AI-Oriented Innovation, Productivity, and Market Structure: Evidence from U.S. Firms

Mahdiyeh Entezarkheir*

24 February 2026

Abstract

This paper estimates the causal effect of artificial intelligence (AI)–oriented innovation on firm productivity and examines how market structure shapes these returns. I construct a panel of publicly traded U.S. firms from 1980 to 2021 by linking Compustat, Discern patent–firm matches, PatentsView citation data, and the USPTO AI Patent Dataset. AI intensity is measured as the share of a firm's patents related to AI technologies; productivity is measured as labour productivity (sales per employee). Exploiting within-firm variation among patenting firms, I distinguish between the extensive margin of AI adoption and the intensive margin of AI orientation. Controlling for R&D, overall patenting, patent stocks, fragmentation, and firm and year fixed effects, I show that sustained increases in AI intensity significantly raise labour productivity, while discrete adoption episodes do not generate gains. The results indicate that productivity improvements reflect persistent reallocation of innovative effort toward AI rather than temporary experimentation or accumulation of AI patent stocks, highlighting the importance of the direction of technical change. To address endogeneity, I implement a shift–share (Bartik) instrument using firms' exposure to exogenous growth in AI technological opportunities. Instrumental-variable estimates confirm that the productivity effects are causal. Returns to AI-oriented innovation are larger for firms with broader AI engagement and are amplified in more concentrated industries and more fragmented technology markets. The findings show that the productivity impact of frontier technologies depends critically on market structure.

Keywords: Artificial Intelligence; Firm Productivity; Innovation; Patents; Market Structure; General-Purpose Technologies

JEL Classification Numbers: O33, O31, D24, L25, L13, O34

*Associate Professor, Department of Economics, Huron at University of Western Ontario, 1349 Western Rd, London, ON N6G 1H3, Canada. Email: menteza@uwo.ca.

1 Introduction

Artificial intelligence (AI) has emerged as one of the most consequential technologies of the modern economy. Rapid advances in machine learning, data analytics, and computing have expanded the range of automatable tasks, raising expectations of substantial productivity gains (Acemoglu et al., 2022; Acemoglu and Restrepo, 2020; Brynjolfsson et al., 2018; Aghion et al., 2017; Brynjolfsson and McAfee, 2014). Yet measured productivity growth has often appeared modest relative to visible AI progress (Brynjolfsson et al., 2021; Syverson, 2017; David, 1990). This "modern productivity paradox" suggests that AI's benefits may be delayed and uneven because they depend on complementary investments and gradual diffusion (Brynjolfsson et al., 2018).

This paper studies AI-oriented innovation, productivity, and market structure using a panel of publicly traded U.S. firms from 1980–2021. The dataset links Compustat financials, Discern patent–firm matches, PatentsView citation data, and the USPTO AI Patent Dataset. AI intensity is measured as the share of a firm's patents related to AI; productivity is measured as sales per employee. By focusing on shifts in AI intensity rather than total innovation scale, the analysis tests whether the direction of innovation predicts higher productivity.

The empirical strategy separates what firms innovate in from how much they innovate. Controlling for patent stock, citation-weighted patent stock, R&D stock per worker, and technological fragmentation (Ziedonis, 2004), and including firm and year fixed effects, identification relies on within-firm changes in AI orientation of innovation. Because AI intensity is defined as the share of AI patents in total patenting, the baseline focuses on patenting firms and identifies the intensive margin—how reallocating innovative effort toward AI affects productivity—rather than the discrete decision to adopt AI. This intensive-margin design aligns directly with models of directed technical change and the general-purpose technology (GPT) view, in which productivity gains arise from sustained technological reorientation and complementary investment rather than from discrete adoption episodes.

Innovation-based growth theory emphasizes that productivity depends not only on the amount of innovation but also on its direction (Acemoglu, 1998, 2002; Acemoglu et al., 2018; Aghion et al., 2016). In this framework, reallocating innovative effort toward AI may generate different productivity outcomes even when total innovation remains unchanged. The general-purpose technology (GPT) view further suggests that broad technologies such as AI yield gains only when deeply embedded through complementary investments in skills, data, and organizational change (Bresnahan and Trajtenberg, 1995; Bloom et al., 2012; Aghion et al., 2017; Brynjolfsson et al., 2018; Cockburn et al., 2018; Brynjolfsson et al., 2021). Together, these perspectives imply that sustained AI orientation—rather than one-time adoption—should drive productivity gains.

Empirical evidence on AI and productivity is expanding but remains incomplete. Survey-based studies generally find positive productivity effects of AI adoption, often contingent on complementary assets and organizational capabilities (Czarnitzki et al., 2023; Calvino and Fontanelli, 2023; Babina et al., 2024). Patent-based studies similarly link AI innovation to higher productivity (Cockburn et al., 2018; Damioli et al., 2021), but typically rely on shorter time horizons or narrower samples. Less is known about how firms' innovation portfolios evolve toward AI over long periods, how productivity effects differ across competitive environments, and how intellectual property structure shapes returns to AI-oriented innovation.

Market structure and intellectual property environments are central to understanding AI returns. Prior research shows that competition and appropriability shape innovation incentives and performance (Aghion et al., 2005; Teece, 1986). AI's scale economies in data and computing may reinforce concentration (Autor et al., 2020; De Loecker and Eeckhout, 2020; Korinek and Vipra, 2024; Gans, 2024). A further, closely related dimension that has been largely absent from the AI–productivity literature is the structure of intellectual property rights in the underlying technology space. A large body of work shows that fragmented patent ownership and "patent thickets" can raise transaction costs, increase bargaining frictions, and slow the commercialization of new technologies (Ziedonis, 2004). For complex technolo-

gies such as AI, which rely on many complementary and overlapping patents, fragmentation may be particularly harmful, as firms must navigate congested IP landscapes before they can effectively deploy new innovations. These considerations suggest that productivity gains from AI-oriented innovation are likely to vary across competitive and technological environments.

The results show that firms whose innovation portfolios tilt more strongly toward AI experience significantly higher labour productivity. When extensive and intensive margins are separated, gains arise from sustained increases in AI intensity rather than discrete adoption episodes. Fragmentation in technology markets lowers productivity overall but does not eliminate AI's benefits. This pattern suggests that AI can deliver robust efficiency gains even when institutional frictions impede access to complementary IP. The effects are also heterogeneous in ways that are consistent with the GPT view of AI and with market-structure considerations. Productivity gains from AI-oriented innovation are larger for firms with persistent AI activity, pointing to the importance of learning and organizational embedding. They increase with the breadth of a firm's AI engagement, measured by the number of distinct AI technology classes in which it patents, highlighting the value of diversified AI capabilities. Finally, the productivity gains associated with AI-oriented innovation are stronger in more concentrated industries, consistent with the role of appropriability, scale-related complementarities, and strategic deployment of AI in shaping competitive advantage.

A key concern in this setting is endogeneity: more productive firms may be more likely to pursue AI, and time-varying managerial or strategic factors may jointly drive AI-oriented patenting and productivity. To address endogeneity, I implement a shift–share (Bartik) instrument based on firms' exposure to exogenous growth in AI technological opportunities. IV estimates confirm a causal effect of AI-oriented innovation on productivity and exceed fixed-effects estimates, consistent with attenuation from measurement error and endogenous adoption decisions along the extensive margin.

This paper contributes to the literature in four ways. First, it provides long-horizon U.S.

evidence linking productivity to the direction of innovation toward AI. Second, it shows that market structure shapes AI returns, integrating the AI and competition literatures. Third, it documents that patent fragmentation lowers productivity but does not neutralize AI's gains. This shows that patent ownership fragmentation is an important drag on performance, but that AI can partially overcome these frictions through efficiency gains. Fourth, it demonstrates that persistence and breadth in AI innovation amplify productivity effects, reinforcing the view of AI as a GPT requiring sustained investment.

The findings have direct policy relevance. If AI is a GPT whose gains depend on persistence, breadth, competitive conditions, and the structure of intellectual property rights, then effective AI policy must go beyond encouraging more innovation. It should also support complementary investments in skills, data infrastructure, and organizational capital, while promoting competition and diffusion mechanisms that prevent AI advantages from becoming excessively concentrated—without undermining incentives for costly, long-run AI investment. In parallel, policies that reduce frictions from patent thickets—through clearer patent boundaries, improved licensing institutions, or mechanisms that lower transaction costs—may raise productivity directly and facilitate the translation of AI innovation into broad-based performance improvements.

2 Empirical Framework

2.1 Production Function and Direction of Innovation

To motivate the empirical specification, I begin from a standard firm-level Cobb-Douglas production function in which output is produced using inputs, following the knowledge-capital tradition in the empirical productivity literature (Griliches, 1979; Hall and Mairesse, 1995; Hall et al., 2010):

$$Y_{it} = A_{it} K_{it}^{\alpha} L_{it}^{1-\alpha}, \tag{1}$$

where Y_{it} denotes firm output, K_{it} capital, L_{it} labour, and A_{it} firm-specific efficiency. Dividing both sides by labour and taking logarithms yields

$$\log\left(\frac{Y_{it}}{L_{it}}\right) = \log A_{it} + \alpha \log\left(\frac{K_{it}}{L_{it}}\right). \tag{2}$$

Thus, labour productivity depends on capital deepening and firm-specific efficiency, which captures differences in technology, organization, and other factors not directly measured by inputs.

The central object of interest is therefore A_{it}. Following the knowledge-capital framework (Griliches, 1979; Hall et al., 2005), productivity is assumed to depend on the firm's accumulated stock of knowledge. However, innovation-based growth theory emphasizes that productivity depends not only on how much firms innovate, but also on the direction of technological change. In models of directed technical change, profit incentives shape the allocation of innovative effort across technological domains, and reallocating innovation toward particular technologies can generate different productivity outcomes even if total innovative effort remains unchanged (Acemoglu, 1998, 2002; Acemoglu et al., 2018; Aghion et al., 2016). In the context of artificial intelligence, this perspective implies that productivity gains depend critically on whether firms reorient their innovation portfolios toward AI technologies rather than merely expanding overall R&D (Acemoglu and Restrepo, 2020).

In parallel, the general-purpose technology (GPT) view argues that broad technologies such as electricity, ICT, and plausibly AI, translate into productivity gains only when firms develop complementary capabilities and embed the technology deeply into their production and organizational processes (Bresnahan and Trajtenberg, 1995; Helpman and Trajtenberg, 1994). Historical evidence from ICT diffusion shows that productivity effects arise primarily through organizational transformation, management practices, and the accumulation of intangible capital rather than through hardware adoption alone (Bresnahan et al., 2002; Brynjolfsson and Hitt, 2000, 2003; Bloom et al., 2012). Recent work argues that AI exhibits

the defining characteristics of a GPT, with potentially far-reaching effects on production and growth, but whose economic impact depends critically on complementary investments in data infrastructure, computing capacity, skills, and organizational capital (Cockburn et al., 2018; Aghion et al., 2017; Brynjolfsson et al., 2021).

Consistent with these perspectives, I model firm productivity as a function of the direction of innovation ($\log(AIShare)_{it}$), the level of innovative effort ($\log(CitePatSock/L)_{it}$ and $\log(R\&DStock/L)_{it}$), and institutional frictions ($\log F_{it}$):

$$\begin{aligned}\log A_{it} = \alpha_1 \log(AIShare)_{it} + \alpha_2 \log(CitePatSock/L)_{it} + \alpha_3 \log(R\&DStock/L)_{it} \\ + \alpha_4 \log F_{it} + X'_{it}\delta + \mu_i + \lambda_t + \varepsilon_{it}.\end{aligned} \tag{3}$$

Substituting equation (3) into equation (2) yields the baseline empirical specification:

$$\begin{aligned}\log(Productivity)_{it} = \alpha_1 \log(AIShare)_{it} + \alpha_2 \log(CitePatSock/L)_{it} + \alpha_3 \log(R\&DStock/L)_{it} \\ + \alpha_4 \log F_{it} + \alpha_5 \log(K/L)_{it} + X'_{it}\delta + \mu_i + \lambda_t + \varepsilon_{it}.\end{aligned} \tag{4}$$

It is important to clarify that equation (4) identifies the intensive margin of AI-oriented innovation—namely, changes in the share of AI patents within patenting firms. Because $AIShare_{it}$ is defined as the share of AI patents in total patenting activity, the coefficient α_1 captures changes in the composition of innovation within patenting firms. The baseline specification therefore identifies the intensive margin of AI-oriented innovation—how deeper AI orientation affects productivity—rather than the discrete decision to begin AI patenting. The extensive margin of AI adoption is examined separately in Section 4.

