on assets and profit margin. Our findings suggest that AI production decreases cash flow risk. Table 8 shows that a 10% increase in AI patent counts relative to its mean decreases the volatilities of return on assets and profit margin by 6% and 5% relative to their respective means. We also examine the volatility of daily stock returns, and we find a confirmatory reduction of 2% relative to its mean. Similarly across all three measures, the increase in AI patent counts corresponds to roughly 3%-4.5% of the respective dependent variable's standard deviation.

7. Mechanisms

We investigate four possible mechanisms underlying the effect of AI production on the value of the producer firm: labor productivity, physical capital intensity, innovation capacity, 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 helps 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). 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.

7.1. 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 9 about here]

We first examine labor productivity, which we capture using profit per employee. Table 9 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). This finding suggests that AI production increases labor productivity.

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 9. 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.

7.2. 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 10 about here]

We therefore examine the capital intensity of AI producer firms along various dimensions. Table 10 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 10 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.

7.3. Innovation Capacity

AI technology also allows firms to imitate the sensory, cognitive, and decision making abilities of humans, and, therefore, to automate such tasks. Applied to innovation activities, AI enables producer firms to reduce their R&D inputs, for example, by enhancing the productivity of high skilled workers in R&D hubs. However, AI can also improve the resulting innovation outputs, for instance, in terms of scale and speed (producing more innovation, and doing so more quickly). (The producer firm can also be affected indirectly, motivated by its customers' AI exposure, for instance, if its AI embedded products stimulate demand for further innovation upgrades in existing or new products.) We have already found that AI production decreases future R&D spending. We now consider future innovation outputs, in absolute terms and relative to the corresponding inputs.

[Insert Table 11 about here]

We begin by examining innovation output. Table 11 Panel A shows that AI production, as measured by successful current AI patent counts (i.e., in year t), increases future innovation output, as captured by future patent counts (i.e., in year t+1). A 10% increase in AI patent counts results in roughly 8% higher total patent counts (AI plus non-AI) relative to its mean, which corresponds to about 3% of its standard deviation. For AI and non-AI patent counts, respectively, the increase is 14% and 5% relative to the mean, or 8% and 2% relative to the standard deviation.

We then turn to examining innovation efficiency, as captured by innovation output per corresponding dollar of R&D input. Specifically, we calculate innovation efficiency as the ratio of patent counts to R&D spending. More precisely, future patents (in year t+1) are scaled by

R&D (in year t-1), lagging R&D by two years to reflect the time it typically takes for patents to be granted. In this analysis only, we add $1,000 (the smallest increment in our data for R&D spending) to the R&D of every firm to avoid zero denominators and reducing the sample size by almost half.

Table 11 Panel B shows that current AI production increases future innovation efficiency. For a 10% increase in AI patent counts, innovation efficiency measured using total patent counts (AI plus non-AI) are higher by 18% relative to the mean, corresponding to 5% of a standard deviation. In the case of AI and non-AI patent counts, respectively, the increase is 10% and 20% relative to the mean, or 4% and 5.5% relative to the standard deviation.

7.4. 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 even increase the stability of the firm's production outputs and inputs, and thereby reduce firm 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 but 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, 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.

8. 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. Overall, AI production appears to increase financial structure aggressiveness.

9. 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. We find that AI producers have higher risk-adjusted stock returns.

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, successful AI production causes higher profitability and lower risk.

We document four mechanisms through which AI production affects firm value. AI production increases the firm's labor productivity; it decreases the firm's physical capital intensity; it increases the firm's capacity for more innovation in absolute terms and relative to innovation costs; and it increases the firm's bargaining power vis-à-vis its business counterparties. We also find that, consistent with the decrease in risk, firms adopt more aggressive financial structures.

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. AI production to date has been value enhancing for producer firms across several operational dimensions. This will likely continue in the near future.

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. 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.[18]

[18] 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).

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.[19] 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.

[19] 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.[20] 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).

[20] 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.

Agrawal, Ajay K., John McHale, and Alexander Oettl, 2023, Artificial intelligence and scientific discovery: A model of prioritized search, working paper.

Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska, 2021, The demand for AI skills in the labor market, *Labour Economics* 71, 102002.

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, 2023, Artificial intelligence, firm growth, and product innovation, forthcoming *Journal of Financial Economics.*

Bena, Jan, and Elena Simintzi, 2022, Machines could not compete with Chinese labor: Evidence from US firms' innovation, working paper.

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.

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).

Derrien, François, Ambrus Kecskés, and Phuong-Anh Nguyen, 2023, Labor force demographics and corporate innovation, *Review of Financial Studies* 36, 2797-2838.

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.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665-712.

Liu, Jianan, Robert F. Stambaugh, and Yu Yuan, 2019, Size and value in China, *Journal of Financial Economics* 134, 48-69.

Rock, Daniel, 2021, Engineering value: The returns to technological talent and investments in artificial intelligence, working paper.

Stoffman, Noah, Michael Woeppel, 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).

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.

