AI-Skilled Labor, Intangible Capital, and Firm Productivity: Evidence from the TensorFlow Shock*

Suleyman Faruk Gozen† Yusuf Ozkara‡ Ruveyda Nur Gozen§

University of Bristol *Boston College* *Cardiff Business School & CEP LSE*

August 2026

Abstract

This paper studies how the complementarity between intangible capital and AI-skilled workers shapes firm-level productivity. We exploit the 2015 release of Google TensorFlow as a plausibly exogenous shock to AI effectiveness. Using a difference-in-differences framework with value-added labor productivity and total factor productivity, we show that firms with a higher AI-worker share increase their productivity. The gains rise with firm size and intangible capital, and the effect is concentrated in the upper part of the productivity distribution. This evidence indicates that the complementarity between intangible capital and AI-skilled labor widens productivity dispersion. We also provide TensorFlow-specific evidence using job postings and worker profiles from Revelio Labs proprietary data.

Keywords: Productivity, Intangible Capital, Artificial Intelligence, Firm Size

JEL codes: D22, D24, E22, J24, O33

*The previously circulated title was "Intangible Capital Meets Skilled Labor: The Implications for Productivity Dynamics". We thank Arpad Abraham, Ufuk Akcigit, Pascal St-Amour, Susanto Basu, Fabio Bertolotti, Tsenguunjav Byambasuren, Andrea Caggese, Andrea Chiavari, Se Kyu Choi, Bruno Conte, Luis Correia, Diane Coyle, Richard Dickens, Pawel Doligalski, Abdul A. Erumban, Francois Fontaine, Lars Peter Hansen, Jonathan Haskel, Iftikhar Hussain, Giammario Impullitti, Peter Ireland, Greg Kaplan, Ezgi Kaya, William R. Kerr, Jin Ho Kim, Per Krusell, Maximilian Konradt, Danial Lashkari, Cecilia Jona-Lasinio, Christos Makridis, Josh Martin, Matthias Mertens, Roxana Mihet, Patrick Minford, Iacopo Morchio, Sergio Feijoo Moreira, Mary O'Mahony, Luisanna Onnis, Wojtek Paczos, Theodore Papageorgiou, Ziho Park, Ahmed Pirzada, Sergey Popov, Lukasz Rachel, Tommaso Reggiani, Jose Victor Rios Rull, Carolina Villegas-Sanchez, Edouard Schaal, Fabio Schiantarelli, Paul Schreyer, Saverio Spinella, Kjetil Storesletten, Chad Syverson, Richard Tol, Harald Uhlig, Haomin Wang, Christoph Weiss, Thomas Winberry, Sezer Yasar, and Hakki Yazici as well as seminar and workshop participants at Boston College, Bristol-Surrey Workshop in Macroeconomics, BSE Summer Forum, Cardiff Business School, Comparative Analysis of Enterprise Data (CAED) Conference (the Pennsylvania State University), ESCoE Conference on Economic Measurement 2026, Global INTAN-Invest Conference, The Productivity Institute Spring Brown Bag Seminar Series 2026, Royal Economic Society 2025 Annual Conference, Scottish Economic Society 2026 Annual Conference, University of Bristol, the University of Chicago, and University of Sussex Business School for helpful comments and suggestions. All errors are our own.

†Corresponding author. School of Economics, University of Bristol. 12 Priory Road, Bristol BS8 1TU, United Kingdom. Email: suleyman.gozen@bristol.ac.uk.

‡Economics Department, Boston College. 140 Commonwealth Avenue, Chestnut Hill, MA 02467, United States. Email: ozkaray@bc.edu.

§Economics Department, Cardiff Business School & Research Associate, CEP LSE. Cardiff Business School: Aberconway Building, Colum Drive, Cardiff CF10 3EU, United Kingdom. CEP LSE: Houghton Street, London WC2A 2AE, United Kingdom. Email: gozenr@cardiff.ac.uk & r.n.gozen@lse.ac.uk.

1 Introduction

Productivity differences across firms have widened markedly over the past several decades, with an increasing share of productivity growth driven by large firms (Andrews et al., 2016; Decker et al., 2020; Akcigit and Ates, 2023). While existing explanations emphasize market power (Gutiérrez and Philippon, 2017), declining competition, or the role of superstar firms (Bessen, 2017; Autor et al., 2020), the causes and mechanisms underlying the rise in productivity dispersion remain contested.

At the same time, two structural transformations have reshaped the firm production. First, intangible capital, such as know-how, proprietary knowledge, and data-driven processes, has become an increasingly important input into production. Second, firms have experienced a pronounced rise in demand for workers with advanced computational and analytical expertise, and in particular for AI-skilled labor. Despite extensive research on each of these trends in isolation, much less is known about how intangible capital and AI-skilled labor interact within firms to shape productivity.

This paper argues that the interaction between intangible capital and AI-skilled labor constitutes a key yet underexplored channel behind rising productivity dispersion. We propose that intangible capital and AI-skilled labor are strongly complementary, in the sense that intangible capital generates productivity gains when firms possess the AI-skilled workforce required to deploy, adapt, and scale them effectively. Crucially, because intangible capital is scalable and non-rivalrous, this complementarity is amplified by firm size, allowing larger firms to exhibit higher productivity from the same technological inputs (Haskel and Westlake, 2017; Crouzet et al., 2022b).

This perspective leads us to three questions. What is the effect of AI-skilled labor on firm-level productivity? What is the role of the complementarity between intangible capital and AI-skilled labor in the productivity channel? How does firm-size heterogeneity shape the magnitude of this complementarity? To address these questions, we document and quantify a new channel, namely the complementarity between intangible capital and AI-skilled labor, as a partial explanation for the increasing productivity dispersion across the firm-size distribution in the U.S. economy.

To answer these questions, we need variation in the effectiveness of AI-skilled labor that is plausibly unrelated to firms' other choices. Following the exogenous variation insight of Rock (2019), we exploit the November 2015 release of Google's TensorFlow, which is an open-source machine learning framework that sharply reduced the cost and complexity of deploying artificial intelligence tools. We conceptualize it as a

technology shock that increased the effectiveness of AI-skilled workers. We measure a firm's exposure to the shock with its share of AI-skilled workers at the time of the release, constructed from the AI measure of Babina et al. (2024). We hypothesize that AI-skilled workers embody the frontier human capital required to develop, deploy, and scale the data, software, and algorithmic tools that constitute intangible capital, so firms with greater AI readiness at the shock were differentially positioned to benefit from it. This differential exposure allows us to identify the causal impact of improvements in AI effectiveness on intangible capital and productivity.

Using a difference-in-differences (DiD) framework, we proceed in three steps. First, we establish the productivity effect. Firms with a higher share of AI-skilled workers at the time of the shock become substantially more productive after the TensorFlow release. We find a value-added labor productivity gain of roughly 3.8 log points for the average exposed firm, and the estimates build gradually from 2017 onward with no evidence of differential pre-trends. We also find similar set of effects on total factor productivity estimated with four alternative and standard production function methods. Second, we document the complementarity. Exposed firms accumulate substantially more intangible capital after the shock, and the productivity response rises systematically with the firm's intangible capital and size. Third, we show that productivity dispersion increases. We conduct a quantile DiD, and the causal estimates rise from 2.6 log points at the 10th percentile of the productivity distribution to 11.7 at the 90th. Moreover, we implement a location-scale model, which shows that the conditional variance of productivity widens only among large and intangible capital rich firms.

We also provide TensorFlow-specific and worker-specific evidence. A distinctive feature of our analysis is the use of proprietary worker-level data from Revelio Labs, which aggregate the near-universe of online professional profiles and job postings and allow us to observe the specific deep-learning frameworks listed in workers' skill profiles. These data discipline the interpretation of our estimates in two ways. First, they tie the workforce response to TensorFlow itself rather than to an undifferentiated post-2015 wave of digital diffusion. We find that after the release, TensorFlow skills in high-exposure firms grow to more than six times all other competing deep-learning frameworks combined, and the TensorFlow share of these firms' deep-learning workforce is 9.4 percentage points higher than in low-exposure firms. Second, it allows us to examine the productivity effect at a finer level of disaggregation. We show that high-exposure firms post higher pay for the same occupation, and that this premium is as large for non-technical occupations as for computing roles. It suggests a firm-wide pattern rather than only being specific to the roles that build AI systems and is

therefore consistent with the productivity effect shown in our baseline results.

An extensive battery of identification and robustness exercises supports these findings. We incorporate state- and industry-specific linear trends, industry-by-year and state-by-year fixed effects, a placebo test, propensity-score matching, a narrower definition of AI skills, two-way clustered inference, alternative measures, external validity through the exclusion of the AI-producing sectors, and a structurally estimated markup control following De Loecker et al. (2020). We find that the results remain robust.

This paper makes three main contributions. First, we provide causal firm-level evidence that intangible capital and AI-skilled labor are strongly complementary for productivity gains, and that this complementarity is amplified by firm size. In doing so, we complement Babina et al. (2024), which documents that AI investments predict firm growth and product innovation. We show that the productivity returns to the AI-skilled workers run through the accumulation of complementary intangible capital, and we investigate their distributional consequences. Second, our findings offer a technology-based explanation for rising productivity dispersion, which complements existing explanations regarding market power and competition. Third, rather than treating AI as a broad and unmeasured technology wave, we exploit a specific technology shock and verify with the Revelio Labs workforce data that firms' responses were directed at that specific technology.

Our findings carry policy implications. Because productivity gains from improved AI effectiveness depend on firms' existing intangible capital and scale, policies aimed at fostering AI adoption are likely to be most effective when combined with support for complementary intangible capital, while uniform AI subsidies may disproportionately benefit large firms and thereby widen productivity dispersion. This insight suggests a role for policies that help smaller and intangible capital poor firms build the complementary assets needed to benefit from AI technologies.

We situate our analysis for the U.S., where detailed firm-level data on intangible capital and productivity are available and where the rise in productivity dispersion has been extensively documented. There are similar patterns of rising intangible capital and skill-biased technological change, which have been observed in Europe and other advanced economies (Van Ark et al., 2013; Andrews et al., 2016; Bartelsman and Wolf, 2018; Adler and Siegel, 2019). Even though our evidence cannot transport beyond U.S. listed firms, these parallels suggest that the mechanism is unlikely to be U.S.-specific.

Related Literature Our paper is related to several strands of literature. The first strand focuses on the declining business dynamism in the U.S. economy. Explanations for this decline include slowing technological diffusion (Akcigit and Ates, 2023), reallocation of factors toward superstar firms (Autor et al., 2020), implementation and restructuring lags of breakthrough technology (Brynjolfsson et al., 2018), structural changes in the cost structure due to intangible capital (De Ridder, 2024), and increased market power driven by intangible capital (Crouzet and Eberly, 2019), among others. Our contribution is to emphasize an additional channel, the complementarity between intangible capital and AI-skilled labor, which benefits large firms and leads to an increase in productivity dispersion.

The second strand examines the secular rise of corporate intangible capital (Corrado et al., 2009; Corrado and Hulten, 2010; McGrattan and Prescott, 2010; Eisfeldt and Papanikolaou, 2014; Corrado et al., 2016; Haskel and Westlake, 2017; McGrattan, 2020). This literature documents that the accumulation of intangible capital affects several dimensions of firm dynamics, including productivity growth (Corrado et al., 2017; McGrattan, 2020), competition (Ayyagari et al., 2019), market power (Crouzet and Eberly, 2019; Zhang, 2019; De Ridder, 2024), markup (Altomonte et al., 2021), rents (Crouzet and Eberly, 2023), and factor inputs (Chiavari and Goraya, 2025). Closest to our emphasis on heterogeneity, Arrighetti et al. (2014) show, using a panel of Italian firms, that the propensity to invest in intangible assets increases with firm size, human capital, and past intangible investment, and that these differences are systematically linked to performance. Our contribution is to argue that, along with the rising share of intangible capital in the U.S. economy, the heterogeneity of intangible capital across different firm sizes can partially explain the increasing productivity dispersion.

Closest to our mechanism, Brynjolfsson et al. (2021) argue that general purpose technologies depress productivity while their complementary intangible capital is being built, and later raise it, and Tambe et al. (2020) show that firm-level digital capital is concentrated in and helps account for the performance of superstar firms. Our contribution is firm-level causal evidence on the complementarity. More specifically, we identify an improvement in the effectiveness of AI-skilled labor and show that its productivity returns are increasing in the firm's intangible capital and scale, and that this is a particular channel behind the increasing productivity dispersion.

The third strand highlights the relationship between skilled workers, firm performance, and technological progress, particularly in the context of intangible capital and AI. One set of studies shows that intangible assets are built by, and reward, specialized human capital. For instance, intangible intensive firms rely on workforces with

specialized skills and knowledge (Crouzet et al., 2022a), intangible assets result from joint investments by firms and skilled workers (Ladika et al., 2025), and these joint investments create deferred wage obligations that constrain firms' debt capacity (Sun and Xiaolan, 2019). A second set provides firm- and worker-level evidence linking skills, new technologies, and performance such that productivity is significantly higher in R&D- and skill-intensive industries (O'Mahony and Vecchi, 2009); upskilling is driven by the co-evolution of technology and organizational restructuring rather than by technology adoption alone (Piva et al., 2005); and AI technologies in particular are associated with higher productivity, employment, and innovation (Yang, 2022; Rammer et al., 2022), while AI exposure reshapes the task and skill content of occupations at the worker level (Engberg et al., 2025). Our contribution is to highlight the role of the complementarity between intangible capital and AI-skilled labor in the firm-level productivity and its impact on productivity dispersion across firms.

The final strand investigates the impact of technical change on labor market dynamics. Several papers explore wage dynamics (Katz and Murphy, 1992; Acemoglu, 1998; Katz et al., 1999; Autor et al., 2008; Violante, 2008), skill-biased technological change (Solow, 1957; Greenwood et al., 1997; Krusell et al., 2000; Acemoglu, 2002a,b; Aghion et al., 2002; Bresnahan et al., 2002; Hornstein et al., 2005), and capital-skill complementarity (Griliches, 1969; Greenwood and Yorukoglu, 1997; Goldin and Katz, 1998; Bresnahan et al., 2002; Autor et al., 2003). We complement this literature in two ways. We emphasize the role of the secular rise in intangible capital as a new form of technical change in the U.S. economy, and we then attribute heterogeneous productivity responses to a specific technology shock rather than to a broad technological trend.

The paper is organized as follows. Section 2 describes the data, measurement and sample. Section 3 presents stylized facts. Section 4 presents the research design and the causal results. Section 5 documents the complementarity between intangible capital and AI-skilled labor, the role of firm size. Section 6 provides the productivity-dispersion implication with distributional methods. Section 7 provides TensorFlow-specific evidence on the mechanism using the Revelio Labs data. Section 8 summarizes an extensive battery of robustness checks. Section 9 concludes. The appendix collects supporting material, including the sample statistics, the full set of robustness checks, and the additional analyses.

2 Data, Measurement, and Sample Composition

2.1 Data and Measurement

We use the U.S. Compustat database to measure firm-level financial balance sheet variables on an annual basis.[1]

In line with standard procedures in the literature, we exclude utilities (SIC codes 4900-4999), financial firms (SIC codes 6000-6999), and government entities (SIC code 9000 and above). We deflate all nominal variables with industry-level producer price indices, and Section A.1 describes the construction of these deflators in detail.

Measurement of Productivity Our baseline productivity measure is value-added labor productivity, defined as the logarithm of operating income before depreciation (Compustat `oibdp`) per employee. Operating income before depreciation equals revenue minus the cost of goods sold and selling, general, and administrative expenses. Therefore, it nets out intermediate inputs and is the closest available proxy for value added in Compustat. Constructing a textbook value-added measure would require a firm-level wage bill, and Compustat's staff expense item is populated for only 8.9% of our estimation sample, so we have a data limitation to implement the textbook version of value-added measure. We accordingly interpret it as operating performance per worker rather than as a pure form of technical efficiency. We winsorize it at the 1st and 99th percentiles within each year. As a complementary measure, we also estimate firm-level total factor productivity using four alternative and standard production function estimators with Olley and Pakes (1996), Levinsohn and Petrin (2003), Wooldridge (2009), and Ackerberg et al. (2015).

Measurement of Intangible Capital We measure firm-level intangible capital using the Total Q database of Peters and Taylor (2017), which captures both external and internal channels of intangible investment. External intangible capital reflects the acquisition of existing technologies, knowledge, and organizational capabilities through mergers and acquisitions. Internal intangible capital encompasses two main

[1]The unit of observation is the consolidated firm as Compustat reports company-level financial statements, so all variables aggregate over a firm's plants and establishments, and we observe no establishment-level variation. Compustat also covers only publicly listed firms, so our sample excludes private firms, startups, and small businesses. Therefore, our analyses should be read as evidence on the listed-firm segment of the economy, and we return to this limitation in the conclusion.

components called as knowledge capital and organizational capital. Knowledge capital includes investments in R&D, software, data infrastructure, and related innovation activities. Organizational capital captures firm-specific managerial practices, coordination routines, incentive structures, and decision-making processes. Both components are largely expensed rather than capitalized under accounting standards and therefore reside off the balance sheet. Therefore, we follow Peters and Taylor (2017) and construct internal intangible capital using a perpetual inventory method, consistent with the view that intangible assets accumulate through sustained investment (Lev and Radhakrishnan 2005; Eisfeldt and Papanikolaou 2014; Ewens et al. 2019).

Measurement of Firm Size Throughout the paper, firm size is proxied by the market value of equity, and this is a common approach in the related literature (e.g., DeAngelo et al., 2006; Linck et al., 2008; Edmans et al., 2009; Kogan et al., 2017). Market value is well suited to our setting for two reasons. First, because intangible capital largely resides off the balance sheet, book-based size measures systematically understate the scale of intangible capital rich firms. However, market value capitalizes the expected returns to both tangible and intangible capital, and provides a size ranking that does not mechanically penalize the firms included in our analysis. Second, it captures a forward-looking firm scale, which is measured consistently across industries with very different asset structures. A potential drawback is that market value also embeds the market's valuation of growth prospects, and so a size split based on it could partly sort firms on expected performance rather than on scale itself. To address this, in Section E.8 we re-estimate our related results with two accounting-based measures using total assets and employment, and show that the baseline pattern carries over.

Measurement of Firm-level AI Workers We use the measure constructed by Babina et al. (2024) for the firm-level share of AI workers, which proxies a novel measure of firms' AI investment by incorporating both the stock of AI-skilled workers currently employed by the firm and the flow of AI talent as reflected in job postings. This measure combines employee resume data from Cognism with job posting data from Burning Glass Technologies to create a comprehensive view of firms' engagement with AI talent. Specifically, the resume data from Cognism provides detailed employee profiles, which includes educational background, job titles, and self-reported skills, and it allows for the identification of existing workers with AI-related expertise. The Burning Glass data captures the demand side through job advertisements to determine firms' intent to hire workers with AI capabilities.

AI workers are identified based on a curated list of AI-related skills, which are classified into two categories. Narrow AI skills represent highly specialized and directly AI-relevant capabilities, and broad AI skills are more widely applicable and may include peripheral technologies. Our baseline measure relies on the broad definition.[2] We treat missing values as missing instead of filling them with zeros. This is a conservative choice that reduces sample size but removes any concern that imputed zeros generate spurious variation.

Two features of this measure discipline its interpretation. First, the AI-worker share measures AI-related human capital (e.g. the firm's capability to develop, deploy, and adapt AI tools), and it may not represent the general skill level of the firm-level workforce. Therefore, we do not interpret it as a measure of workforce quality, education, or human capital in the broad sense, and we refer to it as the firm's AI readiness. Second, the measure is a stock of workers rather than a flow of investment since it records the AI-skilled workforce a firm has in place, and it does not capture firm-level expenditure on AI adoption. This is precisely the object that our design requires because we hypothesize that the TensorFlow release raised the marginal product of the AI-skilled workers a firm already employed.

2.2 Sample Composition

The final sample used in the main analysis spans 2010-2019. We begin in 2010 to keep the Global Financial Crisis and its immediate aftermath out of the pre-treatment period as the sharp contraction of 2008-2009 would otherwise contaminate the pre-trends with a shock unrelated to AI. We end in 2019 to keep the COVID-19 period out of the post-treatment window since the 2020-2021 pandemic produced large swings in measured productivity that are unrelated to the TensorFlow shock and would confound the estimated response.[3] The baseline sample comprises 8,603 firm-year observations for 1,213 firms. The panel is unbalanced; firms enter and exit with their Compustat coverage. Section A.2 reports the full set of sample statistics, including variable definitions, the sample-selection steps, the distribution of observations across years and industries, summary statistics, and pairwise correlations.

[2] Section E.6 shows that our results are robust to the narrow definition.
[3] Section 4.3.1 provides an analysis by including 2008-2009 and post-2020 periods.

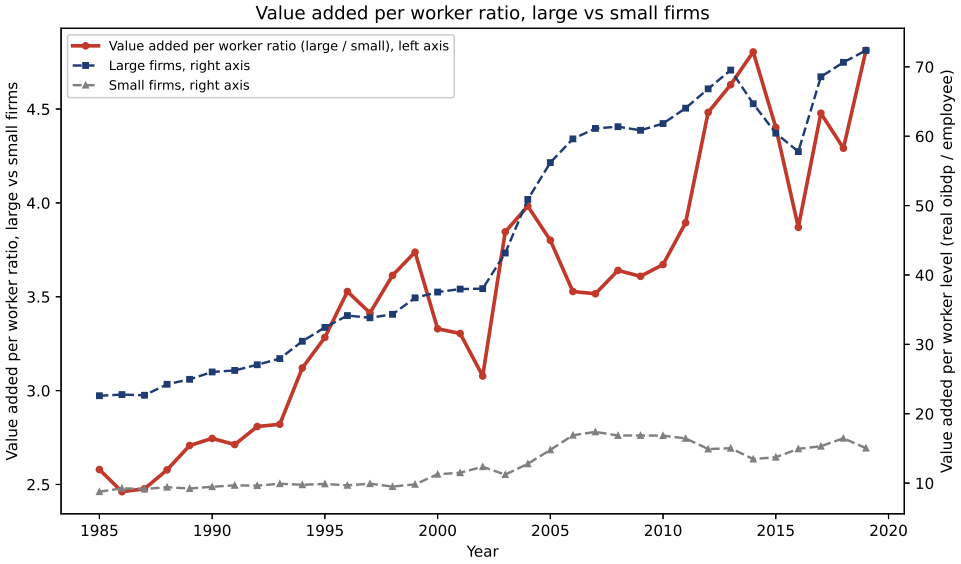
3 Stylized Facts

In this section, we document firm-level stylized facts which present patterns on productivity dispersion, intangible capital, and AI-skilled labor.

Fact 1: The value-added labor productivity gap between large and small firms has widened substantially over the past decades.

Figure 1 shows the average value-added labor productivity ratio between large firms (90th percentile) and small firms (10th percentile) of the firm size distribution within each industry and year, together with the group-level productivity paths.[4]

Figure 1: Value-Added Labor Productivity by Firm Size



Note: The left axis of the figure shows the average value-added labor productivity ratio between large firms and small firms. The right axis shows the average productivity level of large and small firms. Small firms are the ones at the 10th percentile and large firms are the ones at the 90th percentile of market value within each year and industry (*NAICS*). Industry-year cells are weighted by their number of firms.

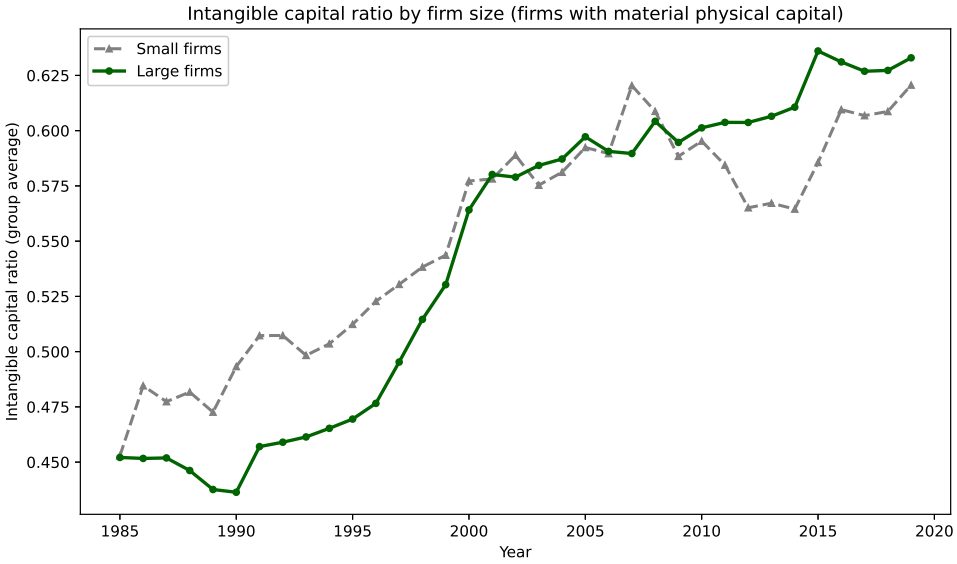
The figure shows that the productivity gap between large and small firms nearly doubled between the mid-1980s and the end of the sample. It rises from a ratio of about 2.5 to close to 5 with the widening especially pronounced after 2000. The right axis indicates that large firms exhibit a much steeper upward trend in productivity than small firms. The real value added per worker roughly triples for large firms, while it rises far more modestly for small firms. This pattern suggests that large firms seem to be the primary drivers of productivity gains, while smaller firms are unable to catch up, which is consistent with evidence in the related literature.

[4]Our value-added productivity measure is constructed from firm-level Compustat data. In Section B we validate it against official aggregate statistics.

Fact 2: Intangible capital rises across the U.S. economy, and large firms accumulate it steadily and persistently, converging to and eventually overtaking smaller firms.

Figure 2 reports the averages of the intangible capital ratio for small and large firms.

Figure 2: Intangible Capital Ratio by Firm Size



Note: This figure shows the annual average of the intangible capital ratio for small and large firms, restricting attention to firms with material physical capital. The intangible capital ratio is intangible capital divided by the sum of intangible and tangible capital. Small (large) firms are those in the bottom (top) decile of market value within industry and year, and industry-year cells are weighted by their number of firms.

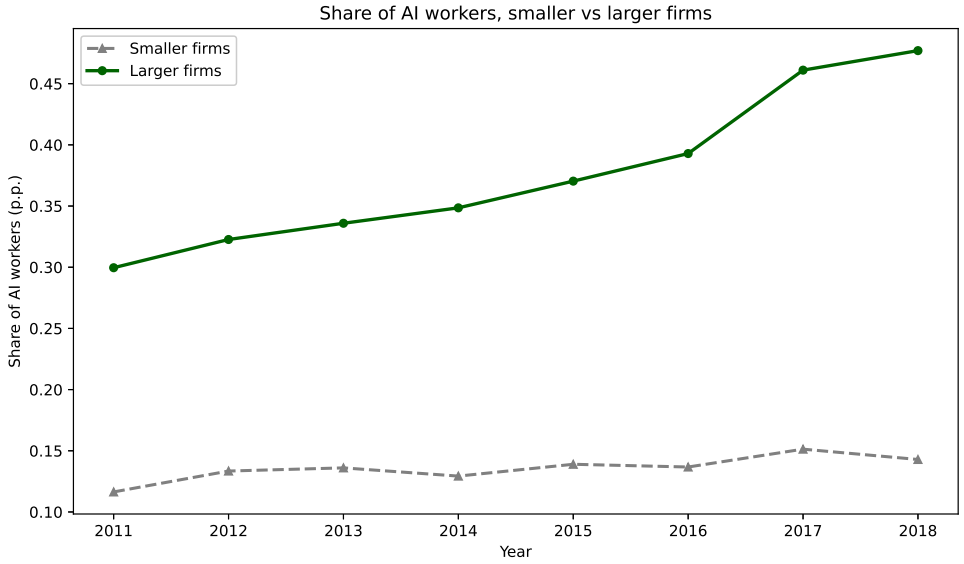
Two observations stand out. First, both groups accumulate intangible capital over the sample period, with the average ratio rising by close to 20 percentage points. This is consistent with the literature which documents that the composition of corporate capital in the U.S. economy has become increasingly weighted toward intangible capital. Second, the pattern is heterogeneous by size such that intangible capital ratio of small firms is on average higher early in the sample, which is partly a composition feature as small firms typically hold little physical capital that mechanically inflates the ratio. We see that large firms accumulate intangible capital steadily and persistently, and they converge to the level of small firms around the year 2000 and, they move above the small firms in the mid-2010s. Therefore, the fact provides suggestive evidence that intangible capital intensity has risen faster among large firms over the past two decades, and it motivates the firm size heterogeneity we examine in the causal analysis.

Fact 3: AI-skilled workers disproportionately rise in larger firms.

Figure 3 shows that larger firms have accumulated more AI-skilled workers than smaller firms. Larger firms hold a higher AI-worker share throughout the sample, and the gap widens markedly over time. However, the share by small firms remains almost flat throughout.

This size gradient is consistent with evidence that AI adoption rises strongly with firm size (McElheran et al., 2024) and that AI-driven growth concentrates among larger firms (Babina et al., 2024). Such economies of scale in technology adoption are also a recurring finding in the literature (Begenau et al., 2018; Autor et al., 2020; Tambe et al., 2020; Sui, 2021; Aghion et al., 2023; De Ridder, 2024; Kwon et al., 2024; Lashkari et al., 2024; Adams et al., 2025; Firooz et al., 2025, 2026; Hubmer and Restrepo, 2026). Therefore, this fact suggests that the accumulation of AI-related human capital would be subject to economies of scale, with larger firms being better positioned to attract, retain, and deploy specialized skills.

Figure 3: Share of AI Workers by Firm Size



Note: This figure plots the average share of AI workers (in percentage points) for larger and smaller firms. Small (large) firms are those below (above) the median of market value within industry and year, and industry-year cells are weighted by their number of firms.

4 Empirical Analysis

In this section, we exploit an exogenous shock to the effectiveness of AI workers to investigate the causal effect of AI effectiveness on firm productivity. We first present the empirical design, then the baseline results for value-added labor productivity and total factor productivity, and finally a series of exercises that address the main identification and specification concerns.

4.1 Empirical Design

We now turn to establishing a causal claim for the mechanism we propose. Specifically, first, we investigate whether greater effectiveness of AI workers leads to improved firm productivity. To do so, following Rock (2019), we exploit the 2015 release of Google TensorFlow, which is a widely adopted open-source machine learning platform, as an exogenous shock. TensorFlow is an open-source, general-purpose machine learning framework that substantially reduced the technical complexity and cost of deploying advanced AI models.

More importantly, its release did not target specific firms, industries, or regions, but instead represented a broad technological improvement that enhanced the productivity of existing AI expertise. Section C discusses the shock in more detail, including evidence that its timing was largely unanticipated, its rapid adoption relative to competing libraries, and the growing body of studies that treats the release as a turning point in the AI diffusion.