In the empirical implementation of equation (4), $Productivity_{it}$ is defined as sales per employee. Because sales equal price times physical output, the logarithm of sales per worker can be decomposed into the logarithm of physical output per worker and the logarithm of prices. Consequently, the dependent variable reflects both technological efficiency and potential variation in mark ups. This interpretation is consistent with the production-function

framework and recognizes that revenue-based productivity measures may incorporate both efficiency improvements and pricing power.[1]

The coefficient on $\log(\text{AIShare}_{it})$ therefore captures the association between reallocating innovative effort toward AI and revenue-based labour productivity, conditional on total innovative intensity, capital deepening, and other firm characteristics. Because $AIShare_{it}$ is defined as the share of AI patents within total patenting activity, this coefficient identifies changes in technological orientation within patenting firms. In this sense, the specification provides a reduced-form test of whether shifting the composition of innovation toward AI predicts productivity differences across firms and over time, isolating the intensive margin of AI-oriented innovation rather than the discrete decision to initiate AI patenting.

2.2 Innovation Measures and Patent Fragmentation

The key explanatory variable, AIShare_{it}, is defined as

$$AIShare_{it} = \frac{\text{Number of AI patents of firm } i \text{ in year } t}{\text{Total number of patents of firm } i \text{ in year } t}. \tag{5}$$

This measure captures the extent to which a firm's innovative activity is directed toward artificial intelligence rather than other technological domains. By construction, it reflects changes in the composition of innovation toward AI rather than the scale of overall patenting activity. The variable is constructed using the 86% probability threshold from the AIPD2023 data set. Alternative thresholds of 50% and 93% are considered in robustness analysis, and the main findings are robust to these alternatives.

AI innovation typically requires complementary investments in specialized technical expertise, structured datasets, computing infrastructure, and organizational redesign. As a result, a higher AI patent share proxies for deeper technological capabilities that are not

[1]Robustness checks using alternative measures of productivity—including total factor productivity constructed from a fixed-effects Solow residual, total factor productivity constructed from a structural Levinsohn–Petrin estimator implemented following Ackerberg et al. (2015), and value added per worker—are reported in Section 4.1. The findings are robust to these different productivity measures.

directly observable but are closely linked to productivity.

To control for overall innovative effort, I include the logarithm of citation-weighted patent stock per worker, $\log(CitePatSock/L)_{it}$, and the logarithm of R&D stock per worker, $\log(R\&DStock/L)_{it}$, as R&D and patent intensities. The citation-weighted patent stock is constructed using a declining-balance formula,

$$\text{CitePatSock}_{it} = (1-\delta)\text{CitePatSock}_{it-1} + flowCi_{it},$$

with a depreciation rate of $\delta = 15\%$ (Hall et al., 2000, 2005, 2007; Noel and Schankerman, 2013).[2] The variables CitePatSock_{it} and flowCi_{it} stand for citation stock and citation counts of firm i at time t, respectively.[3] Other research, such as Hasan et al. (2020) use the count of patents to measure innovation. However, I follow the majority of studies that view innovation as being best measured through weighting patents by citations and then building a depreciation rate that acknowledges the falling value of innovation over time (Hall et al., 2000, 2005, 2007; Lanjouw and Schankerman, 2004; Noel and Schankerman, 2013). Using citation-weighted patent stock allows us to take into account not only the output of the innovative activity but also the quality of the innovation via its citations. Many patents are not necessarily valuable and the incorporation of citations in measuring innovation as a proxy of patent quality partially addresses the problem of low value patents.[4] The R&D stock is constructed analogously using deflated R&D expenditures and a similar depreciation structure (Griliches, 1980; Hall et al., 2005; Bloom et al., 2013).[5]

As a result of a gap between the application and grant date of patents, the data on patents

[2]In the declining balance formula used for building the citation-weighted patent stock measure, most researchers settled on δ=15% (Hall et al. , 2000, 2005, and 2007). Hall and Mairesse (1995) apply alternative depreciation rates and conclude that changing the rate from 15% does not make a difference.

[3]I define the initial stock of citations as the initial sample values of citation counts similar to Noel and Schankerman (2013).

[4]In Section 4, I checked the robustness of main findings using a measure of innovation intensity, which is built upon the ratio of patent stocks rather than citation-weighted patent stocks to the number of employees. The findings are robust to the using of this measure of innovation.

[5]R&D measure is adjusted for inflation with the CPI urban U.S. index for 1992 (Source: http://www.bls.gov).

and their citations are truncated. Patents with an application date close to the end of the sample might be granted out of the reach of the sample. Similarly, patents in the sample might receive further citations outside the sample period. I use the methods employed by Hall et al. (2005) to correct for these truncations in building citation-weighted patent stocks (CitePatSock_{it}). For the correction of truncation procedure in patent and citation counts see the Appendix.

Patent thickets, situations in which multiple overlapping and complementary patents within a technological space are held by different owners, may impede productivity by increasing licensing complexity, transaction costs, and the risk of hold-up, particularly in cumulative technological domains (Shapiro, 2012; Entezarkheir, 2017). These frictions can delay technology implementation, discourage efficient reorganization of production, and divert managerial and technical resources toward legal and contractual activities rather than productive use. To capture these frictions, I include the logarithm of the patent ownership fragmentation index, $\log F_{it}$, following Ziedonis (2004). The fragmentation index is defined as

$$F_{it} = 1 - \sum_{j=1}^{J} \left(\frac{cite_{ijt}}{cite_{it}} \right)^2, \tag{6}$$

where $cite_{ijt}$ denotes the total number of citations made by firm i in all of its patents granted at time t to the patents of firm j, and $cite_{it}$ represents the total number of citations made by firm i in all of its patents granted in year t to patents of all other firms. Each citation in a patent document references a related patent. The lower limit of the index F_{it} equals $1/J$ (when all citations are made to the patents of one firm) and its upper limit equals $1 - (1/J)$ (when each citation references a patent from a different firm). Citations to a firm's own patents, which do not pose fragmentation threats, are excluded from the index calculation, resulting in missing values for F_{it} in such cases. An indicator variable, $MissingF_{it}$, is included to account for observations where the index cannot be computed. Higher values of F_{it} indicate more fragmented ownership in technology market and denser

patent thickets.[6]

2.3 Capital Deepening and Additional Controls

The logarithm of the capital-labor ratio, $\log(K/L)_{it}$, enters directly from the Cobb-Douglas production structure and captures capital deepening in equation (4), a central determinant of labour productivity in neoclassical growth models (Solow, 1956; Mankiw et al., 1992). It is measured as the ratio of property, plant, and equipment to the number of employees.

The vector X_{it} in equation (4) includes additional firm- and market-level characteristics. R&D spillovers can affect labour productivity in both positive and negative ways. On the positive side, knowledge generated by technologically proximate firms can diffuse across organizational boundaries, expanding the set of ideas, techniques, and problem-solving approaches available to the firm. When firms possess sufficient absorptive capacity, these external knowledge flows complement internal R&D, improve production processes, thereby increasing labour productivity (Griliches, 1979; Jaffe, 1986; Cohen and Levinthal, 1990). On the other hand, spillovers may reduce firm-level productivity if external R&D intensifies competitive pressures or weakens a firm's ability to appropriate returns from its own innovation efforts. In particular, R&D conducted by technologically close rivals may erode private gains from innovation, leading to lower measured productivity for incumbent firms even if aggregate efficiency improves (Bloom et al., 2013). I include the logarithm of R&D spillovers received by firm i at time t in X_{it}, defined as

$$\text{R\&DSpillover}_{it} = \sum_{j \neq i} \rho_{ij} \times R\&DStock_{jt}, \tag{7}$$

[6] As Ziedonis (2004) explained, this measure of patent ownership fragmentation has inherent limitations. Patent citations can only be observed if firms choose to patent their innovations and successfully obtain those patents. Many firms, however, opt for alternative methods, such as secrecy, to protect their innovations (Cohen et al., 2000). Additionally, a firm may cite a patent in its documentation, but the patent holder may never enforce it, or the patent itself may not be valid. Furthermore, firms may engage in IP-related negotiations with other patent owners without citing those patents. Nevertheless, as Ziedonis (2004) points out, there is little reason to believe that these limitations introduce systematic biases that favour certain types of firms, such as those with higher capital intensity. While these imperfections exist, they are unlikely to distort the overall purpose of the measure.

where $R\&DStock_{jt}$ denotes the R&D stock of firm j at time t ($j \neq i$) and ρ_{ij} shows the technological proximity between firms i and j, following Jaffe (1986), Bloom et al. (2013), and Noel and Schankerman (2013).

To quantify technological proximity, each firm is represented by a location vector $S_i = \{s_{i1}, s_{i2}, ..., s_{iN}\}$ summarizing the distribution of its citations across N technology classes (Noel and Schankerman, 2013). Specifically, s_{ik} represents the share of all citations made in firm i's patents that correspond to technology class k. I use 4-digit Cooperative Patent Classification (CPC 4) technology classes.[7] Following Bloom et al. (2013) and Noel and Schankerman (2013), technological proximity is measured using the uncentered correlation coefficient between firms' location vectors:

$$\rho_{ij} = \frac{S_i' S_j}{(S_i' S_i)^{1/2} (S_j' S_j)^{1/2}}. \tag{8}$$

This structure is appealing because it weights spillovers according to technological overlap, consistent with the idea that firms learn more from others operating in closely related technological domains (Jaffe, 1986).[8]

Market structure is controlled for using the logarithm of firm market share ($logMShare_{it}$) and the logarithm of the Herfindahl–Hirschman Index ($logHHI_{it}$) constructed at the four-digit SIC level. Changes in a firm's market share can result in a reallocation from less efficient firms to more efficient ones, thus impacting firm productivity (Baldwin and Gu, 2004). Moreover, in more concentrated markets, weaker competitive pressure may allow managerial slack, inefficient production, and slower adoption of cost-reducing technologies, leading to lower productivity. In contrast, greater competitive pressure in less concentrated markets increases the threat of entry and survival pressure, incentivizing firms to operate

[7]The CPC 4 technology classes are defined jointly by the U.S. Patent and Trademark Office and the European Patent Office. My data contains 677 distinct CPC 4 technology classes.

[8]This measure is continuous and bounded between 0 and 1. Values close to 1 indicate that two firms allocate their innovative efforts to similar technology classes (high alignment of vectors), while a value of 0 indicates completely unrelated technological portfolios (orthogonal vectors). The proximity measure is symmetric: $\rho_{ij} = \rho_{ji}$.

more efficiently and adopt productivity-enhancing innovations (Nickell, 1996; Aghion et al., 2005). I measure market share of each firm from the ratio of the firm 's sales in year t to total sales of the primary four-digit standard industry classification (SIC4) that the firm belongs to in year t. This method of calculating market share follows Blundell et al. (1999), who employ sales in the primary SIC3; Giroud and Mueller (2010), who use sales in the primary SIC2 and SIC4; and Duso et al. (2014), who utilize sales in the primary SIC4.[9]

Firm size ($logSize_{it}$), measured as the logarithm of sales, captures scale effects and organizational capacity (e.g., Bloom and Van Reenen, 2007).[10] Firm age ($logFirmAge_{it}$), measured as the logarithm of years since listing (Demirel and Mazzucato, 2012; Brown et al., 2009), captures life-cycle dynamics (Evans, 1987; Dunne and Hughes, 1994).[11]

Firm fixed effects μ_i absorb time-invariant heterogeneity, such as managerial quality and baseline technological capability, while year fixed effects λ_t capture economy-wide shocks and technological trends in equation (4). The variable ε_{it} is an idiosyncratic error term. Standard errors are clustered at the firm level to allow for arbitrary serial correlation. I use a within-firm estimator to estimate equation (4), which helps to account for possible time-invariant, unobserved firm-level heterogeneities (μ_i) inherent in panel data.