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	310.9	397.25	6	1	3
2	367	Electronic components and accessories	139.6	226.50	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	97.9	99.14	8	4	5
5	366	Communications equipment	92.2	87.04	7	5	6
6	384	Surgical, medical, and dental instruments and supplies	138.4	74.36	13	6	8
7	481	Telephone communications	45.7	58.25	4	7	7
8	372	Aircraft and parts	20.4	52.36	1	9	4
9	371	Motor vehicles and motor vehicle equipment	45.5	45.89	9	10	10
10	283	Drugs	270.6	38.89	24	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	28.2	11.50	14	20	19
16	138	Oil and gas field services	30.4	11.36	15	22	24
17	873	Research, development, and testing services	39.4	7.21	20	13	12
18	284	Soap, detergents, and cleaning preparations; perfumes, cosmetics, and other toilet preparations	22.4	6.39	16	30	27
19	421	Trucking and courier services, except air	27.6	5.96	17	33	39
20	738	Miscellaneous business services	37.6	5.36	19	15	16

21	369	Miscellaneous electrical machinery, equipment, and supplies	20.8	5.21	18	14	13
22	799	Miscellaneous amusement and recreation services	29.4	4.68	21	24	21
23	483	Radio and television broadcasting stations	25.2	4.57	22	16	15
24	356	General industrial machinery and equipment	39.6	3.93	26	17	20
25	874	Management and public relations services	24.9	2.68	23	23	22
26	504	Professional and commercial equipment and supplies	30.6	1.68	30	21	23
27	208	Beverages	18.7	1.57	28	31	32
28	451	Air transportation, scheduled, and air courier services	19.9	1.57	27	26	26
29	506	Electrical goods	22.6	1.57	25	29	28
30	281	Industrial inorganic chemicals	16.4	1.50	29	28	25
31	131	Crude petroleum and natural gas	98.4	1.39	33	25	37
32	308	Miscellaneous plastics products	29.2	0.64	32	27	29
33	581	Eating and drinking places	68.5	0.64	37	40	40
34	495	Sanitary services	24.8	0.61	31	32	38
35	809	Miscellaneous health and allied services, not elsewhere classified	18.1	0.25	34	34	34
36	736	Personnel supply services	26.4	0.21	36	37	33
37	331	Steel works, blast furnaces, and rolling and finishing mills	27.3	0.14	35	36	30
38	735	Miscellaneous equipment rental and leasing	15.6	0.07	40	39	35
39	153	Operative builders	22.8	0.04	39	35	31
40	565	Family clothing stores	15.3	0.04	38	38	36

Table 2
Descriptive Statistics

This table presents descriptive statistics for the main variables used in this paper. Variables are defined in Appendix Table 1.

	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)	127.7	493.1	0.0	1.9	40.3
- 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)	2,307	6,830	62	260	1,201
- Firm age (years)	15	15	4	10	21
- Innovation dummy variable	0.33	0.47	0.00	0.00	1.00
Variables used in stock returns analysis					
- Monthly stock return (%)	0.66	16.37	-7.82	-0.22	7.73
- Market capitalization ($ M)	2,314	7,481	46	217	1,075
- Market-to-book of equity	2.7	4.0	0.8	1.4	2.8
- Momentum (%)	9.7	60.6	-26.9	0.6	30.9
- Short-term reversal (%)	0.70	16.31	-7.84	-0.24	7.75
- Stock price ($)	19.9	22.8	4.3	12.0	27.2
Dependent variables: Profitability					
- Return on assets	-0.057	0.279	-0.084	0.026	0.079
- Profit margin	-0.093	0.327	-0.080	0.021	0.067
Dependent variables: Risk					
- Volatility of return on assets	0.03	0.07	0.01	0.01	0.03
- Volatility of profit margin	0.74	3.63	0.02	0.04	0.16
- Volatility of stock returns (%)	4.18	2.70	2.31	3.44	5.18
Dependent variables: Labor productivity					
- Profit per employee ($ 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)	2,423	7,100	63	279	1,293
Dependent variables: Capital intensity					
- PP&E / Total assets	0.26	0.22	0.08	0.19	0.37
- Inventory / Total assets	0.14	0.16	0.01	0.09	0.21
- 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	0.092	0.000	0.000	0.007
Dependent variables: Innovation capacity					
- Total patent counts / Total assets (per $ B)	10.6	31.6	0.0	0.0	4.2
- AI patent counts / Total assets (per $ B)	0.67	3.25	0.00	0.00	0.00
- Non-AI patent counts / Total assets (per $ B)	9.6	29.4	0.0	0.0	3.2
- Total innovation efficiency	117.46	629.27	0.00	0.00	0.13
- AI innovation efficiency	11.50	106.58	0.00	0.00	0.00
- Non-AI innovation efficiency	96.40	519.79	0.00	0.00	0.11

Dependent variables: Bargaining power					
- Volatility of (Sales / Total assets)	0.039	0.046	0.011	0.023	0.049
- Product differentiation (%)	96.9	2.2	96.3	97.5	98.3
- Volatility of (Total costs / Total assets)	0.035	0.044	0.009	0.020	0.042
- Volatility of (SG&A / Total assets)	0.010	0.016	0.002	0.004	0.011
- Volatility of (COGS / Total assets)	0.028	0.037	0.006	0.015	0.034
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	0.268	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. The sample and specifications are described in the text. Variables are defined in Appendix Table 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is ln(0.1+AI patent counts)		
	(1)	(2)	(3)
ln(1+R&D stock) [tax credit induced] × Firm's AI exposure	0.052***		0.024*
	(6.44)		(1.90)
ln(1+R&D stock) [tax credit induced] × Customers' AI exposure		0.057***	0.040***
		(6.58)	(3.01)
ln(1+R&D stock) [tax credit induced]	0.037***	0.038***	0.038***
	(2.61)	(2.61)	(2.62)
ln(1+Non-AI patent counts)	0.381***	0.381***	0.381***
	(12.48)	(12.50)	(12.51)
Innovation dummy variable	0.165***	0.165***	0.165***
	(2.59)	(2.59)	(2.59)
ln(Total assets)	0.051***	0.051***	0.051***
	(6.31)	(6.31)	(6.34)
ln(Firm age)	-0.036***	-0.038***	-0.038***
	(-3.82)	(-4.07)	(-4.07)
Fixed effects			
State × Year?	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes
Firm?	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes
Observations	91,868	91,868	91,868
Adjusted R^2	0.702	0.702	0.702