Empirical Strategy To capture firm-level exposure to the TensorFlow shock, we use the firm's share of AI workers in 2015 (the year of the release) as a proxy for firm-level "AI readiness", at the moment the shock arrived. This exposure variable proxies for AI-related human capital (e.g. the firm's capability to deploy artificial intelligence). Hence, what we want to identify is the differential return to AI readiness. Two considerations motivate measuring exposure in the shock year. First, conceptually, the treatment we study is an improvement in the effectiveness of the AI-skilled workforce that firms had in place when TensorFlow became available.[5] Second, measuring exposure in the shock year minimizes measurement error because the 2015 share captures the AI-skilled workforce a firm actually had in place when TensorFlow arrived, whereas a more distant pre-shock share proxies this readiness stock with additional noise and would therefore attenuate our estimates toward zero. Having said that, we also confirm this is not driving our results as the 2014 share delivers similar conclusions.[6]

[5] TensorFlow was released in November 2015, so the 2015 AI share overwhelmingly reflects hiring decisions taken before the release since workforce shares are typically slow-moving stock variables that would not respond sharply and meaningfully within the last two months of 2015. Moreover, as discussed in Section C, the release was not publicly anticipated since Google's decision to open-source its internal framework came as a surprise to the industry, so 2015 hiring would not have been undertaken in anticipation of the shock.

[6] Section E.9 reports these exercises using the pre-shock 2014 AI share and the mean, median, minimum, and maximum of the 2010-2014 shares, and an approach that projects the 2015 exposure with the 2014 share.

We estimate the following baseline difference-in-differences (DiD) specification:

$$\begin{aligned} Y_{it} =& \beta_1 \text{Post}_t + \beta_2 \text{Share of AI Workers}_{i,2015} + \beta_3 \left(\text{Post}_t \times \text{Share of AI Workers}_{i,2015}\right) \\ &+ X'_{it-1}\gamma + \lambda_t + \mu_i + \varepsilon_{it} \end{aligned} \tag{1}$$

where Y_{it} denotes the outcome of interest for firm i in year t, measured by the logarithm of productivity measures. The indicator variable Post_t equals one for the years following the first full calendar year after the release of TensorFlow (i.e., $t \geq 2016$), and zero otherwise. Because TensorFlow was released in November 2015, fiscal year 2015 is partially exposed to the shock. We therefore treat 2015 as a transition year and exclude it from the pooled DiD estimation sample. $\text{Share of AI Workers}_{i,2015}$ represents the firm's share of AI-related workers in 2015 and proxies for the firm's AI readiness as its capacity to benefit from TensorFlow's release.

The baseline specification includes firm and year fixed effects. Year fixed effects (λ_t) control for macroeconomic shocks and firm fixed effects (μ_i) control for time-invariant unobservable heterogeneity at the firm level. Standard errors are clustered at the firm level.[7] Industry and state affiliations are time-invariant for all firms in our sample and would be absorbed by the firm fixed effects.[8]

We write the two constituent main effects (β_1 and β_2) explicitly in equation (1) for the sake of providing the full estimated model, but neither is separately identified once the fixed effects are included. Post_t is common to all firms in a given year and is therefore absorbed by the year fixed effect λ_t. $\text{Share of AI Workers}_{i,2015}$ is fixed at its 2015 value and is firm-constant, and is absorbed by the firm fixed effect μ_i. Hence, the coefficients β_1 and β_2 drop out of the estimated model, and identification rests on the interaction term $\text{Post}_t \times \text{Share of AI Workers}_{i,2015}$. β_3 is the coefficient of interest, which measures the average effect of the exogenous improvement in AI effectiveness after the shock. Therefore, β_3 is the linear slope relating post-shock changes in the outcome to a firm's initial AI exposure. Interpreting this slope as an average treatment effect on the treated would additionally require a defined no-exposure counterfactual and assumptions under which heterogeneous dose responses aggregate linearly (Callaway et al., 2024); we assess the plausibility of the underlying dose-response pattern in Section E.1 by

[7] Because the TensorFlow release is a single aggregate event, Section E.7 verifies that inference is robust to two-way clustering by firm and year.

[8] Where time-varying sectoral or regional shocks are a concern, the appropriate specification interacts these dimensions with time, and we show in Section E.2 that our results are robust to industry- and state-specific linear trends as well as to industry-by-year and state-by-year fixed effects in Section E.3.

comparing the pooled estimate to a discretized specification.

The vector X_{it-1} includes one-year lagged control variables such as the logarithms of total assets, R&D intensity (R&D expenditures over sales), capital expenditures, the cash-to-asset ratio, firm age, markup, the dividend-to-asset ratio, and lagged sales growth. We include these one-year lagged controls to absorb standard determinants of firm growth and investment that may correlate with both AI readiness and subsequent outcomes. The logarithm of total assets controls for firm operational scale. The logarithm of R&D intensity captures innovation and technological intensity. The logarithm of capital expenditures accounts for investment opportunities and capital intensity. The cash-to-asset ratio proxies for financial slack and the ability to fund new technology adoption. Firm age captures life-cycle differences in growth and technology adoption. Markup controls for market power and pricing ability. The dividend-to-asset ratio reflects firm maturity and payout policy. Finally, lagged sales growth accounts for pre-existing growth momentum. All controls are lagged one year to mitigate simultaneity with the outcome. Following the insight by De Loecker et al. (2020), the markup control is an accounting approximation of the following form: $0.85 \times \frac{\text{Sales}}{\text{Cost of Goods Sold}}$, where the constant proxies for the output elasticity of the variable input.[9] By conditioning on markups we account for the possibility that market power constitutes an alternative channel, while acknowledging that a control of this form cannot fully rule out market power effects.

The validity of our causal interpretation relies on several identification assumptions. First, the key assumption in the DiD design is that, absent the TensorFlow shock, firms with varying shares of AI workers would have followed parallel trends in productivity. We provide support for this assumption through various dynamic DiD analyses. Second, we assume that the 2015 release of Google TensorFlow represents a plausibly exogenous improvement in AI effectiveness that differentially benefits firms based on their AI workforce, but is otherwise unrelated to concurrent shocks to productivity.[10] Third, the firm and year fixed effects, the lagged controls, and the battery of robustness checks in Section E address unobserved heterogeneity and time-varying confounders.

[9] In Section 4.3.5 we estimate a structural markup measure following De Loecker et al. (2020). We substitute the structurally estimated markup, which leaves the baseline results almost unchanged.

[10] This assumption is supported by the timing nature of TensorFlow's release as elaborated in Section C, which was not targeted at specific sectors or firm types, and by the TensorFlow-specific workforce evidence in Section 7.

4.2 Results

4.2.1 Value Added Labor Productivity

We begin with the effect of the shock on our baseline productivity measure, value-added labor productivity.

Table 1: DiD Estimates with Sequential Controls: Value-Added Labor Productivity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post×AI Share$_{2015}$	6.30**	6.54**	6.39**	7.33**	6.41**	6.32**	6.31**	6.59**	6.59**
	(3.09)	(3.00)	(3.04)	(3.14)	(3.19)	(3.18)	(3.18)	(3.15)	(3.15)
L.Sales Growth		0.18***	0.18***	0.16***	0.16***	0.16***	0.16***	0.16***	0.16***
		(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
L.Firm Age			0.07	0.05	0.02	0.02	0.03	0.03	0.03
			(0.10)	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)
L.ln(R&D/Sales)				-0.13**	-0.12**	-0.12**	-0.12**	-0.11**	-0.11**
				(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
L.ln(Assets)					0.05*	0.08**	0.08**	0.09***	0.10***
					(0.03)	(0.04)	(0.04)	(0.04)	(0.04)
L.ln(Capex)						-0.04	-0.04	-0.04	-0.04
						(0.03)	(0.03)	(0.03)	(0.03)
L.Dividends/Assets							0.40*	0.39*	0.40*
							(0.21)	(0.21)	(0.21)
L.Cash/Assets								0.28**	0.28**
								(0.13)	(0.13)
L.Markup									0.00
									(0.00)
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	8,603	8,603	8,603	8,603	8,603	8,603	8,603	8,603	8,603

Note: This table reports pooled DiD estimates of value-added labor productivity with sequential controls. $\text{Post}_t = \mathbf{1}[\text{year} \geq 2016]$. Column (j) adds the $(j-1)$-th lagged control. Standard errors clustered by firm in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 1 reports the pooled DiD estimates of equation (1). To make the role of the control set fully transparent, columns (1)-(9) add the eight lagged controls sequentially, from no controls to the full set. The main coefficient is positive and stable across the control build-up, which moves between roughly 6.3 and 7.3 log points per percentage point increase in the 2015 AI worker share. In the full specification of column (9) the estimate is 6.59, i.e. a one percentage point higher AI worker share is associated with

about 6.8[11] percent higher value-added labor productivity.[12]

The estimated effect is economically meaningful. Our results reveal that a one standard deviation increase in the 2015 AI worker share (0.68 percentage points) is associated with 4.5 log points, or about 4.6 percent, higher value-added labor productivity. When we evaluate at the median firm, where value-added per employee is roughly $38,900, this corresponds to an additional $1,780 of value-added per employee per year, or about $10.5 million for a firm with the median workforce of 5,900 employees. Moving from the 25th to the 75th percentile of 2015 AI worker share (0.22 to 0.70 percentage points) implies a value-added productivity gain of about 3.2 percent.[13]

To explore the dynamics and assess the validity of the parallel trends assumption, we estimate a dynamic DiD specification of the following form:

$$Y_{it} = \sum_{\tau \neq 2015} \beta_\tau \left(\mathbb{1}\{t = \tau\} \times \text{Share of AI Workers}_{i,2015} \right) + X'_{it-1}\gamma + \lambda_t + \mu_i + \varepsilon_{it} \quad (2)$$

where the exposure is interacted with a full set of year indicators. We take 2015 as the partial-exposure year of the TensorFlow release as the omitted reference period. This specification traces the differential trajectory of outcomes across firms with varying AI readiness, before and after the shock.

Figure 4 plots the estimates for value-added labor productivity. Three features stand out. First, the pre-treatment coefficients for 2010-2014 are individually small and statistically insignificant. Second, the response is not immediate such that the 2016 coefficient is positive but modest, and the effect becomes large from 2017 onward. This gradual build-up is consistent with what the productivity J-curve insight of Brynjolfsson et al. (2021) predicts for a general purpose technology, which argues that firms first accumulate the technology, and measured productivity responds only once that accumulation begins to pay off.[14] Third, the effect is persistent through remaining at around 11 log points through the end of the sample.

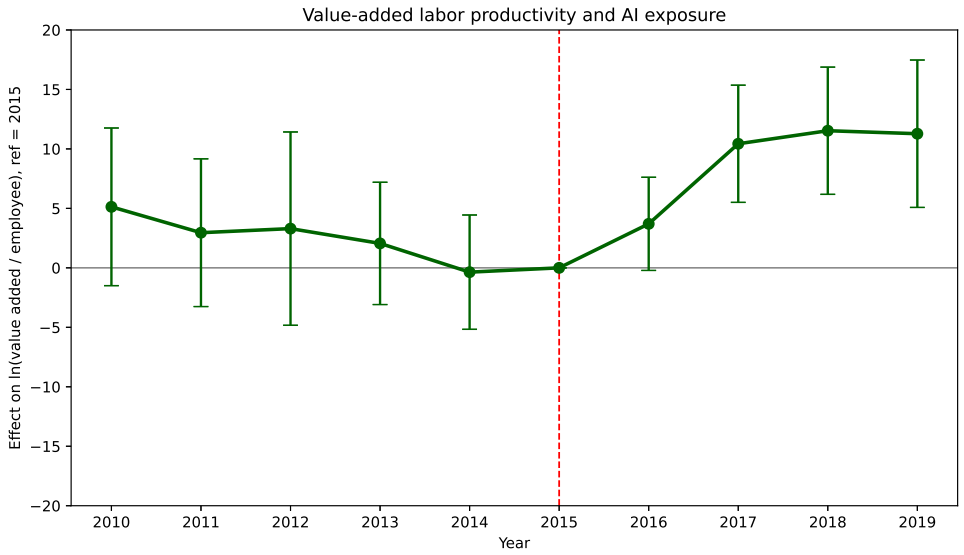
[11] Since the exposure variable is measured as a fraction, a one percentage point increase corresponds to $0.01 \times 6.59 = 0.0659$, which implies a percent change of $e^{0.0659} - 1 = 0.068$, or about 6.8 percent.

[12] Throughout the paper, estimates reported in log points refer to a one percentage point change in the exposure variable.

[13] Section E.1 provides evidence on the dose-response relation between exposure and productivity by replacing the continuous share with exposure dummies.

[14] Section 5 provides evidence that exposed firms build intangible capital after the shock, and that the productivity response is concentrated in firms that entered the shock intangible capital rich.

Figure 4: Dynamic DiD Estimates: Value-Added Labor Productivity and AI Workers



Note: This figure plots the estimated coefficients from a dynamic DiD regression of the impact of the TensorFlow shock on the logarithm of value-added labor productivity. The coefficients represent the interaction between the firm's 2015 share of AI workers and year indicators, with 2015 (the year of the November TensorFlow release) as the omitted reference year, marked by the dashed line. Control variables include the one-year lag of the firm-level logarithm of total assets, logarithm of R&D intensity, logarithm of capital expenditures, the cash-to-asset ratio, firm age, markup, the dividend-to-asset ratio, and lagged sales growth. The specification includes firm and year fixed effects. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

4.2.2 Total Factor Productivity

Value-added labor productivity nets out intermediate inputs, but it does not control for capital deepening. Therefore, exposed firms could raise value added per worker simply by accumulating more capital per worker. Therefore, to isolate efficiency, and to also ensure that the conclusions do not only hinge on the labor productivity, we construct firm-level total factor productivity (TFP) and re-estimate the empirical design accordingly. Because our productivity concept is value added, we estimate a value-added production function as follows:

$$\ln \mathrm{VA}_{it} = \beta_L \ln L_{it} + \beta_K \ln K_{it} + \omega_{it} + \epsilon_{it}$$

where VA_{it} is real (industry-PPI-deflated) operating income before depreciation, L_{it} is employment, K_{it} is net property, plant, and equipment, and ω_{it} is unobserved productivity. Because input choices respond to ω_{it}, OLS estimates of the elasticities are biased. We therefore estimate the production function with four standard alternative control-function estimators that address this simultaneity through different identifying assumptions and proxy variables. We follow i) Olley and Pakes (1996) (OP), which proxies unobserved productivity with investment, ii) Levinsohn and Petrin (2003) (LP), which uses intermediate inputs, iii) Wooldridge (2009) (WRDG), which implements

the control-function moments in a single GMM step, and iv) Ackerberg et al. (2015) (ACF), which addresses the functional-dependence critique of the earlier two-stage procedures. Table G.1 reports the estimated labor and capital elasticities for each method. Firm-year TFP is then computed as the production-function residual, $\widehat{\text{TFP}}_{it} = \ln \text{VA}_{it} - \hat{\beta}_L \ln L_{it} - \hat{\beta}_K \ln K_{it}$, winsorized at the 1st and 99th percentiles within year.

Table 2: Pooled DiD Estimates: Total Factor Productivity

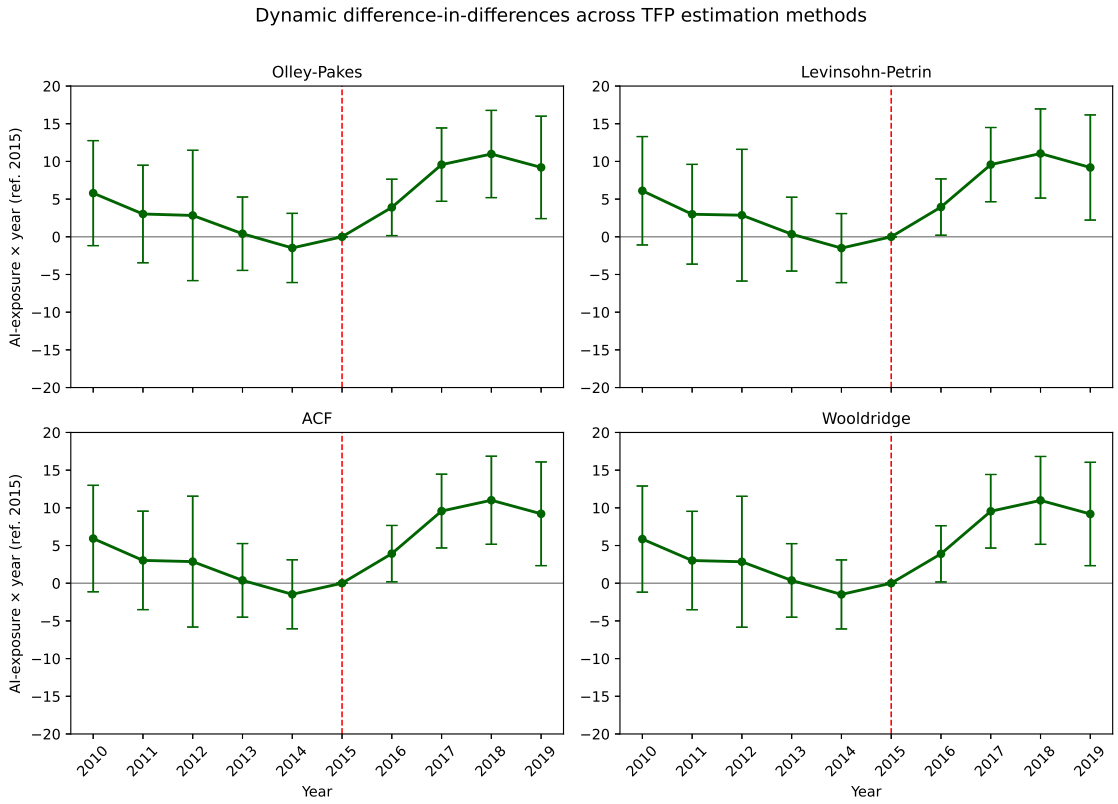
	(1) OP	(2) LP	(3) ACF	(4) WRDG
Post×AI Share$_{2015}$	6.626*	6.608*	6.609*	6.614*
	(3.530)	(3.572)	(3.546)	(3.545)
Controls	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	8,603	8,603	8,603	8,603

Note: This table reports pooled DiD estimates of the effect of pre-treatment AI exposure on total factor productivity. The outcome is ln(TFP) from a value-added production function estimated by Olley and Pakes (1996) (OP), Levinsohn and Petrin (2003) (LP), Ackerberg et al. (2015) (ACF), Wooldridge (2009) (WRDG); $\text{Post}_t = \mathbf{1}[\text{year} \geq 2016]$. Standard errors clustered by firm in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 2 reports the pooled DiD estimates of equation (1) with TFP as the outcome, one column per estimator (OP, LP, ACF, and WRDG). Across all four methods, the coefficient of interest implies an increase in TFP by around 6.6 log points, which is almost identical to the value-added labor productivity estimate. The close agreement between the two outcomes, and across the four estimation methods, suggests that the productivity response is not driven by capital deepening or by the choice of production function estimator.

Figure 5 presents the corresponding dynamic estimates for each TFP estimator. The pattern mirrors the value-added event study such that pre-treatment coefficients are flat and jointly insignificant for every method, and the effect emerges from 2017 onward through reaching roughly 9 to 11 log points by 2018-2019.

Figure 5: Dynamic DiD Estimates across TFP Estimation Methods



Note: This figure plots the estimated coefficients from dynamic DiD regressions of the impact of the TensorFlow shock on the logarithm of total factor productivity, estimated with four alternative production-function methods (OP: Olley and Pakes (1996); LP: Levinsohn and Petrin (2003); ACF: Ackerberg et al. (2015); WRDG: Wooldridge (2009)). The coefficients represent the interaction between the firm's 2015 share of AI workers and year indicators, with 2015 as the reference year. The specification includes firm and year fixed effects. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

What the productivity measures do and do not net out. Both of our outcomes are revenue-based, which basically recovers revenue productivity (TFPR) but not physical productivity (TFPQ) (Foster et al., 2008). Therefore, any deviation of a firm's own output price from its industry index remains in the outcome, and three potential channels such as market power, demand shifts, and changes in product mix or quality can generate such a deviation. We address market power directly in Section 4.3.5, and the treatment effect is almost unchanged by controlling structurally estimated De Loecker et al. (2020) markup. We cannot rule out the other two due to the data limitation. If AI-exposed firms faced stronger demand or upgraded their product mix, part of the estimated response reflects what those firms sell rather than how efficiently they produce it. Separating the two requires firm- or product-level prices and quantities, which Compustat does not contain, and which is the measurement issue that Klette and Griliches (1996) and De Loecker (2011) identify. As De Loecker and Goldberg (2014) emphasize, this is a pervasive constraint rather than a feature of our data alone. Due to the absence of firm-level prices, revenue-based measures

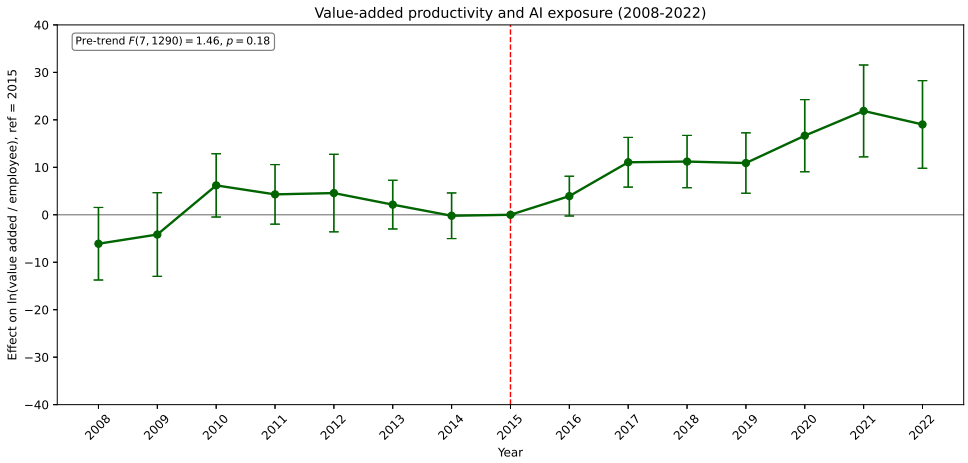
would conflate technical efficiency with pricing, demand, and product quality, and the literature has proceeded by being explicit about that content, which we also adopt the same posture here. Two considerations bound the concern. First, the resulting bias need not be upward because if improved AI effectiveness lowered marginal cost and was passed through to firm prices, an industry-level deflator attributes too little real output to exposed firms, and our estimates would understate the physical productivity gain. Second, our central argument concerns the distribution of the returns rather than their precise composition. We argue that whether greater AI readiness raised technical efficiency, demand, or product quality, the finding that the returns accrue to large and intangible capital rich firms, and therefore widen productivity dispersion, is not dramatically affected.

4.3 Refinements on Identification and Specification

4.3.1 Including More Leads in the Event Study

Our baseline estimation window is 2010-2019 to isolate a plausibly clean pre- and post-shock comparison. For the purpose of the rigorous pre-trend test, however, longer lead periods are informative.

Figure 6: Dynamic DiD Estimates on the Extended 2008-2022 Window - Value Added

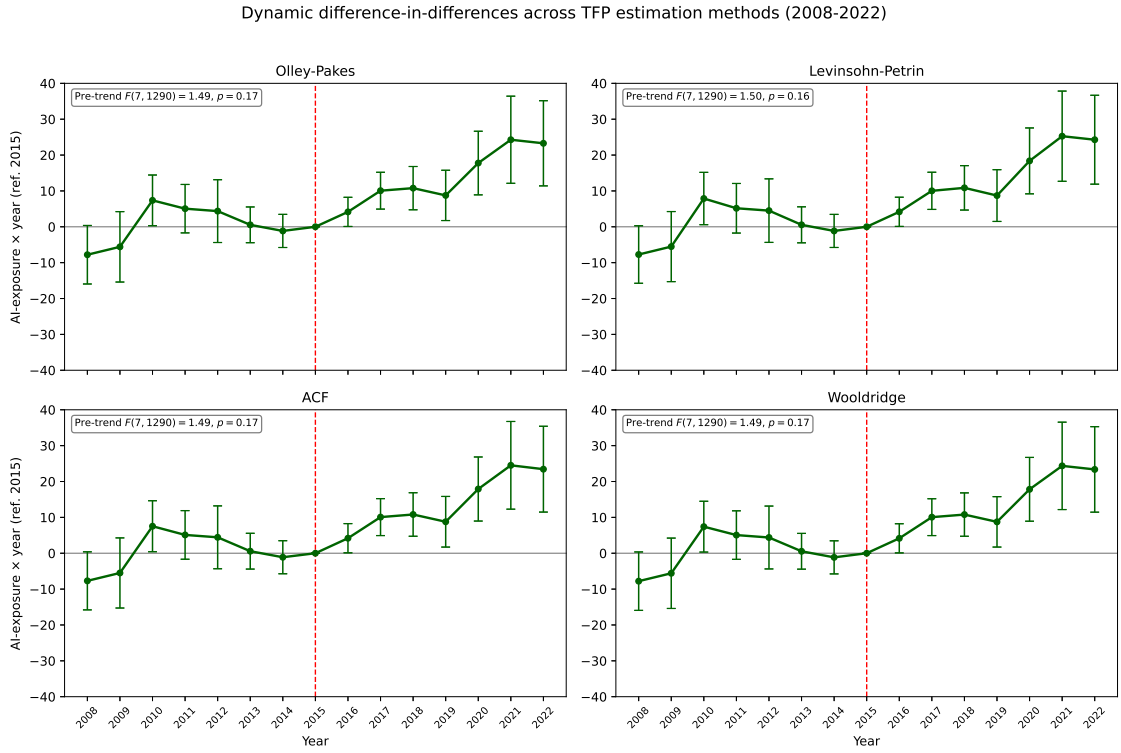


Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity on the extended 2008-2022 panel. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

As a robustness check, we extend the panel from 2008 through 2022. Table G.2 reports the pooled DiD results for both value-added labor productivity and TFP, and we find that the effect is still positive and statistically significant. Figure 6 shows that

even with the crisis years added to the pre-period, the extended set of leads remains jointly insignificant, and the baseline effect for value-added labor productivity not only persists but continues to grow after the baseline window closes through reaching roughly 17 to 22 log points in 2020-2022. Moreover, the gradual strengthening of the estimated effect over seven post-shock years is consistent with the slow diffusion of a general-purpose technology (Brynjolfsson et al., 2021). We do not, however, interpret the 2020-2022 coefficients as a direct evidence of continued TensorFlow diffusion since the pandemic responses would still correlate with firms' 2015 AI exposure. We read the extended window as demonstrating robustness of the association and as adding pre-treatment leads.

Figure 7: Dynamic DiD Estimates on the Extended 2008-2022 Window - TFP



Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for TFP on the extended 2008-2022 panel. The coefficients represent the interaction between the firm's 2015 share of AI workers and year indicators, with 2015 as the reference year. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

Figure 7 repeats this extended-window analysis for TFP across all four estimators. The overall finding is almost identical to the one with value-added. The extended set of pre-treatment leads remains jointly insignificant for every method, so the parallel-trends assumption survives the more demanding test even after the 2008-2009 crisis years are included into the pre-period. We find that the response emerges from 2016 onward and then continues to strengthen beyond the baseline window through reaching around 18

log points in 2020 and 24 to 25 in 2021-2022.

4.3.2 Fiscal-Year Timing and Partial Exposure

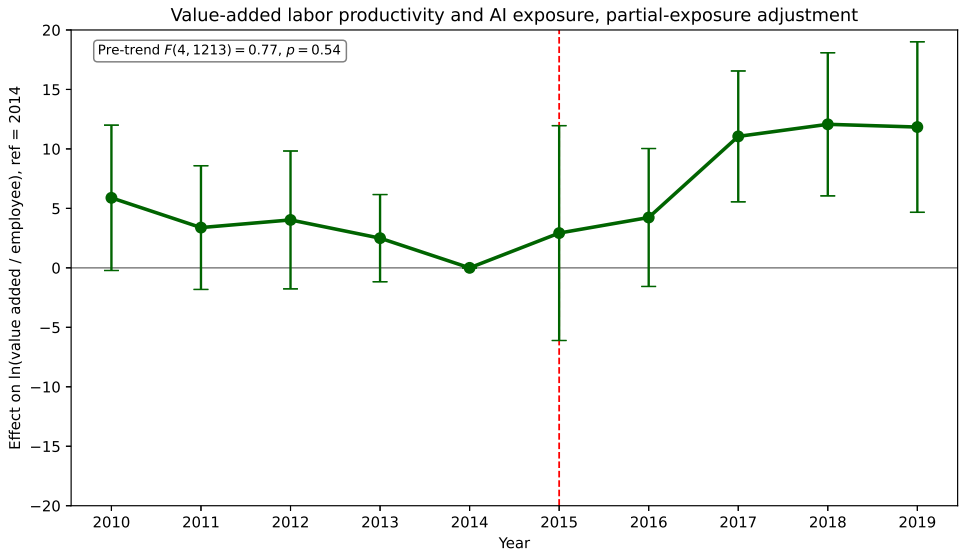
A natural concern with any calendar-time shock is that Compustat fiscal years can be firm-specific, and hence the same nominal "fiscal 2015" can correspond to very different windows of calendar exposure to the November 2015 TensorFlow release. For instance, a firm with a December fiscal-year end reports fiscal 2015 variables that close just weeks after the release, whereas a firm with a June year end reports fiscal 2015 variables that were entirely settled months before it. Therefore, treating all fiscal 2015 observations alike would misattribute the timing of the shock.

To address this concern, our baseline specification takes the conservative path. Since the release lands close to the end of the calendar year, fiscal 2015 variables are neither cleanly pre- nor cleanly post-treatment, so we exclude 2015 from the pooled DiD sample and consider it as the omitted reference period in the dynamic DiD specification. As a sharper robustness check, we evaluate the timing information at the level of the firm calendar-year observation. We classify each fiscal 2015 observation by each firm's fiscal-year-end date relative to the November 2015 release. More specifically, fiscal years ending June-October 2015 closed before the release are pre-treatment, and we include them as pre-period observations for 2015, while fiscal years ending November 2015-May 2016[15] straddle the release, are partially exposed, and we drop them. The treatment window is unchanged ($\text{Post}_t = \mathbf{1}[\text{fiscal year} \geq 2016]$). Table G.3 reports the pooled DiD sequential-controls estimates under this adjustment. We find that the main coefficient of interest is stable between 6.31 and 7.32 log points across all nine columns, and the full specification yields 6.64 that is almost identical to the baseline estimate of 6.59.

Figure 8 documents the dynamic DiD estimates. Because clean 2015 observations are now part of the sample and hence the dynamic DiD specification can assign 2015 its own coefficient, we move the reference period to 2014. The estimates show a flat pre-period and a 2015 coefficient of 2.9 that is statistically indistinguishable from zero. Therefore, there is no effect in the release year itself, which provides a direct evidence against an implausibly instantaneous response. Similar to the baseline result, we find that the effect emerges from 2017 onward.