2.4 Addressing Endogeneity with a Shift–Share (Bartik) Instrument for AI Innovation

A central concern in estimating the productivity effects of AI-oriented innovation is that a firm's AI patent share may be endogenous. More productive firms may be better positioned to invest in AI, and unobserved managerial ability, organizational quality, or forward-looking strategies may simultaneously drive both productivity and the shift toward AI. To address this issue, I employ a shift–share (Bartik) instrumental-variable strategy following the frame-

[9] I would have liked to build the market share variable at a more disaggregate level, but my data only provide information by SIC4.

[10] I do not measure firm size with the number of employees to avoid mechanical overlap with the capital–labour ratio ($log(K/L)_{it}$), in which the number of employees appears in its denominator.

[11] The sample analyzed in this paper consists of publicly traded firms from Compustat.

work formalized by Goldsmith-Pinkham et al. (2020) and Bartik (1991).

The Bartik instrument combines predetermined firm-level exposure to different AI technology components with aggregate, economy-wide shocks to those components. In my setting, I exploit the fact that firms differ in their historical exposure to eight fine-grained AI technology classes c which are machine learning, evolutionary computation, natural language processing, vision, speech, knowl;edge processing, planning and control, and AI hardware, while the growth of these technologies at the aggregate level reflects common technological progress that is plausibly exogenous to any individual firm's productivity.

Formally, the instrument is constructed as:

$$Z_{it} = \sum_{c} Exposure_{ic}^{pre} \times AIGrowth_{ct}, \tag{9}$$

where $Exposure_{ic}^{pre}$ is firm i's share of patents in AI technology class c during a pre-period baseline window (1980–2009), and $AIGrowth_{ct}$ is the aggregate growth of AI patenting in class c at time t, constructed using a leave-one-out procedure that excludes firm i. The exposure shares are predetermined and fixed before the modern AI boom, while the aggregate growth rates capture common technological opportunities that are outside the control of individual firms. Therefore, my instrumental variable approach will have the following first and second stages, respectively:

$$\log(AIShare)_{it} = \pi Z_{it} + X_{it}'\gamma + \mu_i + \lambda_t + \varepsilon_{it}, \tag{10}$$

$$\begin{aligned}\log(Productivity_{it}) = \beta_1 \log(\widehat{AIShare})_{it} + \beta_2 \log(CitePatSock/L)_{it} + \beta_3 \log(R\&DStock/L)_{it} \\ + \beta_4 \log F_{it} + \beta_5 \log(K/L)_{it} + X_{it}'\delta + \mu_i + \lambda_t + \varepsilon_{it},\end{aligned} \tag{11}$$

This design is particularly well suited to the historical evolution of AI over my sam-

ple period. Prior to 2010, AI was a niche technology, diffusing slowly and applied mainly in specialized or experimental contexts. As a result, aggregate shocks to AI technologies were economically small, and the Bartik instrument is empirically weak in this early period. This weakness is itself informative: it is consistent with AI not yet constituting a major technological opportunity that could meaningfully shape firm behaviour. After 2010, however, breakthroughs in deep learning, the expansion of large-scale datasets, and advances in computing power transformed AI into a general-purpose technology with broad commercial relevance. In this later period, the Bartik instrument becomes strong and highly predictive of firm-level AI innovation.[12] This sharp contrast between the pre- and post-2010 periods supports the interpretation that the instrument captures exogenous shifts in technological opportunities rather than endogenous firm responses to productivity shocks.

Intuitively, a firm that historically specialized in, say, vision or natural language processing technologies experiences a larger exogenous increase in its predicted AI orientation when those fields expand rapidly at the aggregate level. The resulting variation in AI patent share is driven by common technological progress interacting with predetermined technological portfolios, not by firm-specific productivity dynamics. Conditional on firm fixed effects, year fixed effects, and rich controls for R&D, patenting activity, patent quality, and market structure, identification relies on the assumption that aggregate AI trends are orthogonal to firm-level productivity shocks.

This approach also mitigates concerns about measurement error. AI patent shares may be noisy at the firm-year level, especially when patent counts are small. The Bartik instrument replaces this potentially noisy variation with a predicted component based on long-run exposure and large-scale aggregate trends, which is typically less contaminated by firm-specific reporting or timing issues. In this sense, the IV strategy addresses both endogeneity and attenuation bias arising from measurement error in AI patenting.

Although Bartik instruments are more common in labor and regional economics, they are

[12] https://insightfulai.co.uk/a-brief-history-of-ai-key-milestones-and-innovations/ (Last accessed: Jan 23, 2026).

increasingly used in innovation and technology-related contexts. Examples include Aghion et al. (2019), Acemoglu et al. (2020), Bloom et al. (2013), Kogan et al. (2017), and Autor et al. (2014). These studies share the same core logic as my approach: firms differ in predetermined exposure to external forces, and common shocks generate plausibly exogenous variation in outcomes. In my context, AI is the external technological force, and firms' historical AI technology portfolios determine their sensitivity to its rise.

Overall, the Bartik design allows me to separate the causal impact of shifts toward AI-oriented innovation from endogenous firm behavior. The fact that the instrument is weak before 2010 but strong afterward is not a weakness of the design; rather, it is precisely what one would expect if AI only became an economically meaningful general-purpose technology in the modern era. Importantly, the IV strategy instruments for variation in AI intensity among patenting firms, thereby identifying the causal effect of the intensive margin of technological reorientation rather than the adoption decision itself.

3 Data

I construct a panel dataset of publicly traded U.S. firms from 1980 to 2021 by combining four complementary data sources: Compustat North America Fundamentals Annuals, the latest Discern patent–firm linkage dataset,[13] PatentsView citation data, and the USPTO Artificial Intelligence Patent Dataset (AIPD2023). Compustat provides annual firm-level financial information for U.S. public firms.[14] I restrict attention to firms headquartered in the United States with at least three consecutive years of activity in Compustat, following Arora et al. (2024). Over 1980–2021, this yields an unbalanced panel of 23,619 firms and 252,360 firm-year observations.

Patent data are obtained from Discern, which links U.S. utility patents to Compustat firms using standardized firm identifiers and dynamic ownership tracking. Discern improves

[13] I use Discern 2.0.1, which is available at https://discern-project.com.

[14] Publicly traded firms include those listed on the New York Stock Exchange, American Stock Exchange, and regional exchanges, as well as over-the-counter firms traded on NASDAQ.

matching accuracy relative to earlier patent–firm linkages by incorporating name changes, subsidiary information, and merger activity, ensuring that patents are dynamically reassigned when firms restructure. After merging Compustat with Discern, 1,865,424 patents are linked to firms over the sample period.[15]

I augment these data with citation information from PatentsView, which provides 149,301,482 citations associated with 8,596,058 patents from 1980 to 2022.[16] Citation data are used to construct citation-weighted patent stocks as a measure of innovative output as described in Section 2. Because of the lag between patent application and grant, and because citations accrue gradually over time, both patent counts and citation counts are subject to truncation. I therefore apply truncation corrections following Hall et al. (2005). The details of the correction procedures are provided in the Appendix.

In addition to citation information, PatentsView provides detailed Cooperative Patent Classification (CPC) codes for each patent. I use the CPC subclass-level classification to characterize the technological domain of each patent. The merged dataset spans 677 distinct CPC technology classes over the sample period. These technology classes form the basis for constructing measures of technological relatedness in R&D spillovers. After merging Compustat, Discern, and PatentsView, the resulting dataset contains 8,036 firms and 111,082 firm-year observations over the period 1980–2021.

AI-related patents are identified using AIPD2023.[17] AIPD2023 provides detailed information on both AI granted patents and pre-grant publications from 1976 to 2023; I restrict attention to granted patents in order to ensure consistency with the Discern data.[18] AIPD2023 assigns each patent document a predicted probability between 0 and 100 percent of belonging to the AI category, using an enhanced multi-step deep learning approach based

[15]Detailed descriptions of the Discern dataset are provided in Arora et al. (2024).

[16]PatentsView is available at: http://patentsview.org.

[17]AIPD2023 is available at: https://www.uspto.gov/ip-policy/economic-research/research-datasets/artificial-intelligence-patent-dataset.

[18]The choice of AIPD is motivated by Giczy et al. (2022), who benchmarked AIPD against several alternative approaches in the literature, including the classification and keyword methods proposed by Cockburn et al. (2019) and WIPO (2019), and show that AIPD substantially outperforms these alternatives in identifying AI patents.

on the automated patent landscaping methodology of Abood and Feltenberger (2018) and Giczy et al. (2022). The dataset proposes three probability thresholds — 50%, 86%, and 93% — to classify patents as AI-related. For example, under the 86% threshold, a patent is classified as AI if its predicted probability is greater than or equal to 86%; otherwise, it is classified as non-AI. Lower thresholds classify more patents as AI, while higher thresholds increase precision but reduce recall by excluding patents with intermediate probabilities. Pairolero et al. (2025) show that the 86% threshold provides the best balance between precision and recall. Accordingly, I adopt the 86% cut-off as my baseline definition of AI patents and conduct robustness checks using the 50% and 93% thresholds in Section 4.1. After merging Compustat, Discern, PatentsView, and AIPD2023, the combined dataset contains 7,842 firms and 104,456 firm-year observations over 1980–2021.

The key explanatory variable in the empirical analysis is the share of a firm's patents that are AI-related, defined as the number of AI patents divided by total patents in a given firm-year. By construction, this measure is defined only for firm-years with positive patenting activity. Accordingly, the analysis focuses on patenting firms. This restriction is economically motivated: the empirical objective is to identify how shifts in the direction of innovation affect productivity, holding constant overall innovative effort. For firms without patenting activity, the technological composition of innovation is unobserved, and the AI share is mechanically undefined. Restricting the sample to patenting firms ensures that the identifying variation arises from within-firm changes in the composition of innovation rather than from differences between innovative and non-innovative firms.

Within the patenting sample, firms differ in whether and when they engage in AI-related patenting. The baseline specifications exploit within-firm variation in the AI share of patents, effectively isolating the intensive margin of AI-oriented innovation. This design captures changes in the allocation of innovative effort toward AI conditional on participation in patenting. Conceptually, this approach aligns with the directed technical change framework developed in Section 2, which emphasizes that productivity depends not only on the

level of innovation but on its direction. It also reflects the view of AI as a general-purpose technology whose productivity effects materialize when it becomes embedded in firms' technological capabilities through sustained reallocation of innovative effort. Specifications that include both AI adopters and non-adopters and explicitly separate extensive and intensive margins are reported in Section 4, confirming that the positive productivity effects are driven by deeper AI orientation rather than by initial adoption episodes.

Requiring the availability of financial variables, patent measures, AI classifications, and market-structure controls defines the regression sample used in the empirical analysis. The resulting estimation sample includes 2,086 firms and 12,227 firm-year observations over 1980–2021. As is typical in long-horizon panels linking financial and patent data, variation in sample size across specifications reflects the availability of detailed innovation and market-structure measures. The panel remains unbalanced, and identification relies on within-firm changes over time.

Firms in the sample are R&D-intensive, which means average R&D stock per worker is 165, and citation-weighted patent stock per worker displays considerable dispersion. Market concentration varies widely across industries, with a mean HHI of 0.331, suggesting substantial heterogeneity in competitive environments.

Descriptive statistics for the estimation sample are reported in Table 1. On average, firms exhibit a high degree of patent ownership fragmentation, with a mean index of 0.84 (median 0.90), indicating that innovation frequently occurs in technologically congested environments characterized by dispersed patent ownership. AI-oriented innovation accounts for 30 percent of total patenting activity on average, with a median share of 0.17, suggesting a right-skewed distribution in which some firms are highly specialized in AI. Firms in the sample are innovation-intensive: mean R&D stock per employee equals 165, and citation-weighted patent stock per employee averages 20,909, though both measures display substantial dispersion. Market structure also varies considerably across industries, with a mean HHI of 0.33 (median 0.25), indicating meaningful heterogeneity in competitive environments.