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. Firms are sorted into three quasi-terciles: zero, low, and high AI patents (T1, T2, and T3). The zero AI patents tercile is further sorted into two groups: zero innovation (T1a) and non-zero innovation (T1b). 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: Equally Weighted Portfolios

	AI = 0 & Innovation = 0 (T1a)	AI = 0 & Innovation > 0 (T1b)	Low AI (T2)	High AI (T3)	T3 – T1a	T3 – T1b
MKT	0.99***	1.10***	1.15***	1.17***	0.18***	0.07*
	(23.02)	(33.22)	(24.90)	(23.78)	(3.96)	(1.82)
SMB	0.83***	0.84***	0.62***	0.35***	-0.48***	-0.50***
	(12.48)	(16.25)	(7.36)	(3.89)	(-7.60)	(-8.86)
HML	0.19***	-0.05	-0.16**	-0.07	-0.26***	-0.02
	(2.76)	(-0.96)	(-2.08)	(-0.84)	(-4.45)	(-0.34)
RMW	-0.14*	-0.28***	-0.43***	-0.32**	-0.18	-0.04
	(-1.77)	(-5.13)	(-3.64)	(-2.19)	(-1.46)	(-0.35)
CMA	-0.17	0.12	0.02	0.12	0.29*	-0.01
	(-1.21)	(1.38)	(0.13)	(0.72)	(1.87)	(-0.04)
Alpha	0.07	0.20	0.36**	0.49***	0.42***	0.29**
	(0.49)	(1.56)	(2.20)	(3.71)	(3.28)	(2.17)
Observations	336	336	336	336	336	336
Adjusted R^2	0.868	0.917	0.873	0.847	0.295	0.314

Panel B: Value Weighted Portfolios

	AI = 0 & Innovation = 0 (T1a)	AI = 0 & Innovation > 0 (T1b)	Low AI (T2)	High AI (T3)	T3 – T1a	T3 – T1b
MKT	1.11***	1.03***	0.97***	1.01***	-0.09**	-0.02
	(35.89)	(39.83)	(31.30)	(27.46)	(-2.04)	(-0.35)
SMB	0.25***	0.02	-0.17***	-0.20***	-0.46***	-0.22***
	(5.08)	(0.41)	(-2.91)	(-3.75)	(-6.94)	(-3.39)
HML	0.09	-0.09**	-0.17***	-0.16**	-0.25***	-0.07
	(1.41)	(-2.24)	(-3.16)	(-2.51)	(-3.57)	(-0.78)
RMW	0.20***	0.17***	0.00	-0.15	-0.35***	-0.31***
	(2.79)	(3.42)	(0.03)	(-1.51)	(-5.27)	(-3.39)
CMA	-0.22**	0.16**	0.18*	-0.13	0.09	-0.30
	(-2.30)	(2.10)	(1.77)	(-0.91)	(0.78)	(-1.55)
Alpha	-0.19**	-0.09	0.01	0.23**	0.42***	0.32**
	(-2.42)	(-1.18)	(0.07)	(2.07)	(3.40)	(2.11)
Observations	336	336	336	336	336	336
Adjusted R^2	0.907	0.892	0.804	0.864	0.306	0.180

Panel C: Size Neutral Portfolios						
	AI = 0 & Innovation = 0 (T1a)	AI = 0 & Innovation > 0 (T1b)	Low AI (T2)	High AI (T3)	T3 – T1a	T3 – T1b
MKT	1.10***	1.09***	1.07***	1.12***	0.02	0.03
	(37.47)	(46.86)	(34.00)	(19.48)	(0.36)	(0.53)
SMB	0.51***	0.45***	0.33***	0.38***	-0.13	-0.07
	(9.81)	(9.52)	(5.65)	(4.18)	(-1.63)	(-0.92)
HML	0.14**	-0.05	-0.21***	-0.18*	-0.32***	-0.12
	(2.26)	(-1.33)	(-4.65)	(-1.92)	(-4.52)	(-1.44)
RMW	0.13	0.01	-0.26***	-0.26*	-0.39***	-0.28**
	(1.61)	(0.22)	(-3.12)	(-1.68)	(-3.63)	(-2.48)
CMA	-0.23**	0.11	0.05	0.04	0.27*	-0.07
	(-2.14)	(1.47)	(0.42)	(0.19)	(1.69)	(-0.38)
Alpha	-0.17**	-0.03	0.21**	0.39***	0.56***	0.41***
	(-2.10)	(-0.34)	(2.11)	(2.65)	(4.32)	(2.85)
Observations	336	336	336	336	336	336
Adjusted R^2	0.925	0.931	0.889	0.810	0.215	0.117