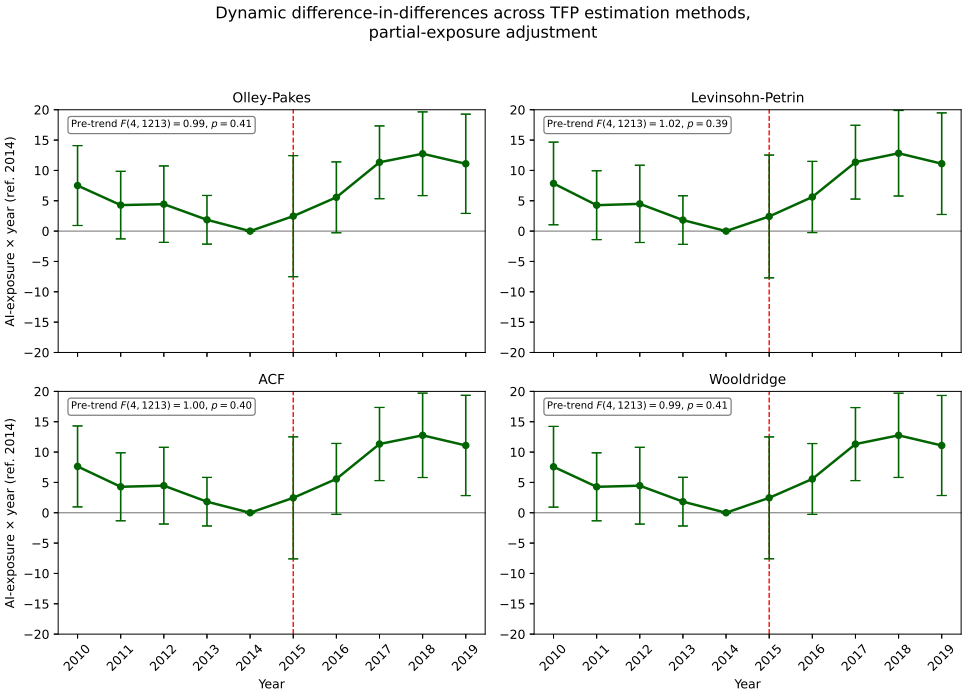
[15] Based on Compustat's fiscal-year convention, fiscal years ending between June 2016 and May 2017 are assigned to fiscal year 2016. This is the reason why May 2016 is the end point for this consideration.

Figure 8: Dynamic DiD Estimates for Value-Added Labor Productivity: Partial-Exposure Adjustment



Note: This figure plots the estimated coefficients from the dynamic DiD regression for value-added labor productivity under the observation-level partial-exposure adjustment: fiscal-2015 observations ending June-October 2015 (closed before the November 2015 TensorFlow release) are kept as pre-treatment observations, while fiscal-2015 observations ending November 2015-May 2016 are dropped as partially exposed. The coefficients represent the interaction between the firm's 2015 share of AI workers and year indicators, with 2014 as the reference year; the dashed line marks the 2015 release. Standard errors clustered at the firm level. Error bars show the 90% confidence interval.

Figure 9: Dynamic DiD Estimates for TFP: Partial-Exposure Adjustment



Note: This figure plots the estimated coefficients from the dynamic DiD regression for TFP under the observation-level partial-exposure adjustment: fiscal-2015 observations ending June-October 2015 (closed before the November 2015 TensorFlow release) are kept as pre-treatment observations, while fiscal-2015 observations ending November 2015–May 2016 are dropped as partially exposed. The coefficients represent the interaction between the firm's 2015 share of AI workers and year indicators, with 2014 (the last fully pre-shock year) as the reference year; the dashed line marks the 2015 release. Standard errors clustered at the firm level. Error bars show the 90% confidence interval.

We also extend the partial-exposure adjustment to the TFP analysis, and Table G.4

reports the results. The main coefficient is around 6.71 log points across the four methods. The estimates are almost identical to the value-added result under the same adjustment and to the baseline TFP estimates on the unadjusted sample. Therefore, we conclude that the timing convention does not dramatically affect the pooled DiD effect.

The dynamic DiD estimates in Figure 9 deliver a similar insight as their value-added counterpart. A coefficient of 2015 is around 2.5 log points which is statistically indistinguishable from zero. Hence, we find no significant efficiency response in the release year itself. The pre-treatment coefficients for 2010-2013 are jointly insignificant for all four methods, and the effect emerges from 2017 onward through reaching around 11 log points in 2017, 12 in 2018, and 11 in 2019.

4.3.3 Quarterly Evidence

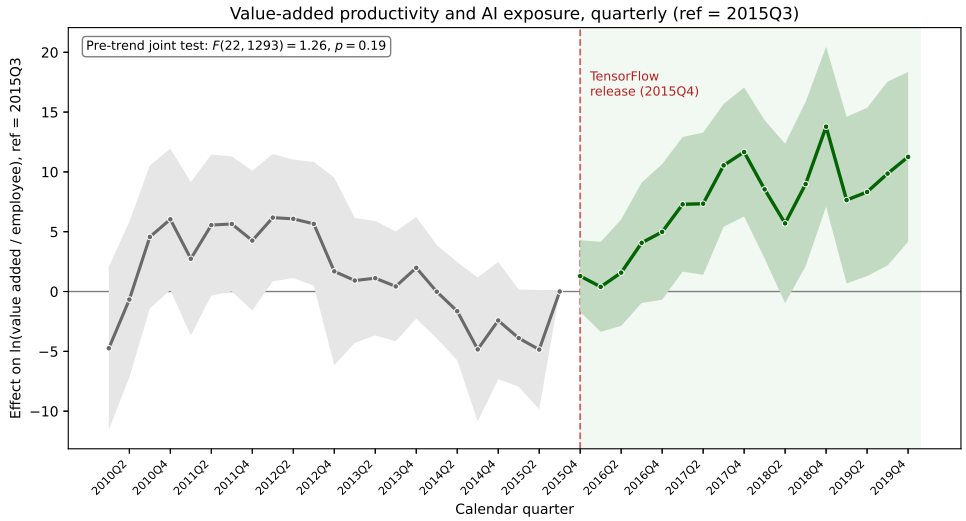
The annual design aggregates outcomes to the calendar year. To pin down the timing more precisely and obtain a more granular view of the dynamics, we re-estimate the empirical design at the quarterly frequency using Compustat quarterly data.

First of all, quarterly operating income is strongly seasonal and frequently negative in individual quarters. Therefore, we implement a trailing four-quarter (TTM) sum of operating income before depreciation, which removes seasonality while preserving quarterly variation in timing.[16] We compute the quarterly outcome as the logarithm of real TTM operating income per employee, and we winsorize at the 1st and 99th percentiles within each calendar quarter. We include firm and year-quarter fixed effects, and cluster the standard errors by firm.

Table G.5 reports the pooled DiD estimates. We find that the main coefficient of the full specification is around 6.08 log points, which is in line with the annual baseline estimate. Figure 10 plots the quarterly dynamic DiD estimates. With twenty-two pre-treatment quarters, we find that the joint pre-trend test is statistically insignificant, and the overall path shows the effect emerging gradually during 2016, and strengthens through 2017, and reaches about 11 log points by the end of 2019.

[16] The TTM sum at quarter q aggregates `oibdpq` over quarters $q-3$ through q and is defined only when all four quarters are observed.

Figure 10: Quarterly Dynamic DiD Estimates: Value-Added Labor Productivity



Note: This figure plots the estimated coefficients from a quarterly dynamic DiD regression of the impact of the TensorFlow shock on the logarithm of value-added labor productivity (trailing four-quarter operating income before depreciation). The coefficients represent the interaction between the firm's 2015 share of AI workers and year-quarter indicators, with 2015Q3 as the reference period. The specification includes firm and year-quarter fixed effects, and the lagged controls. Standard errors clustered at the firm level. The shaded ribbon shows the 90% confidence interval.

We also implement the same quarterly design for the TFP analysis.[17] Table G.6 shows that the pooled Post×AI estimates are about 6.5 log points for all four methods, and Figure H.2 shows the corresponding dynamic DiD estimates with flat and statistically insignificant pre-trends and the same gradual post-2016 build-up as in the value added.

Taking together, the quarterly evidence reinforces both the identification and the interpretation such that the effect appears after the release quarter with a delayed and persistent efficiency response rather than an instantaneous adjustment.

4.3.4 Dynamic Panel (Nickell) Bias

Our baseline outcome is value-added labor productivity, measured as value added per employee, and our baseline specification can be considered as a static two-way fixed-effects model. Although the specification includes lagged control variables (X_{it-1}), the outcome does not depend on its own past. Therefore, the classic bias of Nickell (1981) does not arise mechanically.

The residual concern would be that firm productivity would be persistent such that a firm's value added per worker in a particular year could be strongly related to its value

[17] We apply the annual production-function elasticities of Section 4.2.2 to quarterly inputs (TTM real operating income, quarterly net property, plant, and equipment, and annual employment) to obtain quarterly TFP for each of the four estimators.

in the previous year, and hence omitting this dynamic term could in principle distort the estimated treatment effect. To confront the concern, we re-estimate a dynamic version of the baseline model by adding the lagged dependent variable Y_{it-1} so that we can model persistency explicitly. This is the case in which the Nickell (1981) bias bites in a way that once a lagged dependent variable meets the firm fixed effects, the estimated persistence parameter is biased toward zero in a short panel, and the distortion carries over to the coefficients of interest.

In the dynamic setting, we attempt to correct this bias using the split-panel jackknife of Dhaene and Jochmans (2015). The logic of the correction is that the Nickell (1981) bias shrinks as the panel lengthens at a known first-order rate. The estimator therefore splits the time dimension into two consecutive half-panels (2010-2014 and 2015-2019 in our context), and re-estimates the model on each half. Then it combines the full-panel estimate with the half-panel estimates so that the leading $O(1/T)$ bias term (the component declining at rate $1/T$, where T denotes the number of time periods) cancels. Because the treatment turns on mid-sample and is not identified in the pre-only first half, the jackknife is applied to the persistence parameter ρ (i.e. the locus of the Nickell (1981) bias), and identified in both halves while the impact coefficient is carried from the dynamic specification.[18]

Table 3 reports the results. Column (1) is the static baseline version. Column (2) reports the uncorrected dynamic fixed-effects model in which the impact effect, the response of productivity within the same period, is 3.18, and the estimated persistence is $\rho = 0.41$. Taken at face value, this dynamic-FE persistence would imply a long-run effect of $\hat{b}/(1-\hat{\rho}) = 5.4$ log points, however ρ here is exactly the quantity biased toward zero due to the Nickell (1981) bias, and hence it understates the true dynamics. Column (3) applies the split-panel jackknife due to Dhaene and Jochmans (2015) such that correcting the incidental-parameter bias lifts the persistence parameter from 0.41 to $\rho = 0.68$, and the implied long-run effect, $\hat{b}/(1-\hat{\rho})$, rises to approximately 10 log points. What we mainly learn from this analysis is that the static baseline estimate of 6.59 lies between the dynamic impact and the bias-corrected long-run effect, and this is what we expect when a persistent process is summarized by a static model. Moreover, we confirm that the treatment effect remains positive and significant throughout. As a result, a potential dynamic-panel bias does not drive our results, and if anything, accounting for persistence implies a larger cumulative effect than the static estimate.

[18] It is important to note that the split-panel jackknife corrects the persistence parameter.

Table 3: Dynamic-Panel (Nickell) Bias Robustness: Value-Added Labor Productivity

	(1) Static two-way FE	(2) Dynamic FE	(3) Half-panel jackknife
Post×AI Share$_{2015}$ (impact b)	6.594**	3.182*	3.182*
	(3.154)	(1.735)	(1.735)
Persistence (ρ)	—	0.407***	0.682
		(0.031)	
Implied long-run effect $b/(1-\rho)$	6.594	5.370	10.008
Firm FE	✓	✓	✓
Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	8,603	8,360	8,360

Note: This table reports the value-added DiD with and without a dynamic-panel bias correction. Column (1) is the static baseline (a strictly exogenous treatment in a static fixed-effects model carries no Nickell (1981) bias). Column (2) adds the lagged dependent variable; its persistence coefficient is biased toward zero by the $O(1/T)$ incidental-parameter bias. Column (3) applies the split-panel jackknife of Dhaene and Jochmans (2015), $\hat{\theta}_{\text{SPJ}} = 2\hat{\theta} - \frac{1}{2}(\hat{\theta}_{S_1} + \hat{\theta}_{S_2})$, with time halves S_1 = 2010-2014 and S_2 = 2015-2019, to correct the persistence parameter; the impact coefficient is carried from column (2) because the mid-sample treatment onset leaves it unidentified in S_1. The implied long-run effect is $\hat{b}/(1-\hat{\rho})$. * $p < .10$, ** $p < .05$, *** $p < .01$.

4.3.5 Markup Estimation

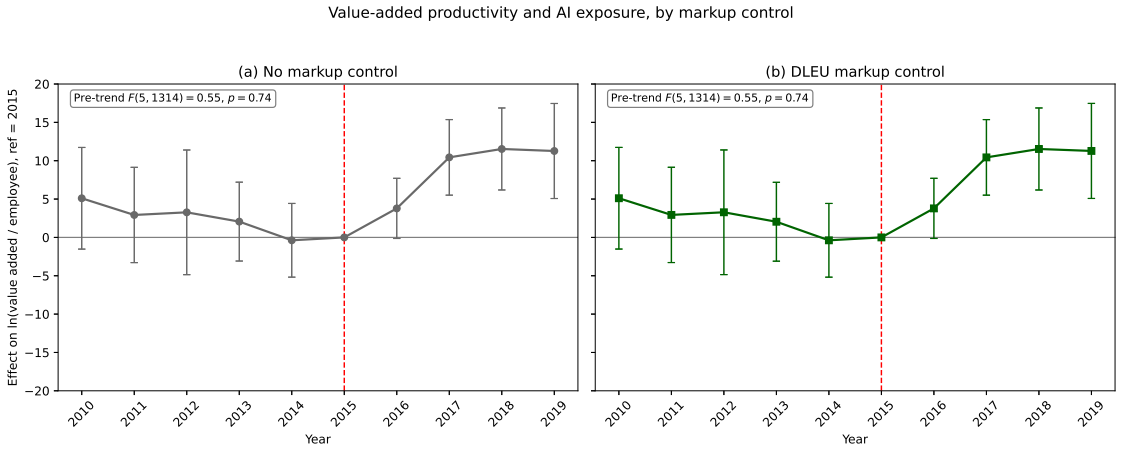
Our object of interest is the productivity gain from AI-skilled labor, but a revenue-based productivity measure can reflect more than efficiency alone. A firm that raises its prices, or that operates in a less competitive market would record higher value added per worker even without a real improvement in how efficiently it converts inputs into output. Improvements in AI effectiveness could plausibly move markups as well as efficiency, so without accounting for market power we could not tell whether the estimated response reflects firms producing more with the same inputs or simply charging more. We therefore condition on a firm-level markup in the baseline specification, which potentially absorbs the market power component of measured productivity and isolates the efficiency channel we wish to identify.

Our baseline design conditions on an accounting markup proxy of the form $0.85 \times$ (*sales per cost of goods sold*) following the insight of De Loecker et al. (2020), in which the constant 0.85 proxies the output elasticity of the variable input. Because this proxy imposes a common elasticity on every firm, a natural concern would be that it mismeasures market power, or that, if the true elasticity differs systematically across

industries, it leaves a residual market power channel in the outcome. To assess whether this proxy approach could be an issue, we replace the imposed constant with a structurally estimated, industry-specific output elasticity following the production function approach of De Loecker et al. (2020), and use it to construct a firm-level structural markup.[19]

Table G.7 incorporates this structural markup into the baseline DiD. We find that the baseline effect is almost similar. Figure 11 confirms that the dynamic DiD estimates are indistinguishable across markup controls, in which we find the same flat and statistically insignificant pre-period coefficients and the same post-2017 build-up whether the markup is omitted or measured structurally. Beside the refined results, we nonetheless interpret the markup control cautiously. It conditions on a revenue-based markup and cannot by itself fully rule out all other potential market power channels.

Figure 11: Dynamic DiD Estimates under Alternative Markup Controls



Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity under two alternative markup controls: panel (a) omits the markup control, and panel (b) uses the structurally estimated De Loecker et al. (2020) markup. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

4.3.6 Controlling for Other Technologies

A potential concern would be that the estimated response may not be specific to AI but rather broad wave of advanced technologies evolving over time. As an example, firms that employed AI-skilled workers in 2015 would be technology intensive, and continuously adopt a range of advanced technologies over the same time period. Hence, the post-2015 treatment effect could reflect a broad technology-adoption profile rather than the improved effectiveness of AI-skilled labor specifically.

[19]The construction of this structural markup and its related details are discussed in Section D.

To isolate AI-specific effects from broader technological progress, we control for the firm's exposure to other advanced technologies using the related shares of Babina et al. (2024) such as the firm's IT-, data-, data-analysis-, and robotics-worker shares. We incorporate these additional control variables in our baseline specification contemporaneously since we investigate whether concurrent digital adoption accounts for the response.[20]

Table 4: Controlling for Other Technology Exposure: Value-Added Labor Productivity

	(1)	(2)	(3)	(4)	(5)	(6)
Post×AI Share$_{2015}$	6.594**	7.582**	7.502**	7.509**	7.447**	7.342**
	(3.154)	(3.400)	(3.362)	(3.368)	(3.398)	(3.395)
Share of IT Workers	—	0.000	—	—	—	-0.001
		(0.001)				(0.001)
Share of Data Workers	—	—	0.005	—	—	0.007
			(0.004)			(0.006)
Share of Data-Analysis Workers	—	—	—	0.006	—	-0.002
				(0.007)		(0.010)
Share of Robotics Workers	—	—	—	—	0.023	0.023
					(0.021)	(0.021)
Firm FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
Observations	8,603	7,292	7,292	7,292	7,292	7,292

Note: This table reports the pooled value-added DiD controlling for the firm's exposure to other advanced technologies. Column (1) is the baseline; columns (2)-(5) add the firm's contemporaneous IT-, data-, data-analysis-, and robotics-worker shares from Babina et al. (2024) one at a time, and column (6) includes all four jointly. The technology-share data end in 2018, so columns (2)-(6) effectively cover 2010-2018. Standard errors, clustered by firm, in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 4 reports the results. Adding the technology controls separately leaves the baseline effect between 7.4 and 7.6 log points, and including them jointly yields 7.3 log points.[21] Notably, if the productivity response reflected a broader technology-adoption profile, these other technology exposures should predict productivity gains themselves. However, what we find is that each of the four technology-worker shares enters statistically insignificant both one at a time and jointly, while the baseline AI

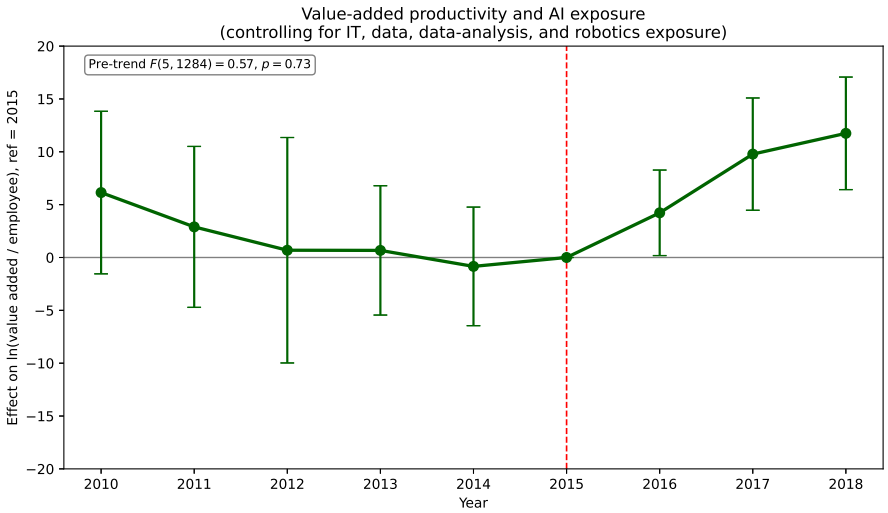
[20] Because the availability of these data on the additional variables ends in 2018, the resulting estimation sample covers 2010-2018.

[21] The slightly larger magnitude relative to the baseline estimate of 6.6 reflects the change in sample rather than the controls themselves as the technology shares are not available for 2019.

effect is almost unchanged.

The dynamic DiD estimates in Figure 12 with all four additional technology controls are likewise almost unaffected. The pre-treatment coefficients remain jointly insignificant and the effect becomes positive and significant in 2017 and 2018. As a result, the estimated response is therefore attributable to AI-related workforce rather than to a generic technology adoption profile.

Figure 12: Dynamic DiD Estimates Controlling for Other Technology Exposure



Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity, additionally controlling for the firm's contemporaneous IT-, data-, data-analysis-, and robotics-worker shares from Babina et al. (2024). Because these data end in 2018, the event study covers 2010-2018, with 2015 as the omitted reference year. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

5 Complementarity and Size

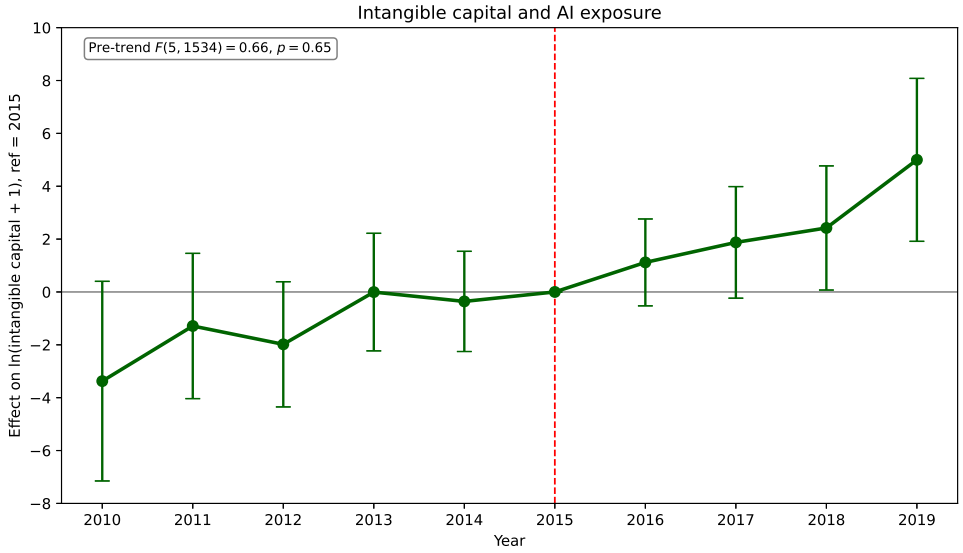
5.1 Intangible Capital Accumulation

Having established the productivity effects, we turn to the mechanism through which they arise in our context, which is through the accumulation of intangible capital. Our hypothesis is that if improvements in AI effectiveness raise the return to intangible capital, exposed firms in our sample analysis should respond by building more of them.

We re-estimate the baseline design with the firm's intangible capital as the outcome. Figure 13 presents the corresponding dynamic DiD estimates. We find that the pre-treatment coefficients are flat and jointly insignificant, and the stock builds steadily after the shock by reaching about 5 log points by 2019. The pattern is also consistent with the gradual accumulation of intangible capital rather than a one-off jump (Brynjolfsson

et al. (2021)). In light of the results in this section, we show that firms with a higher share of AI workers at the time of the shock experience greater intangible capital accumulation following the 2015 TensorFlow release.

Figure 13: Dynamic DiD Estimates: AI Workers and Intangible Capital



Note: This figure plots the estimated coefficients from a dynamic DiD regression of the impact of the TensorFlow shock on the logarithm of intangible capital. The coefficients represent the interaction between the firm's 2015 share of AI workers and year indicators, with 2015 as the reference year. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

5.2 Complementarity Effect on Productivity

The central prediction of our mechanism is that the productivity response to improved AI effectiveness would increase in the firm's intangible capital and, since intangible capital is scalable and non-rivalrous (Haskel and Westlake, 2017; Crouzet et al., 2022a), also in firm size.

To examine the joint role of the two moderators, Table 5 splits firms at the median of their 2015 intangible intensity and, separately, at the median of their 2015 size. The estimates for each of the resulting groups' effect are relative to the small and low intangible capital baseline group.[22] Column (1) reports each group's effect difference from this baseline, and column (2) the implied per-group effect. Based on the column (1), we find that the productivity effect is positive and statistically significant in firms that are both large and intangible capital rich.

[22] We fix both moderators at their 2015 value since intangible capital would be itself an endogenous response to the shock (as we established in the previous Section 5.1). Fixing the moderators at their 2015 level makes them predetermined with respect to the shock, so that the interaction measures how the treatment effect varies with a firm's initial intangible capital and size.

Moreover, it is also the group whose effect differs significantly from the baseline group. Importantly, based on the Column (1), the remaining groups show no statistically significant response. Hence, our results suggest that improved AI effectiveness raises productivity in firms with larger size and higher intangible capital. We read this pattern as a size-intangible gradient. Based on this insight, in Section 6.1, we further investigate the productivity amplification of the complementarity effect more formally using a quantile DiD analysis.

Table 5: DiD Estimates by Intangible Intensity and Firm Size: Value-Added Labor Productivity

	(1) Difference vs. baseline	(2) Implied effect
Post×AI Share$_{2015}$ (small × low intangible, baseline)	—	0.65
		(2.56)
× Small × high intangible	12.44	13.09
	(15.56)	(15.63)
× Large × low intangible	6.99	7.64*
	(4.37)	(4.22)
× Large × high intangible	8.22*	8.87**
	(4.33)	(4.11)
Firm FE	✓	✓
Year FE	✓	✓
Controls	✓	✓
Observations	8,603	8,603

Note: This table reports the value-added AI effect by intangible intensity and firm size, estimated as four median-split groups in one pooled DiD with small, low-intangible firms as the omitted baseline. Firms are split at the sample median of 2015 intangible intensity and 2015 (log) market value. The Post×AI Share$_{2015}$ is interacted with the three non-baseline group dummies; the specification is otherwise the baseline value-added DiD. Column (1) reports the baseline-group effect (first row) and each group's effect difference from the baseline; column (2) reports the implied per-group effect (baseline + difference). * $p < .10$, ** $p < .05$, *** $p < .01$.

5.3 Refinements on the Complementarity

A natural concern with our headline complementarity result is that it may be mechanical. As discussed in Section 2.1, the Peters and Taylor (2017) intangible capital measure is built from capitalized R&D and a fraction of SG&A, which potentially embodies human capital. At the same time, our treatment intensity is the share of AI workers, which could proxy human capital as well. If intangible capital is in part a function of human capital, then interacting it with the AI-worker share in our complementarity analysis could mechanically capture a nonlinearity in human capital (e.g. a form of

$f(\text{skills}) \times \text{skills})$ rather than a complementarity channel.

Table 6: Refinements on the Complementarity

	(A) AI share square controlled	(B) Orthogonalized intangible split
Post×AI Share$_{2015}$ (small × low intangible, baseline)	6.55	3.29
	(8.77)	(14.85)
× Small × high intangible	16.56	15.49
	(16.48)	(20.65)
× Large × low intangible	10.30*	6.52*
	(5.74)	(3.46)
× Large × high intangible	11.93**	19.12**
	(5.96)	(7.82)
Post×AI Share$^2_{2015}$	−81.03	—
	(105.54)	
Firm FE	✓	✓
Year FE	✓	✓
Controls	✓	✓
Observations	8,603	7,883

Note: This table probes whether the intangible×size value-added AI effect reflects a nonlinearity in skills. Both columns report the implied per-group AI effect (the omitted baseline small, low-intangible effect, first row, plus each group's difference from it) in the pooled four-group difference-in-differences; the specification is otherwise the baseline value-added DiD. Column (A) adds $\text{Post}_t \times \text{AI}^2$ (the treatment interacted with the squared 2015 AI-worker share), so the group pattern cannot merely capture curvature in skills. Column (B) forms the high/low-intangible split from the 2015 log intangible intensity after residualizing it on the 2015 AI-worker share at the firm level. * $p < .10$, ** $p < .05$, *** $p < .01$.

We address this concern in two distinct ways as we document in Table 6.[23] First, we directly control for the suspected nonlinearity by re-estimating the pooled DiD specification through adding the interaction term with respect to the squared AI exposure so that any curvature in the response to AI-skilled labor is controlled. What we find in the column (A) of Table 6 that the pattern is unaffected as the implied AI effect for large and high intangible firms is 11.93 log points.[24] Importantly, the coefficient of the interaction term with the squared of AI exposure is statistically insignificant, which suggests that curvature in exposure does not account for the productivity response.[25] This evidence suggests that a potential nonlinearity concern does not seem to be driving our baseline results.

Second, we implement an approach to purge the intangible measure of its potential human capital component. We residualize the 2015 intangible intensity on the 2015 AI-worker share at the firm level and then re-form the high and low intangible split

[23] As a simple diagnostic, the firm-level correlation between the 2015 intangible intensity and the 2015 AI-worker share is 0.21, which is far from the near-collinearity.

[24] Once the quadratic term is included, the fitted marginal effect at exposure is interpreted differently than a constant slope, so the comparison with the baseline (Table 5) should not be read as an unconditional increase in the implied effect.

[25] Its point estimate is large in absolute terms because the squared AI exposure is small with its square in of order 10^{-5}. Therefore, it is two orders of magnitude below the AI exposure measure itself, and any coefficient on the interaction term with the squared of AI exposure is correspondingly large to move the outcome.

on the residual, which conceptually make the intangible capital split orthogonal to its human capital aspect. Column (B) shows that the large size and high intangible capital effect not only survives but instead strengthens to 19.12 log points compared to the baseline complementarity estimate. Therefore, removing the component of intangible capital that is linearly associated with the AI-worker share does not eliminate the complementarity pattern.[26] This is the opposite of what a mechanical interpretation would predict. Both results indicate that the complementarity reflects an interaction between intangible capital and AI-skilled labor rather than a nonlinearity in skills.

5.4 More Granular Complementarity: ICT Capital

Another potential concern would be that the Peters and Taylor (2017) intangible capital measure may omit AI-relevant investments such as information and communications technology (ICT) capital including internal software development, databases, and data curation. These would be the kind of relevant capital through which AI-skilled labor is likely to operate. For instance, Brynjolfsson and Hitt (2000) argue that the excessive returns to ICT capital reflect its dependence on complementary organizational assets. Compustat, however, contains no precise firm-level measure of ICT capital that would let us capture these inputs directly from the reported accounting items.

To provide evidence on the complementarity between ICT capital and AI-skilled labor, we adopt an imputation approach following the strategy of Falato et al. (2022). More specifically, we construct a firm-level ICT capital stock from the BEA detailed fixed-assets tables and compute, for each industry and year, the ratio of the current-cost net stock of information equipment and software (computers, peripherals, and prepackaged, custom, and own-account software) to the tangible capital stock, and we apply this industry multiple to the firm's deflated net property, plant, and equipment.[27] The resulting variable proxies the software and data-related assets that the Peters and Taylor (2017) measure may miss.[28]

A second advantage of the ICT measure relates to a potential nonlinearity concern in

[26] Residualizing on the AI-worker share removes only that linear association. Hence, it may not be able to fully purge the human capital component of intangible capital if there is an additional remaining nonlinear dependence or a measurement error.