Figure 1 illustrates the diffusion and distribution of AI-oriented innovation in the regression sample. Panel A shows a pronounced upward trend in the average AI patent share over time, particularly after the late 2000s, consistent with the rise of modern machine learning and the emergence of AI as a general-purpose technology. The interquartile range also widens over time, indicating increasing heterogeneity in firms' AI orientation. Panel B displays the cross-sectional distribution of AI patent shares, revealing substantial dispersion and right-skewness, with a mass of observations at lower values and a nontrivial fraction of firm-years exhibiting high AI specialization. Together, the figure underscores both the dynamic diffusion of AI and the meaningful within-sample variation that underpins the empirical identification strategy.

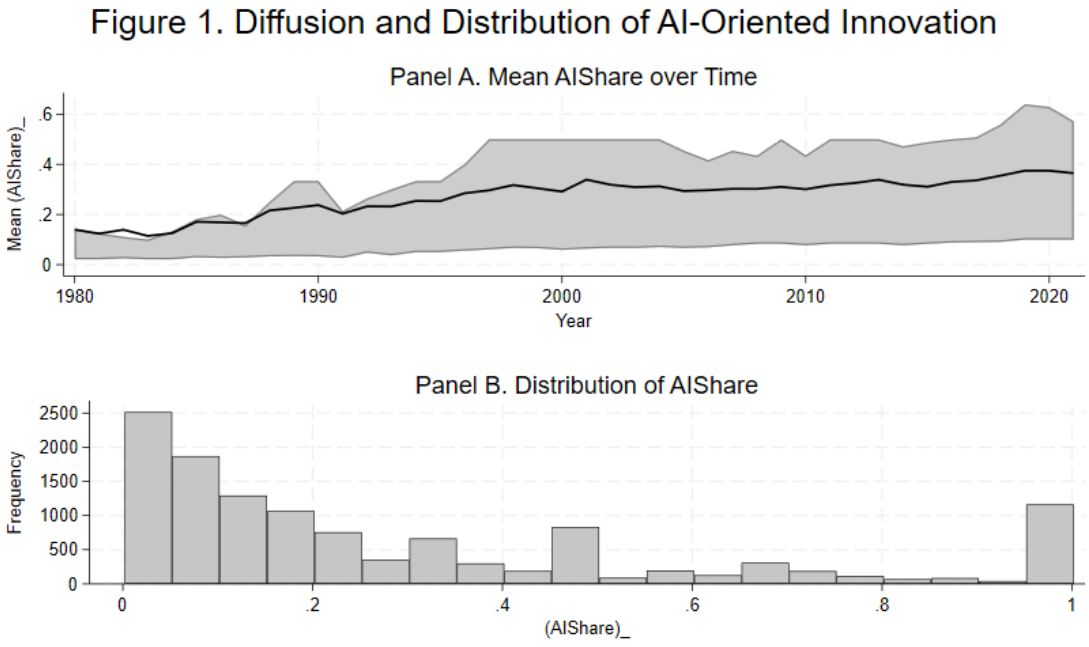


Figure 1: **Diffusion and Distribution of AI-Oriented Innovation.** Panel A plots the annual mean of $(AIShare)_{it}$ among AI-adopting patenting firms between 1980 and 2021. The shaded region represents the interquartile range. Panel B displays the cross-sectional distribution of $(AIShare)_{it}$ across firm-year observations in the regression sample.

Table 1: Descriptive Statistics: AI-Adopting Patenting Firms

Variable	Description	Obs	Mean	Median	Std. Dev.	Min	Max
$Productivity_{it}$	Sales per employee	12,227	331.09	256.92	402.96	0.02	12,302.89
$AIShare_{it}$	Share of AI patents in total patents	12,227	0.30	0.17	0.31	0.002	1.00
$(CitePatStock/L)_{it}$	Citation-weighted patent stock per employee	12,227	20,908.74	3,246.01	145,062.50	0.60	6,370,407
$(PatentStock/L)_{it}$	Patent stock per employee	12,227	83.52	23.31	315.50	0.01	14,650
$(R\&DStock/L)_{it}$	R&D stock per employee	12,227	165.38	90.20	368.04	0.36	15,843.69
$R\&DSpillover_{it}$	Technological R&D spillover	12,227	100,675	74,250	82,416	371	468,162
F_{it}	Fragmentation index	11,841	0.84	0.90	0.19	0.00	0.99
$(K/L)_{it}$	Capital–labor ratio	12,227	173.77	93.35	479.48	2.25	13,258.94
$Size_{it}$	Sales (million USD)	12,227	8,125.58	991.03	25,335.91	0.01	469,822
$MShare_{it}$	Market Share	12,227	0.19	0.04	0.28	0.00	1.00
HHI_{it}	Herfindahl index	12,227	0.33	0.25	0.23	0.05	1.00
$FirmAge_{it}$	Years since listing	12,227	15.38	13.00	10.70	1	42

Notes: This table reports descriptive statistics for the regression sample of AI-adopting patenting public U.S. firms observed between 1980 and 2021. The sample consists of 12,227 firm-year observations. $AIShare_{it}$ is constructed using the 86% probability threshold from AIPD2023. Citation-weighted patent stock and R&D stock are constructed using a perpetual-inventory method with a 15% depreciation rate. Fragmentation is the Ziedonis (2004) patent ownership index. Market share and HHI are computed at the SIC4-year level.

4 Results

Table 2 reports the baseline and selected sensitivity estimates of the effect of shifting a firm's innovation portfolio toward AI, measured by $\log(AIShare)_{it}$, on labor productivity, $logProductivity_{it}$. The baseline specifications are estimated on the regression sample of AI-adopting patenting firms described in Section 3, and therefore identify the intensive margin of AI-oriented innovation within firms that engage in patenting activity. Columns (1)–(5) are estimated using panel data methods, and the Hausman specification test supports the use of fixed-effects estimators. Across all specifications, I control for firms' market share, capital–labour ratio, size, age, and market concentration, all expressed in natural logarithms, as well as year fixed effects. Firms' innovative activity is further accounted for in Columns (2) and (3) by including citation-weighted patent stock intensity, $log(CitePatSock/L)_{it}$, and R&D stock intensity, $log(R\&DStock/L)_{it}$. In Column (4), citation-weighted patent stock intensity is replaced by the patent stock per employee, $log(PatentStock/L)_{it}$, while maintaining the R&D stock control. Column (5) replaces firm fixed effects with industry fixed effects (SIC4). Standard errors are clustered at the firm level to allow for arbitrary serial correlation in productivity shocks within firms over time. The results are robust to clustering at the industry level.

The coefficient on AI patent share, $\log(AIShare)_{it}$, is positive and statistically significant across all specifications in Table 2. Because the model exploits within-firm variation among patenting firms, the estimated effect captures sustained shifts in the technological orientation of the firm toward AI rather than the discrete decision to begin patenting in AI. The results therefore indicate that, among firms with comparable R&D intensity and overall patenting activity, those whose innovation portfolios become more AI-intensive experience higher labour productivity.

This pattern is consistent with AI-oriented innovation enhancing efficiency through AI-embedded process improvements—such as prediction, optimization, and automation—that raise the productivity of knowledge work and R&D activities. It also aligns with the presence

of complementarities between AI technologies and co-investments in data infrastructure, organizational restructuring, and skilled human capital that are typically required for effective development and deployment. Although the data do not directly observe ICT investment, the inclusion of firm fixed effects, time fixed effects, and detailed innovation controls mitigates concerns that the estimates merely reflect broad digitalization trends. Instead, the evidence supports an interpretation centered on the direction of technological change toward AI.

The coefficients on the standard firm-level controls behave in a manner consistent with economic intuition and prior productivity studies. The capital–labour ratio ($log(K/L)_{it}$) is positive and highly significant across all specifications, indicating that capital deepening is an important driver of labor productivity. Firm size, measured by $logSize_{it}$, also enters positively, reflecting economies of scale, organizational advantages, and greater scope for adopting productivity-enhancing technologies in larger firms. Firms with higher market shares ($logMShare_{it}$), exhibit higher productivity in the firm fixed-effects specifications (Columns 1–4), consistent with more productive firms gaining larger market shares through competitive selection. When industry fixed effects replace firm fixed effects in Column (5), the coefficient on market share becomes statistically insignificant, suggesting that much of the variation in market share operates at the firm level rather than across firms within industries. Market concentration, measured by $log(HHI)_{it}$, is generally negative and statistically significant in the firm fixed-effects models, indicating that greater competitive pressure is associated with higher productivity. This pattern is consistent with theories in which competition disciplines managerial slack and accelerates efficiency-enhancing innovation. Firm age, $logFirmAge_{it}$, enters with a negative and statistically significant coefficient throughout, suggesting that older firms tend to have lower productivity growth conditional on firm fixed effects. This result is consistent with organizational inertia, slower technology adoption, or vintage effects in production technologies.

Measures of innovative activity display strong and economically meaningful effects. Both

citation-weighted patent stock intensity, $log(CitePatSock/L)_{it}$, and R&D stock intensity, $log(R\&DStock/L)_{it}$, are positive and highly significant in Columns (2)–(4), confirming that general innovative capacity independently raises firm productivity. Replacing citation-weighted patents with truncation-corrected patent stock intensity ($log(PatentStock/L)_{it}$) in Column (4) yields very similar conclusions, indicating that the productivity effect of innovation is not driven by the specific construction of the patent measure.

The fragmentation measure, $logF_{it}$, is negative and highly significant in Columns (3)–(5), indicating that more fragmented patent ownership within a technological space is associated with lower productivity, consistent with higher transaction costs, coordination failures, and weaker appropriability. The R&D spillover measure, $logR\&DSpillover_{it}$, enters with a negative coefficient in most specifications. This suggests that external R&D activity by technologically proximate firms may intensify competitive pressures or reduce appropriability, thereby lowering measured productivity for incumbent firms, even though it may benefit aggregate efficiency. The positive coefficient on $MissingF_{it}$ reflects systematic differences between firms for which fragmentation is observed and those for which it is missing.

Finally, Column (5) demonstrates that the estimated impact of AI-oriented innovation remains positive and statistically significant when industry fixed effects replace firm fixed effects. This confirms that the productivity gains associated with AI are not driven solely by persistent firm-level heterogeneity but are also present in comparisons across firms within the same industry. The positive and statistically significant impact of AI on productivity is also robust to when I control for industry-by-year fixed effects (coefficient = 0.034, standard error = 0.008).

4.1 Robustness Checks

This section evaluates the robustness of the baseline relationship between AI-oriented innovation ($\log(AIShare)_{it}$) and firm productivity ($logProductivity)_{it}$) to a wide range of alternative specifications and identification concerns. I assess sensitivity to (i) alternative

Table 2: AI Innovation and Labor Productivity: Baseline Results

$\log(Productivity)_{it}$	(1)	(2)	(3)	(4)	(5)
$\log(AIShare)_{it}$	0.044***	0.059***	0.053***	0.055***	0.089***
	(0.008)	(0.008)	(0.007)	(0.008)	(0.012)
$\log(CitePatStock/L)_{it}$		0.074***	0.078***		0.065***
		(0.013)	(0.013)		(0.010)
$\log(PatentStock/L)_{it}$				0.057***	
				(0.017)	
$\log(R\&DStock/L)_{it}$		0.308***	0.306***	0.316***	0.178***
		(0.029)	(0.028)	(0.029)	(0.016)
$\log F_{it}$			-0.173***	-0.161***	-0.312***
			(0.035)	(0.034)	(0.046)
$\log R\&DSpillover_{it}$		-0.099***	-0.085***	-0.085***	-0.093***
		(0.013)	(0.013)	(0.013)	(0.016)
$\log(K/L)_{it}$	0.249***	0.064***	0.065***	0.069***	0.090***
	(0.026)	(0.025)	(0.024)	(0.025)	(0.021)
$\log Size_{it}$	0.213***	0.292***	0.294***	0.284***	0.241***
	(0.029)	(0.031)	(0.031)	(0.031)	(0.021)
$\log MShare_{it}$	0.077***	0.114***	0.117***	0.118***	-0.007
	(0.025)	(0.026)	(0.026)	(0.027)	(0.020)
$\log(HHI)_{it}$	-0.071***	-0.059**	-0.058**	-0.057*	-0.022
	(0.026)	(0.029)	(0.029)	(0.030)	(0.024)
$\log FirmAge_{it}$	-0.099***	-0.133***	-0.129***	-0.117***	-0.103***
	(0.025)	(0.027)	(0.027)	(0.029)	(0.016)
$MissingF_{it}$			0.103***	0.083***	0.198***
			(0.030)	(0.029)	(0.039)
Firm Fixed Effects	Yes	Yes	Yes	Yes	No
Industry Fixed Effects	No	No	No	No	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations	14,511	12,227	12,227	12,249	12,227
R-squared	0.603	0.699	0.702	0.697	0.673

Notes: Standard errors are clustered at the firm level and reported in parentheses. Columns (1)–(4) are estimated using firm fixed effects; Column (5) uses industry fixed effects (SIC4) instead. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

definitions of AI exposure based on different patent-classification thresholds, (ii) timing assumptions and dynamic effects using lags and leads, (iii) alternative measures of AI investment capturing innovation direction versus accumulated knowledge, (iv) placebo tests distinguishing experimentation from substantive AI commitment, (v) sample restrictions and trimming of extreme observations, (vi) alternative productivity measures including structural TFP estimators, and (vii) potential multicollinearity among controls. Across all exercises, the core result remains robust in comparison to the baseline results in Column (1) of Table 3: productivity gains arise when firms reallocate innovative effort toward AI, rather than from simple experimentation, entry into AI patenting, or the accumulation of AI-related patent stocks. These exercises confirm that the productivity gains are driven by sustained increases in AI intensity within firms, rather than by discrete adoption episodes or mechanical scaling effects.