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. Firms are sorted into two groups based on R&D capital stock: zero R&D stock ("low", L), and non-zero R&D stock ("high", H). Independently, firms are sorted into quintiles based on AI exposure 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: Equally Weighted Portfolios						
	Exposure Q1 (low)	Exposure Q2	Exposure Q3	Exposure Q4	Exposure Q5 (high)	Exposure (Q5 – Q1)
Low R&D stock (L)	-0.28**	-0.43**	-0.11	-0.23	0.15	0.42***
	(-1.98)	(-2.43)	(-0.71)	(-1.07)	(0.85)	(2.89)
High R&D stock (H)	-0.14	0.03	0.60**	0.55**	0.86***	1.00***
	(-1.17)	(0.30)	(2.47)	(2.46)	(4.01)	(4.02)
R&D stock (H – L)	0.13	0.47***	0.70***	0.79***	0.71***	0.58***
	(1.37)	(2.88)	(2.68)	(3.23)	(3.75)	(2.72)
Panel B: Value Weighted Portfolios						
	Exposure Q1 (low)	Exposure Q2	Exposure Q3	Exposure Q4	Exposure Q5 (high)	Exposure (Q5 – Q1)
Low R&D stock (L)	-0.33***	-0.30**	-0.43***	-0.55***	-0.04	0.29**
	(-3.22)	(-2.50)	(-3.34)	(-3.05)	(-0.34)	(2.15)
High R&D stock (H)	-0.30***	0.05	0.10	0.15	0.53***	0.83***
	(-2.64)	(0.45)	(0.91)	(0.93)	(3.60)	(3.74)
R&D stock (H – L)	0.03	0.36*	0.52***	0.69***	0.57***	0.54**
	(0.32)	(1.89)	(3.45)	(3.09)	(3.15)	(2.47)
Panel C: Size Neutral Portfolios						
	Exposure Q1 (low)	Exposure Q2	Exposure Q3	Exposure Q4	Exposure Q5 (high)	Exposure (Q5 – Q1)
Low R&D stock (L)	-0.35***	-0.39***	-0.36***	-0.55***	0.03	0.37***
	(-3.53)	(-3.13)	(-3.11)	(-2.97)	(0.21)	(2.89)
High R&D stock (H)	-0.31***	-0.05	0.24*	0.20	0.61***	0.93***
	(-2.68)	(-0.49)	(1.85)	(1.40)	(3.75)	(4.31)
R&D stock (H – L)	0.03	0.34**	0.61***	0.76***	0.59***	0.55***
	(0.43)	(2.27)	(3.70)	(3.59)	(4.10)	(2.98)

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. 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, turnover, and firm age. Variables are defined in Appendix Table 1. 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.

	Dependent variable is monthly stock return			
	(1)	(2)	(3)	(4)
ln(0.1+AI patent counts)	0.002	0.016		
	(0.11)	(0.93)		
ln(1+R&D stock) [tax credit induced]			0.031***	0.028***
× AI exposure			(3.07)	(2.85)
AI exposure			0.009	0.008
			(0.15)	(0.15)
ln(1+R&D stock) [tax credit induced]			0.209***	0.116***
			(5.83)	(4.93)
ln(1+Non-AI patent counts)	-0.081**		-0.066*	
	(-2.08)		(-1.74)	
ln(1+R&D spending)	0.138***		-0.097**	
	(3.78)		(-2.43)	
ln(0.1+Non-AI patent counts/TA)		0.037*		0.047**
		(1.83)		(2.23)
ln(0.001+R&D spending/TA)		0.112***		0.024
		(3.36)		(0.63)
Innovation dummy variable	0.125	-0.130*	0.078	-0.228***
	(1.60)	(-1.95)	(1.04)	(-3.21)
Other control variables?	Yes	Yes	Yes	Yes
FF48 industry fixed effects?	Yes	Yes	Yes	Yes
Observations	993,135	993,135	993,135	993,135
R^2	0.090	0.091	0.091	0.092

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. The sample and specifications are described in the text. Variables are defined in Appendix Table 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is	
	Return on assets	Profit margin
ln(0.1+AI patent counts) [instrumented]	0.075**	0.075**
	(2.18)	(2.08)
ln(1+R&D stock) [tax credit induced]	0.006**	0.010***
	(2.51)	(3.43)
ln(1+Non-AI patent counts)	-0.038***	-0.029**
	(-2.89)	(-2.06)
Control variables?	Yes	Yes
Fixed effects		
State × Year?	Yes	Yes
SIC3 industry?	Yes	Yes
Firm?	Yes	Yes
SIC2 industry × Year?	Yes	Yes
Observations	91,868	90,366
F-statistic for instrument	24.9	25.2
Hansen J-statistic	0.00	0.62
p-value of Hansen J-statistic	0.979	0.430

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

This table shows the results of regressions of cash flow volatility 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 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is ln(Volatility of •)		
	Return on assets	Profit margin	Stock returns
ln(0.1+AI patent counts) [instrumented]	-0.610***	-0.523***	-0.231***
	(-3.78)	(-2.64)	(-4.20)
ln(1+R&D stock) [tax credit induced]	-0.005	-0.049***	-0.002
	(-0.33)	(-2.70)	(-0.38)
ln(1+Non-AI patent counts)	0.243***	0.230***	0.083***
	(3.89)	(3.00)	(3.94)
Control variables?	Yes	Yes	Yes
Fixed effects	.	.	.
State × Year?	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes
Firm?	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes
Observations	90,763	88,384	91,457
F-statistic for instrument	25.0	25.0	25.1
Hansen J-statistic	0.13	0.41	0.10
p-value of Hansen J-statistic	0.719	0.524	0.754

Table 9
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 1. ***, **, 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]	0.075*	-0.081	-0.033
	(1.95)	(-0.90)	(-0.79)
ln(1+R&D stock) [tax credit induced]	0.005**	0.020**	-0.002
	(2.36)	(2.23)	(-0.53)
ln(1+Non-AI patent counts)	-0.027*	0.044	0.022
	(-1.93)	(1.30)	(1.34)
Control variables?	Yes	Yes	Yes
Fixed effects	.	.	.
State × Year?	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes
Firm?	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes
Observations	90,494	90,335	91,914
F-statistic for instrument	24.4	25.3	24.9
Hansen J-statistic	0.01	0.13	0.03
p-value of Hansen J-statistic	0.934	0.721	0.871

Table 10
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 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable is ln(• / Total assets)				
	Property, plant, and equipment	Inventory	Capex	R&D spending	Acquisitions expenditures
ln(0.1+AI patent counts) [instrumented]	-0.237**	-0.344**	-0.392***	-0.685***	0.355*
	(-2.39)	(-2.11)	(-2.85)	(-3.65)	(1.70)
ln(1+R&D stock) [tax credit induced]	-0.001	0.040***	-0.025*	0.167***	0.012
	(-0.09)	(2.68)	(-1.96)	(8.84)	(0.57)
ln(1+Non-AI patent counts)	0.099**	0.148**	0.121**	0.285***	-0.105
	(2.52)	(2.30)	(2.24)	(3.88)	(-1.27)
Control variables?	Yes	Yes	Yes	Yes	Yes
Fixed effects	.	.	.	.	.
State × Year?	Yes	Yes	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes	Yes	Yes
Firm?	Yes	Yes	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes	Yes	Yes
Observations	91,453	91,914	92,088	92,106	92,109
F-statistic for instrument	25.3	24.9	24.9	24.9	24.9
Hansen J-statistic	2.65	0.00	0.86	0.16	0.06
p-value of Hansen J-statistic	0.104	0.984	0.355	0.693	0.803