[27] Since the ICT capital measure is an industry-year ratio scaled by firm property, plant, and equipment, its within-industry variation is driven by tangible capital. Hence, it captures a proxy for predicted ICT intensity rather than firm-specific precise software and data holdings.

[28] Its correlation with the Peters and Taylor (2017) stock is 0.37, confirming that the two measures are related but distinct (Figure H.1).

our complementarity analysis as discussed in Section 5.3. By construction, ICT measure contains no human-capital or labor component as it is imputed from industry-level BEA stocks of information equipment and software (e.g. physical investment goods), scaled to the firm's tangible capital. Therefore, it captures ICT and data assets rather than the skills of the workforce.[29]

Based on our measure of ICT capital, Table 7 repeats our complementarity analysis by splitting firms by their 2015 ICT capital and their 2015 size. We find that the pattern mirrors the baseline results such that the productivity effect is positive and statistically significant only for firms that are both large and ICT capital rich, while all other groups are statistically indistinguishable from zero. Hence, we argue that the baseline complementarity effect extends to a capital measure explicitly designed to capture software and data assets. This additional evidence supports the interpretation that our benchmark intangible capital measure proxies the relevant AI-complementary capital.

Table 7: Value-Added Productivity Effect by ICT Capital and Firm Size

	Small firms	Large firms
Low ICT capital	6.81	-0.87
	(21.02)	(5.62)
High ICT capital	4.22	10.59***
	(2.95)	(3.37)

Note: Each cell reports the Post×AI Share_{2015} coefficient from a separate baseline DiD regression of value added per employee estimated within groups defined by median splits of the firm's 2015 informational (ICT) capital stock and 2015 market value. All regressions include the full controls, firm and year fixed effects, and firm-clustered standard errors (in parentheses). * $p < .10$, ** $p < .05$, *** $p < .01$.

6 Productivity Dispersion

The effects we document so far show that AI-related gains are unevenly distributed across firms with different size and intangible capital. Although unequal gains are consistent with a widening of the cross-firm productivity distribution, we have not systematically tested whether the mechanism can directly account for the observed productivity dispersion between small and large firms shown in Figure 1. Because rising productivity dispersion is a central part of our contribution, in this section we

[29] The two are close to orthogonal such that the firm-level correlation between the 2015 ICT capital and the 2015 AI-worker share is 0.06.

test the distributional implication directly with two plausible approaches.

6.1 Quantile Difference-in-Differences

The first approach estimates how the treatment effect varies across the conditional distribution of productivity. Our conceptual hypothesis is that if the effect is larger at higher productivity quantiles, the TensorFlow shock would be productivity dispersion enhancing, whereas the reverse pattern would compress the dispersion. We estimate it using the two-step quantile DiD estimator of Canay (2011), which accommodates firm and year fixed effects before recovering quantile treatment effects. Because this two-step procedure understates standard errors, we additionally implement a firm-level bootstrap correction.

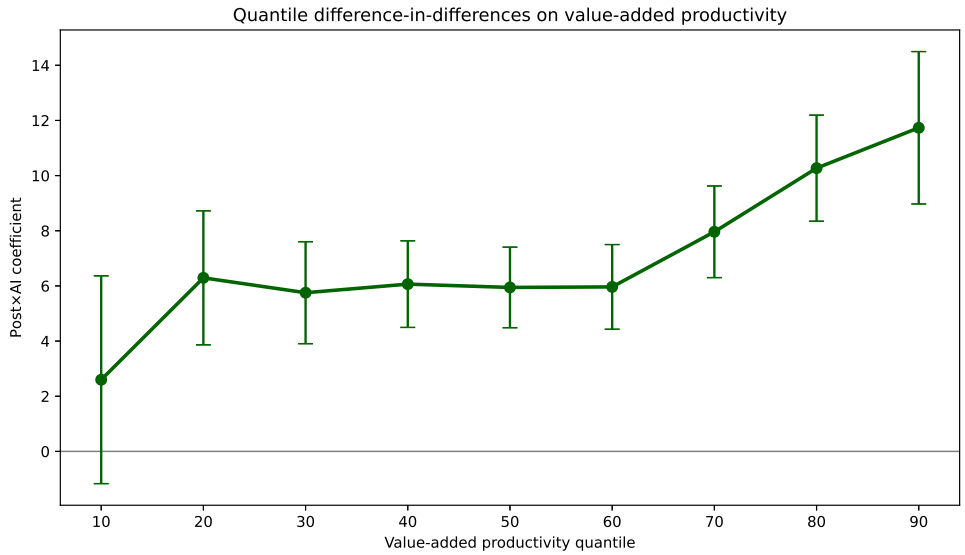
Table 8: Quantile DiD Estimates: Productivity Dispersion

Quantile of value-added prod.	Post×AI Share$_{2015}$	Second-step SE	Bootstrap SE	90% bootstrap interval
0.10	2.60	(2.29)	(4.64)	[−1.03, 13.22]
0.25	7.15**	(1.28)	(3.30)	[1.41, 12.40]
0.50	5.94***	(0.89)	(2.24)	[2.98, 10.08]
0.75	8.84***	(1.07)	(3.20)	[4.22, 15.02]
0.90	11.73***	(1.68)	(3.77)	[7.33, 19.31]
$\beta_{q90} - \beta_{q10}$	9.14**	—	(3.64)	[1.22, 12.83]

Note: This table reports quantile difference-in-differences estimates. Coefficients are in log points per percentage point of 2015 AI exposure, estimated across conditional quantiles with firm and year fixed effects removed by the two-step procedure of Canay (2011). The second-step standard error is the one returned by the quantile regression; it conditions on the firm fixed effect estimated in the first step and therefore understates the sampling variability of the two-step estimator. The bootstrap standard error resamples firms with replacement and re-runs both steps within each of 499 replications, assigning a fresh identifier to each drawn firm so that a firm drawn twice contributes two separate fixed effects; the reported intervals are percentile intervals. Stars are based on the bootstrap standard errors. * $p < .10$, ** $p < .05$, *** $p < .01$.

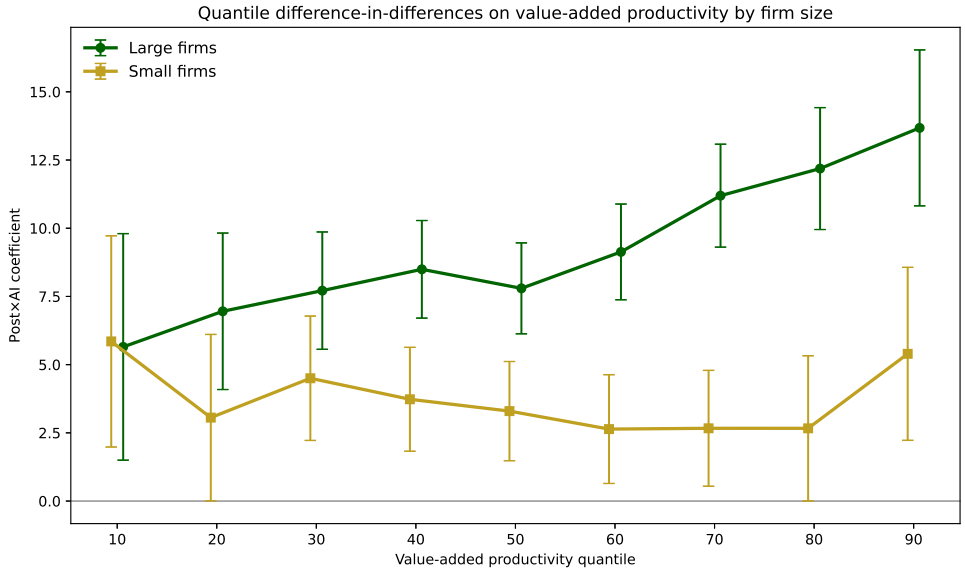
Table 8 reports the results for the key quantiles, and Figure 14 plots the estimated treatment coefficient in each decile of the value-added productivity distribution. We find that the profile rises over the upper half of the distribution. The effect is small at the 10th percentile and statistically indistinguishable from zero, and is around 6 log points in the middle of the distribution. However, it increases steadily above the 60th percentile to 11.7 log points at the 90th percentile. Hence, we find that the shock stretches the productivity distribution to the right at the top. We also find that the overall pattern survives the bootstrap correction such that the coefficient remains bounded away from zero at every quantile above the 10th with the 10th percentile effect remains indistinguishable from zero, which is in line with our hypothesis.

Figure 14: Quantile DiD Estimates: Value-Added Labor Productivity



Note: This figure plots the main coefficient from quantile DiD regressions of the logarithm of value-added labor productivity, estimated at each decile of the conditional productivity distribution using the two-step fixed-effects quantile estimator of Canay (2011). All regressions include the lagged controls, firm and year fixed effects (removed in the first step), and standard errors clustered at the firm level. Error bars show the 90% confidence interval.

Figure 15: Quantile DiD Estimates by Firm Size



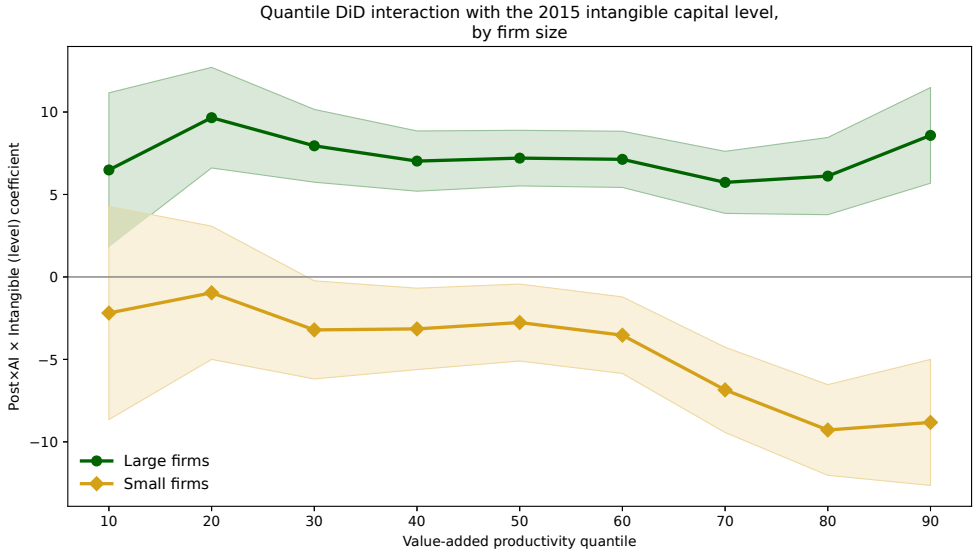
Note: This figure plots the main coefficient from quantile DiD regressions of the logarithm of value-added labor productivity, estimated at each decile separately for large and small firms (median 2015 market value split), using the two-step fixed-effects quantile estimator of Canay (2011) with the full lagged controls. Error bars show the 90% confidence interval.

We also test the role of firm size. Figure 15 estimates the quantile DiD separately for large and small firms based on the median 2015 size split. We find that among large firms the treatment effect rises steadily across the distribution from about 6 log points at the bottom deciles to 14 at the top. However, among small firms the profile is flat at

a low level. Exposure therefore raises productivity by more for large firms than for small ones across most of the distribution, and by progressively more toward the top of the large-firm distribution.[30] Hence, this evidence suggests that our mechanism can account for part of the increasing productivity dispersion documented in Figure 1.

Figure 16 pushes this insight one step further and interacts the treatment with the firm's 2015 intangible capital within each size group. We find that for large firms the intangible capital interaction is positive at every decile, while for small firms it is zero at the bottom of the distribution and turns negative at the top. It suggests that the complementarity between the intangible capital and AI-skilled labor operates across the entire productivity distribution for large firms and is absent for small firms. This evidence is in line with the pattern of our mechanism which widens productivity dispersion by large firms. This result also supports the complementarity evidence documented in Table 5 such that the complementarity effect on productivity is amplified with firm size.

Figure 16: Quantile DiD Interaction with the 2015 Intangible Capital Level, by Firm Size



Note: This figure plots the coefficient on Post × AI × Intangible Capital, the interaction of the treatment with the firm's standardized 2015 intangible capital stock, from quantile DiD regressions of the logarithm of value-added labor productivity, estimated at each decile separately for large and small firms (median 2015 market value split) using the two-step estimator of Canay (2011). Shaded bands show the 90% confidence interval.

[30] Because the two profiles come from separate subsample regressions, we do not test the difference of the large-small coefficient formally.

6.2 Location-Scale Maximum Likelihood

The quantile approach describes where in the distribution the effect falls. In this section, we apply the location-scale maximum likelihood approach that directly models the spread of the distribution. A standard regression assumes that all firms share a common error variance and asks only how the treatment shifts the mean of the outcome. We follow a location-scale model by Gould et al. (2010), which relaxes this assumption by letting the variance of the outcome depend on covariates, so that the treatment can shift both the location (the conditional mean) and the scale (the conditional variance) of our outcome variable of productivity. This is what a productivity dispersion hypothesis requires as we want to know not only whether AI-exposed firms become more productive on average, but whether the shock pushes firms further apart.

More specifically, we do an estimation by maximum likelihood the model $y_{it} \sim N(X_{it}\beta, \exp(Z_{it}\gamma))$. In this specification, the mean equation ($X_{it}\beta$) is the baseline productivity regression, and the scale equation ($\exp(Z_{it}\gamma)$) models the standard deviation of the outcome as a function of the same variables where the exponential form guarantees a positive variance. Here, the coefficient of interest is the treatment coefficient in the scale equation such that a positive value means that the shock widens the conditional variance of productivity across firms, i.e., it makes exposed firms more unequal. On the other hand, a negative value means it compresses them.

Table 9 reports the corresponding results. We find that, on average, the treatment leaves the conditional variance essentially unchanged (a scale coefficient of 0.010, statistically indistinguishable from zero). Moreover, the scale effect rises with the firm's intangible capital intensity, and in the size-by-intangible capital split the conditional variance falls for the baseline group of small and low intangible capital firms, while the interactions with high intangible capital intensity and large size are both positive.[31] Therefore, the results suggest that the implied variance effect is positive only for firms that are both large and intangible capital intensive. In other words, the shock compresses outcomes among firms with low intangible capital, and spreads them among firms with high intangible capital. This is also exactly where the baseline complementarity effects are concentrated.

[31] The reported significance tests apply to the individual scale coefficients, so Table 9 indicates the sign pattern across different groups.

Table 9: Location-Scale Estimates: Productivity Dispersion

	Coefficient	(SE)
Pooled, baseline		
Post×AI Share$_{2015}$	0.010	(0.015)
Pooled, complementarity		
Post×AI Share$_{2015}$	0.024	(0.015)
Post×AI Share$_{2015}$ × L.ln(Intangible)	0.015*	(0.009)
Size × intangible split		
Post×AI Share$_{2015}$ (base: low int., small)	−0.065***	(0.020)
Post×AI Share$_{2015}$ × high intangible	0.028**	(0.014)
Post×AI Share$_{2015}$ × large	0.062***	(0.016)

Note: This table reports location-scale MLE estimates of the effect of AI exposure on productivity dispersion. It reports the scale ($\log \sigma$) equation of a location-scale normal MLE, $y \sim N(X\beta, \exp(Z\gamma))$; and the regressors are standardized. In the split specification the group dummies partition firms by 2015 intangible intensity and 2015 size at their medians; the base group is low-intangible, small firms. Standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Taken together, the quantile and location-scale evidence implies that the complementarity between intangible capital and AI-skilled labor raises productivity most at the top of the productivity distribution and widens the conditional variance precisely among large and intangible capital rich firms, which contributes to rising productivity dispersion between small and large firms.

6.3 How Much of the Productivity Dispersion Accounted For?

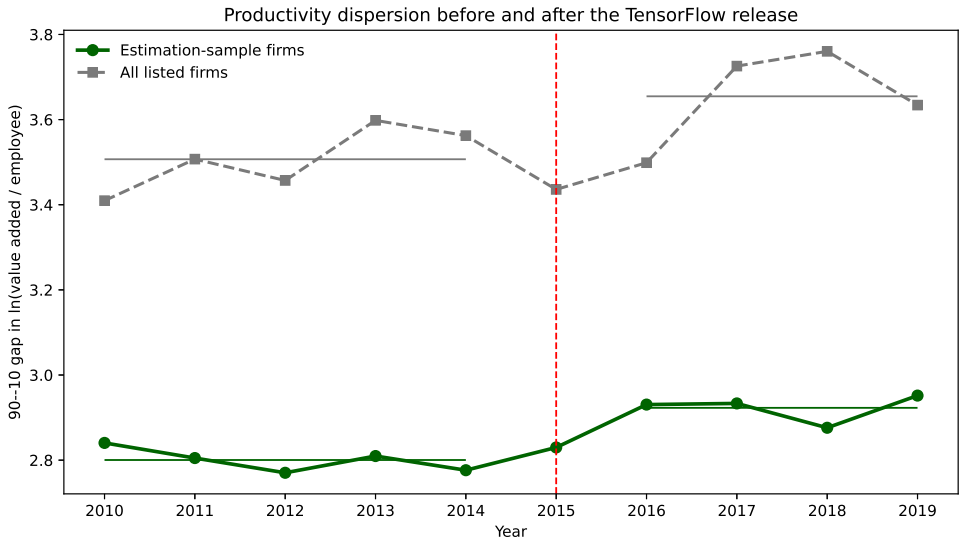
We also apply a back-of-the-envelope calculation to provide an illustrative magnitude on how much of the observed productivity dispersion between small and large firms can be accounted by our mechanism.

As Figure 17 suggests, the size 90-10 gap in log value added per worker rises from an average of 2.80 over 2010-2014 to 2.92 over 2016-2019 among the firms of our estimation sample, which represents an increase of about 12 log points, and it corresponds to 14.8 log points across all listed Compustat firms in our overall panel.[32] As documented in Table 8, the quantile estimates imply a differential effect of $\beta_{q90} - \beta_{q10} = 9.1$ log points per percentage point of exposure, and our mechanism implies a widening of roughly 5.3 log points given the mean 2015 AI-worker share of 0.58 percentage

[32] We compare period averages rather than end-points because the gap is noisy from year to year.

points. Therefore, our back-of-the-envelope calculation is on the order of two-fifths of the observed rise in productivity dispersion. We present this as a heuristic scale comparison rather than a formal statistical calculation since multiplying a difference between conditional quantile slopes by mean exposure does not in general recover the change in an unconditional quantile gap, which depends on the joint distribution of exposure, covariates, and treatment effects. Hence, the comparison requires firm ranks to be stable enough that effects estimated on the conditional distribution carry over to the cross-firm distribution. Even with this limitation, we think that it is still informative about order of magnitude, which indicates that the complementarity we document is large enough to be a material part of the productivity dispersion trend.

Figure 17: Productivity Dispersion Before and After the TensorFlow Release



Note: This figure plots the 90-10 gap in the logarithm of value-added labor productivity by year, separately for the firms of the pooled estimation sample and for all listed firms in the 2010-2019 panel. Horizontal segments mark the pre-period (2010-2014) and post-period (2016-2019) averages of each series with respect to the 2015 TensorFlow release. The dashed vertical line marks the 2015 release of Google TensorFlow.

7 TensorFlow-Specific and Worker-Specific Evidence

Our identification interprets the post-2015 productivity effect driven by AI-exposed firms as the effect materialized by the TensorFlow release. Although we show in Section 4.3.6 that controlling alternative firm-level technologies do not affect our baseline results, a natural objection would still be that the estimates might instead capture a broader post-2015 wave of AI and digital diffusion (e.g. cloud computing, data infrastructure, and open-source ecosystems), in which case the results would still document heterogeneous returns to AI readiness, but the mechanism would be less

precisely identified. To confront this concern, in this section we provide direct evidence that the workforce response of exposed firms is specifically directed at TensorFlow.

We use the detailed resume, LinkedIn profiles and job posting data from Revelio Labs' proprietary database, a comprehensive workforce data, which allow us to observe firm-year level the number of workers whose profiles list specific deep-learning frameworks. Based on the Revelio Labs' skill taxonomy regarding the deep-learning frameworks, we handle TensorFlow and its bundled high-level API Keras, as well as the competing and legacy frameworks PyTorch, Caffe, Torch, and Apache MXNet.[33]

Our approach is in line with a general convention in the related literature in which job-posting databases have become a standard source of firm-level evidence on technology adoption. As an example, Alekseeva et al. (2021), Acemoglu et al. (2022), Babina et al. (2024), and Hampole et al. (2025) (which also uses Revelio Labs database) use online vacancies to measure firm-level AI exposure. Our analysis using Revelio Labs database also closely follows that approach. We need to note that we have a certain data limitation such that the framework-specific skills in our sample are essentially almost absent from worker profiles before 2015 and begin to appear from 2017 onward, which can be driven by natural resume updating lags. Due to this data limitation, our approach in this section delivers a descriptive analysis rather than a causal inference.

7.1 TensorFlow-Specific Evidence

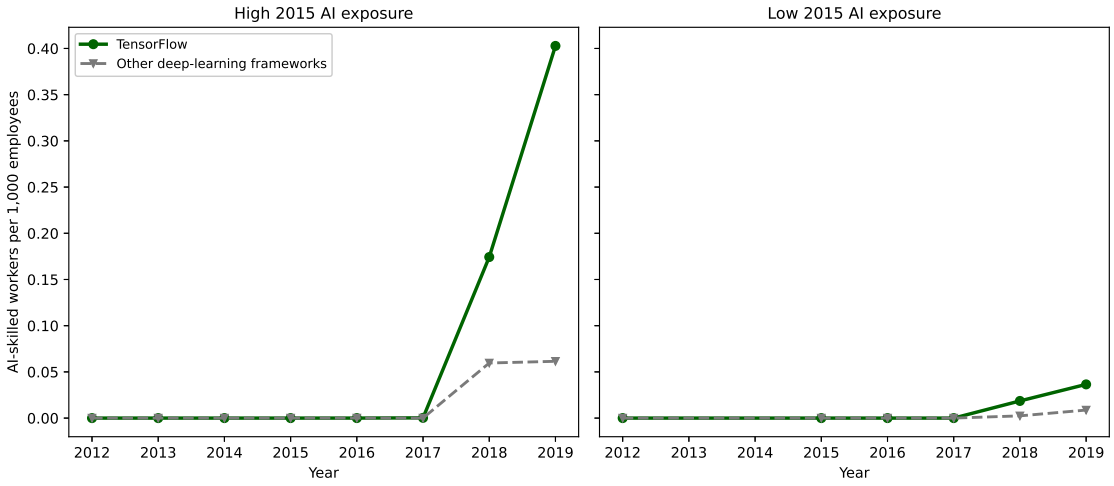
We first investigate whether our running variable of AI exposure is mainly associated with TensorFlow or also with alternative deep learning frameworks. This analysis helps us determine whether our baseline analysis captures a technological wave specific to TensorFlow or a broader one.

First, Figure 18 plots the average number of TensorFlow workers and of workers using all other alternative deep-learning frameworks, per 1,000 employees, separately for firms with high and low 2015 AI exposure. We find a pattern such that in high exposure firms, TensorFlow skills surge after 2017 through reaching about 0.40 workers per 1,000 employees by 2019, more than six times the corresponding count for all other frameworks combined. However, in low exposure firms both series barely move. These levels are small by construction. These deep-learning framework skills are listed by a narrow set of workers, and hence we interpret that these numbers capture where

[33] In Section F, we provide details on the data structure and additional related evidence.

AI capability is being built rather than the size of the AI workforce. In this analysis, we focus on the contrast between high and low exposure firms but not the absolute magnitude of either series since the identifying variation in our baseline empirical analysis is mainly driven by the difference across firms with different AI exposure.

Figure 18: TensorFlow versus Other Deep-Learning Frameworks, by 2015 AI Exposure

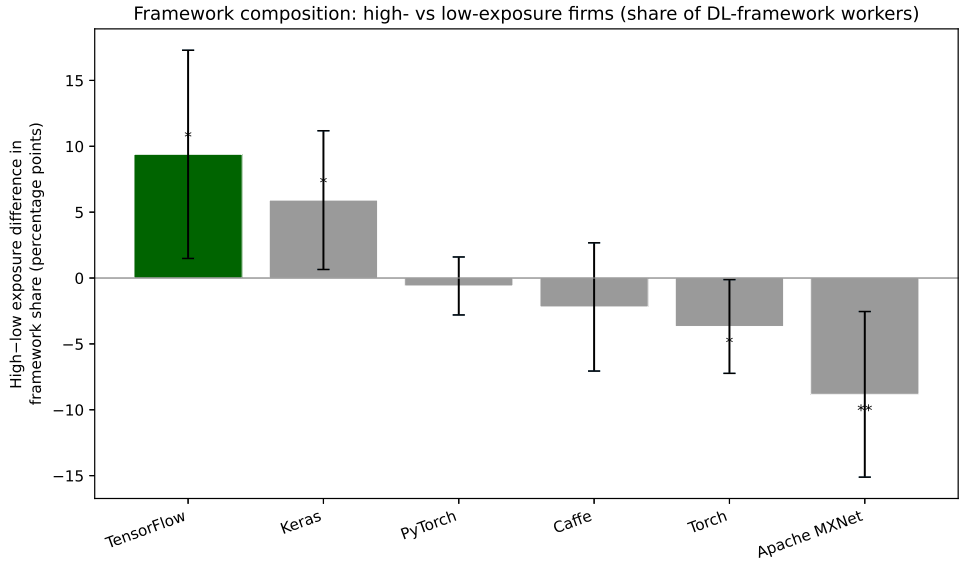


Note: This figure plots the average number of workers with TensorFlow skills and with skills in other deep-learning frameworks (Keras, PyTorch, Caffe, Torch, Apache MXNet), per 1,000 employees, separately for firms with high (left panel) and low (right panel) 2015 AI-worker exposure, using the Revelio Labs framework-level skill data.

Second, Figure 19 formalizes the previous observed pattern with a formal composition test. We restrict our sample to firm-years that hire any deep-learning framework workers, and evaluate how high and low exposure firms adjust their composition of the deep-learning. In more detail, we compute each deep-learning framework's share of the firm-year's total deep-learning framework workers, and regress on a high versus low 2015 AI exposure indicator. We find that the TensorFlow share is 9.4 percentage points higher in high exposure firms (and the share of Keras, which is TensorFlow's bundled high-level interface, 5.9 points higher), while the shares of the other frameworks such as Torch and MXNet are lower, and PyTorch is flat.

Hence, we find that the workforce response of more exposed firms is TensorFlow directed in composition (Figure 19), which suggests that our running variable of AI exposure does not seem to capture a broader deep-learning or technology waves. Together with the analysis where we control for other firm-level technology exposures in Section 4.3.6, this further evidence supports our interpretation that the TensorFlow release is a key driver behind the improvements in AI effectiveness evaluated by our baseline analysis.

Figure 19: Deep-Learning Workforce Composition by Framework



Note: This figure plots, for each deep-learning framework, the difference in the framework's workforce share between firms with high and low 2015 AI exposure. The sample is restricted to firm-years that employ at least one worker with any deep-learning framework skill in the Revelio Labs skill data. The plotted bar is the coefficient on an indicator for above-median 2015 AI exposure from a regression that includes year fixed effects and the full set of lagged controls. High and low exposure denote the two halves of the distribution of the firm's 2015 share of AI workers. Error bars show the 90% confidence interval, based on standard errors clustered by firm.

Two features of this evidence should be stated. First, as we noted before, Revelio Labs' framework-level skill records are almost empty before the release and its postings coverage begins in 2016. Therefore, it does not provide us a plausible pre-period. Therefore, we are not able to conduct a difference-in-differences methodology in this section. Second, workers would update their resumes or LinkedIn profiles with a lag, so we may not pin down the timing of the adoption of skills precisely. Therefore, what this section establishes is the composition of the workforce response but not its timing or its causal magnitude.[34]

7.2 Worker-Specific Evidence on the Productivity Channel

Our baseline productivity outcome, value-added labor productivity, is a firm-level aggregate. The Revelio Labs data allow us to look inside that aggregate along two margins such as what firms offer to pay, and whom they hire. First, we argue that posted pay would serve as a worker side proxy for productivity in a way that the pay would track the value of what workers produce, and the way it is distributed across occupations would help us understand any productivity gain accrues narrowly or across the firm more equally. Second, whom firms hire would indicate whether firms

[34]Section F sets out the limitations in more detail.

also differ in the workforce they employ. We are aware that neither margin substitutes for our firm-level productivity measures because posted pay would reflect productivity under certain conditions in the labor market that we cannot test here, and, due to the data limitation, neither has a plausible pre-period so that we cannot provide a causal inference in these both margins. That said, cross-section descriptive analyses in these margins would still help us discipline the interpretations of the baseline result.

Posted pay. A raw comparison of average posted salaries would be uninformative, because high and low AI exposure firms would advertise different jobs in different places. For instance, in Section F.3, we find that high AI exposure firms post far more computing and AI-titled roles, and they do so within industries. Therefore, we construct an occupation-adjusted measure by residualizing the logarithm of the posted salary on occupation-by-year and metropolitan area effects.[35] We then evaluate how this adjusted posted pay measure differs with respect to firms' AI exposure.

Table 10: Posted Pay and 2015 AI Exposure

	(1) All postings	(2) Computer/math	(3) All other
2015 AI-worker share (p.p.)	0.0518***	0.0778***	0.0665***
	(0.0096)	(0.0119)	(0.0123)
Year effects	✓	✓	✓
Occupation×year, metro absorbed	✓	✓	✓
Firm-years	2,606	2,078	2,569
Firms	725	653	722

Note: The dependent variable is the firm-year mean residual of log posted salary after absorbing occupation×year and metropolitan-area effects, so the comparison holds occupation, location and year fixed. Columns (2) and (3) split the postings by occupation group: the association is of similar magnitude in both, so the pay difference is firm-wide rather than specific to technical roles. Column (1) averages over all of a firm-year's postings and therefore weights occupations by the firm's own posting mix; since exposed firms post relatively more computing roles, that mix is itself a function of the exposure measure, so column (1) is a composition-weighted blend and need not lie between columns (2) and (3). Standard errors clustered by firm. * $p < .10$, ** $p < .05$, *** $p < .01$.

Column (1) of Table 10 reports the result.[36] We find that one percentage point higher 2015 AI-worker share is associated with a 5.2 log point higher pay premium.[37] To the extent that posted pay would track the value of what workers produce, we interpret

[35] Averaging the residual within the firm-year yields a posted-pay premium that compares firms' pay offers in the same occupation, metropolitan area and year.

[36] Since not all firms in our DiD sample appear in the Revelio Labs database, we have a lower number of firms and observations in this particular analysis.

[37] Figure F.3 shows the relationship across the whole exposure distribution; it is positive throughout, with a mild concavity at the top, so the linear coefficient should be read as an average rather than as a constant marginal effect.

that this pay premium would be indicative of more AI exposed firms operating with productive workers, which is in line with our firm-level evidence in the baseline setting.