I begin by examining whether the results are sensitive to the threshold used to classify AI-related patents. Columns (2) and (3) of Table 3 re-estimate the baseline specification in equation (4) using alternative AI share measures constructed under more inclusive (50 percent) and more restrictive (93 percent) classification cutoffs, denoted $\log(\text{AIShare})_{it50}$ and $\log(\text{AIShare})_{it93}$. These cutoff points are described in Section 3. Across both definitions, the estimated coefficients remain positive, highly statistically significant, and nearly identical in magnitude to the baseline estimate reported in Column (1). This stability indicates that the productivity effects of AI-oriented innovation are not driven by a particular classification rule and reflect a robust relationship between innovation direction and firm performance. Similar insensitivity to technology-classification thresholds has been documented in studies of digital and general-purpose technologies, such as Bloom et al. (2013) and Bessen et al. (2022).

To address concerns about simultaneity and delayed adjustment, I next replace contemporaneous AI patent share ($\log(AIShare)_{it}$) with its one-year lag, $\log(AIShare)_{it-1}$. The estimated coefficient on lagged AI share remains positive and statistically significant (0.007,

Std.Error=0.043), with a magnitude comparable to the baseline estimate. This finding suggests that shifts toward AI-oriented innovation precede productivity improvements rather than responding contemporaneously to them. Using lagged technology measures to mitigate reverse causality and capture delayed effects is standard in the innovation and productivity literature (Aghion et al., 2019; Acemoglu et al., 2020), and the stability of the coefficient reinforces the causal interpretation of the baseline results.

Following the event-study/distributed-lag approach (e.g., Clarke, 2021; Miller et al., 2023; Schmidheiny and Siegloch, 2019), I further estimate a dynamic specification that augments equation (4) with two leads and two lags of the AI intensity variable, ($\log(AIShare)_{it}$), to assess pre-trends/anticipation and allow for delayed effects. The estimated coefficients on future AI shares are small and statistically insignificant ($\log(AIShare)_{it+1}$: 0.002, Std.Error=0.008; $\log(AIShare)_{it+2}$: 0.008, Std.Error=0.009), providing no evidence of pre-trends or anticipatory productivity effects. In contrast, the coefficient on the second lag of AI intensity is positive and statistically significant (0.017, Std.Error=0.007), while the first lag is smaller and imprecisely estimated (0.006, Std.Error=0.007). This pattern indicates that productivity gains from AI-oriented innovation materialize with a delay of roughly two years, consistent with organizational adjustment costs, learning, and complementary investments required for AI adoption to translate into higher productivity.

To distinguish between the direction of innovation and the scale of accumulated AI investments, Column (4) of Table 3 replaces AI patent share ($\log(AIShare)_{it}$) with an AI patent knowledge stock, $\log(AIStock)_{it}$ in equation (4), constructed using the perpetual-inventory method explained in Section 2 with a 15 percent depreciation rate. In contrast to the positive and significant share effect reported in Column (1), the estimated coefficient on AI patent stock is negative and statistically significant. This contrast is informative rather than contradictory. While AI share captures a reallocation of innovative effort toward AI technologies holding total innovation constant, AI stock reflects cumulative investments that may involve substantial experimentation, organizational restructuring, and integration

costs. The negative coefficient is consistent with short-run adjustment frictions and diminishing returns to accumulated AI investments, whereas the positive share effect highlights the productivity gains associated with strategic reorientation toward AI. Similar distinctions between innovation direction and accumulated knowledge are emphasized in the broader innovation literature (Aghion et al., 2005; Brynjolfsson et al., 2021).

To test whether merely entering AI-related innovation is sufficient to generate productivity gains, Column (5) of Table 3 replaces AI patent share with a binary indicator equal to one if the firm files at least one AI-related patent in a given year (AI_Dummy_{it}). The estimated coefficient on this AI adoption dummy is negative and highly statistically significant. This result indicates that initial experimentation with AI technologies is associated with short-run productivity costs rather than gains. The dummy captures minimal engagement in AI-related patenting rather than meaningful deployment or integration, and the negative coefficient is consistent with learning costs, adjustment frictions, and opportunity costs during early experimentation. In contrast, measures capturing the intensity of AI-oriented innovation— such as AI patent share— yield positive productivity effects, suggesting that gains emerge only once firms move beyond experimentation and commit to AI at scale (Brynjolfsson et al., 2021; Tambe, 2014).

I also examine whether the results are driven by specific time periods or extreme observations. Restricting the sample to the post-2010 period— when AI technologies became economically consequential— yields a positive and statistically significant elasticity of labor productivity with respect to AI patent share (0.018, Std. Error=0.009), confirming that within-firm shifts toward AI-intensive innovation remain productivity-enhancing in the modern AI era (Agrawal et al., 2019). Excluding observations with AI patent shares above the 99th percentile likewise leaves the estimated coefficient positive, statistically significant, and quantitatively similar to the baseline (0.053, Std. Error=0.007), indicating that the results are not driven by a small number of extreme innovators (Bloom et al., 2013; De Loecker et al., 2016).

Table 3: Robustness Checks

$\log(Productivity)_{it}$	(1)	(2)	(3)	(4)	(5)
$\log(AIShare)_{it}$	0.053***				
	(0.007)				
$\log(AIShare)_{it50}$		0.052***			
		(0.008)			
$\log(AIShare)_{it93}$			0.050***		
			(0.007)		
$\log(AIStock)_{it}$				-0.003***	
				(0.001)	
AI_Dummy_{it}					-0.142***
					(0.009)
Controls	Yes	Yes	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations	12,227	15,187	11,783	19,527	36,146
R-squared	0.702	0.704	0.704	0.695	0.744

Notes: The dependent variable is $\log(Productivity)_{it}$, measured as the logarithm of sales per employee. Standard errors are clustered at the firm level and reported in parentheses. Columns are estimated using firm fixed effects. $\log(AIShare)_{it}$ is the log share of AI patents in the firm-year. $\log(AIShare)_{it50}$ and $\log(AIShare)_{it93}$ apply alternative cut off thresholds in constructing AI-related patents. $\log(AIStock)_{it}$ is the logarithm of the stock of AI patents. AI_Dummy_{it} equals one if the firm has at least one AI patent in year t. All controls in 4 are included. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

To assess sensitivity to the choice of productivity measure, Table 4 re-estimates the baseline specification in equation (4) using alternative proxies for productivity. Columns (2) and (3) employ total factor productivity measures constructed using, respectively, a fixed-effects Solow residual and a structural Levinsohn–Petrin estimator implemented following Ackerberg et al. (2015).[19] Column (4) uses log value added per worker.[20] Across all measures, the coefficient on AI patent share ($\log(\text{AIShare})_{it}$) remains positive and statistically significant, with magnitudes comparable to the baseline labour-productivity specification in Column (1). These results indicate that the estimated effects are not driven by scale effects, labor intensity, or a particular productivity metric.

As a placebo test, I replace AI patent share in equation (4) with the firm's share of non-AI patents, defined as one minus the AI patent share. The estimated coefficient on non-AI patent share is small and statistically insignificant (0.022, Std. Error=0.021), while the effects of R&D intensity, citation-weighted patent stock, and other firm controls remain stable. This null result indicates that productivity gains are not associated with shifts toward non-AI innovation and reinforces the interpretation that the baseline findings reflect the direction of technological change toward AI rather than general patenting activity per se (Acemoglu et al., 2016; Aghion et al., 2019).

[19]First, I construct a simple Solow-style total factor productivity (TFP) proxy by estimating a Cobb–Douglas production function with firm fixed effects, where output is measured by logarithm of sales, capital by logarithm of net property, plant, and equipment, and labour by logarithm of employment. The estimating equation is $logSale_{it} = \alpha_1\ logK_{it} + \alpha_2 logL_{it} + e_{it}$. The residual from this regression represents firm-level logarithm of TFP. This residual-based approach has been widely used in the empirical productivity and innovation literature as a transparent proxy for firm-level efficiency (e.g., Griliches and Mairesse, 1995; Bloom et al., 2013).

For the other total factor productivity measure, I employ a structural TFP measure based on the Levinsohn–Petrin (LP) approach, which uses intermediate inputs as proxies to control for unobserved productivity shocks and address simultaneity in input choices (Levinsohn and Petrin, 2003). In implementing this approach, I follow the timing and identification assumptions emphasized by Ackerberg et al. (2015) by treating labour as a freely adjustable input and capital as a state variable, thereby addressing potential identification concerns in production-function estimation. This Levinsohn–Petrin / Ackerberg et al style implementation is widely regarded as a preferred method in the productivity literature and is commonly used in studies of technology adoption and innovation (e.g., Bloom et al, 2013; De Loecker et al., 2016).

[20]I construct value added per worker, defined as firm-level sales net of intermediate inputs (the direct costs of producing the goods or services that a firm sells during the year) and operating expenses (All operating expenses not directly tied to production), divided by employment (number of employees) following standard practice in the productivity literature (e.g., Hall, Mairesse, and Mohnen, 2009; De Loecker and Warzynski, 2012).

Table 4: AI-Oriented Innovation and Alternative Productivity Measures

	(1) Labor Prod.	(2) TFP (FE)	(3) TFP (LP/ACF)	(4) Value Added
$\log(AIShare)_{it}$	0.053***	0.053***	0.040***	0.056***
	(0.007)	(0.007)	(0.007)	(0.016)
Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	12,227	12,225	12,217	9,499
R-squared				

Notes: Each column reports the coefficient on the firm-level share of AI-related patents from a fixed-effects regression using a different productivity measure in the logarithm format as the dependent variable. Column (1) reports the baseline labour productivity specification. All controls in 4 are included. Standard errors are clustered at the firm level and reported in parentheses. Column (2) uses a Solow-style total factor productivity (TFP) proxy obtained as the residual from a Cobb–Douglas production function estimated with firm fixed effects. Column (3) reports results using a structural TFP measure estimated via the Levinsohn–Petrin approach with timing assumptions following Ackerberg et al. (2015). Column (4) uses log value added per worker, defined as sales net of intermediate inputs and operating expenses divided by employment. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Finally, I assess whether the estimated productivity effects of AI-oriented innovation ($\log(AIShare)_{it}$) are sensitive to multicollinearity among controls by re-estimating the baseline specification excluding firm market share, industry concentration, and patent fragmentation measures following Autor et al. (2014) and Bloom et al. (2013). Although these variables capture distinct mechanisms, - market position, industry structure, and technological fragmentation- they exhibit limited within-firm variation once firm and year fixed effects are included and may absorb overlapping residual variation. Reassuringly, removing these controls leaves the estimated coefficient on AI patent share virtually unchanged in magnitude and statistical significance (0.060, Std.Error=0.008), confirming that the baseline results are not driven by any single control or by collinearity among measures of market structure and technological environment.