Table 11

Mechanisms Underlying the Effect of AI Production: The Producer Firm's Innovation Capacity

This table shows the results of regressions of innovation 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 1. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Innovation Output

	Dependent variable is ln(Patent counts / Total assets)		
	Total patents	AI patents	Non-AI patents
ln(0.1+AI patent counts) [instrumented]	0.774***	1.420***	0.485***
	(4.29)	(11.18)	(2.81)
ln(1+R&D stock) [tax credit induced]	0.031	0.001	0.029
	(1.46)	(0.05)	(1.50)
ln(1+Non-AI patent counts)	0.346***	-0.007	0.474***
	(4.73)	(-0.14)	(6.89)
Control variables?	Yes	Yes	Yes
Fixed effects	.	.	.
State × Year?	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes
Firm?	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes
Observations	89,773	89,773	89,773
F-statistic for instrument	25.2	25.2	25.2
Hansen J-statistic	2.03	0.12	2.15
p-value of Hansen J-statistic	0.155	0.727	0.142

Panel B: Innovative Efficiency

	Dependent variable is ln(• innovation efficiency)		
	Total innovation	AI innovation	Non-AI innovation
ln(0.1+AI patent counts) [instrumented]	1.822***	0.956***	1.973***
	(3.68)	(3.11)	(3.91)
ln(1+R&D stock) [tax credit induced]	-0.552***	-0.137***	-0.523***
	(-8.58)	(-4.55)	(-8.07)
ln(1+Non-AI patent counts)	-0.001	0.327***	-0.009
	(-0.01)	(2.62)	(-0.05)
Control variables?	Yes	Yes	Yes
Fixed effects	.	.	.
State × Year?	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes
Firm?	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes
Observations	89,647	89,647	89,647
F-statistic for instrument	25.3	25.3	25.3
Hansen J-statistic	2.98	1.55	4.33
p-value of Hansen J-statistic	0.084	0.212	0.038

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 1. ***, **, 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]	-0.344***	0.506*	-0.351***	-0.260**	-0.282***
	(-3.28)	(1.88)	(-3.29)	(-2.48)	(-2.65)
ln(1+R&D stock) [tax credit induced]	0.011	-0.024	0.012	0.009	0.003
	(1.07)	(-0.81)	(1.07)	(0.86)	(0.33)
ln(1+Non-AI patent counts)	0.134***	-0.156	0.132***	0.105**	0.110***
	(3.10)	(-1.52)	(3.15)	(2.44)	(2.59)
Control variables?	Yes	Yes	Yes	Yes	Yes
Fixed effects	.	.	.	.	.
State × Year?	Yes	Yes	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes	Yes	Yes
Firm?	Yes	Yes	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes	Yes	Yes
Observations	89,203	83,753	88,043	80,446	87,635
F-statistic for instrument	25.1	22.2	25.1	28.6	25.7
Hansen J-statistic	0.04	1.30	0.13	0.39	0.05
p-value of Hansen J-statistic	0.834	0.254	0.723	0.533	0.825

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 1. ***, **, 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]	0.287*	-0.251**	0.023*	-0.083***	0.008
	(1.71)	(-2.13)	(1.67)	(-2.58)	(1.57)
ln(1+R&D stock) [tax credit induced]	0.021	0.002	0.002	-0.008***	0.000
	(1.24)	(0.19)	(1.55)	(-2.82)	(0.33)
ln(1+Non-AI patent counts)	-0.046	0.095**	-0.004	0.045***	-0.001
	(-0.70)	(2.06)	(-0.64)	(3.48)	(-0.31)
Control variables?	Yes	Yes	Yes	Yes	Yes
Fixed effects	.	.	.	.	.
State × Year?	Yes	Yes	Yes	Yes	Yes
SIC3 industry?	Yes	Yes	Yes	Yes	Yes
Firm?	Yes	Yes	Yes	Yes	Yes
SIC2 industry × Year?	Yes	Yes	Yes	Yes	Yes
Observations	91,914	91,913	92,109	92,109	92,109
F-statistic for instrument	24.9	24.9	24.9	24.9	24.9
Hansen J-statistic	0.63	2.13	1.16	0.11	0.12
p-value of Hansen J-statistic	0.427	0.144	0.281	0.742	0.726