Which workers? If high exposure firms pay more because they compete for scarce AI talent, then we can argue that the pay premium may not reflect that firms compensate higher productive workers with higher pay, but instead it can reflect that the pay premium would be concentrated in such technical roles. We evaluate this argument in Columns (2) and (3) of Table 10, and we do not find such evidence in favor of the pay premium concentrated in technical roles. We find that the corresponding association is 7.8 log points (column (2)) for computing and mathematical occupations and around 6.7 log points (column (3)) for every other occupation.[38] Hence, this result provides suggestive evidence that the pay premium would be an overall feature within the firm (i.e., productivity channel) but would not be only specific to the AI-related roles.

Workforce composition. A natural concern for the evidence on posted pay would be that high exposure firms pay more because they employ different people in their workforce. To evaluate this potential channel, we use workers' education profiles in the Revelio Labs database, and integrate it into our descriptive analysis on the posted pay premium.

First of all, we investigate whether firms with different AI exposure employ different composition of workers with respect to their education degrees. Panel A of Table 11 and Figure F.4 show that high exposure firms significantly differs than low exposure firms in terms of education profiles of their workers. What we find is that the share of workers holding at least a bachelor's degree is 93.1% at high exposure firms against 84.7% at low-exposure firms. The graduate share that covers master's degrees and doctorates is 42.4% against 31.4%; and the doctorate share is 6.17% against 3.74%, for high exposure and low exposure firms respectively. To bring more systematic evidence, when we regress the education shares on the AI exposure, we find that a one percentage point higher 2015 AI-worker share is associated with a 7.55 percentage point higher graduate share, conditional on firm size.[39]

[38] Both exceed the all-postings coefficient of column (1), which is not a contradiction but a composition effect since column (1) weights occupations by each firm's own posting mix, and that mix is itself part of what the exposure measure captures. Hence, the all-postings average is a blend whose weights move with exposure. As a validation check, holding the occupation mix fixed at the sample-wide share of computing postings instead of the firm's own gives the coefficient of 0.077 that is between the two occupation groups.

[39] Figure F.4 plots the same relationship across the exposure distribution, drawn on the wider set of

Table 11: Workforce Education, 2015 AI Exposure and Posted Pay

Panel A: Firm workforce education and 2015 AI exposure			
Share of the workforce (p.p.)	Low exposure	High exposure	Coef. on 2015 AI share
Bachelor's or above	84.7	93.1	5.377***
			(0.664)
Graduate (master's or doctorate)	31.4	42.4	7.551***
			(1.117)
Doctorate	3.74	6.17	1.362***
			(0.315)
Panel B: Posted-pay premium, exposure and education jointly			
	(1)	(2)	(3)
2015 AI-worker share (p.p.)	0.0518***		0.0218**
	(0.0096)		(0.0090)
Graduate share of workforce (p.p.)		0.0043***	0.0040***
		(0.0004)	(0.0004)
Year effects	✓	✓	✓
Firm-years	2,606	2,606	2,606
Firms	725	725	725

Note: Panel A is a firm-level cross-section and controls for lagged log assets. Panel B is a firm-year panel whose dependent variable is the occupation-adjusted posted-pay premium defined in Table 10; the graduate share is firm-constant. Every column of Panel B conditions on lagged log assets and year effects, so columns differ from one another only through the graduate share, and column (1) is the same regression, on the same rows, as column (1) of Table 10. Standard errors clustered by firm. * $p < .10$, ** $p < .05$, *** $p < .01$.

Panel B of Table 11 brings the AI exposure and workforce education profile together to evaluate whether the AI exposure can still explain the posted pay by controlling workforce education.[40] First of all, as documented previously, the 2015 AI-worker share is itself associated with a 5.2 log point posted-pay premium. Importantly, when we include it together with the graduate share of the workforce, even though the coefficient of AI exposure falls by about three fifth and the graduate share itself contributes significantly to the posted pay premium, the association between AI exposure and the posted pay is still positive and statistically significant even after controlling the graduate share of workforce. Therefore, even though the education profiles would be one of the key driver behind the posted pay, we find that AI exposure itself can contribute to it in a positive and statistically significant manner.

Taken together, the two margins line up with the mechanism of the paper, and they let

firms with usable education records. Most of the difference is accumulated among firms with quite modest AI workforces rather than driven by a handful of technology firms at the extreme.

[40] Every column conditions on firm size and year. Column (1) is the same regression as column (1) of Table 10.

posted pay serve as a worker-side proxy for productivity. To the extent that pay tracks the value of what workers produce, we assume that a firm that offers more for the same occupation, metropolitan area, and year is a more productive firm. If improved AI effectiveness raises productivity, then the pay difference should appear across the firm rather than in the roles that operate the technology, which is what Table 10 shows, and workforce composition alone should not explain it, which is what Table 11 shows.

8 Robustness Checks

We perform an extensive battery of additional exercises for our baseline results, which we report in Section E in detail. As a summary, first of all, we relax the linear dose-response that the continuous exposure measure imposes by implementing above-median and quartile exposure indicators (Section E.1). Moreover, we address potential confounding dynamics by adding state-specific and industry-specific linear time trends (Section E.2) and by incorporating industry-by-year and state-by-year fixed effects (Section E.3). We also address a potential concern that the response would reflect pre-existing differences of AI-intensive firms rather than the shock through a placebo test (Section E.4) and propensity-score matching (Section E.5). Furthermore, we use alternative measurement of the treatment with a narrower definition of AI skills (Section E.6). We assess inference with two-way clustering by firm and year (Section E.7). We investigate alternative firm size measures with respect to firm assets and employment (Section E.8). We also measure exposure with the firm's pre-shock 2014 AI-worker share (Section E.9). Finally, we perform an external validity check by dropping the AI-intensive industries (Section E.10).

9 Conclusion

This paper studies how the complementarity between intangible capital and AI-skilled labor shapes firm-level productivity dynamics. We measure a firm's AI readiness by its share of AI-skilled workers, which captures AI-related human capital. We exploit the November 2015 release of Google's TensorFlow as a plausibly exogenous improvement in the effectiveness of AI-skilled workers. Using a difference-in-differences framework with value-added labor productivity and total factor productivity as outcomes, we show that firms with a higher AI-worker share at the time of the shock become more productive, which is pronounced with the firm's intangible capital and size. We

perform direct distributional tests and show that the shock raised productivity most at the top of the distribution, and the observed productivity dispersion between small and large firms would be disciplined through our mechanism.

We interpret these results as heterogeneous returns to AI readiness during the diffusion of deep learning, with the TensorFlow release as the trigger through which the effectiveness of AI-skilled labor improved. We provide a direct evidence to support this insight using the proprietary Revelio Labs data. We show that, after the release, the deep learning workforce of higher exposed firms is TensorFlow dominant in magnitude and TensorFlow directed in composition.

Our findings provide three key implications for innovation and technology policy. First, policies promoting AI adoption are unlikely to generate uniform gains, since in our estimates the returns to AI readiness are concentrated among firms that already hold complementary intangible capital and operate at scale. Second, because those gains rise with firm size and intangible capital, broad-based AI support may amplify productivity dispersion by disproportionately favoring large and intangible capital rich firms, which would yield a key policy trade-off. For instance, policies that incentivize frontier AI adoption in superstar firms would raise productivity for those firms relative to small firms, and would therefore widen relative productivity dispersion. Third, workforce policies that expand advanced digital and AI skills together with support for intangible capital accumulation such as data access, managerial training, and organizational upgrading would raise the returns to intangible capital and firms' capacity to convert technological advance into sustained productivity growth.

Our analysis also highlights an important avenue for future research. Our reliance on Compustat data confines the analysis to publicly traded firms. Extending the framework to private firms including startups and small businesses through micro-level census or administrative data would offer a more comprehensive view of how intangible capital and AI-skilled labor interact across the firm-size distribution.

References

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

_ , "Directed technical change," *The Review of Economic Studies*, 2002, *69* (4), 781–809.

_ , "Technical change, inequality, and the labor market," *Journal of Economic Literature*, 2002, *40* (1), 7–72.

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

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer, "Identification properties of recent production function estimators," *Econometrica*, 2015, *83* (6), 2411–2451.

Adams, Jonathan J, Min Fang, Zheng Liu, and Yajie Wang, "The rise of AI pricing: Trends, driving forces, and implications for firm performance," *Journal of Monetary Economics*, 2025, p. 103875.

Adler, David and Laurence B Siegel, *The productivity puzzle: Restoring economic dynamism*, CFA Institute Research Foundation, 2019.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li, "A theory of falling growth and rising rents," *Review of Economic Studies*, 2023, *90* (6), 2675–2702.

_ **, Peter Howitt, and Giovanni L Violante**, "General purpose technology and wage inequality," *Journal of Economic Growth*, 2002, *7* (4), 315–345.

Akcigit, Ufuk and Sina T Ates, "What happened to US business dynamism?," *Journal of Political Economy*, 2023, *131* (8), 2059–2124.

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

Altomonte, Carlo, Domenico Favoino, Monica Morlacco, Tommaso Sonno et al., "Markups, intangible capital and heterogeneous financial frictions," Technical Report, Centre for Economic Performance, LSE 2021.

Andrews, Dan, Chiara Criscuolo, and Peter N Gal, "The best versus the rest: the global productivity slowdown, divergence across firms and the role of public policy," 2016.

Ark, Bart Van, Vivian Chen, and Kirsten Jäger, "European productivity growth since 2000 and future prospects," *International Productivity Monitor*, 2013, (25), 65.

Arrighetti, Alessandro, Fabio Landini, and Andrea Lasagni, "Intangible assets and firm heterogeneity: Evidence from Italy," *Research Policy*, 2014, *43* (1), 202–213.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen, "The fall of the labor share and the rise of superstar firms," *The Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Autor, David H, Frank Levy, and Richard J Murnane, "The skill content of recent technological change: An empirical exploration," *The Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

_ , Lawrence F Katz, and Melissa S Kearney, "Trends in US wage inequality: Revising the revisionists," *The Review of Economics and Statistics*, 2008, *90* (2), 300–323.

Ayyagari, Meghana, Asli Demirguc-Kunt, and Vojislav Maksimovic, *The rise of star firms: intangible capital and competition*, The World Bank, 2019.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial intelligence, firm growth, and product innovation," *Journal of Financial Economics*, 2024, *151*, 103745.

Bartelsman, Eric J and Zoltan Wolf, "Measuring productivity dispersion," 2018.

Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp, "Big data in finance and the growth of large firms," *Journal of monetary economics*, 2018, *97*, 71–87.

Bessen, James, "Information technology and industry concentration," 2017.

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

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

_ , **Daniel Rock, and Chad Syverson**, "Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics," in "The Economics of Artificial Intelligence: An Agenda," University of Chicago Press, 2018, pp. 23–57.

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

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna, "Difference-in-differences with a continuous treatment," Technical Report, National Bureau of Economic Research 2024.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Robust inference with multiway clustering," *Journal of Business & Economic Statistics*, 2011, *29* (2), 238–249.

Canay, Ivan A, "A simple approach to quantile regression for panel data," *The econometrics journal*, 2011, *14* (3), 368–386.

Chen, Sheng-Syan, Yan-Shing Chen, and Shu-Cing Peng, "The Double-Edged Role of Artificial Intelligence in Corporate Decarbonization," Working Paper, SSRN August 2025. Available at SSRN: https://ssrn.com/abstract=5362708.

Chen, Wilbur Xinyuan, Terrence Tianshuo Shi, and Suraj Srinivasan, "The Value of AI Innovations," Working Paper 24-069, Harvard Business School May 2024.

Chiavari, Andrea and Sampreet Singh Goraya, "The rise of intangible capital and the macroeconomic implications," 2025.

Corrado, Carol A and Charles R Hulten, "How do you measure a" technological revolution"?," *American Economic Review*, 2010, *100* (2), 99–104.

Corrado, Carol, Charles Hulten, and Daniel Sichel, "Intangible capital and US economic growth," *Review of Income and Wealth*, 2009, *55* (3), 661–685.

_ , **Jonathan Haskel, and Cecilia Jona-Lasinio**, "Knowledge spillovers, ICT and productivity growth," *Oxford Bulletin of Economics and Statistics*, 2017, *79* (4), 592–618.

_ , _ , _ , **and Massimiliano Iommi**, "Intangible investment in the EU and US before and since the Great Recession and its contribution to productivity growth," Technical Report, EIB Working Papers 2016.

Crouzet, Nicolas and Janice C Eberly, "Understanding weak capital investment: The role of market concentration and intangibles," Technical Report, National Bureau of Economic Research 2019.

_ **and Janice Eberly**, "Rents and intangible capital: A q+ framework," *The Journal of Finance*, 2023, *78* (4), 1873–1916.

_ **, Janice C Eberly, Andrea L Eisfeldt, and Dimitris Papanikolaou**, "The economics of intangible capital," *Journal of Economic Perspectives*, 2022, *36* (3), 29–52.

_ **, Janice Eberly, Andrea Eisfeldt, and Dimitris Papanikolaou**, "Intangible capital, non-rivalry, and growth," *Available at SSRN*, 2022.

DeAngelo, Harry, Linda DeAngelo, and René M Stulz, "Dividend policy and the earned/contributed capital mix: a test of the life-cycle theory," *Journal of Financial economics*, 2006, *81* (2), 227–254.

Decker, Ryan A, John Haltiwanger, Ron S Jarmin, and Javier Miranda, "Changing business dynamism and productivity: Shocks versus responsiveness," *American Economic Review*, 2020, *110* (12), 3952–3990.

Dhaene, Geert and Koen Jochmans, "Split-panel jackknife estimation of fixed-effect models," *The Review of Economic Studies*, 2015, *82* (3), 991–1030.

Edmans, Alex, Xavier Gabaix, and Augustin Landier, "A multiplicative model of optimal CEO incentives in market equilibrium," *The Review of Financial Studies*, 2009, *22* (12), 4881–4917.

Eisfeldt, Andrea L and Dimitris Papanikolaou, "The value and ownership of intangible capital," *American Economic Review*, 2014, *104* (5), 189–94.

Engberg, Erik, Michael Koch, Magnus Lodefalk, and Sarah Schroeder, "Artificial intelligence, tasks, skills, and wages: Worker-level evidence from Germany," *Research Policy*, 2025, *54* (8), 105285.

Ewens, Michael, Ryan H Peters, and Sean Wang, *Acquisition prices and the measurement of intangible capital*, National Bureau of Economic Research, 2019.

Falato, Antonio, Dalida Kadyrzhanova, Jae Sim, and Roberto Steri, "Rising intangible capital, shrinking debt capacity, and the US corporate savings glut," *The Journal of Finance*, 2022, *77* (5), 2799–2852.

Firooz, Hamid, Sylvain Leduc, and Zheng Liu, "Will AI Intensify or Weaken Market Competition?," Working Paper 2024-15, Federal Reserve Bank of San Francisco 2026.

_ , Zheng Liu, and Yajie Wang, "Automation and the rise of superstar firms," *Journal of Monetary Economics*, 2025, *151*, 103733.

Foster, Lucia, John Haltiwanger, and Chad Syverson, "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?," *American Economic Review*, 2008, *98* (1), 394–425.

Goldin, Claudia and Lawrence F Katz, "The origins of technology-skill complementarity," *The Quarterly Journal of Economics*, 1998, *113* (3), 693–732.

Gould, William, Jeffrey Pitblado, and Brian Poi, *Maximum Likelihood Estimation with Stata*, 4 ed., College Station, TX: Stata Press, 2010.

Greenwood, Jeremy and Mehmet Yorukoglu, "1974," in "Carnegie-Rochester conference series on public policy," Vol. 46 Elsevier 1997, pp. 49–95.

_ , Zvi Hercowitz, and Per Krusell, "Long-run implications of investment-specific technological change," *The American Economic Review*, 1997, pp. 342–362.

Griliches, Zvi, "Capital-skill complementarity," *The Review of Economics and Statistics*, 1969, pp. 465–468.

Gutiérrez, Germán and Thomas Philippon, "Declining Competition and Investment in the US," Technical Report, National Bureau of Economic Research 2017.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller, "Artificial intelligence and the labor market," Technical Report, National Bureau of Economic Research 2025.

Haskel, Jonathan and Stian Westlake, *Capitalism without capital*, Princeton University Press, 2017.

Hornstein, Andreas, Per Krusell, and Giovanni L Violante, "The effects of technical change on labor market inequalities," in "Handbook of Economic Growth," Vol. 1, Elsevier, 2005, pp. 1275–1370.

Hubmer, Joachim and Pascual Restrepo, "Not a Typical Firm: Capital–Labor Substitution and Firms' Labor Shares," *American Economic Journal: Macroeconomics*, 2026, *18* (2), 34–71.

Katz, Lawrence F and Kevin M Murphy, "Changes in relative wages, 1963–1987: supply and demand factors," *The Quarterly Journal of Economics*, 1992, *107* (1), 35–78.

_ et al., "Changes in the wage structure and earnings inequality," in "Handbook of Labor Economics," Vol. 3, Elsevier, 1999, pp. 1463–1555.

Klette, Tor Jakob and Zvi Griliches, "The inconsistency of common scale estimators when output prices are unobserved and endogenous," *Journal of applied econometrics*, 1996, *11* (4), 343–361.

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

Krusell, Per, Lee E Ohanian, José-Víctor Ríos-Rull, and Giovanni L Violante, "Capital-skill complementarity and inequality: A macroeconomic analysis," *Econometrica*, 2000, *68* (5), 1029–1053.

Kwon, Spencer Y, Yueran Ma, and Kaspar Zimmermann, "100 years of rising corporate concentration," *American Economic Review*, 2024, *114* (7), 2111–2140.

Ladika, Tomislav, Robin Döttling, and Enrico C Perotti, "Creating intangible capital," *Available at SSRN 3479152*, 2025.

Lashkari, Danial, Arthur Bauer, and Jocelyn Boussard, "Information technology and returns to scale," *American Economic Review*, 2024, *114* (6), 1769–1815.

Lev, Baruch and Suresh Radhakrishnan, "The valuation of organization capital," in "Measuring Capital in the New Economy," University of Chicago Press, 2005, pp. 73–110.

Levinsohn, James and Amil Petrin, "Estimating production functions using inputs to control for unobservables," *The Review of Economic Studies*, 2003, *70* (2), 317–341.

Linck, James S, Jeffry M Netter, and Tina Yang, "The determinants of board structure," *Journal of financial economics*, 2008, *87* (2), 308–328.

Loecker, Jan De, "Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity," *Econometrica*, 2011, *79* (5), 1407–1451.

_ and Pinelopi Koujianou Goldberg, "Firm performance in a global market," *Annual Review of Economics*, 2014, *6* (1), 201–227.

_ , **Jan Eeckhout, and Gabriel Unger**, "The rise of market power and the macroeconomic implications," *The Quarterly journal of economics*, 2020, *135* (2), 561–644.

McElheran, Kristina, J Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas, "AI adoption in America: Who, what, and where," *Journal of Economics & Management Strategy*, 2024, *33* (2), 375–415.

McGrattan, Ellen R, "Intangible capital and measured productivity," *Review of Economic Dynamics*, 2020, *37*, S147–S166.

_ **and Edward C Prescott**, "Unmeasured Investment and the Puzzling US Boom in the 1990s," *American Economic Journal: Macroeconomics*, 2010, *2* (4), 88–123.

Nickell, Stephen, "Biases in dynamic models with fixed effects," *Econometrica: Journal of the econometric society*, 1981, pp. 1417–1426.

Olley, G Steven and Ariel Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica: Journal of the Econometric Society*, 1996, pp. 1263–1297.

O'Mahony, Mary and Michela Vecchi, "R&D, knowledge spillovers and company productivity performance," *Research Policy*, 2009, *38* (1), 35–44.

Peters, Ryan H and Lucian A Taylor, "Intangible capital and the investment-q relation," *Journal of Financial Economics*, 2017, *123* (2), 251–272.

Piva, Mariacristina, Enrico Santarelli, and Marco Vivarelli, "The skill bias effect of technological and organisational change: Evidence and policy implications," *Research Policy*, 2005, *34* (2), 141–157.

Rammer, Christian, Gastón P Fernández, and Dirk Czarnitzki, "Artificial intelligence and industrial innovation: Evidence from German firm-level data," *Research Policy*, 2022, *51* (7), 104555.

Ridder, Maarten De, "Market power and innovation in the intangible economy," *American Economic Review*, 2024, *114* (1), 199–251.

Rock, Daniel, "Engineering value: The returns to technological talent and investments in artificial intelligence," *Available at SSRN 3427412*, 2019.

Solow, Robert M, "Technical change and the aggregate production function," *The Review of Economics and Statistics*, 1957, pp. 312–320.

Sui, Xiaomei, "Uneven firm growth in a globalized world," *Available at SSRN 4208473*, 2021.

Sun, Qi and Mindy Z Xiaolan, "Financing intangible capital," *Journal of Financial Economics*, 2019, *133* (3), 564–588.

Tambe, Prasanna, Lorin Hitt, Daniel Rock, and Erik Brynjolfsson, "Digital capital and superstar firms," Technical Report, National Bureau of Economic Research 2020.

Violante, Giovanni L, "Skill-biased technical change," *The New Palgrave Dictionary of Economics*, 2008, *2*, 1–6.

Wooldridge, Jeffrey M, "On estimating firm-level production functions using proxy variables to control for unobservables," *Economics letters*, 2009, *104* (3), 112–114.

Yang, Chih-Hai, "How artificial intelligence technology affects productivity and employment: Firm-level evidence from Taiwan," *Research Policy*, 2022, *51* (6), 104536.

Zhang, Lichen, "Intangible-Investment-Specific Technical Change, Concentration and Labor Share," 2019.

Appendices

A Data and Sample Statistics

A.1 Industry-Level Deflation

We deflate all nominal variables with industry-level producer price indices. First, we use the Bureau of Labor Statistics (BLS) Producer Price Index industry data, which report the net-output price index of each 3- to 6-digit NAICS industry. We use the industry total series taken as annual averages and rebased to 2012 = 100. Second, we match each Compustat firm to its industry hierarchically as we first attempt a 6-digit NAICS match and then progressively coarsen to the 5-, 4-, and 3-digit level.[41] Third, for the small set of service industries for which the BLS does not publish a producer price index, we use the aggregate CPI rebased to the same base year.[42,43] The nominal control variables are deflated with the same industry index.

A.2 Sample Statistics

Table A.1 presents the firm-level constructed variables and their definitions, together with the Compustat labels from which each is built. Table A.2 summarizes the sample-selection procedure, which reports the firm-years and distinct firms surviving each step in the order in which the steps are applied. The resulting 8,603 firm-years on 1,213 firms are the sample on which the pooled DiD specification for the value-added labor productivity is estimated.

Table A.3 reports the distribution of observations across years and industries. The sample is relatively balanced over time such that each year accounts for between 10.1% and 12.0% of the observations, so no single year drives the estimates. Industry composition reveals that Business Equipment is the largest group with 30.8% of firm-years, which is

[41] Because BLS publishes these series on the 2022 NAICS vintage while Compustat carries earlier vintages, we apply a concordance for the reclassified ICT industries (e.g., software publishers and internet publishing).

[42] Over our 2010-2019 estimation window, 85.6% of firm-year observations are deflated with their own-industry producer price index and the remainder with the aggregate measure, and hence no observation is left undeflated.

[43] The economic difference is substantial. For example, the net-output price index of software publishers falls over the sample period while the aggregate CPI rises by roughly 18%, so industry-level deflation reveals real output growth in ICT-intensive industries that CPI deflation absorbs into prices.

expected given that AI-skilled labor is concentrated in technology-producing industries. Manufacturing, Wholesale & Retail and Healthcare each contribute around 14-17%, and hence the estimation sample is not only confined to the technology-intensive sectors.

Table A.1: Firm-Level Constructed Variables and Descriptions

Variable	Definition	Reference
Panel A: Outcome variables		
Value-added labor productivity	$\ln\left(\frac{\text{Operating income before depreciation } (oibdp)}{\text{Employees } (emp)}\right)$	
Intangible capital	$\ln(\text{Intangible capital } (K^{int}) + 1)$	Peters and Taylor (2017)
Panel B: Treatment variables		
AI Share_{2015}	The firm's share of AI workers in 2015, the year of the TensorFlow release	
Post	$\mathbf{1}[\text{year} \geq 2016]$; 2015 is the base year and is excluded from the baseline pooled DiD specifications	
Post $\times$ AI Share_{2015}	Interaction of Post with AI exposure	
Panel C: Firm characteristics and controls		
ln(R&D/Sales)	$\ln\left(\frac{\text{R\&D expenditure } (xrd) + 1}{\text{Sales } (sale) + 1}\right)$	
ln(Capex)	$\ln(\text{Capital expenditure } (capx) + 1)$	
Cash/Assets	$\frac{\text{Cash and short-term investments } (ch)}{\text{Assets } (at)}$	
Firm age	$\ln(\text{year} - \text{founding year} + 1)$	Founding years are from Jay R. Ritter's Website
Markup	$0.85 \frac{\text{Sales } (sale)}{\text{Cost of goods sold } (cogs)}$, an accounting approximation in which the constant proxies for the output elasticity of the variable input; the structurally estimated counterpart is described in Section D	De Loecker et al. (2020)
Dividends/Assets	$\frac{\text{Common dividends } (dvc)}{\text{Lagged assets } (at_{t-1})}$	
Sales Growth	Year-on-year growth rate of *sale*	
Panel D: Deflation and classification		
Industry price deflator	Bureau of Labor Statistics producer price index for the firm's own NAICS industry (net output, annual average, 2012 = 100), matched hierarchically from the 6- to the 3-digit level, with the aggregate CPI as a fallback; see Appendix A.1	
Industry classification	Fama-French 12-industry groups constructed from 4-digit SIC codes	

Note: This table summarizes the firm-level constructed variables used in the analysis and their definitions. Compustat labels are given in italics. All nominal dollar variables are deflated with the industry-level producer price index described in Panel D and are expressed in 2012 prices.

Table A.2: Sample Selection

Selection step	Firm-years	Firms	Dropped
Panel A: Construction of the firm-year panel			
Compustat North America, annual industrial file	516,431	39,360	
Retain observations denominated in U.S. dollars	516,431	39,360	0
Retain the industrial-format record	516,431	39,360	0
Drop financial firms (SIC 6000–6999)	384,052	27,172	132,379
Drop utilities (SIC 4900–4999)	360,485	26,305	23,567
Drop public administration (SIC $\geq$ 9000)	353,517	25,534	6,968
Drop observations with a missing firm identifier	353,517	25,534	0
Drop duplicate firm-year records	352,636	25,534	881
Panel B: From the estimation window to the estimated sample			
Restrict to the 2010–2019 estimation window	52,638	8,671	299,998
Require non-missing value added and employment (outcome ln(oibdp/emp), positive values only)	30,392	5,170	22,246
Require a non-missing 2015 AI-worker share (treatment exposure)	15,969	1,982	14,423
Drop the 2015 base year	14,274	1,974	1,695
Require non-missing values for all eight lagged controls	8,705	1,304	5,569
Drop observations with non-finite values in the outcome or controls	8,692	1,302	13
Drop firms observed only once (absorbed by the firm fixed effect)	8,603	1,213	89
Final pooled difference-in-differences sample	8,603	1,213	

Note: This table reports the number of firm-year observations and distinct firms remaining after each step of the sample-selection procedure, in the order in which the steps are applied. Panel A starts from the raw Compustat North America annual industrial file and constructs the firm-year panel; the "Dropped" column reports the firm-years removed by each step. Panel B begins from the 2010-2019 estimation window and applies the data requirements of the baseline specification for the value-added labor productivity. Because we do not impute missing values, the requirements in Panel B bind jointly and the resulting sample is the one on which every pooled specification is estimated. The final two steps are imposed by the estimator: 13 observations carry non-finite values (nine in the outcome and four in lagged firm age) and cannot enter the regression, and 89 firms observed exactly once are absorbed by the firm fixed effect and contribute no identifying variation. The final row corresponds to the estimation sample reported in Table 1.

Table A.3: Distribution of Observations Across Years and Industries

	Firm-years	% of sample
Panel A: Distribution across years		
2010	919	10.68
2011	955	11.10
2012	971	11.29
2013	997	11.59
2014	1,032	12.00
2016	996	11.58
2017	944	10.97
2018	923	10.73
2019	866	10.07
Panel B: Distribution across Fama-French 12 industries		
Business Equipment	2,646	30.76
Manufacturing	1,453	16.89
Wholesale & Retail	1,361	15.82
Healthcare	1,214	14.11
Consumer Durables	465	5.41
Other	422	4.91
Chemicals	395	4.59
Consumer Non-Durables	352	4.09
Telephone & Television	152	1.77
Oil, Gas & Coal	143	1.66
Total	8,603	100.00

Note: This table shows the distribution of observations across years and industries. Industries follow the Fama-French 12-industry classification constructed from 4-digit SIC codes. Percentages are shares of the 8,603 firm-year observations.

Table A.4 presents summary statistics for the main variables. Firms exhibit substantial heterogeneity in size, investment, productivity, intangible capital, and the share of AI-skilled workers. The interquartile range of total assets spans $0.5 billion to $6.3 billion, and log value added per employee has a standard deviation of 1.19, which indicates that the estimation sample contains wide variation in the outcome. The share of AI workers averages 0.50 percentage points but is strongly right-skewed such that the median firm employs 0.31% AI workers while the 75th percentile is twice that. It is precisely this dispersion in pre-shock AI exposure that we use for identification.

Table A.5 reports pairwise correlations among key firm characteristics and highlights the joint variation across firm size, investment, productivity, and intangible capital.

Table A.4: Summary Statistics for the Main Variables

Variable	*N*	Mean	S.D.	p25	Median	p75
Panel A: Outcome and treatment						
ln(value added / employee)	8,603	3.63	1.19	2.95	3.66	4.34
Share of AI Workers (p.p.)	8,603	0.50	0.74	0.15	0.31	0.63
AI Exposure: 2015 Share of AI Workers (p.p.)	8,603	0.58	0.68	0.22	0.37	0.70
Panel B: Intangible Capital						
ln(Intangible Capital)	8,594	7.03	1.86	5.75	6.97	8.25
ln(Intangible Capital / L.Assets)	8,594	-0.40	0.69	-0.70	-0.30	0.00
Panel C: Lagged controls						
L.ln(Assets)	8,603	7.44	1.93	6.08	7.36	8.68
L.ln(Capital Expenditure)	8,603	4.00	1.96	2.63	3.93	5.25
L.ln(R&D/Sales)	8,603	-4.04	2.25	-5.17	-3.55	-2.28
L.Cash / Assets	8,603	0.14	0.12	0.05	0.10	0.19
L.ln(Firm Age)	8,603	3.44	0.69	3.00	3.47	3.91
L.Markup	8,603	2.33	8.03	1.19	1.43	2.08
L.Dividends / Assets	8,603	0.02	0.05	0.00	0.00	0.02
L.Sales Growth	8,603	0.10	0.34	-0.02	0.06	0.15
Panel D: Firm size (levels)						
Total Assets ($m, 2012 prices)	8,603	11,320	36,377	489	1,693	6,261
Employees (000s)	8,603	23.34	50.80	1.60	5.90	19.80
Market Value ($m, 2012 prices)	7,960	12,627	49,377	594	1,905	6,903

Note: The table documents the summary statistics for the main variables.