4.2 Shift–Share (Bartik) IV Evidence on the Productivity Effects of AI-Oriented Innovation

This section reports instrumental-variables estimates of the causal effect of $\log(AIShare)_{it}$ on labour productivity ($\log Productivity_{it}$) using a Bartik (shift–share) design in Table 5. The variable $\log(AIShare)_{it}$ is instrumented with a post-2010 AI opportunity shock interacted with predetermined, pre-2010 technological exposure. All specifications across columns (1) to (3) of Table 5 include all the right hand side variables in equation (4) as well as firm and year fixed effects. All standard errors are clustered at the firm level. Estimation method across columns is based on Bartik IV method explained in Section 2.4.

Column (1) of Table 5 presents the main IV results (2010–2021; baseline exposure 1980–2009). The first stage is strong and precisely estimated: the coefficient on the Bartik instrument ($bartikIV_{it}$) is large and highly significant (0.590, Std.Error=0.119), and the first-stage F-statistic (24.78) comfortably exceeds conventional weak-instrument thresholds. This indicates that the instrument is highly predictive of firm-level AI patent intensity and that the model is well identified. Underidentification is rejected (Kleibergen–Paap LM p=0.000). In the second stage, the instrumented coefficient on $\log(AIShare)_{it}$, reported in column (1) of Table 5, is 0.377 (Std.Error = 0.157), statistically significant at the 5 percent level. This implies that a 1 percent increase in the share of AI patents in a firm's innovation portfolio causes, on average, a 0.38 percent increase in labour productivity. The magnitude of this estimate is substantially larger than the corresponding fixed-effects estimates in Table 2??, suggesting that conventional within-firm regressions may understate the true causal effect of AI because of attenuation from measurement error and endogenous adoption decisions. The IV estimate isolates variation driven by exogenous shifts in AI technological opportunities rather than firm-specific productivity shocks, and therefore captures a causal productivity effect of AI-oriented innovation.

Column (2) of Table 5 reports the first robustness check, where AI intensity ($\log(AIShare)_{it}$) is reconstructed using a truncation-adjusted denominator for total patent counts. The esti-

mated effect of AI ($\log(AdjAIShare)$) remains virtually unchanged: the coefficient on the instrumented $\log(AIShare)_{it}$ is 0.379 (Std.Error= 0.158), again significant at the 5 percent level. This result demonstrates that the main findings are not sensitive to right-truncation in patent grants or to the precise measurement of patent shares, reinforcing the interpretation that AI-oriented innovation has a robust and economically meaningful impact on productivity.

Column (3) of Table 5 implements a second robustness check by restricting the IV sample to 2010–2018, thereby excluding the most recent years in which patent grant data are most severely right-truncated. The estimated coefficient rises to 0.736 (Std.Error = 0.327) and remains statistically significant. The larger magnitude is economically intuitive: this subsample concentrates on the period in which modern deep-learning methods diffused rapidly and AI became tightly integrated into core business processes. In this environment, firms induced into AI adoption by the technological shock are likely to experience especially large productivity gains, so the local average treatment effect captured by the IV naturally increases.

Taken together, the results in Table 5 provide strong evidence that AI-oriented innovation has a causal and quantitatively important effect on firm productivity. The strength of the first stage after 2010, combined with the consistency of the second-stage estimates across alternative constructions of AI share and alternative sample windows, supports the interpretation that the Bartik instrument captures exogenous shifts in technological opportunities rather than endogenous firm behaviour. The fact that the instrument is weak prior to 2010 but becomes powerful in the modern AI era further reinforces this interpretation, as it mirrors the historical transition of AI from a niche technology to a general-purpose input with broad economic relevance.

While the baseline results focus on the intensive margin of AI-oriented innovation within patenting firms, it is important to assess whether the productivity gains arise from deeper AI orientation or from the initial decision to enter AI patenting. To distinguish these mecha-

nisms, I extend the analysis to specifications that include both AI adopters and non-adopters and explicitly separate the extensive and intensive margins of AI innovation.

Table 5: Bartik IV Estimates of AI-Oriented Innovation and Labour Productivity

	(1) Main IV 2010–2021	(2) Robustness 1 2010–2021	(3) Robustness 2 2010–2018
Endogenous Variable	$\log(AIShare)_{it}$	$\log(AdjAIShare)_{it}$	$\log(AIShare)_{it}$
Second Stage:			
Endogenous AI Share	0.377***	0.380***	0.736**
	(0.157)	(0.158)	(0.327)
First Stage:			
Excluded IV Coef.	0.590***	0.587***	0.611***
(bartikIV_{it})	(0.119)	(0.118)	(0.193)
First-Stage F	24.78	24.55	10.03
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	3,655	3,655	2,847

Notes: Clustred standard errors are at the firm level in parentheses.
***, ** denote significance at the 1% and 5% levels, respectively.

4.3 Extensive and Intensive Margin Effects of AI-Oriented Innovation

The baseline specification in equation (4) isolates the intensive margin of AI-oriented innovation by restricting the sample to firms that have already adopted AI patenting. Within this group, identification relies on within-firm changes in the composition of the patent portfolio. Column (1) of Table 6 shows that reallocating innovative effort toward AI is associated with significantly higher labour productivity. The coefficient on $\log(AIShare)_{it}$ implies that a 10% increase in AI orientation raises labour productivity by approximately 0.5 percent. Because the specification controls for citation-weighted patent stock, R&D stock, capital deepening, fragmentation, market structure, and firm and year fixed effects, this estimate captures the effect of changing the direction of innovation, holding constant the scale and quality of innovative effort.

To assess whether this result reflects the deepening of AI capabilities or merely the act of adoption, I expand the sample to include all patenting firms and explicitly separate the extensive and intensive margins. Column (2) of Table 6 estimates the full-sample specification using $\log(1 + AIShare)_{it}$, allowing firms with zero AI patenting to enter the regression. The coefficient becomes negative and statistically significant. This reversal does not contradict the baseline result; rather, it reflects a change in the estimand. In the full sample, the coefficient combines two economically distinct effects: (i) the discrete transition from zero to positive AI engagement (the extensive margin), and (ii) marginal increases in AI orientation among adopting firms (the intensive margin).

Column (3) of Table 6 isolates the extensive margin by replacing AI share with an indicator for AI adoption. The negative and statistically significant coefficient on AI_Dummy_{it} indicates that within firms, years in which AI patenting begins are associated with lower productivity relative to non-AI years. This negative association is consistent with the presence of adjustment costs, organizational frictions, and complementary investment requirements emphasized in the general-purpose technology (GPT) literature. AI adoption is not a frictionless technology switch; it requires data infrastructure, computing capital, new human capital, and reorganization. In early stages, these investments may depress measured revenue productivity.

Column (4) in Table 6 includes both the adoption indicator and the intensive margin among adopters. The results show that the negative coefficient in the pooled specification is driven entirely by the extensive margin. Once adoption is controlled for separately, increases in AI orientation among adopters are positive and economically meaningful. In other words, initiating AI patenting is associated with short-run productivity reductions, but sustained reallocation toward AI — conditional on adoption — is productivity-enhancing.

This pattern aligns closely with both directed technical change models (Acemoglu, 1998; Acemoglu, 2002) and the GPT framework (Bresnahan and Trajtenberg, 1995). Directed technical change emphasizes that reallocating innovative effort across technologies can gen-

erate different productivity paths even when total innovation remains constant. GPT theory emphasizes that productivity gains materialize only after complementary investments and organizational embedding occur. The evidence here suggests that AI behaves precisely in this way: productivity gains arise from persistent and deep integration of AI into the innovation portfolio, not from sporadic experimentation.

Column (5) in Table 6, which measures AI intensity as the logarithm of AI patents per employee ($log(1 + AIpatents/emp)_{it}$), yields consistent results, reinforcing that the economic mechanism operates through sustained AI capability building rather than through scale alone. Taken together, the margin decomposition clarifies that the productivity returns to AI are not driven by discrete adoption events but by directional specialization and capability accumulation. The negative extensive-margin coefficient captures transitional costs, while the positive intensive-margin effect reflects realized productivity gains once AI becomes embedded within the firm's technological trajectory.

The distinction between extensive and intensive margins also sharpens the interpretation of the instrumental-variable estimates in Section 4.2. The Bartik instrument exploits predetermined exposure to AI technology classes interacted with aggregate AI growth. Importantly, this design generates exogenous variation in firms' predicted AI orientation that operates primarily through shifts in technological opportunity rather than firm-specific productivity shocks. Because the instrument strengthens markedly after 2010—precisely when AI begins to diffuse as a general-purpose technology—it identifies variation in sustained AI reorientation rather than one-time adoption episodes. The IV estimates therefore capture the causal effect of deepening AI orientation conditional on complementary capability building, consistent with the intensive-margin mechanism documented above. The larger magnitude of the IV coefficient relative to fixed-effects estimates is consistent with attenuation bias arising from measurement error in AI shares and with endogenous timing of adoption decisions (Goldsmith-Pinkham et al., 2020). In this sense, the IV results reinforce the interpretation that AI-driven productivity gains arise from exogenous expansions in tech-

nological opportunity interacting with firms' pre-existing technological trajectories, rather than from endogenous selection of high-productivity firms into AI.

Table 6: Extensive-Intensive Margin Effects of AI-Oriented Innovation

$\log(Productivity)_{it}$	(1)	(2)	(3)	(4)	(5)
$\log(AIShare)_{it}$	0.053***				
	(0.007)				
$\log(1 + AIShare)_{it}$		-0.068**			
		(0.030)			
AI_Dummy_{it}			-0.142***	-0.207***	
			(0.009)	(0.013)	
$\log(1 + AIShare)_{it}$				0.353***	
$\times$ AI_Dummy_{it}				(0.042)	
$log(1 + AIpatents/emp)_{it}$					-0.040***
					(0.008)
Controls	Yes	Yes	Yes	Yes	Yes
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
Observations	12,227	36,146	36,146	36,146	36,146
R-squared	0.702	0.739	0.744	0.745	0.739

Notes: The dependent variable is $\log(Productivity)_{it}$, measured as the logarithm of sales per employee. Standard errors are clustered at the firm level and reported in parentheses. All controls in 4 are included. All Columns are estimated using firm fixed effects. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

4.4 Heterogeneous Impacts of AI-Oriented Innovation

To examine when AI-oriented innovation generates larger productivity gains, I estimate a series of fixed-effects regressions that interact AI-oriented innovation ($\log(AIShare)_{it}$) with firm- and market-level characteristics capturing organizational commitment, technological scope, and competitive conditions in Table 7. Each interaction is introduced in a separate specification, holding constant the full set of controls from equation (4) as well as firm and year fixed effects.

Column (1) in Table 7 reports the baseline estimate of a 1 percent increase in the AI patent share is associated with a 0.053 percent increase in labour productivity, measured as sales per employee. Column (2) allows the productivity effect of AI ($\log(\text{AIShare})_{it}$) to vary with

persistence in AI activity ($Persistence_{it}$), which is defined as an indicator variable equal to 1 if a firm has at least three consecutive years of AI patenting. Among firms without persistent AI engagement, the elasticity of labour productivity with respect to AI patent share is 0.038. For firms with persistent AI activity, the total effect rises to 0.057 (Std. Error=0.008), with the interaction term being positive and statistically significant. This pattern indicates that AI generates larger productivity gains when adoption is sustained rather than episodic, consistent with learning-by-doing, organizational adjustment costs, and the accumulation of AI-specific human capital. These results align with evidence that productivity gains from advanced technologies depend critically on complementary organizational investments and experience (Bresnahan et al., 2002).

Column (3) in Table 7 examines heterogeneity with respect to AI diversity ($Diversity_{it}$), measured as the number of distinct AI technology fields in which a firm patents. There are eight fine-grained AI technology classes in the data, which are machine learning, evolutionary computation, natural language processing, vision, speech, knowledge processing, planning and control, and AI hardware. The interaction between AI intensity and AI diversity is positive and statistically significant. For firms with minimal AI scope, the elasticity of labor productivity with respect to AI patent share is 0.056. Each additional AI technology field increases this elasticity by approximately 0.014, yielding a total effect of 0.071 (Std. Error=0.008). This finding suggests that AI-oriented innovation is more productive when deployed across a broader set of technological domains, consistent with AI functioning as a general-purpose technology whose returns are amplified by scope, recombination, and complementarities across tasks (Bresnahan and Trajtenberg, 1995; Cockburn et al. , 2018).