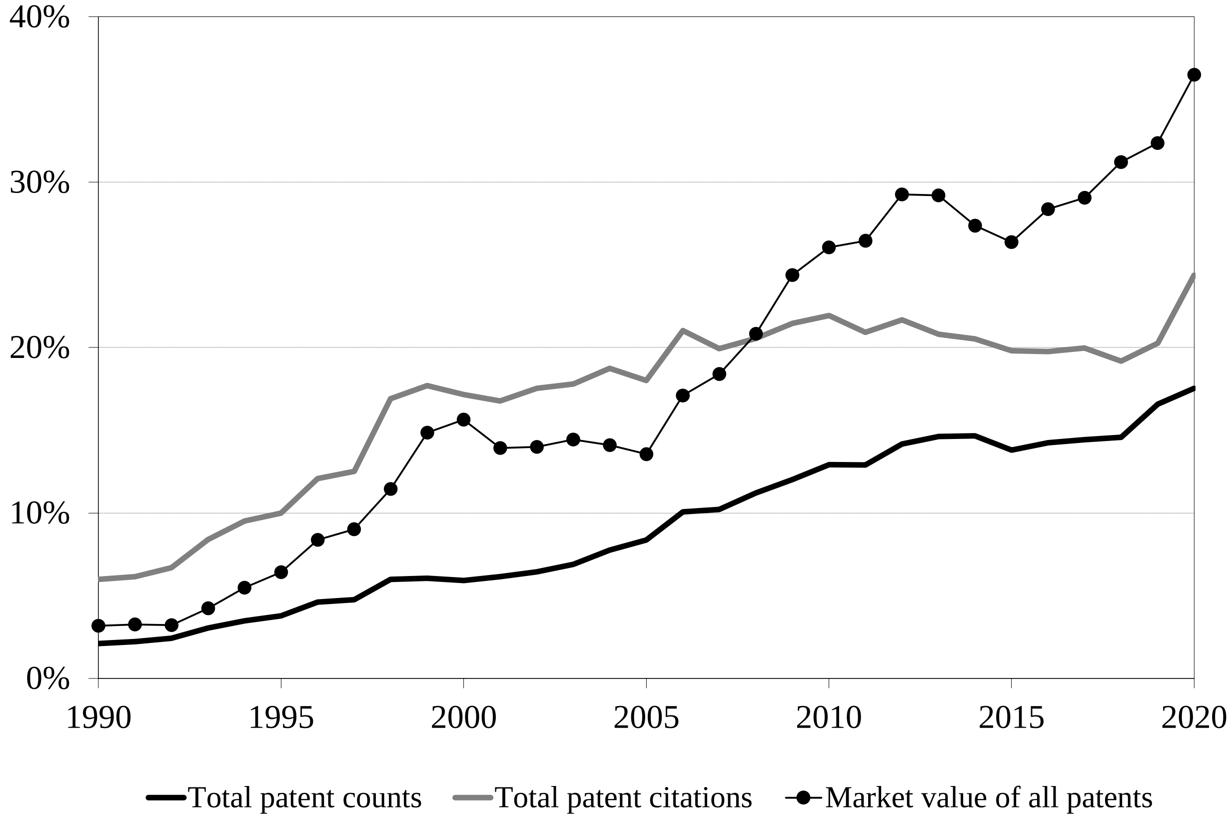

Figure 1. Share of AI innovation in aggregate innovation activity. This figure shows the annual share of AI patent grants in all patent grants (AI and non-AI). Innovation activity is measured variously as patent counts, forward citations to patents, and the market value of patents.