Table A.5: Pairwise Correlations Among Key Firm Characteristics

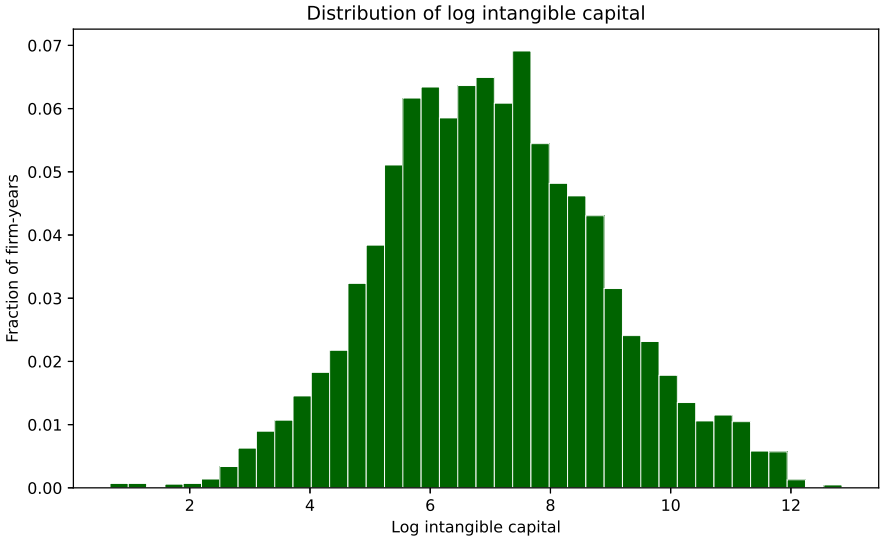
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) ln(VA / Employee)	1.00								
(2) Share of AI Workers	0.14***	1.00							
(3) ln(Intangible Capital)	0.37***	0.06***	1.00						
(4) L.ln(Assets)	0.39***	0.02*	0.94***	1.00					
(5) L.ln(Capex)	0.29***	-0.02*	0.82***	0.92***	1.00				
(6) L.ln(R&D/Sales)	0.27***	0.34***	-0.08***	-0.19***	-0.26***	1.00			
(7) L.Cash/Assets	0.01	0.22***	-0.24***	-0.31***	-0.32***	0.37***	1.00		
(8) L.Markup	0.10***	0.04***	0.00	-0.03***	-0.06***	0.12***	0.10***	1.00	
(9) L.Sales Growth	0.11***	0.06***	-0.07***	-0.09***	-0.10***	0.10***	0.09***	0.05***	1.00

Note: The table presents the pairwise correlations among the key firm characteristics. * $p < .10$, ** $p < .05$, *** $p < .01$.

Figure A.1 shows the histogram of the logarithm of intangible capital. The distribution is approximately symmetric and wide, which spans 4.7 log points between the 10th

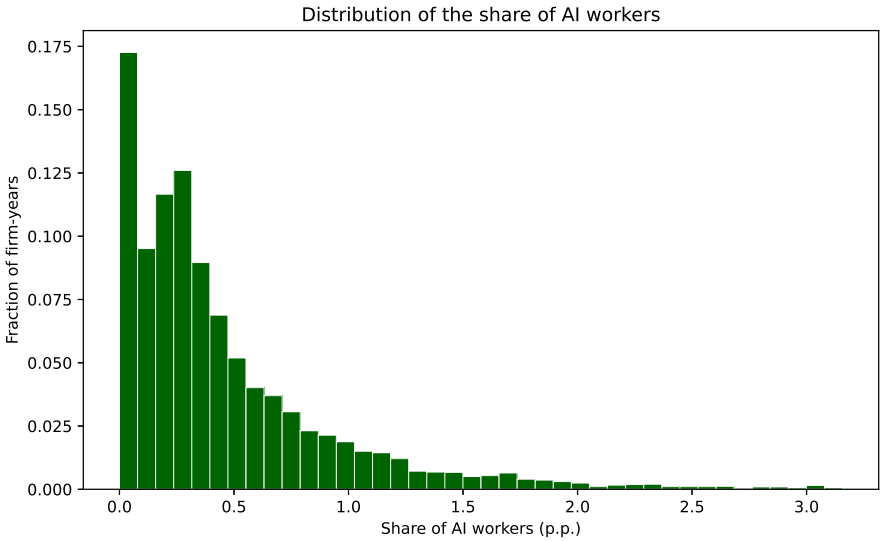
and 90th percentiles. Therefore, the sample is not concentrated at either extreme of intangible capital.

Figure A.1: Distribution of Log Intangible Capital



Note: This figure plots the distribution of the logarithm of intangible capital ($\ln(K^{\text{int}}+1)$), where K^{int} is the Peters and Taylor (2017) measure of firm-level intangible capital. Bar heights are fractions of firm-year observations.

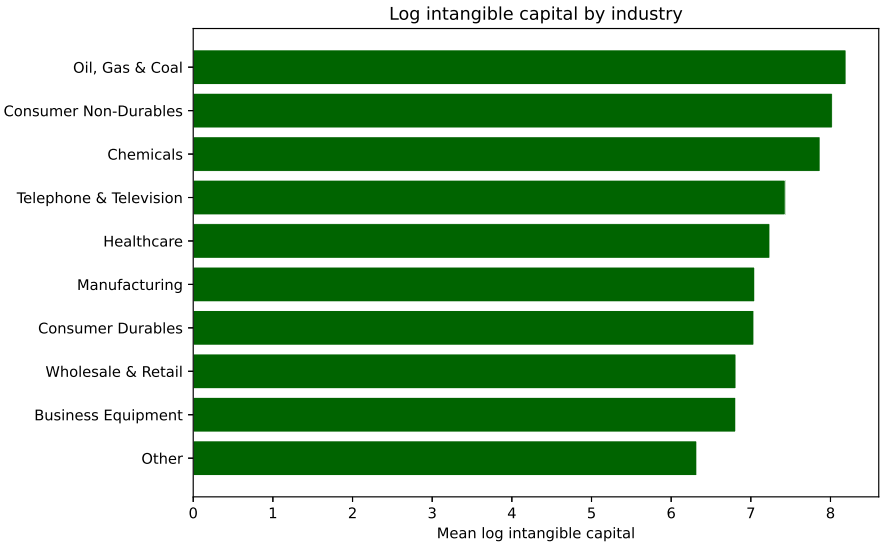
Figure A.2: Distribution of the Share of AI Workers



Note: This figure plots the distribution of the firm-level share of AI-skilled workers (in percentage points) from Babina et al. (2024). Bar heights are fractions of the displayed firm-year observations.

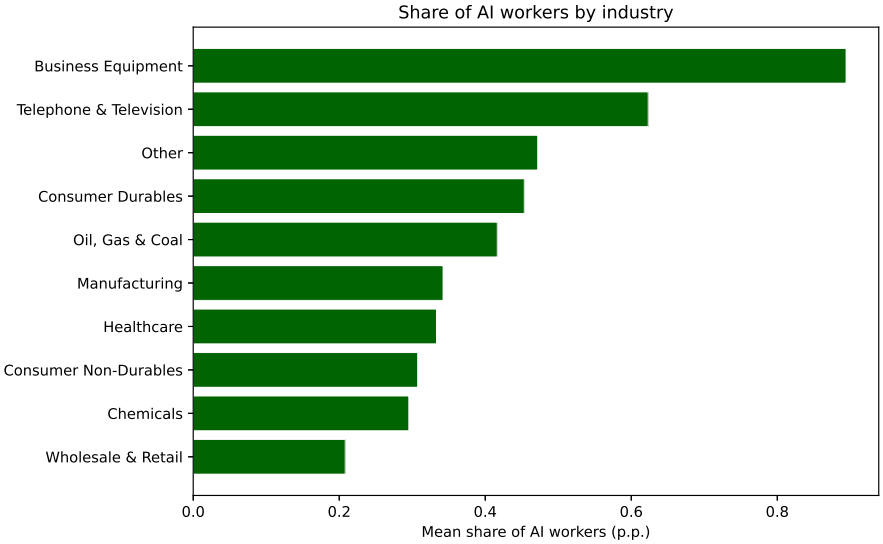
Figure A.2 presents the histogram of the firm-level share of AI workers and reveals a pronounced right skew such that the modal firm has a very small but nonzero share, and the median firm with 0.31%. A long upper tail extends well beyond 3%.

Figure A.3: Log Intangible Capital by Industry



Note: This figure plots the mean of $\ln(K^{\text{int}} + 1)$ by Fama-French 12 industry, constructed from 4-digit SIC codes.

Figure A.4: Share of AI Workers by Industry



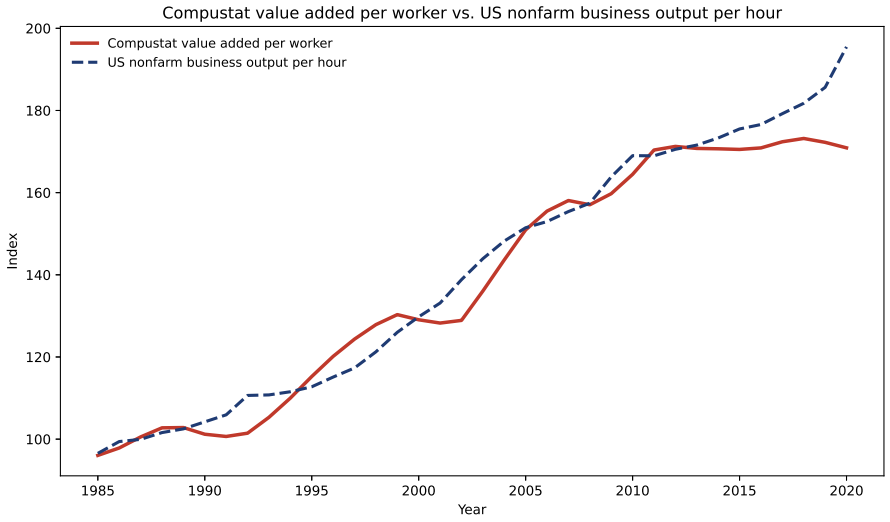
Note: This figure plots the mean firm-level share of AI-skilled workers (in percentage points) by Fama-French 12 industry, constructed from 4-digit SIC codes.

Figures A.3 and A.4 report industry-level averages of the logarithm of intangible capital and of the share of AI workers. Both display pronounced cross-industry heterogeneity.

B Comparison with Aggregate Productivity Data

Because our productivity measures are constructed from Compustat firm-level data, a natural question is whether they are able to track the productivity dynamics in aggregate statistics. To answer it, we compare our value-added productivity measure with the official U.S. labor productivity series over 1985-2020.

Figure B.1: Compustat Value Added per Worker and US Nonfarm Business Labor Productivity



Note: This figure plots the chained median-firm index of Compustat value added per worker (the yearly median of within-firm log changes in deflated operating income before depreciation per employee, cumulated and smoothed with a three-year centered moving average; full Compustat universe) and real output per hour in the U.S. nonfarm business sector (BLS, FRED series *OPHNFB*, annual averages). Both series are indexed to their 1985-1989 average.

For the comparison, on the Compustat side, we use the full universe of U.S. firms (excluding utilities, financial firms, and government entities, as in our estimation sample). In line with our baseline measure, we compute value added per employee exactly as operating income before depreciation, deflated to constant dollars and divided by employment, which is restricted to positive values. We then aggregate these firm-level observations into an annual index in three steps. First, we compute each firm's within-firm log change in value added per worker between consecutive years, so that the index reflects the productivity growth of continuing firms and removes the direct effect of entry, exit, and the changing level composition of listed firms.[44] Second, we take the median of these within-firm changes in each year to make the index robust to the extreme observations, and we chain the median into a level index. Third, we smooth the index with a three-year centered moving average to remove

[44] Each annual median is nonetheless computed over a potentially different set of firm pairs, so selective entry, exit, missingness, or the positive-value restriction can still affect the index.

residual year-to-year sampling noise. On the aggregate side, we use real output per hour in the U.S. nonfarm business sector (Bureau of Labor Statistics, and we retrieve it from FRED series *OPHNFB*). We annualize it by averaging quarterly observations. To avoid dependence on a single base year, we index both series based on their 1985-1989 average.

Figure B.1 shows that the two series move together closely. We find that the correlation between the index levels is 0.98, and the major turning points over time (the early-1990s slowdown, the late-1990s IT boom, the early-2000s IT burst, and the post-2003 surge) appear in both series. The paths also coincide almost exactly until the mid-2010s and diverge only at the very end of the window, which can be the case since the 2019 COVID episode is likely to depress the operating income of listed firms. Therefore, this comparison provides suggestive evidence that the value-added productivity measure in our analysis is likely to capture similar productivity dynamics as the official aggregate statistics, which supports both its construction and the external relevance of our firm-level results.

C Discussion on Google TensorFlow Shock

In this section, we want to discuss the nature and importance of TensorFlow shock in more detail and highlight its two features which reinforces its interpretation as an identifying shock.

First, its release was largely unanticipated. Figure C.1, Panel (a), shows that public attention to TensorFlow was essentially flat before its release in November 2015 and jumped sharply at the moment of the open-source release. This pattern is consistent with a discrete and unexpected shock rather than a gradual and pre-announced diffusion.[45] It also supports the argument that firms' 2015 AI worker intensity was predetermined and not chosen in anticipation of the release.[46]

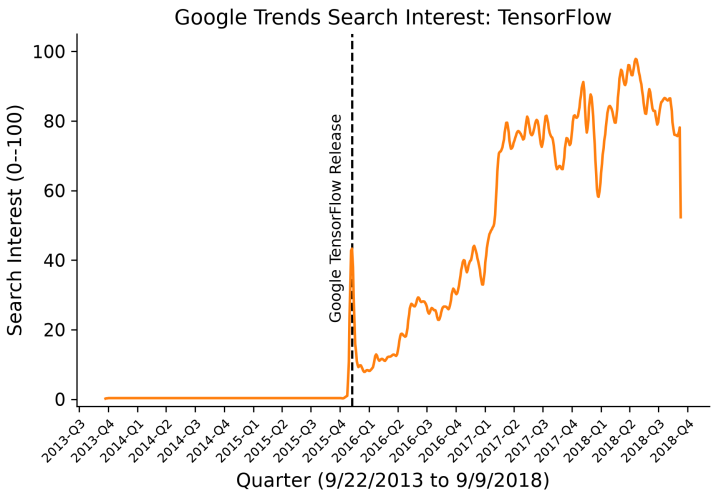
Second, TensorFlow rapidly became prominent among machine-learning frameworks. Figure C.1, Panel (b), reports cumulative GitHub star counts, which rose quickly relative to alternative libraries. This measures developer interest rather than framework usage

[45] We note that exact-name search interest establishes the timing and salience of the public release, but cannot by itself rule out anticipation of comparable open-source AI tools before the TensorFlow name became public.

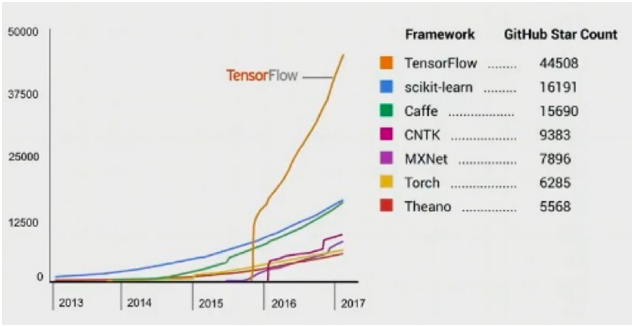
[46] The stronger support for predetermination comes from the specifications in Section E.9, which measure exposure entirely before the release.

or firm adoption, but together with the workforce evidence of Section 7 it indicates that the release was salient and rapidly taken up.

Figure C.1: Google TensorFlow: Public Attention and Machine-Learning Library Usage



(a) Public attention



(b) Usage relative to other libraries

Note: Panel (a) is constructed from Google Trends data, following Daniel Rock's presentation slides here, and plots weekly search interest for the term "TensorFlow" (index 0-100); the vertical dashed line marks the open-source release of Google TensorFlow in November 2015. Panel (b) reports GitHub star counts across major machine-learning frameworks and is adapted from the same slides, which draws on Jeff Dean's 2018 talk at MIT. Further details are available in Daniel Rock's presentation slides.

A growing body of work highlights the TensorFlow as a major point in the AI diffusion. Using patent-level data, Chen et al. (2024) show that the economic value of AI-related innovations rises dramatically after the TensorFlow release, with particularly strong effects in sectors where shared machine-learning frameworks enhance knowledge spillovers. Related evidence from Chen et al. (2025) documents a pronounced post-2015 surge in firms' demand for AI-related skills such that references to TensorFlow are virtually nonexistent prior to 2015 (only a single mention between 2010 and 2015) but increase sharply to more than 26,000 mentions between 2016 and 2021. Leveraging the timing of TensorFlow's release in a DiD design, they provide causal evidence

that AI adoption leads to improvements in climate-related innovation, environmental management, and firms' disclosure practices.

In our context, this shock is expected to raise the marginal product of AI-skilled workers by facilitating faster model development, easier implementation, and more efficient experimentation. Because TensorFlow primarily improves the effectiveness of pre-existing AI talent, we argue that firms with higher 2015 share of AI-skilled workers are plausibly more exposed to the shock. This feature motivates the use of AI worker intensity as a measure of firm-level exposure.

Importantly, we do not interpret the TensorFlow release as generating a purely mechanical increase in labor productivity. If the shock operated solely through task automation, the returns would depend primarily on a firm's task composition rather than on its stock of complementary capital.[47] Instead, this framework suggests that any productivity response should depend on firms' ability to deploy AI technologies within their existing production structure. In particular, intangible capital such as data assets, software, and organizational infrastructure provides the platform through which improvements in AI effectiveness can be integrated into firm operations. This reasoning motivates our focus on heterogeneity across firms with different levels of intangible capital and scale in the empirical analysis.

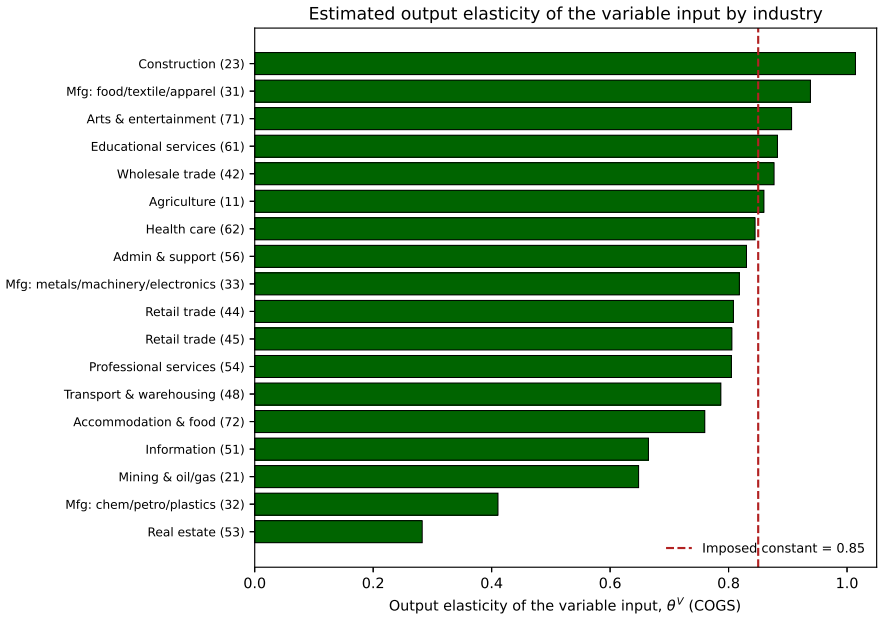
D Structural Markup Estimation

Estimation. In the baseline analysis, following the insight by De Loecker et al. (2020), our markup control is the accounting approximation $0.85 \times (\text{sales per cost of goods sold})$, in which the constant 0.85 proxies for the output elasticity of the variable input. To assess whether this proxy matters, we implement the production function approach of De Loecker et al. (2020). For each 2-digit NAICS industry we estimate a gross output production function in the variable input (COGS), capital, and labor using a control function specification with investment as the productivity proxy. We define the structural markup as the estimated industry output elasticity times the firm's sales-to-cost of goods sold ratio. Figure D.1 plots the estimated elasticities by industry. We find that the industry-level elasticities cluster around the imposed value with a cross-industry mean of 0.78 and median of 0.81, although they display material dispersion across

[47] We do not claim that a pure automation channel would produce uniform firm-level effects. Automation returns can vary with task composition, labor costs, workflows, fixed implementation costs, and scale, and automation may itself require complementary data and software.

industries. This evidence therefore supports the central tendency of the accounting proxy.

Figure D.1: Estimated Output Elasticities of the Variable Input by Industry



Note: This figure plots the output elasticity of the variable input (COGS) estimated by 2-digit NAICS industry from a gross-output production function using a control-function specification with investment as the productivity proxy, following De Loecker et al. (2020). The dashed line marks the value 0.85 imposed by the accounting markup proxy.

Markups as an outcome. In the baseline analysis, we use the structurally estimated markup as an alternative control variable, and we show that the main coefficient of interest in our baseline result is not driven by how markup is calculated. However, this does not still eliminate a potential channel in which markup can also respond to our shock. In this part, we implement a test that puts the markup as an outcome variable. We conjecture that if improved AI effectiveness raises productivity by letting exposed firms charge more, then exposed firms should also show a post-2015 increase in their markups relative to less exposed firms.

Table D.1 reports this analysis. Panel A re-estimates the baseline pooled specification with the markup as the outcome. The coefficient is statistically insignificant for either structurally estimated markup or markup measure with accounting proxy. Hence, a markup response is not supported by our analysis. Figure D.2 shows the dynamic DiD counterpart and confirms our pooled DiD evidence. We find that the pre-treatment leads are jointly insignificant and there are no post-release statistically significant estimates.

Table D.1: Markups as an Outcome and Heterogeneity by Industry Concentration

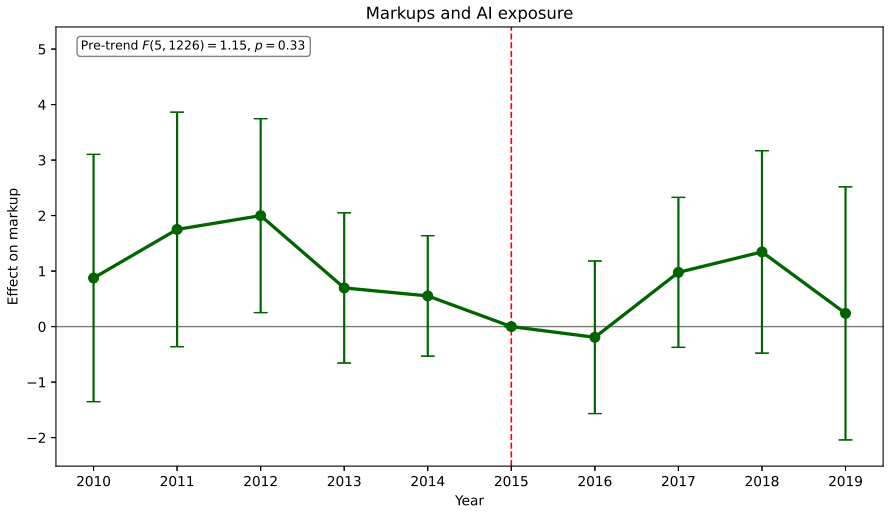
Panel A: the markup as the outcome		
	(1) ln(structural DLEU markup)	(2) ln(accounting 0.85 markup)
Post×AI Share$_{2015}$	-0.507	-0.480
	(1.276)	(1.282)
Observations	8,602	8,602

Panel B: productivity effect by 2015 industry concentration			
	Split sample		Interacted
	(1) Low HHI	(2) High HHI	(3) All
Post×AI Share$_{2015}$	6.21*	7.32	8.95**
	(3.73)	(5.06)	(3.61)
× high HHI			-4.67
			(4.46)
Observations	4,290	4,170	8,460

Note: Panel A re-estimates the baseline pooled difference-in-differences with the markup as the outcome rather than as a control. Column (1) uses the logarithm of the structural De Loecker et al. (2020) markup, the industry output elasticity of the variable input times the firm's sales-to-COGS ratio; column (2) uses the accounting proxy that the paper employs as a control. The markup control is necessarily dropped from the right-hand side. The sample is the pooled estimation sample of Table 1 less one firm-year, for which capital expenditure is missing and the production function underlying the structural markup therefore cannot be estimated. Panel B splits firms at the median 2015 Herfindahl–Hirschman index (HHI) of their 3-digit NAICS industry, computed from Compustat sales over industries with at least 5 firms. Columns (1) and (2) estimate the baseline specification within each half; column (3) pools them and interacts the treatment with an indicator for the concentrated half. Standard errors clustered by firm. * $p < .10$, ** $p < .05$, *** $p < .01$.

Panel B approaches the same question from a different angle. We argue that if the baseline analysis response incorporates a pricing response, then it would be larger where firms have more pricing power. To test this argument, we split firms at the median 2015 Herfindahl-Hirschman index (HHI) of their 3-digit NAICS industry, and re-run our baseline DiD specification for productivity as an outcome variable for low and high HHI industry firms. We find that the main coefficient of interest is 6.21 and statistically significant in low HHI and 7.32 and statistically insignificant in high HHI industries, and the interaction term between our baseline running variable and high HHI dummy variable (column 3) is statistically insignificant, which suggests that the main interest of coefficient is statistically indistinguishable between low and high HHI industries. Hence, we do not find suggestive evidence that our baseline result is driven by industries with different degree of market concentration.

Figure D.2: Markups and AI Exposure: Dynamic Difference-in-Differences



Note: This figure plots the estimated coefficients from a dynamic DiD regression of the logarithm of the structural De Loecker et al. (2020) markup on the interaction between the firm's 2015 share of AI workers and year indicators, with 2015 as the omitted reference year, marked by the dashed line. Coefficients are in log points per percentage point of 2015 exposure. Control variables are the lagged controls of the baseline specification excluding the markup, which is the outcome here. The specification includes firm and year fixed effects. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

We provide these results as suggestive evidence that markup measures which can proxy market power does not seem to drive our baseline results. However, we do not claim that our data would fully rule out this channel. Our calculation of HHI is computed from Compustat that covers listed firms only, so it measures concentration among the firms we observe rather than concentration in the product market.

E Additional Robustness Checks

To further assess the robustness of our causal findings, we conduct an extensive series of exercises to address a distinct threat to underlying identification or measurement.

E.1 Binary DiD Specification

Our baseline specification utilizes AI exposure continuously, and this approach imposes that productivity responds linearly to the 2015 AI-worker share. There can be a natural consideration that if the estimated effect reflects a causal dose-response relation, replacing the continuous treatment with binary exposure groups should deliver coefficients that increase monotonically with exposure.

We therefore re-estimate our baseline model (1) by substituting exposure dummies for the continuous share in two distinct ways. First, we interact Post_t with an indicator for

above-median 2015 AI-worker share. Second, we sort firms into quartiles of the 2015 AI-worker share and interact Post_t with indicators for the second, third, and fourth quartiles, and we take the bottom quartile as the omitted reference group.

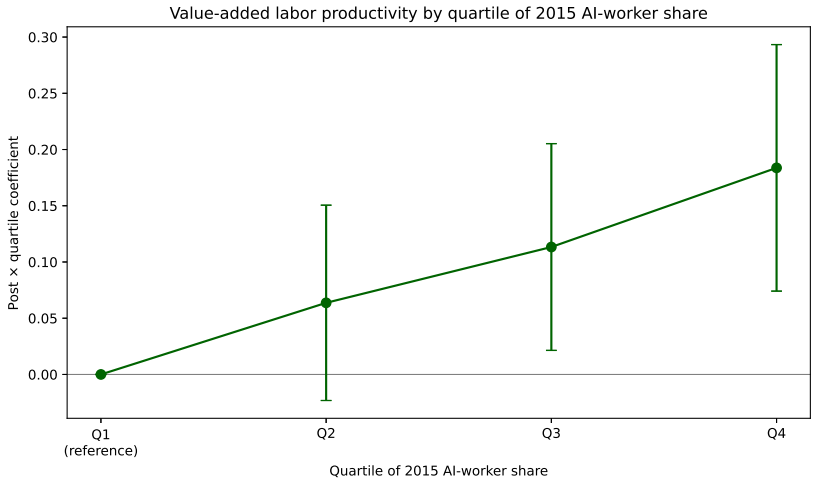
Table E.1: Pooled DiD Estimates: Value-Added Labor Productivity, AI-Exposure Dummies

	(1)
Panel A: High vs. low AI exposure (median split)	
Post×High AI Share_{2015}	0.111***
	(0.039)
Panel B: Quartiles of AI exposure (Q1 = reference)	
Post×Q2	0.064
	(0.053)
Post×Q3	0.113**
	(0.056)
Post×Q4	0.184***
	(0.067)
Controls	✓
Firm FE	✓
Year FE	✓
Observations	8,603

Note: This table reports pooled DiD estimates of the baseline effect with the continuous treatment replaced by exposure dummies. Panel A interacts $\text{Post}_t = \mathbf{1}[\text{year} \geq 2016]$ with an indicator for above-median exposure; Panel B interacts Post_t with indicators for the second, third, and fourth exposure quartiles, with the bottom quartile as the omitted reference. Standard errors clustered by firm in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table E.1 reports the results for value-added labor productivity. In Panel A, we find that firms with above-median exposure have higher productivity by 11.1 log points relative to below-median firms after the shock. Panel B shows that relative to the bottom quartile, the post-shock effect rises from 6.4 log points in the second quartile to 11.3 in the third and 18.4 in the top quartile. Figure E.1 complements it by showing that the implied dose-response profile is broadly monotone in the quartile rank. Because the quartile ranks are equally spaced while the underlying AI-worker share is strongly right-skewed, this exercise reduces concern about outliers and supports an increasing dose response, but it does not test the functional-form restrictions implied by the level-linear specification.