Column (4) in Table 7 interacts AI intensity with product market concentration, measured by the logarithm of the market HHI ($logHHI_{it}$). This variable is explained in Section 2. While the interaction term itself is not statistically significant, evaluating the total effects at different points in the HHI distribution reveals meaningful heterogeneity. AI-oriented innovation raises productivity in both competitive and concentrated markets. However, the

effect is stronger in more concentrated industries. In relatively competitive industries (25th percentile of HHI), a 1 percent increase in AI patent share raises labour productivity by 0.053 percent (Std. Error=0.010). In more concentrated industries (75th percentile), the elasticity rises to 0.063 (Std. Error=0.009). The stronger productivity gains by about 19% observed in concentrated markets suggest that AI enables firms to reinforce market power. Therefore, the weaker competitive pressure may increase the appropriability of AI-driven gains or allow firms to leverage AI strategically to reinforce market power. This interpretation is consistent with work emphasizing the interaction between innovation, market structure, and rent extraction (Aghion et al., 2005; Autor et al., 2020).

Finally, Column (5) in Table 7 considers whether fragmentation of patent ownership in the technology market moderates the productivity returns to AI by interacting $log(AIShare)_{it}$ with patent fragmentation in the technology space ($logF_it$). The construction of the variable $logF_it$ is explained in Section 2. The interaction term is small and statistically insignificant, indicating that the marginal productivity effect of AI is remarkably stable across different degrees of fragmentation. Evaluated at both low and high levels of fragmentation, the total effect of labour productivity with respect to AI patent share at 25th percentile is 0.059 (Std. Error=0.009) and at 75th percentile is 0.057 (Std. Error=0.008). While fragmentation itself is associated with lower productivity levels (-0.200, Std. Error=0.047), it does not attenuate the productivity gains from AI-oriented innovation. This suggests that AI delivers robust efficiency improvements even in complex and fragmented intellectual property environments, consistent with the view that general-purpose technologies can generate gains despite institutional frictions (Ziedonis, 2004; Hall et al., 2014).

5 Conclusion

This paper examines how shifting the direction of firm innovation toward artificial intelligence affects productivity, and how these effects depend on market structure and the organization

Table 7: Heterogeneous Effects of AI-Oriented Innovation on Labor Productivity

$\log(Productivity)_{it}$	(1)	(2)	(3)	(4)	(5)
$\log(AIShare_{it})$	0.053***	0.038***	0.056***	0.069***	0.056***
	(0.007)	(0.011)	(0.008)	(0.016)	(0.008)
$\log(AIShare)_{it}$		0.018**			
$\times$Persistence$_{it}$		(0.008)			
$\log(AIShare)_{it}$			0.014***		
$\times$Diversity$_{it}$			(0.002)		
$\log(AIShare)_{it}$				0.010	
$\times\log$HHI$_{it}$				(0.012)	
$\log(AIShare)_{it}$					-0.011
$\times\log$F$_{it}$					(0.020)
Total effect		0.057***	0.071***	0.063***	0.057***
		(0.008)	(0.008)	(0.010)	(0.008)
Controls	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	12,227	12,227	12,227	12,227	12,227
R^2	0.702	0.697	0.704	0.698	0.700

Notes: Standard errors clustered at the firm level are reported in parentheses. Total effects are computed as the sum of the coefficient on $\log(AIShare)_{it}$ and the interaction term, evaluated at one unit of the interacting variable or at the relevant percentile where applicable. *** , ** , * denote 1%, 5%, 10% levels of significance, respectively. All controls in 4 are included.

of intellectual property rights. Using a newly constructed panel of publicly traded U.S. firms from 1980 to 2021 that links Compustat financials to patent and citation data, I measure AI-oriented innovation as the share of a firm's patents that are AI-related. The empirical framework distinguishes what firms innovate in from how much they innovate by controlling for R&D and patent-based innovation intensity, capital deepening, and a rich set of firm and market characteristics, while exploiting within-firm variation through firm and year fixed effects. Because AI intensity is defined as the share of AI patents within total patenting activity, the baseline estimates identify the intensive margin of AI-oriented innovation among patenting firms rather than the discrete decision to begin AI patenting.

The results provide consistent evidence that reorienting innovation toward AI is associated with higher productivity. In baseline fixed-effects specifications, increases in AI patent share predict positive and statistically significant impacts on labour productivity, and this relationship reflects sustained increases in AI intensity rather than one-time adoption episodes. This relationship is robust to alternative constructions of AI exposure, alternative innovation controls, sample restrictions, and a broad set of alternative productivity measures (including TFP and value added per worker). Dynamic and placebo-style exercises are consistent with the idea that productivity gains reflect substantive reorientation toward AI rather than generic patenting activity or short-run experimentation.

A second central finding is that intellectual-property frictions matter. Greater patent ownership fragmentation in a firm's relevant technology space is associated with lower productivity, consistent with the idea that congested IP environments raise transaction costs and slow effective commercialization and deployment. However, fragmentation does not meaningfully attenuate the marginal productivity returns to AI-oriented innovation: the interaction between AI intensity and fragmentation is small and statistically insignificant. This pattern suggests that AI can deliver efficiency gains even in the presence of patent-thicket frictions, though firms operating in more fragmented technology spaces remain less productive on average.

The paper also documents systematic heterogeneity consistent with AI as a general-purpose technology and with market-structure mechanisms. Productivity gains from AI-oriented innovation are larger for firms with persistent AI activity, indicating that learning, organizational embedding, and complementary investments are important for translating AI innovation into performance. Gains are also larger when firms engage across a broader set of AI technology classes, consistent with scope-based complementarities and recombination across tasks and applications. Finally, returns to AI-oriented innovation are stronger in more concentrated industries, consistent with greater appropriability, scale-related complementarities, and the possibility that AI strengthens competitive advantage in environments where firms can better capture rents.

To address endogeneity concerns, the paper complements fixed-effects evidence with instrumental-variables estimates using a shift–share (Bartik) design that exploits predetermined firm exposure to fine-grained AI technology classes interacted with post-2010 aggregate growth in those classes. The instrument is strongly predictive after 2010—when AI becomes an economically meaningful technological opportunity—and IV estimates imply that exogenous increases in AI-oriented innovation raise productivity. The larger IV magnitude relative to fixed-effects estimates is consistent with attenuation from measurement error in firm-year AI shares and with endogenous selection into AI adoption, particularly along the extensive margin.

Together, these findings contribute to a clearer empirical picture of how AI shapes firm performance. They suggest that the productivity payoff from AI is fundamentally about the direction of innovation and the organizational capacity to sustain and broaden AI engagement, not merely about adopting AI in a limited or experimental way. They also imply that institutional frictions in technology markets—especially patent fragmentation—can depress productivity levels, but do not eliminate the productivity benefits associated with AI-oriented innovation.

The policy implications follow directly. Policies aimed at raising productivity through

AI should emphasize complements—skills, data infrastructure, computing capacity, and organizational capital—because persistence and breadth in AI innovation are strongly linked to larger productivity gains. At the same time, because returns to AI appear stronger in concentrated industries, policymakers should consider diffusion and competition policies that encourage broader access to AI capabilities while preserving incentives for long-run investment. Finally, reforms that reduce transaction costs in fragmented patent landscapes may yield productivity benefits directly and may further facilitate the translation of AI innovation into widespread performance improvements.

Looking forward, several extensions are natural. First, disentangling efficiency gains from markup changes would sharpen the interpretation of revenue-based productivity effects, for example by combining the current framework with demand/markup estimation or alternative quantity-based outcomes where feasible. Second, deeper analysis of mechanism—such as the roles of organizational restructuring, complementary intangible investment, or changes in product-market strategy—would help explain why persistence and breadth matter so strongly. These directions would further clarify when and why AI delivers productivity gains, and how policy can foster broad-based benefits from this general-purpose technology.

References

Acemoglu, Daron. 1998. "Why do new technologies complement skills? Directed technical change and wage inequality." *The quarterly journal of economics*, 113(4): 1055–1089.

Acemoglu, Daron. 2002. "Directed technical change." *The review of economic studies*, 69(4): 781–809.

Acemoglu, Daron, and Pascual Restrepo. 2020. "The wrong kind of AI? Artificial intelligence and the future of labour demand." *Cambridge Journal of Regions, Economy and Society*, 13(1): 25–35.

Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price. 2016. "Import Competition and the Great US Employment Sag of the 2000s." *Journal of Labor Economics*, 34(S1): S141–S198.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. "Artificial intelligence and jobs: Evidence from online vacancies." *Journal of Labor Economics*, 40(S1): S293–S340.

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, reallocation, and growth." *American Economic Review*, 108(11): 3450–3491.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica*, 83(6): 2411–2451.

Aghion, Philippe, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen. 2016. "Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry." *Journal of Political Economy*, 124(1): 1–51.

Aghion, Philippe, Antonin Bergeaud, Mathieu Lequien, and Marc J. Melitz. 2019. "The Impact of Exports on Innovation: Theory and Evidence." *American Economic Review*, 109(7): 2534–2573.

Aghion, Philippe, Benjamin F Jones, and Charles I Jones. 2017. "Artificial intelligence and economic growth." National Bureau of Economic Research.

Aghion, Philippe, Nicholas Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. "Competition and Innovation: An Inverted-U Relationship." *Quarterly Journal of Economics*, 120(2): 701–728.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2019. *Prediction Machines: The Simple Economics of Artificial Intelligence.* Boston, MA:Harvard Business Review Press.

Arora, Ashish, Sharon Belenzon, L Cioaca, L Sheer, HM Shin, and D Shvadron. 2024. "DISCERN 2.0: Duke innovation & scientific enterprises research network [dataset]." *Zenodo (CERN European Organization for Nuclear Research). https://doi.org/10.5281/zenodo*, 3594642.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics*, 135(2): 645–709.

Autor, David H., David Dorn, Gordon H. Hanson, and Jae Song. 2014. "Trade Adjustment: Worker-Level Evidence from the China Shock." *Quarterly Journal of Economics*, 129(4): 1799–1860.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024*a*. "Artificial intelligence, firm growth, and product innovation." *Journal of Financial Economics*, 151: 1–26.

Babina, Tania, Anastassia Fedyk, Alex He, James Hodson, and Geoffrey Tate. 2024*b*. "Artificial Intelligence, Firm Growth, and Industry Concentration." *Journal of Finance*. forthcoming.

Baldwin, John, and Wulong Gu. 2004. "Industrial Competition, Shifts in Market Share and Productivity Growth." Economic Analysis Research Paper Series of Statistics Canada No.021.

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI:W. E. Upjohn Institute for Employment Research.

Bessen, James E., Maarten Goos, Anna Salomons, and Wouter Van den Berge. 2022. "Automatic Reaction—What Happens to Workers at Firms That Automate?" *Review of Economics and Statistics*, 104(4): 675–692.

Bloom, Nicholas, and John Van Reenen. 2007. "Measuring and explaining management practices across firms and countries." *Quarterly Journal of Economics*, 122(4): 1351–1408.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying technology spillovers and product market rivalry." *Econometrica*, 81(4): 1347–1393.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans do IT better: US multinationals and the productivity miracle." *American Economic Review*, 102(1): 167–201.

Blundell, Richard, Rachel Griffith, and John VanReenen. 1999. "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms." *The Review of Economic Studies*, 66(3): 529–554.

Bresnahan, Timothy F, and Manuel Trajtenberg. 1995. "General purpose technologies 'Engines of growth'?" *Journal of econometrics*, 65(1): 83–108.

Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt. 2002. "Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence." *The quarterly journal of economics*, 117(1): 339–376.

Brown, James R, Steven M Fazzari, and Bruce C Petersen. 2009. "Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom." *Journal of Finance*, 64(1): 151–185.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies.* WW Norton & company.

Brynjolfsson, Erik, and Lorin M Hitt. 2000. "Beyond computation: Information technology, organizational transformation and business performance." *Journal of Economic perspectives*, 14(4): 23–48.