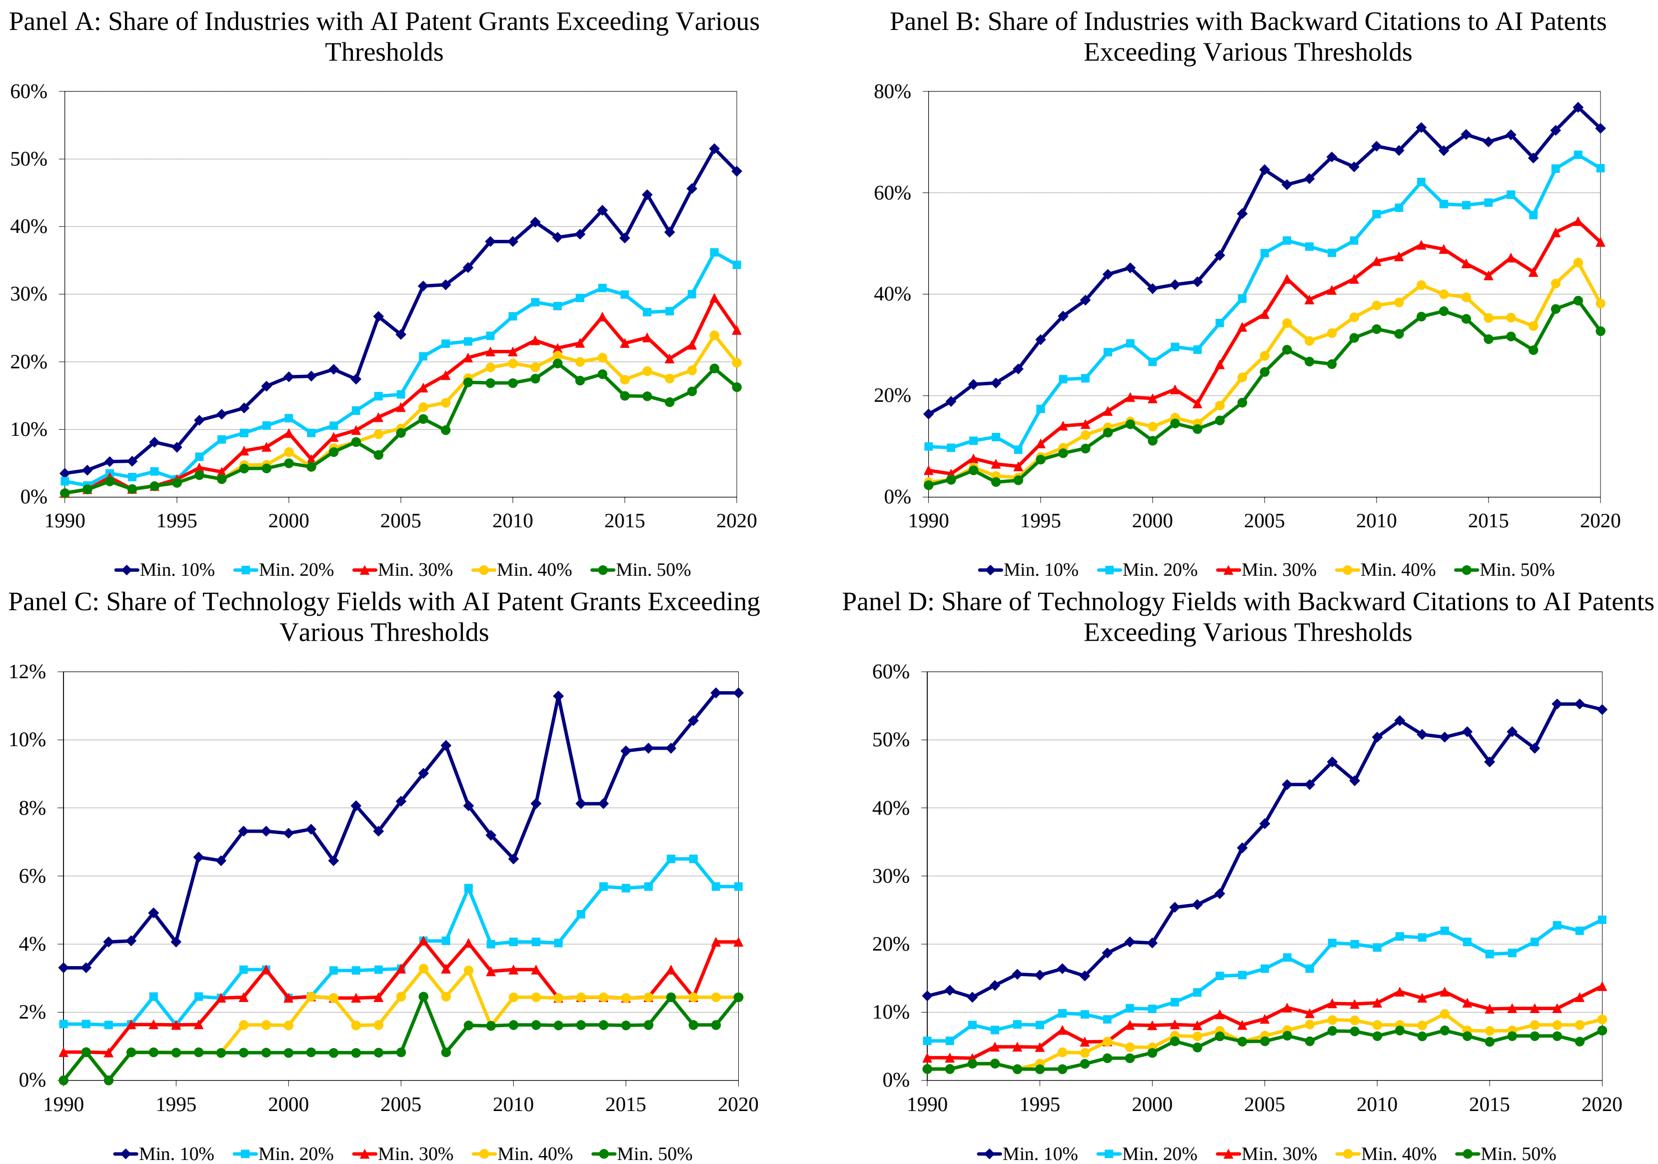

Figure 2. Diffusion of AI innovation across industries and technology fields. This figure shows the diffusion of AI innovation across industries (SIC3s) (Panels A and B) and technology fields (section and class of CPCs) (Panels C and D).