Figure E.1: Pooled DiD Estimates by AI-Exposure Quartile: Value-Added Labor Productivity



Note: This figure plots the estimated coefficients on the interaction between Post_t and indicators for the quartiles of the firm's 2015 share of AI workers, from a pooled DiD regression for the logarithm of value-added labor productivity, with the bottom quartile as the omitted reference group. Quartiles are assigned at the firm level among the firms in the baseline estimation sample. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

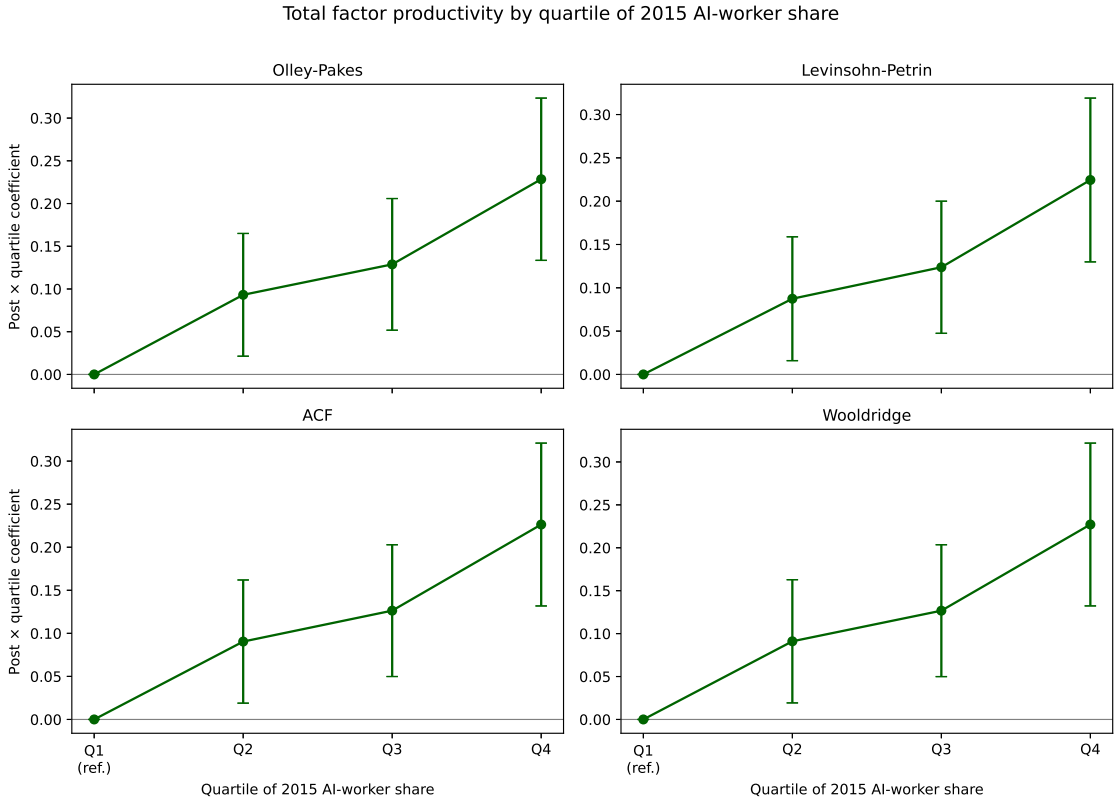
Table E.2: Pooled DiD Estimates: Total Factor Productivity, AI-Exposure Dummies

	(1) OP	(2) LP	(3) ACF	(4) WRDG
Panel A: High vs. low AI exposure (median split)				
Post×High AI Share$_{2015}$	0.124***	0.123***	0.123***	0.124***
	(0.035)	(0.034)	(0.034)	(0.035)
Panel B: Quartiles of AI exposure (Q1 = reference)				
Post×Q2	0.093**	0.087**	0.090**	0.091**
	(0.044)	(0.043)	(0.043)	(0.044)
Post×Q3	0.129***	0.124***	0.126***	0.127***
	(0.047)	(0.046)	(0.046)	(0.047)
Post×Q4	0.228***	0.224***	0.226***	0.227***
	(0.058)	(0.057)	(0.057)	(0.058)
Controls	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	8,603	8,603	8,603	8,603

Note: This table reports pooled DiD estimates of the baseline effect on TFP, with the continuous treatment replaced by exposure dummies. The outcome is $\ln(\text{TFP})$ from a value-added production function estimation. Panel A interacts $\text{Post}_t = \mathbf{1}[\text{year} \geq 2016]$ with an indicator for above-median exposure; Panel B interacts Post_t with indicators for the second, third, and fourth exposure quartiles, with the bottom quartile as the omitted reference. Standard errors clustered by firm in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table E.2 and Figure E.2 repeat the exercise for TFP. The pattern is almost same. We find that the above-median effect is about 12 log points for every TFP estimator, and the quartile coefficients rise monotonically from roughly 9 (second quartile) to 13 (third quartile) and 23 log points (top quartile).

Figure E.2: Pooled DiD Estimates by AI-Exposure Quartile: Total Factor Productivity



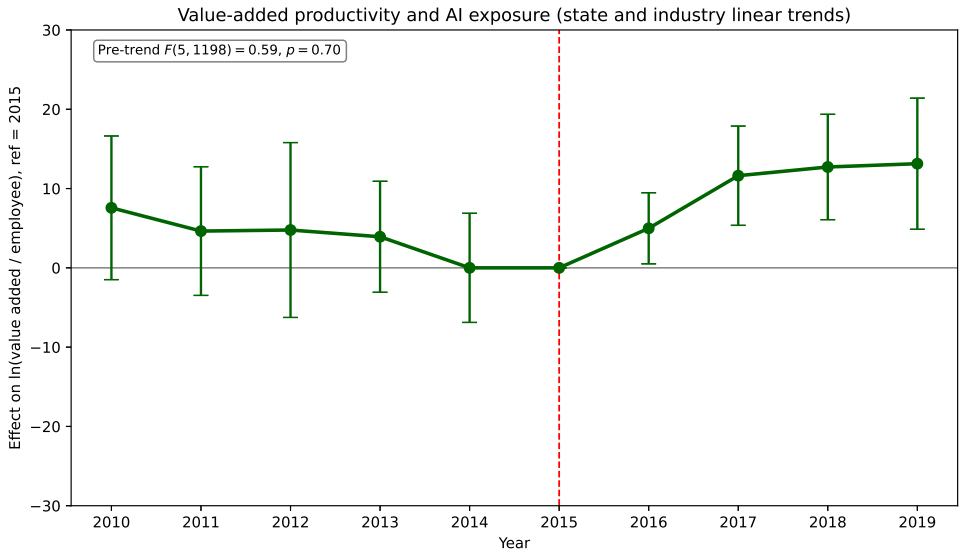
Note: This figure plots the estimated coefficients on the interaction between Post_t and indicators for the quartiles of the firm's 2015 share of AI workers, from pooled DiD regressions for the logarithm of total factor productivity, estimated with four alternative production-function methods. The bottom quartile is the omitted reference group. Quartiles are assigned at the firm level among the firms in the estimation sample. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

These results show that the productivity response increases monotonically, which suggests that our baseline results are not driven by the linear functional form.

E.2 State and Industry Linear Trends

One concern would be that firms with higher AI exposure may be located in regions or industries on systematically different secular trajectories. To address this concern, we augment the benchmark specification with state-specific and industry-specific linear time trends. Figure E.3 shows that the dynamic DiD estimates remain strong with post-2016 effects of 11 to 13 log points and a flat and statistically insignificant pre-period coefficients.

Figure E.3: Dynamic DiD Estimates with State and Industry Linear Trends

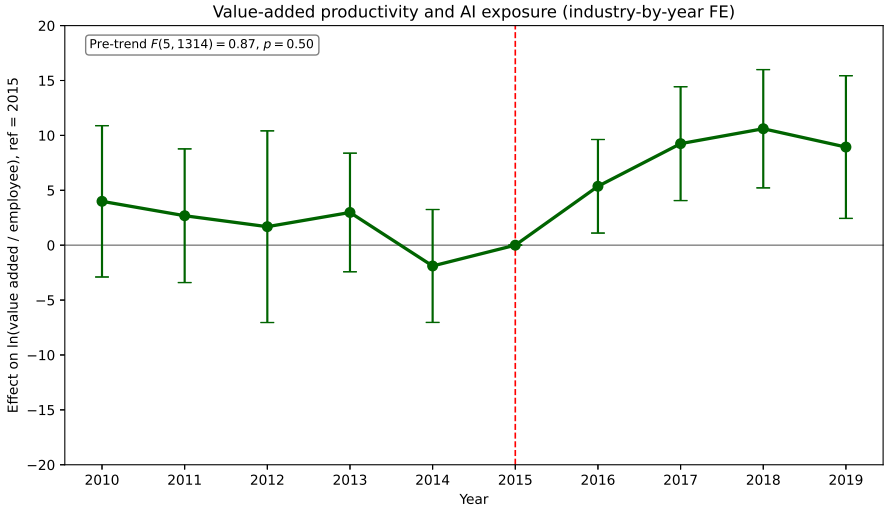


Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity, augmented with state-specific and industry-specific linear time trends. The coefficients represent the interaction between the firm's 2015 share of AI workers and year indicators, with 2015 as the reference year. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

E.3 Industry-by-Year and State-by-Year Fixed Effects

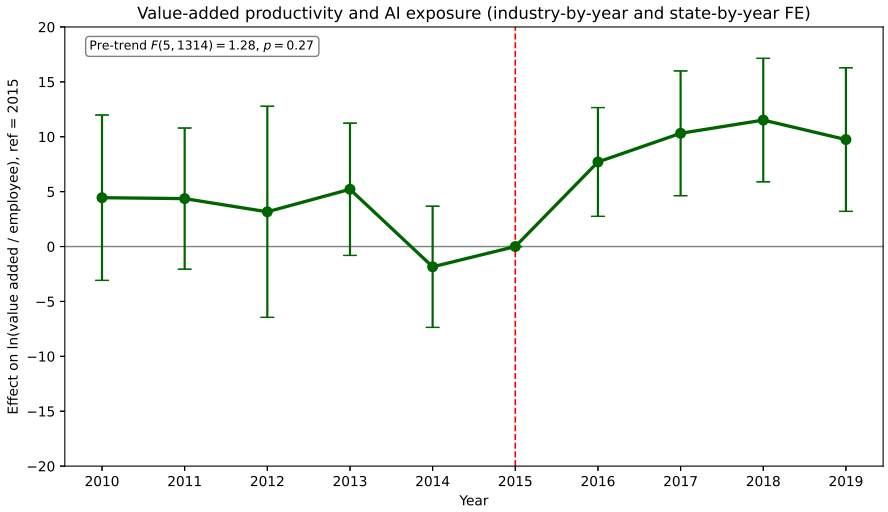
To absorb time-varying sectoral and regional shocks, in this section, we replace the year fixed effects with industry-by-year fixed effects, and then additionally with state-by-year fixed effects. The dynamic DiD estimates in Figures E.4 and E.5 retain the statistically insignificant pre-period and the post-period significant response.

Figure E.4: Dynamic DiD Estimates with Industry-by-Year Fixed Effects



Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity, replacing year fixed effects with industry-by-year fixed effects. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

Figure E.5: Dynamic DiD Estimates with Industry-by-Year and State-by-Year Fixed Effects

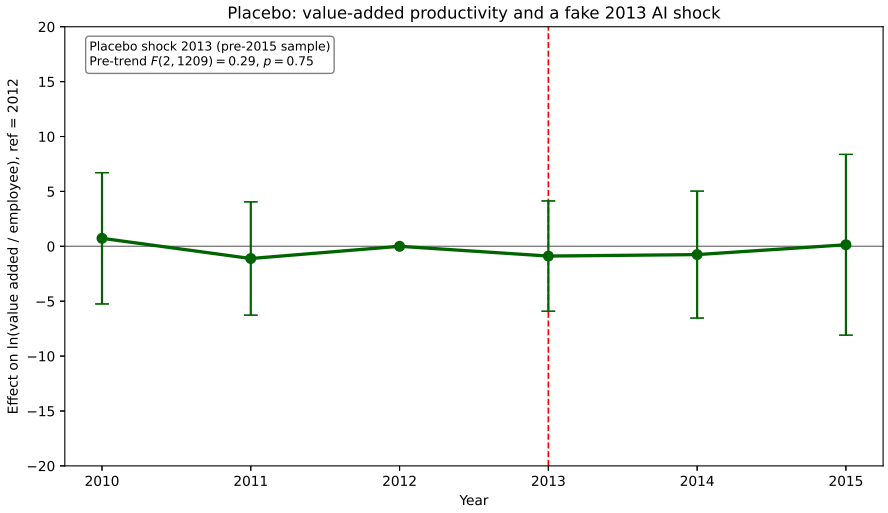


Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity, including firm, industry-by-year, and state-by-year fixed effects. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

E.4 Placebo Test

We conduct a placebo test by redefining the shock year as 2013 and restrict the sample to the pre-TensorFlow period as 2010-2015 so that the true shock cannot contaminate the placebo post-period. Figure E.6 shows that the placebo dynamic DiD estimates are statistically insignificant. This evidence supports that the post-2015 response does not reflect pre-existing differential dynamics of AI-intensive firms.

Figure E.6: Dynamic DiD Estimates: Placebo Shock in 2013

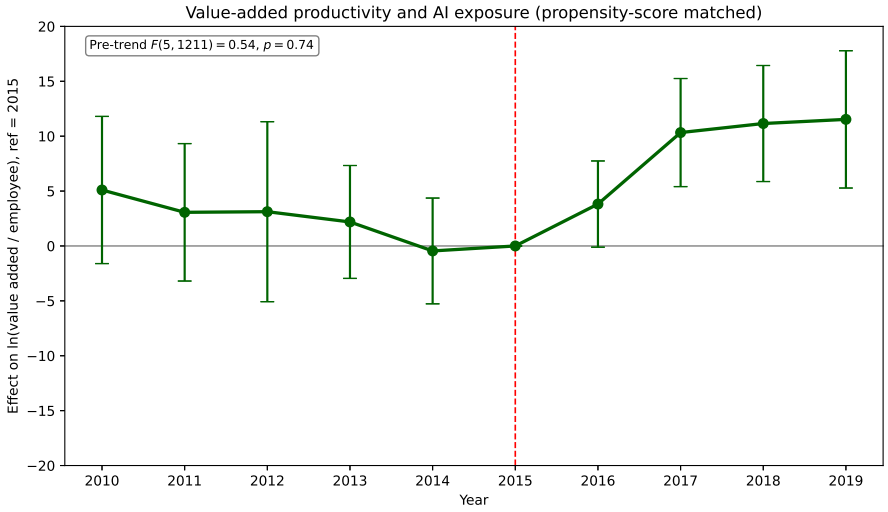


Note: This figure plots the estimated coefficients from a placebo dynamic DiD regression for value-added labor productivity, assigning a counterfactual shock year of 2013 and restricting the sample to 2010-2015. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

E.5 Propensity Score Matching

One potential concern would be that firms with high AI exposure may differ from other firms along observable dimensions that also predict productivity. To address this concern, using a logit propensity score estimated on the full set of pre-treatment firm characteristics, we match each firm with above-median 2015 AI exposure to its nearest neighbor among low-exposure firms.[48] On the matched sample, Figure E.7 shows flat and statistically insignificant pre-trends and similar post-shock significant response. This evidence suggests that potential selection on observables does not seem to explain the baseline results.

Figure E.7: Dynamic DiD Estimates on the Propensity-Score-Matched Sample



Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity on a one-to-one nearest-neighbor propensity-score-matched sample, where firms with above-median 2015 AI exposure are matched to low-exposure firms on the pre-treatment control variables. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

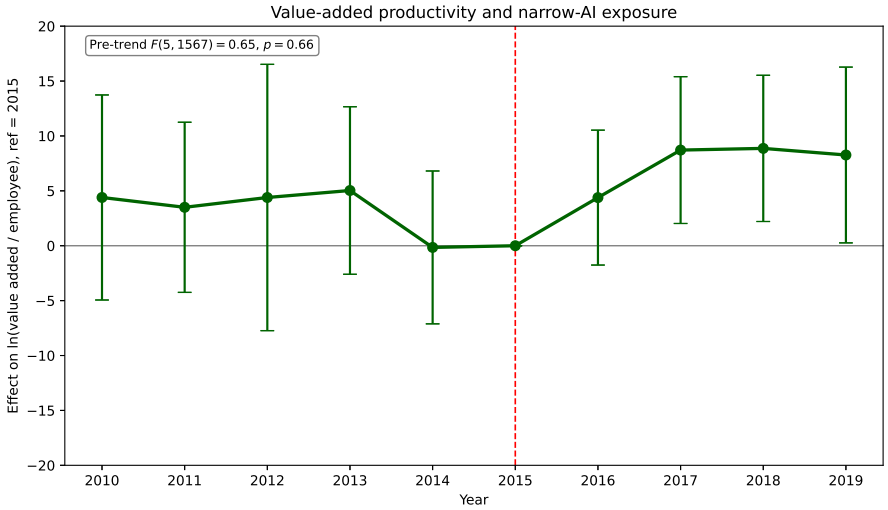
E.6 Alternative AI Measure

We examine whether the results hold under a more stringent definition of AI workers. We replace the broad AI worker share with the narrow AI share of Babina et al. (2024). This alternative definition counts only highly specialized and directly AI-relevant skills. An important note is that the narrow measure is populated for a much smaller set of firms (roughly one-third of the baseline measure). Figure E.8 shows the dynamic DiD estimates with respect to the alternative measure such that the pre-period is flat and statistically insignificant, and we see an overall significant response in the post-period.

[48] The matched sample pairs 693 treated with 693 control firms.

Hence, the underlying pattern in the baseline result is not changed dramatically under the stricter definition of AI skills.

Figure E.8: Dynamic DiD Estimates: Narrow AI Exposure

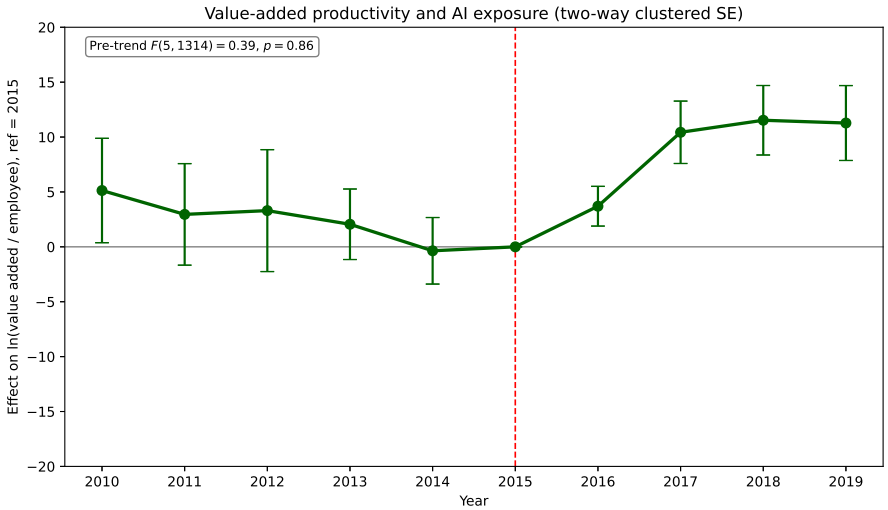


Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity, replacing the broad 2015 AI-worker share with the narrow AI-worker share of Babina et al. (2024). The coefficients represent the interaction between the firm's 2015 narrow-AI share and year indicators, with 2015 as the reference year. Standard errors are clustered at the firm level. Error bars show the 90% confidence interval.

E.7 Two-Way Clustered Standard Errors

Our baseline clusters standard errors at the firm level. Because the shock is a single aggregate event, a concern would be to check arbitrary within year correlation. Regarding this concern, we instead cluster standard errors two-way by firm and year.

Figure E.9: Dynamic DiD Estimates with Two-Way Clustered Standard Errors



Note: This figure plots the estimated coefficients from the baseline dynamic DiD regression for value-added labor productivity with standard errors clustered two-way by firm and year, applying the Cameron et al. (2011) eigenvalue correction to the variance matrix. Error bars show the 90% confidence interval.

Figure E.9 shows the dynamic DiD estimates accordingly with the statistically insignificant joint pre-trends and the overall positive and significant post-period effects.[49]

E.8 Alternative Size Measures

Table E.3: DiD Estimates by Intangible Intensity and Firm Size: Alternative Size Definitions

	(1) Difference vs. baseline	(2) Implied effect
Panel A. Size measured by 2015 total assets		
Post×AI Share$_{2015}$ (small × low intangible, baseline)	—	-5.77
		(5.91)
× Small × high intangible	11.71	5.93
	(8.28)	(6.55)
× Large × low intangible	13.00**	7.23**
	(6.14)	(3.31)
× Large × high intangible	18.37***	12.59***
	(6.75)	(4.44)
Panel B. Size measured by 2015 employment		
Post×AI Share$_{2015}$ (small × low intangible, baseline)	—	0.76
		(5.00)
× Small × high intangible	4.20	4.96
	(7.00)	(5.71)
× Large × low intangible	5.24	6.00*
	(5.30)	(3.13)
× Large × high intangible	16.12***	16.87***
	(6.18)	(4.58)
Firm FE	✓	✓
Year FE	✓	✓
Controls	✓	✓
Observations	8,603	8,603

Note: This table reports the value-added AI effect by intangible intensity and firm size, estimated as four median-split groups in one pooled difference-in-differences with small, low-intangible firms as the omitted baseline, measuring firm size by 2015 total assets (Panel A) and 2015 employment (Panel B) instead of 2015 market value. Firms are split at the estimation-sample median of 2015 intangible intensity and of the (log) 2015 size variable. The Post×AI Share$_{2015}$ treatment (Post$_t$ × 2015 AI-worker share) is interacted with the three non-baseline group dummies. Column (1) reports the baseline-group effect (first row) and each group's effect difference from the baseline; column (2) reports the implied per-group effect (baseline + difference). * $p < .10$, ** $p < .05$, *** $p < .01$.

[49] We apply the eigenvalue correction of Cameron et al. (2011) for the reported confidence intervals.

The complementarity analysis in our baseline specification such as Table 5 splits firms by their 2015 market value. Market value is a natural proxy of firm scale in our context as discussed in Section 2.1, but one potential concern would be that it embeds the market's valuation of the firm's growth prospects such that the size split would partly sort firms on expected performance rather than on scale itself. To address this concern, we use two alternative size measures based on firms' total assets and employment.

Table E.3 re-estimates the baseline design of our complementarity analysis with the alternative size measures of 2015 total assets (Panel A) and 2015 employment (Panel B).[50] The pattern of our baseline Table 5 carries over to both size measures. The small and low intangible baseline group shows no statistically significant effect under either size definition. However, the main coefficient of interest is positive and statistically significant for large and high intangible capital firms with respect to both alternative definitions of firm size. Therefore, we argue that the shown complementarity effect is not driven by the definition of firm size.

E.9 Alternative Exposure Variable

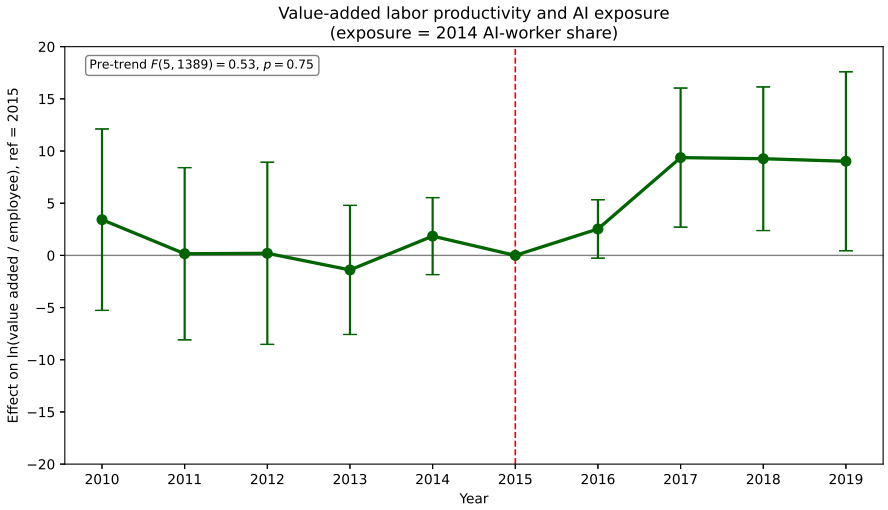
Our baseline exposure measure is the firm's AI worker share in 2015. Because the TensorFlow release occurred in November 2015 and although we discuss in Section C that this release was unexpected, one may still worry that part of the 2015 exposure would reflect firms' hiring in response to the shock. To address this concern, we first use 2014 AI worker share and the mean, median, minimum, and maximum of the firm's AI worker share over 2010-2014 as exposure variables. Second, we project 2015 AI worker share on the firm's 2014 AI worker share.

E.9.1 Prior Exposure Variables

We re-estimate the baseline design with exposure measured before the shock by using the firm's 2014 AI worker share. Figure E.10 plots the corresponding dynamic DiD specification and shows that the pre-treatment coefficients are individually and jointly indistinguishable from zero, the 2016 coefficient remains small, and the significant response emerges in 2017 and persists through 2019.

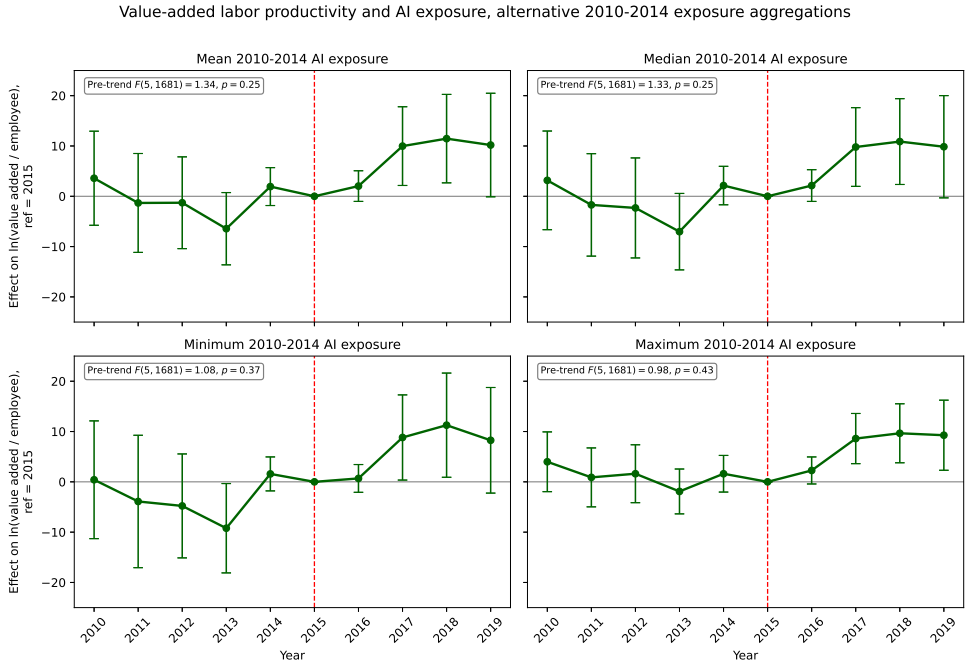
[50] Unlike the market value split, the median is computed within the estimation sample because firms that satisfy the lagged control requirements are systematically larger so that a panel wide median would leave the small firm cells nearly empty.

Figure E.10: Dynamic DiD Estimates with the 2014 AI-Worker Share as Exposure



Note: This figure plots the estimated coefficients from the dynamic DiD regression for value-added labor productivity, measuring exposure by the firm's fully pre-shock 2014 AI-worker share. The reference year is 2015. Standard errors are clustered by firm. Error bars show the 90% confidence interval.

Figure E.11: Dynamic DiD Estimates with Alternative 2010-2014 Exposure Aggregations



Note: This figure plots the estimated coefficients from the dynamic DiD regression for value-added labor productivity, measuring exposure as the mean, median, minimum, and maximum of the firm's AI-worker share over 2010-2014. The reference year is 2015. Standard errors are clustered by firm. Error bars show the 90% confidence interval.

Because a single year exposure would be a noisy snapshot of a firm's AI workforce, we also aggregate the AI worker share over the pre-shock window by measuring exposure as the mean, median, minimum, and maximum over the 2010-2014 period.

Figure E.11 reports the corresponding dynamic DiD estimates. We find that pre-treatment coefficients are jointly statistically insignificant, and we observe positive and statistically significant responses from 2017 onward. As a result, we conclude that the productivity response is not driven by a particular timing choice of AI exposure.

E.9.2 Projected Baseline Exposure Variable

As an alternative approach, we isolate the component of the 2015 exposure that was already in place before the release. We project our running variable on 2014 AI worker share, and re-estimate the design using only the predicted component so that no part of the variation can reflect a response to the release itself.[51] Table E.4 shows that the resulting main coefficient of interest is still positive and statistically significant.

Table E.4: Pooled DiD Estimates for Value-Added Labor Productivity: 2015 Exposure Projected by the 2014 AI-Worker Share

	(1) Baseline Exposure	(2) Predicted Exposure
Post×AI Share$_{2015}$	6.59**	9.21**
	(3.15)	(4.01)
First stage (dependent variable: Post×AI Share$_{2015}$)		
Post×AI Share$_{2014}$	—	0.749***
		(0.140)
First-stage F	—	29
Firm FE	✓	✓
Year FE	✓	✓
Controls	✓	✓
Observations	8,603	8,381

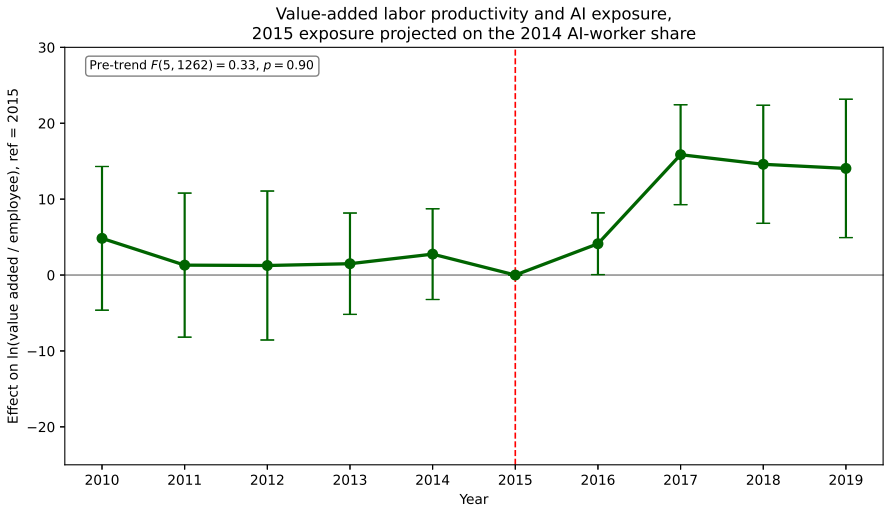
Note: This table reports pooled difference-in-differences estimates for value-added labor productivity. Column (1) is the baseline OLS with the treatment Post$_t$ × 2015 AI-worker share (Post$_t$ = **1**[year ≥ 2016]). Column (2) instruments Post$_t$ × 2015 share with Post$_t$ × the firm's fully pre-shock 2014 AI-worker share, so that only the component of the 2015 exposure predictable from the 2014 share is used for identification. The middle block reports the corresponding first-stage regression of Post$_t$ × 2015 share on Post$_t$ × 2014 share, with the same controls and fixed effects. * $p < .10$, ** $p < .05$, *** $p < .01$.