Brynjolfsson, Erik, and Lorin M Hitt. 2003. "Computing productivity: Firm-level evidence." *Review of economics and statistics*, 85(4): 793–808.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2018. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In *The Economics of Artificial Intelligence: An Agenda.* , ed. Ajay Agrawal, Joshua Gans and Avi Goldfarb, 23–57. Chicago, IL:University of Chicago Press.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The productivity J-curve: How intangibles complement general purpose technologies." *American Economic Journal: Macroeconomics*, 13(1): 333–372.

Calvino, F, and L Fontanelli. 2023. "A Portrait of AI Adopters across Countries: Firm Characteristics, Assets." *Complementarities and Productivity', OECD Working Paper*, 2.

Clarke, Damian. 2021. "Implementing the Panel Event Study." *The Stata Journal*, 21(4): 853–884.

Cockburn, Iain M, Rebecca Henderson, and Scott Stern. 2018. "The impact of artificial intelligence on innovation: An exploratory analysis." In *The economics of artificial intelligence: An agenda.* 115–146. University of Chicago Press.

Cohen, Wesley M., and Daniel A. Levinthal. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation." *Administrative Science Quarterly*, 35(1): 128–152.

Czarnitzki, Dirk, Gastón P Fernández, and Christian Rammer. 2023. "Artificial intelligence and firm-level productivity." *Journal of Economic Behavior & Organization*, 211: 188–205.

Damioli, Giacomo, Vincent Van Roy, and Daniel Vertesy. 2021. "The impact of artificial intelligence on labor productivity." *Eurasian Business Review*, 11(1): 1–25.

David, Paul A. 1990. "The dynamo and the computer: an historical perspective on the modern productivity paradox." *The American economic review*, 80(2): 355–361.

DeLoecker, Jan, and Chad Syverson. 2021. "An industrial organization perspective on productivity." In *Handbook of industrial organization.* Vol. 4, 141–223. Elsevier.

De Loecker, Jan, and Frederic Warzynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review*, 102(6): 2437–2471.

De Loecker, Jan, and Jan Eeckhout. 2020. "Global Market Power." *Quarterly Journal of Economics*, 135(2): 561–644.

De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. 2016. "Prices, Markups, and Trade Reform." *Econometrica*, 84(2): 445–510.

Demirel, Pelin, and Mariana Mazzucato. 2012. "Innovation and Firm Growth: Is R&D Worth It?" *Industry and Innovation*, 19(1): 45–62.

Dunne, Paul, and Alan Hughes. 1994. "Age, Size, Growth and Survival: UK Companies in the 1980s." *Journal of Industrial Economics*, 115–140.

Duso, Tomaso, LarsHendrik Roller, and Jo Seldeslachts. 2014. "Collusion through Joint R&D: An Empirical Assessment." *Review of Economics and Statistics*, 96(2): 349–370.

Entezarkheir, Mahdiyeh. 2017. "Patent Thickets, Defensive Patenting, and Induced R&D: An Empirical Analysis of the Costs and Potential Benefits of Fragmentation in Patent Ownership." *Empirical Economics*, 1–36.

Evans, David. 1987. "Tests of Alternative Theories of Firm Growth." *Journal of Political Economy*, 95(4): 657–674.

Fontanelli, Luca, Mattia Guerini, Raffaele Miniaci, and Angelo Secchi. 2025. "Predictive AI and Productivity Growth Dynamics: Evidence from French Firms." *Journal of Economic Behavior & Organization*, 240: 1–25.

Gans, Joshua S. 2024. "Market Power in Artificial Intelligence." *NBER Working Paper 32408*.

Giczy, Alexander V, Nicholas A Pairolero, and Andrew A Toole. 2022. "Identifying artificial intelligence (AI) invention: A novel AI patent dataset." *The Journal of Technology Transfer*, 47(2): 476–505.

Giroud, Xavier, and Hogler Mueller. 2010. "Does Corporate Governance Matter in Comeptitive Industries?" *Journal of Financial Economics*, 95(1): 312–331.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik instruments: What, when, why, and how." *American Economic Review*, 110(8): 2586–2624.

Griliches, Zvi. 1979. "Issues in Assessing the Contribution of Research and Development to Productivity Growth." *Bell Journal of Economics*, 10(1): 92–116.

Griliches, Zvi, and Ariel Pakes. 1980. "Patents and R&D at the Firm Level: A First Look." National Bureau of Economic Research Working Paper No. 561.

Griliches, Zvi, and Jacques Mairesse. 1995. "Production Functions: The Search for Identification." National Bureau of Economic Research NBER Working Paper 5067.

Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg. 2001. "The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools." National Bureau of Economic Research Working Paper No. 8498.

Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg. 2005. "Market Value and Patent Citations." *RAND Journal of Economics*, 36(1): 16–38.

Hall, Bronwyn H., Christian Helmers, Mark Rogers, and Vania Sena. 2014. "Intellectual property protection and the role of patent intermediaries." *Review of Economics and Statistics*, 96(4): 682–699.

Hall, Bronwyn H., Jacques Mairesse, and Pierre Mohnen. 2009. "Measuring the Returns to R&D." In *Handbook of the Economics of Innovation*. Vol. 2, , ed. Bronwyn H. Hall and Nathan Rosenberg, 1033–1082. Elsevier.

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda. 2013. "Who creates jobs? Small versus large versus young." *Review of Economics and Statistics*, 95(2): 347–361.

Heller, Michael A., and Rebecca S. Eisenberg. 1998. "Can patents deter innovation? The anticommons in biomedical research." *Science*, 280(5364): 698–701.

Helpman, Elhanan, and Manuel Trajtenberg. 1994. "A time to sow and a time to reap: Growth based on general purpose technologies."

Huergo, Elena, and Jordi Jaumandreu. 2004. "How does probability of innovation change with firm age?" *Small Business Economics*, 22(3): 193–207.

Jaffe, Adam. 1986. "Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value." *American Economic Review*, 76(5): 984–1001.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological innovation, resource allocation, and growth." *The quarterly journal of economics*, 132(2): 665–712.

Korinek, Anton, and Joseph E. Vipra. 2024. "Artificial Intelligence Production Networks." *Review of Economic Studies*. forthcoming / working paper version cited.

Lanjouw, Jean O., and Mark Schankerman. 2004. "Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators." *Economic Journal*, 114(495): 441–465.

Lanjouw, Jean O, Ariel Pakes, and Jonathan Putnam. 1998. "How to count patents and value intellectual property: The uses of patent renewal and application data." *The journal of industrial economics*, 46(4): 405–432.

Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, and Sidney G. Winter. 1987. "Appropriating the Returns from Industrial Research and Development." *Brookings Papers on Economic Activity*, , (3): 783–831.

Levinsohn, James, and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Econometrica*, 71(2): 317–341.

Lucas, Robert E. 1978. "On the size distribution of business firms." *Bell Journal of Economics*, 9(2): 508–523.

Luttmer, Erzo G.J. 2007. "Selection, growth, and the size distribution of firms." *Quarterly Journal of Economics*, 122(3): 1103–1144.

Mankiw, N. Gregory, David Romer, and David N. Weil. 1992. "A contribution to the empirics of economic growth." *Quarterly Journal of Economics*, 107(2): 407–437.

Mihet, Raluca. 2025. "Artificial Intelligence, Data, and Market Power." *Journal of Monetary Economics*. forthcoming / working paper version cited.

Miller, Conrad, Nathaniel Johnson, and Laura Wherry. 2023. "An Introductory Guide to Event Study Models." *Journal of Economic Perspectives*, 37(2): 203–230.

Nickell, Stephen J. 1996. "Competition and corporate performance." *Journal of Political Economy*, 104(4): 724–746.

Pairolero, Nicholas A, Alexander V Giczy, Gerard Torres, Tisa Islam Erana, Mark A Finlayson, and Andrew A Toole. 2025. "The artificial intelligence patent dataset (AIPD) 2023 update." *The Journal of Technology Transfer*, 1–24.

Schmidheiny, Kurt, and Sebastian Siegloch. 2019. "On Event Study Designs and Distributed-Lag Models." IZA IZA DP No. 12079.

Siedschlag, Iulia, and Juan David Duran Vanegas. 2025. "Artificial Intelligence and Firm-Level Productivity: Early Evidence from a Small Open Economy." *Available at SSRN 5375462*.

Solow, Robert M. 1956. "A contribution to the theory of economic growth." *Quarterly Journal of Economics*, 70(1): 65–94.

Syverson, Chad. 2017. "Challenges to mismeasurement explanations for the US productivity slowdown." *Journal of Economic Perspectives*, 31(2): 165–186.

Tambe, Prasanna. 2014. "Big Data Investment, Skills, and Firm Value." *Management Science*, 60(6): 1452–1469.

Teece, David J. 1986. "Profiting from Technological Innovation: Implications for Integration, Collaboration, Licensing and Public Policy." *Research Policy*, 15(6): 285–305.

WIPO, WIPO. 2019. "WIPO technology trends 2019: Artificial intelligence." *Geneva: World Intellectual Property Organization*.

Ziedonis, Rosemarie. 2004. "Don't Fence Me In: Fragmented Markets for Technology and the Patent Acquisition Strategies of Firms." *Management Science*, 50(6): 804–820.

Appendix A Correcting Truncation in Patent and Citation Counts

Because of the lag between patent application and grant, and because citations accrue over time, both patent counts and citation counts are subject to truncation. Patents applied for near the end of the sample may be granted after the sample period, and patents granted within the sample may receive citations after the end of the citation window. To address this issue, I follow the correction procedure proposed by Hall et al. (2005).

Using the PatentsView data, I observe all citations made by U.S. granted patents (the *citing* patents) to earlier U.S. granted patents (the *cited* patents). For each citation, I define the citation age as

$$a = \text{GrantYear}_{\text{citing}} - \text{GrantYear}_{\text{cited}},$$

that is, the number of years elapsed between the grant year of the cited patent and the grant year of the citing patent.

Let $g(a)$ denote the empirical distribution of citation ages. Formally, $g(a)$ is the probability that a forward citation to a patent occurs at age a years after the patent is granted. Its cumulative distribution function is

$$S(a) = \sum_{\tau=0}^{a} g(\tau),$$

which represents the share of a patent's lifetime citations that are expected to have occurred within a years of its grant.

For each patent j, let

$$A_j = T_{\text{last}} - \text{GrantYear}_j$$

be its observable citation window, where T_{last} is the final year of the citation sample and GrantYear_j is the grant year of patent j. Thus, A_j measures how many years of potential citation activity we are able to observe for patent j.

Let $\text{cit}_j^{\text{obs}}$ denote the number of forward citations to patent j that are actually observed within the sample period. Following Hall et al. (2005), the truncation-adjusted citation count for patent j is defined as

$$\text{cit}_j = \frac{\text{cit}_j^{\text{obs}}}{S(A_j)},$$

where $S(A_j)$ rescales the observed citations by the fraction of the citation life cycle that is observable for that patent. For patents with a sufficiently long observable window such that their citation life cycle is essentially complete, I impose $S(A_j) = 1$.

I then merge these patent-level corrected citation counts into the Discern patent–firm–year panel, so that each patent j assigned to firm i in year t carries a corrected forward-citation measure cit_{ijt}.

To construct firm-level citation flows, I aggregate corrected citations across all patents granted to firm i in year t:

$$\text{flowCi}_{it}^{\text{raw}} = \sum_{j \in \mathcal{P}_{it}} \text{cit}_{ijt},$$

where $\mathcal{P}_{it}$ is the set of patents granted to firm i in year t. This is the raw citation-weighted patent flow.

Because patent grants themselves are also right-truncated near the end of the sample, I additionally construct a truncation-adjusted version of the citation flow, $\text{flowCi}_{it}^{\text{adj}}$, which scales $\text{flowCi}_{it}^{\text{raw}}$ to be consistent with the corresponding truncation-adjusted patent counts, following the same logic as Hall et al. (2005).

Finally, I build the citation-weighted patent stock, explained in Section 2. This construction ensures that CitPatent_{it} reflects both the quantity and the quality of a firm's innovative activity while appropriately correcting for the two key sources of truncation: incomplete observation of patent grants and incomplete observation of forward citations.