Panel A: Share of Publicly Traded Firms in AI vs. Non-AI Patents

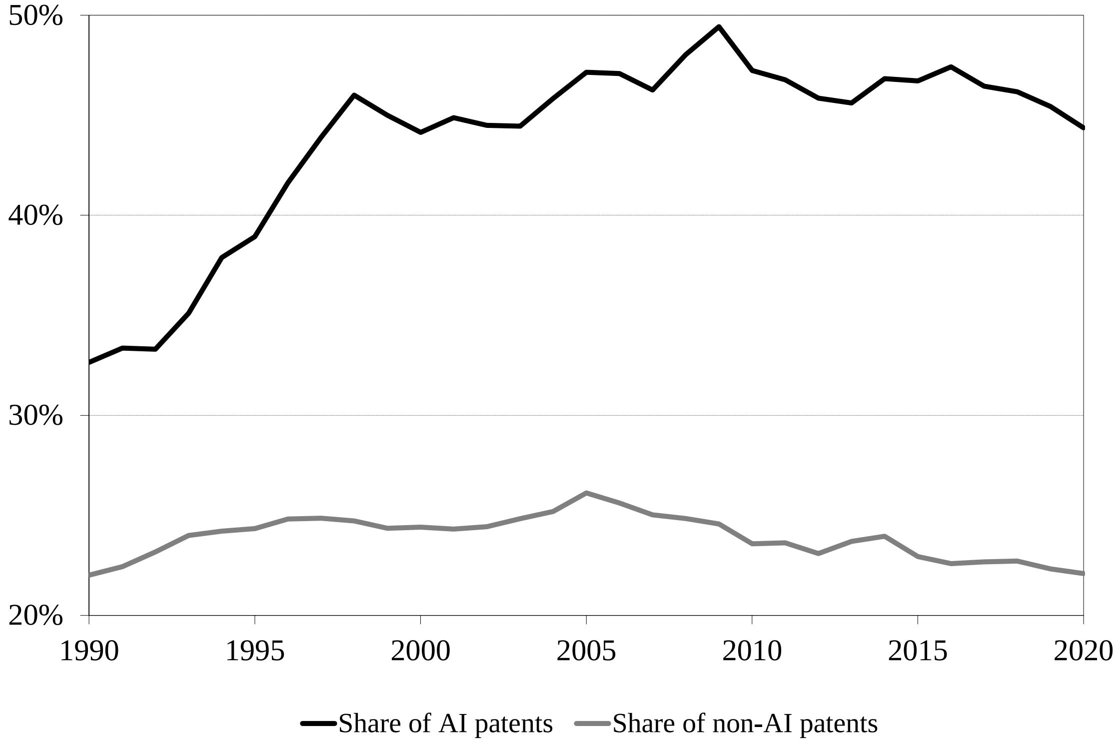

Panel B: Share of Innovative Publicly Traded Firms With At Least One AI Patent

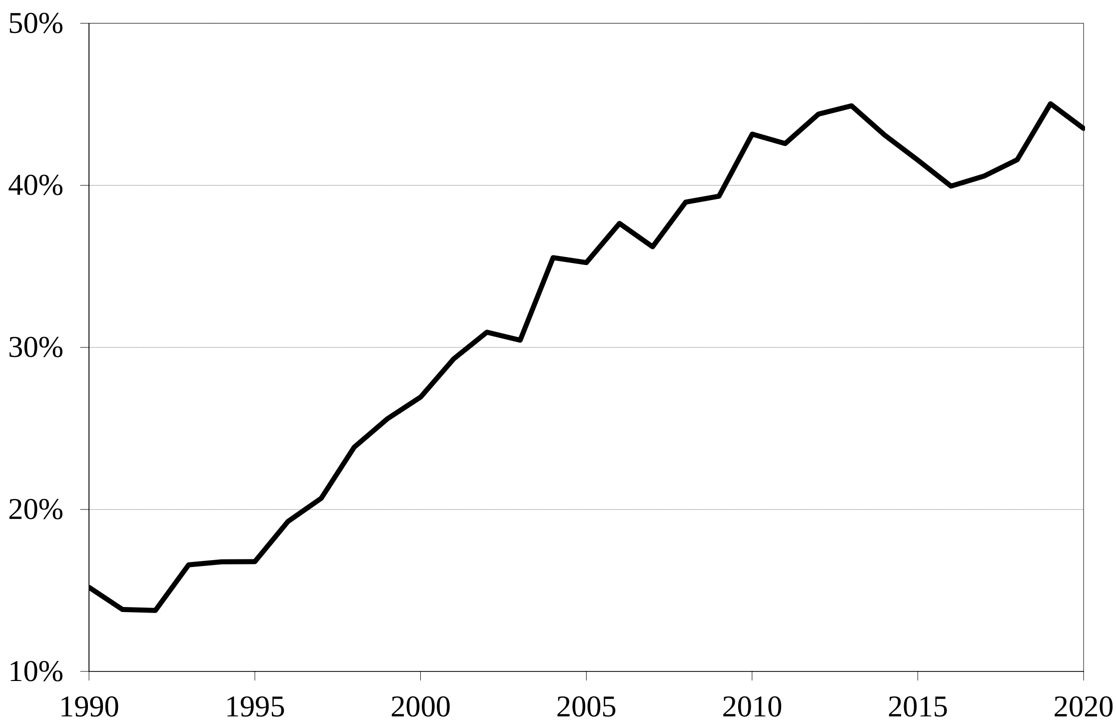

Figure 3. The importance of publicly traded firms in AI innovation. This figure shows the share of publicly traded firms in AI patent grants separately from their share in non-AI patent grants (Panel A). The figure also shows, with the sample of publicly traded firms with at least one patent, the share of firms with at least one AI patent (Panel B).

Appendix Table 1
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 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 3.3 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 3.3 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 3.3 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
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
- Profit margin	NI/SALE from Compustat. Values less than -1 are replaced with -1.
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 profit margin	Standard deviation of quarterly NIQ/SALEQ 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.
Innovation capacity	
- Total patent counts / Total assets	Total patent (AI and non-AI) grants during the 12 months before the fiscal yearend date, scaled by AT from Compustat (in $ billions). Small constant added before taking logarithm: 0.1.
- AI patent counts / Total assets	AI patents granted during the 12 months before the fiscal yearend date, scaled by AT from Compustat (in $ billions). Small constant added before taking logarithm: 0.01.
- Non-AI patent counts / Total assets	Non-AI patent grant during the 12 months before the fiscal yearend date, scaled by AT from Compustat (in $ billions). Small constant added before taking logarithm: 0.1.
- Total innovation efficiency	Total patent (AI and non-AI) grants during the 12 months before the fiscal yearend date. Scaled by XRD from Compustat (in $ millions) plus $1,000. Scalar is lagged by two years. Small constant added before taking logarithm: 0.001.
- AI innovation efficiency	AI patent grants during the 12 months before the fiscal yearend date. Scaled by XRD from Compustat (in $ millions) plus $1,000. Scalar is lagged by two years. Small constant added before taking logarithm: 0.0001.
- Non-AI innovation efficiency	Non-AI patent grants during the 12 months before the fiscal yearend date. Scaled by XRD from Compustat (in $ millions) plus $1,000. Scalar is lagged by two years. 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 2
Illustrative Examples

This table shows the top publicly traded firms by AI patent grants (Panel A), the top occupations in industry SIC 737 (Panel B), and the top customer industries of industry SIC 737 (Panel C). Industry SIC 737 is chosen because it has the most AI patent grants.

Panel A: Top 20 Firms by AI Patent Grants

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
703	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
189	Xerox	357	Computer and office equipment
186	Oracle	737	Computer programming, data processing, and other computer related
169	AT&T	481	Telephone communications
157	Apple	366	Communications equipment
134	Lucent	737	Computer programming, data processing, and other computer related
130	Sun	357	Computer and office equipment
129	Qualcomm	367	Electronic components and accessories
101	Cisco	357	Computer and office equipment
88	Yahoo	737	Computer programming, data processing, and other computer related
76	Adobe	737	Computer programming, data processing, and other computer related
75	Verizon	481	Telephone communications

Panel B: Top 20 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 C: Top 20 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