Figure E.12 reports the corresponding dynamic DiD results. We find that the pre-treatment coefficients are jointly indistinguishable from zero, and the post-treatment

[51] The 2014 share is a strong predictor of the 2015 exposure where the first-stage coefficient is 0.75 with an F-statistic of 29, as documented in Table E.4.

path is similar to the baseline results. Therefore, this analysis provides suggestive evidence that our baseline result is not driven by any component of the exposure measure that could have responded to the shock itself.

Figure E.12: Dynamic DiD Estimates with the 2015 Exposure Projected by the 2014 AI-Worker Share

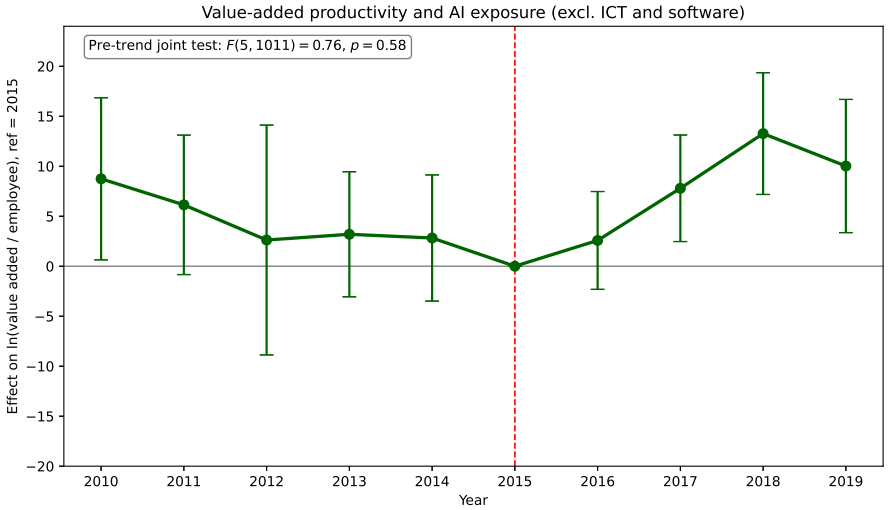


Note: This figure plots two-stage least squares estimates of the dynamic DiD coefficients for value-added labor productivity, instrumenting each year interaction of the 2015 AI-worker share with the corresponding year interaction of the firm's 2014 AI-worker share. The reference year is 2015. Error bars show the 90% confidence interval.

E.10 Excluding AI Intensive Sectors

A natural concern would be that our results might be driven by AI-intensive industries so that the baseline effect would not remain without them.

Figure E.13: Excluding AI Intensive Sectors



Note: The estimation sample excludes computer and electronic product manufacturing (NAICS 334) and software publishers (NAICS 5112 and 5132). Standard errors are clustered by firm, and bars show 90% confidence intervals.

To address this concern, we re-estimate our baseline result after excluding the industries in which AI is most directly produced such as computer and electronic product manufacturing and software publishers.

Figure E.13 presents the dynamic DiD estimates by excluding such industries, and we find that the baseline effect survives the exclusion. Hence, this evidence suggests that the baseline result is not only driven by the AI intensive industries, and hence supports the external relevance of our estimates.

F TensorFlow-Specific and Worker-Specific Evidence

F.1 Revelio Labs

The framework-level workforce evidence in Section 7 is built from Revelio Labs proprietary data. Revelio Labs assembles a large and individual-level database of workers by aggregating public professional (e.g. LinkedIn) profiles and resumes, and provides each profile with a standardized set of skills, positions, employers, and dates. Two features make it well suited to our purpose. First, it records skills at the worker-level so that we observe not only that a firm employs technical workers but which specific tools and frameworks these workers report. Second, workers' employer is mapped to a firm identifier that we can link to Compustat through the firm identifier `gvkey`, which lets us aggregate worker skills up to the firm-year level and merge them with our Compustat panel.

Table F.1: TensorFlow - Revelio Labs Skill Lookup

Skill code	Category
`skill_k15`	Data Analytics (Broad)
`skill_k50`	Data Analytics (Narrow)
`skill_k150`	AI Data Processing
`skill_k500`	Artificial Intelligence
`skill_k1500`	AI Modelling
`skill_k5000`	Deep Learning
`skill_k15000`	Deep Learning Frameworks
`skill_k35000`	TensorFlow

Note: This table lists the Revelio Labs skill-taxonomy codes used to construct the worker-level skill measures.

We extract the AI- and deep-learning-related skills from the category of Revelio Labs'

skill taxonomy. Revelio Labs organizes skills into categories of increasing specificity, and Table F.1 lists the codes we particularly use for deep learning skills among which TensorFlow belongs to. This category ranges from broad data-analytics skills through AI and deep learning down to the most specific deep learning framework-level skill of TensorFlow. In this specific deep learning framework-level skill definition, Revelio Labs classifies TensorFlow and its bundled high-level API Keras, together with PyTorch, Caffe, Torch, and Apache MXNet. For workers we flag whether their profile lists a skill in each of these categories with particular attention to the deep learning framework-specific skills. We then aggregate these individual indicators to the firm-year level by counting, for each firm and year, the number of workers reporting each skill, and we normalize these counts by firm employment (workers per 1,000 employees) so that they are comparable across firms of different sizes. It is an important note that because workers would update their profiles with a lag, framework-specific skills are almost absent before the TensorFlow release and begin to appear from 2017 onward in our Revelio Labs sample.

F.2 Job Postings and Education Records

Beyond the skill records used above, we also draw on further Revelio Labs' data files that let us observe what kinds of jobs they advertise, what they offer to pay for them, and how educated their workforce is.

The first is Revelio Labs' database of United States job postings. Each observation is a single advertised vacancy, which is collected from employer career pages and job boards. It includes information on the occupation codes from the O*NET, the location at state and metropolitan-area level, an advertised salary, and the employer's identifier, which Revelio Labs also maps to the Compustat firm identifier `gvkey`.

Two features of the posted salary information govern how we use it. Revelio Labs reports a salary for nearly every posting, but for most it is not the figure written in the advertisement. It is imputed by a model from the job title, the location and the employer, and a flag distinguishes the two cases. Because an imputed salary is by construction largely a function of the occupation mix, and because our treatment variable is itself a measure of occupation mix, regressing an imputed salary on AI exposure would be close to circular. We therefore discard imputed salaries entirely and use only salaries stated in the posting.

The other data file we utilize is Revelio Labs' education records. It contains one observation per degree per worker by recording the degree level, the field, the awarding

institution and the dates of study. Because a worker would hold several education records, counting records would over weight those with many degrees, so we reduce each worker to a single highest degree on the ladder high school, associate, bachelor's, master's or MBA, doctorate, and then aggregate to the firm. We drop firms with fewer than 100 workers whose degree is recorded.

Data limitations. One key limitation is that Revelio Labs' job postings coverage effectively begins in 2016. Therefore, for our context, there is no pre-treatment period, and we are not able to conduct a DiD design for the sample we use from Revelio Labs. We therefore make no causal claim in Section 7. Instead, we provide post-period cross sectional comparisons between high and low 2015 AI exposure firms.

F.3 Hiring Composition by AI Exposure

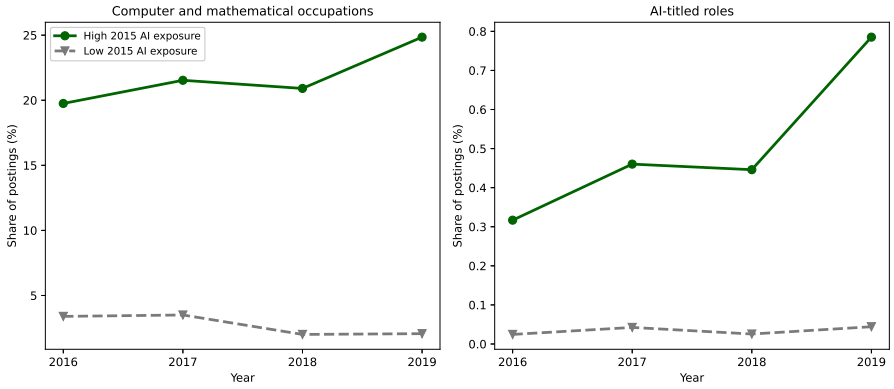
We provide a set of descriptive evidence using the corresponding Revelio Labs related data files. First, we analyze how high and low AI exposure firms differ with respect to what they advertise in their job postings. Table F.2 and Figure F.1 show that the difference is large. We find that among firms with low 2015 AI exposure, the share of postings in Computer and Mathematical occupations falls over the window, from 3.4% in 2016 to 2.1% in 2019. Among firms with high exposure it rises, from 19.8% to 24.9%. The divergence is still sharper for the narrower measure of roles whose title names an AI task. It increases from 0.32% to 0.79% among high-exposure firms against 0.02% to 0.04% among the rest. It suggests that by 2019 the high exposure group advertises AI roles at roughly eighteen times the rate of the low exposure group.

Table F.2: Hiring Composition and Posted Pay by 2015 AI Exposure, 2016–2019

	Low 2015 AI exposure				High 2015 AI exposure			
Year	Postings	Comp./math %	AI-titled %	Adj. pay	Postings	Comp./math %	AI-titled %	Adj. pay
2016	77,666	3.39	0.024	-0.046	34,087	19.75	0.317	+0.087
2017	233,612	3.50	0.042	-0.043	91,687	21.53	0.460	+0.081
2018	286,068	2.01	0.026	-0.049	67,935	20.91	0.446	+0.150
2019	188,005	2.07	0.044	-0.042	43,174	24.85	0.785	+0.130

Note: "Comp./math" is the share of postings in Computer and Mathematical occupations (O*NET major group 15); "AI-titled" is the share whose role title contains an artificial-intelligence term. "Adj. pay" is the residual of log posted salary after absorbing occupation×year and metropolitan-area effects, averaged within the exposure group, so it compares pay for the same occupation in the same metropolitan area and year; only salaries observed in the posting are used. High and low 2015 AI exposure are the two halves of the distribution of the firm's 2015 share of AI workers.

Figure F.1: Hiring Composition by 2015 AI Exposure



Note: The left panel plots the share of job postings in Computer and Mathematical occupations (O*NET major group 15); the right panel the share of postings whose role title contains an artificial-intelligence term. High and low 2015 AI exposure are the two halves of the distribution of the firm's 2015 share of AI workers.

Table F.3 puts this perspective on a regression analysis. We find that a one percentage point higher 2015 AI-worker share is associated with a 16.9 percentage point higher share of postings in computer and mathematical occupations and a 0.85 percentage point higher share of AI-titled postings, conditional on firm size and year.

Table F.3: Hiring Composition and 2015 AI Exposure

	Computer/math share		AI-titled share	
	(1)	(2)	(3)	(4)
2015 AI-worker share (p.p.)	16.852***	10.945***	0.849***	0.918***
	(1.754)	(1.612)	(0.111)	(0.131)
Year fixed effects	✓	✓	✓	✓
Industry fixed effects (FF12)		✓		✓
Firm-years	1,845	1,845	1,845	1,845
Firms	583	583	583	583

Note: The dependent variable is the firm-year share of postings in Computer and Mathematical occupations (columns 1-2) or in roles whose title contains an artificial-intelligence term (columns 3-4), in percentage points. Firm-years with fewer than 20 postings are excluded. Standard errors clustered by firm. * $p < .10$, ** $p < .05$, *** $p < .01$.

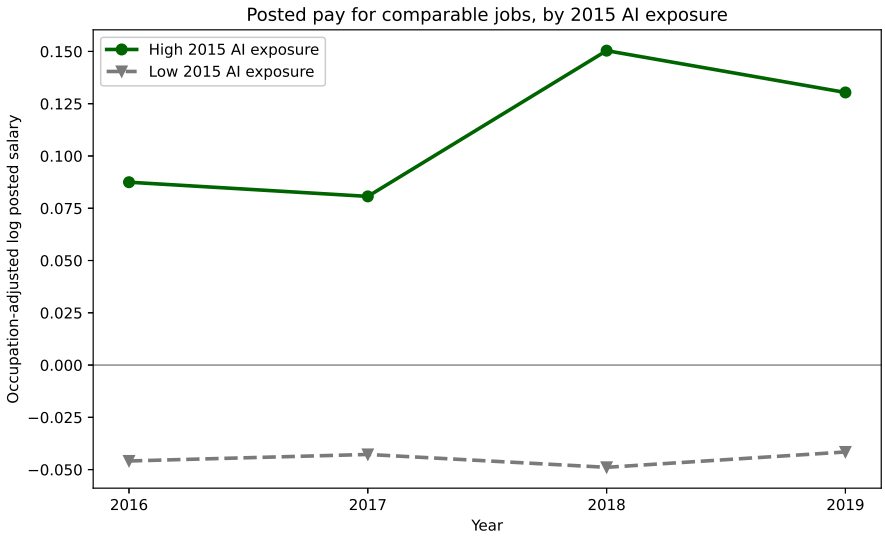
F.4 Posted Pay and Workforce Education

We next ask what the two groups offer to pay. A raw comparison of average posted salaries would be misleading since we have just shown in the previous section that the two groups advertise different jobs in different places. Therefore, we construct an occupation-adjusted measure. We take the logarithm of the stated salary and

residualize it on occupation-by-year and metropolitan-area effects.

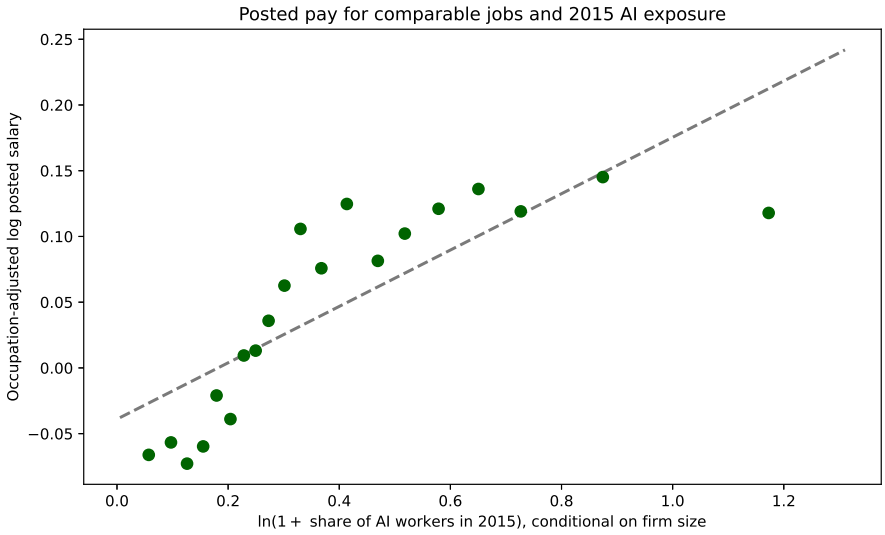
Figure F.2 traces the adjusted posted pay by the two groups. The pay premium is positive for high exposure firms and negative for low exposure firms.

Figure F.2: Posted Pay for Comparable Jobs, by Year and 2015 AI Exposure



Note: Each series is the mean, within an exposure group, of the firm-year residual of log posted salary after absorbing occupation-by-year and metropolitan-area effects, weighted by the number of postings. Only salaries stated in the posting are used. The horizontal line at zero is the sample mean of the residual by construction. High and low 2015 AI exposure are the two halves of the distribution of the firm's 2015 share of AI workers.

Figure F.3: Posted Pay for Comparable Jobs and 2015 AI Exposure

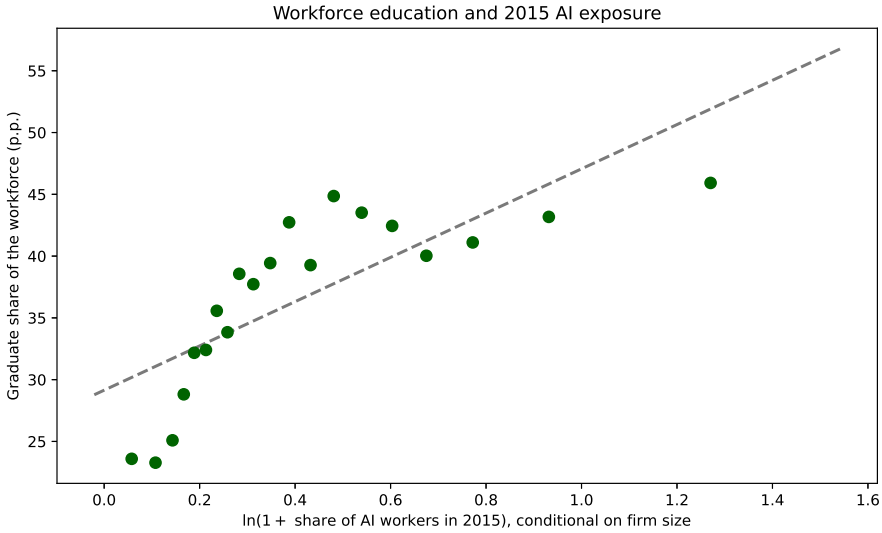


Note: The vertical axis is the firm's mean residual of log posted salary after absorbing occupation-by-year and metropolitan-area effects; only salaries stated in the posting are used. The horizontal axis is $\ln(1 + 2015 \text{ share of AI workers})$, winsorized at the ninety-ninth percentile. Both axes are residualized on lagged log assets, with the unconditional mean added back so the units remain readable. The fitted slope is 0.214.

Figure F.3 shows the relationship between the posted pay and the AI worker share.[52] We see that the relationship is monotone such that moving from the twenty-fifth to the seventy-fifth percentile of exposure is associated with 7.9 log points of additional posted pay for comparable work.

Figure F.4 traces the relationship between the 2015 AI exposure and graduate share of the workforce. We observe that it is positive, and the graduate share increases steadily from about 24% in the lowest bins to roughly 44% around the seventieth percentile of exposure.

Figure F.4: Workforce Education and 2015 AI Exposure



Note: The vertical axis is the share of the firm's workers whose highest recorded degree is a master's, an MBA or a doctorate. Both axes are residualized on lagged log assets, with the unconditional mean added back. The fitted slope is 17.9.

[52] Because the AI worker share is strongly right-skewed, we plot it as $\ln(1 + \text{share})$, and we residualize both axes on log assets. We also restrict the figure to firms with at least 100 postings, since a firm's premium is an average and is noisy when computed from few advertisements.

G Additional Tables

Table G.1: Production-Function Elasticities by Estimation Method

	OP	LP	ACF	WRDG
Labor $\hat{b}_L$	0.428	0.241	0.337	0.364
Capital $\hat{b}_K$	0.644	0.694	0.664	0.669

Note: This table reports the labor and capital elasticities $(\hat{b}_L, \hat{b}_K)$ used to construct the total factor productivity measures, one column per structural estimator. Output is value added (industry-deflated operating income before depreciation, positive values); labor is employment; capital is net property, plant and equipment; the proxy is cost of goods sold (Levinsohn and Petrin (2003), Ackerberg et al. (2015), Wooldridge (2009)) or capital expenditure (Olley and Pakes (1996)). $\ln(\text{TFP}) = \ln(\text{value added}) - \hat{b}_L \ln(\text{employment}) - \hat{b}_K \ln(\text{capital})$.

Table G.2: DiD Estimates: Extended Window

	(1) VA	(2) OP	(3) LP	(4) ACF	(5) WRDG
Post$\times$AI Share$_{2015}$	11.32***	12.41***	12.61***	12.42***	12.43***
	(3.47)	(4.38)	(4.47)	(4.39)	(4.39)
Controls	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Observations	12,732	12,732	12,732	12,732	12,732

Note: This table reports pooled DiD estimates over the extended 2008-2022 window. The outcome is ln(value added/employee) (= ln(oibdp/emp)) in column (1) and ln(TFP) in columns (2)-(5) (OP: Olley and Pakes (1996); LP: Levinsohn and Petrin (2003); ACF: Ackerberg et al. (2015); WRDG: Wooldridge (2009)). Treatment = Post$_t$ $\times$ 2015 AI-worker share (Post$_t$ = $\mathbf{1}$[year $\geq$ 2016]). Standard errors clustered by firm in parentheses. $^{*}p < .10$, $^{**}p < .05$, $^{***}p < .01$.

Table G.3: Pooled DiD Estimates: Partial-Exposure Adjustment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post×AI Share$_{2015}$	6.31**	6.55**	6.40**	7.32**	6.44**	6.36**	6.34**	6.64**	6.64**
	(3.03)	(2.95)	(2.98)	(3.09)	(3.13)	(3.12)	(3.12)	(3.09)	(3.09)
L.Sales Growth		0.17***	0.18***	0.16***	0.16***	0.16***	0.16***	0.16***	0.16***
		(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
L.Firm Age			0.07	0.05	0.02	0.02	0.03	0.03	0.03
			(0.10)	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)
L.ln(R&D/Sales)				-0.13**	-0.12**	-0.12**	-0.12**	-0.11**	-0.11**
				(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
L.ln(Assets)					0.05	0.08**	0.08**	0.09***	0.09***
					(0.03)	(0.04)	(0.04)	(0.04)	(0.04)
L.ln(Capex)						-0.04	-0.04	-0.04	-0.04
						(0.03)	(0.03)	(0.03)	(0.03)
L.Dividends/Assets							0.40*	0.39*	0.40*
							(0.21)	(0.21)	(0.21)
L.Cash/Assets								0.29**	0.29**
								(0.13)	(0.13)
L.Markup									0.00
									(0.00)
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	8,779	8,779	8,779	8,779	8,779	8,779	8,779	8,779	8,779

Note: This table reports pooled DiD estimates of value-added labor productivity under the firm-level partial-exposure adjustment: instead of excluding the entire 2015 calendar year, fiscal-2015 observations ending June-October 2015 are kept as pre-treatment, while those ending November 2015-May 2016 are dropped as partially exposed. Treatment = Post$_t$ × 2015 AI-worker share, Post$_t$ = $\mathbf{1}$[fiscal year $\geq$ 2016]. Columns (1)-(9) add the controls sequentially. Standard errors clustered by firm in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table G.4: Pooled DiD Estimates for TFP: Partial-Exposure Adjustment

	(1) OP	(2) LP	(3) ACF	(4) WRDG
Post×AI Share$_{2015}$	6.72*	6.71*	6.71*	6.71*
	(3.45)	(3.49)	(3.47)	(3.47)
Controls	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	8,779	8,779	8,779	8,779

Note: The outcome is the logarithm of TFP from a value-added production function estimated with four alternative methods. The treatment is Post$_t$ × 2015 AI-worker share (Post$_t$ = $\mathbf{1}$[fiscal year $\geq$ 2016]). Fiscal-2015 observations ending June–October 2015 (closed before the November 2015 TensorFlow release) are kept as pre-treatment observations; those ending November 2015-May 2016 are dropped as partially exposed. Standard errors are clustered by firm. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table G.5: DiD Estimates: Quarterly Value Added per Employee

	(1)	(2)
Post×AI Share$_{2015}$	4.467**	6.080**
	(2.183)	(2.784)
Firm FE	✓	✓
Year-Quarter FE	×	✓
Controls	✓	✓
Observations (firm-quarters)	37,098	37,098

Note: This table reports the quarterly value-added productivity effect of AI exposure. The outcome is ln(value added/employee) = ln(TTM oibdp/emp), where TTM oibdp is the trailing four-quarter sum of quarterly operating income (de-seasonalized, annual scale); industry-deflated, winsorized 1/99 within calendar quarter. Treatment = Post$_q$ × AI-exposure$_i$, where Post$_q$ = **1**[quarter ≥ 2015Q4] (the TensorFlow-release quarter) and AI-exposure is the firm's 2015 share of AI workers. Standard errors clustered by firm in parentheses. 2015Q3 and earlier are pre-treatment and the post period starts in 2015Q4. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table G.6: DiD Estimates: Quarterly TFP across Production-Function Estimators

	OP	LP	ACF	WRDG
Post×AI Share$_{2015}$	6.500**	6.567**	6.517**	6.514**
	(3.268)	(3.323)	(3.291)	(3.285)
Firm FE	✓	✓	✓	✓
Year-Quarter FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations (firm-quarters)	37,098	37,098	37,098	37,098

Note: This table reports the quarterly TFP effect of AI exposure across production-function estimators. Each column is a separate TFP measure: output is the trailing four-quarter sum of quarterly operating income (de-seasonalized, industry-deflated, positive values); inputs are quarterly net PP&E and annual employment; the labor and capital elasticities come from the annual production-function estimation of the corresponding method, and the resulting TFP residual is winsorized 1/99 within calendar quarter. Treatment = Post$_q$ × AI-exposure$_i$, where Post$_q$ = **1**[quarter ≥ 2015Q4] (the TensorFlow-release quarter) and AI-exposure is the firm's 2015 share of AI workers. Standard errors clustered by firm in parentheses. 2015Q3 and earlier are pre-treatment and the post period starts in 2015Q4. * $p < .10$, ** $p < .05$, *** $p < .01$.

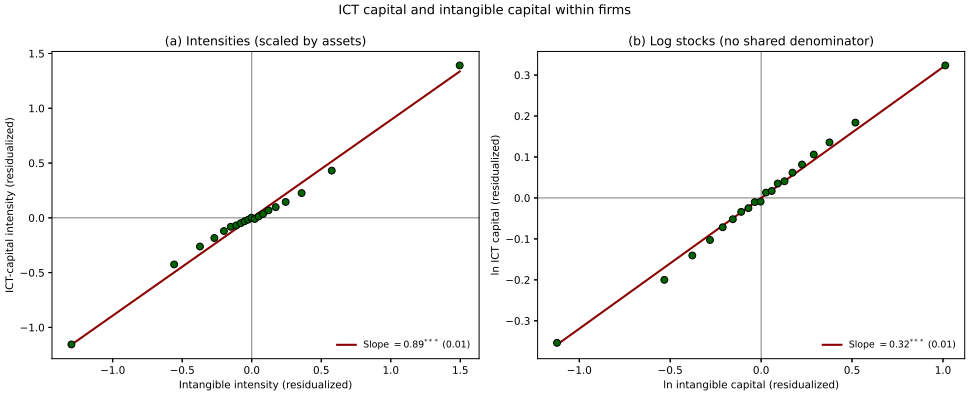
Table G.7: Production-Function-Based Markup

Panel A: estimated output elasticity of COGS, θ^V			
	Pooled OLS	Pooled LP	Pooled FE
θ^V	0.773	0.705	0.456
Panel B: value-added DiD, markup control swapped			
	0.85 proxy	DLEU markup	No markup
Post$\times$AI Share$_{2015}$	6.638**	6.637**	6.636**
	(3.156)	(3.156)	(3.156)
Observations	8,602	8,602	8,602

Note: Panel A reports the output elasticity of the variable input (COGS), θ^V, estimated from a gross-output production function $q = \theta^V v + \theta^K k + \theta^L l + \omega + \varepsilon$ (q = ln sales, v = ln COGS, k = ln PPE, l = ln employment) by pooled OLS, Levinsohn and Petrin (2003) (LP, investment proxy), and firm fixed effects. Panel B re-estimates the value-added pooled DiD with the lagged markup control set to the 0.85 proxy, De Loecker et al. (2020) markup $\theta^V_s \times$Sales/COGS, or omitted. Standard errors clustered by firm in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

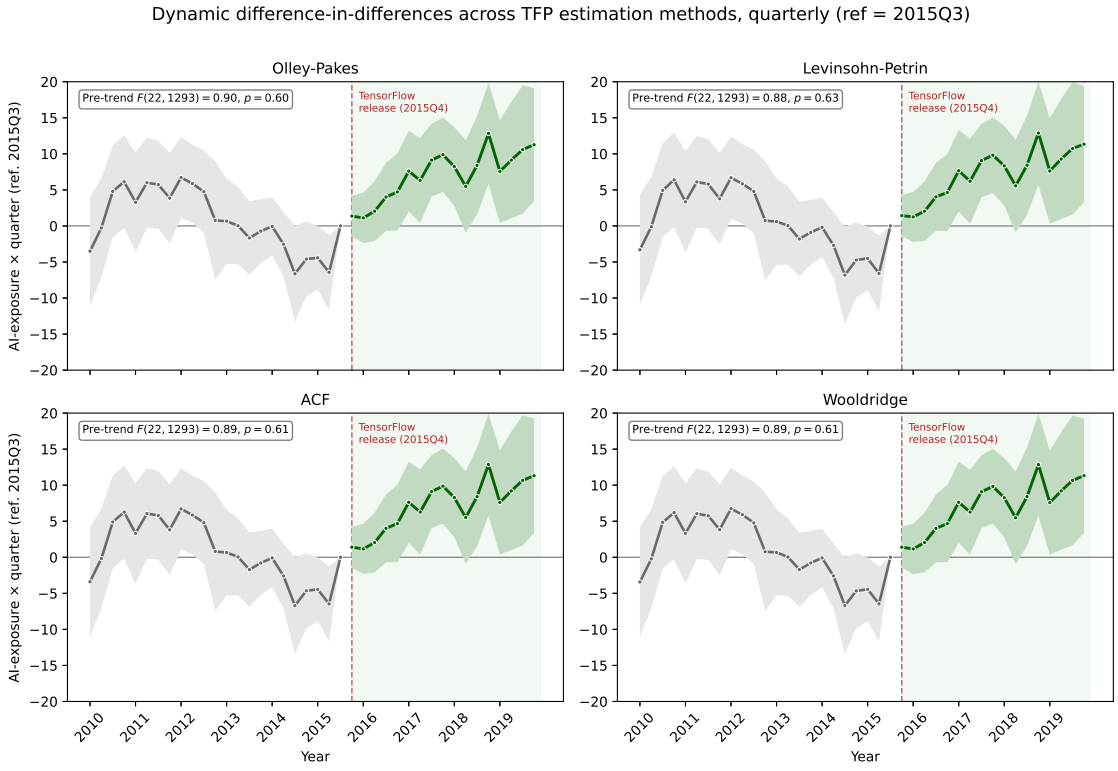
H Additional Figures

Figure H.1: Informational Capital and Intangible Capital



Note: This figure presents binned scatterplots of firm-level informational (ICT) capital against Peters and Taylor (2017) intangible capital, both residualized on firm and year fixed effects. Panel (a) uses capital intensities (stocks normalized by lagged assets); panel (b) uses log stocks. Informational capital is constructed following Falato et al. (2022) by applying the industry-year ratio of BEA information equipment and software stocks to tangible stocks to the firm's deflated net property, plant, and equipment. Each panel reports the within-firm slope with firm-clustered standard errors.

Figure H.2: Quarterly Dynamic DiD Estimates across TFP Estimation Methods



Note: This figure plots the estimated coefficients from quarterly dynamic DiD regressions of the impact of the TensorFlow shock on the logarithm of total factor productivity, estimated with four alternative production-function methods. Quarterly TFP applies the annual value-added production-function elasticities to trailing four-quarter real operating income, quarterly net property, plant, and equipment, and annual employment. The coefficients represent the interaction between the firm's 2015 share of AI workers and year-quarter indicators, with 2015Q3 as the reference period. The specification includes firm and year-quarter fixed effects, and the lagged annual controls. Standard errors clustered at the firm level. The shaded ribbons show the 90% confidence interval