AI Unbound:
Digital Infrastructure, AI Adoption, and Firm Performance*

Nuriye Melisa Bilgin
Koç University

Gianmarco Ottaviano
Bocconi University
BAFFI, GREEN, IGIER, CEP, CEPR

April 13, 2026

Abstract

We study how digital infrastructure relaxes constraints on the diffusion and economic impact of artificial intelligence (AI). Using administrative data and a nationally representative enterprise survey from Turkey (2021–2024), we document significant disparities in AI adoption. Adoption is concentrated among large firms and in regions with high-speed broadband and proximity to data centers, particularly for software-intensive and cloud-based applications. To identify causal effects, we exploit the staggered expansion of Turkey's national natural gas pipeline network, which serves as a conduit for fiber-optic deployment. Because pipeline routing is determined by energy distribution priorities rather than digital demand, it provides plausibly exogenous variation in connectivity. Difference-in-differences estimates show that improved connectivity significantly increases AI adoption, particularly for software-intensive technologies and among small and medium-sized enterprises. Instrumental-variable estimates indicate that infrastructure-driven AI adoption raises labor productivity and export intensity while shifting labor composition toward ICT-related roles. These findings highlight digital infrastructure as a primary determinant of both the pace of AI diffusion and its resulting economic returns.

Keywords: *Artificial Intelligence, Digital Infrastructure, Broadband, Technology Diffusion, Firm Productivity, Cloud Computing*

JEL Codes: *O33, L86, D24, J24, O14, R12*

*Contacts: mebilgin@ku.edu.tr; gianmarco.ottaviano@unibocconi.it. We thank Pamela Bombarda, Giuseppe Berlingieri, Maggie Chen, Julian Hinz, Sebastien Laffitte, Thepthida Sopraseuth, Rachel Sun, Davide Suverato, and participants at conferences, seminars and workshops for helpful comments.

1 Introduction

Artificial intelligence (AI) is rapidly becoming a core input in modern production, management, and trade. Firms increasingly rely on AI systems to analyze large datasets, automate routine processes, and improve decision-making. Yet the diffusion of AI remains highly uneven across firms and regions. Large firms adopt AI at substantially higher rates than small and medium-sized enterprises (SMEs), and adoption is concentrated in digitally advanced locations. These disparities raise a fundamental question: to what extent does local digital infrastructure shape which firms are able to adopt AI technologies and benefit from them?

This paper argues that broadband connectivity is a key enabling input for the deployment of modern AI systems. Many AI applications rely on cloud computing, large-scale data transmission, and remote storage and processing. Faster and more reliable broadband therefore lowers the cost of accessing machine-learning models, uploading large datasets, and running AI inference in real time. Conversely, limited connectivity can represent a binding constraint, especially for software-intensive and cloud-dependent AI applications such as deep learning, natural language processing, and robotic process automation. As a result, improvements in digital infrastructure may play a central role in determining which firms adopt AI and how they reorganize production in response.

To study this question, we assemble a novel firm-level dataset combining multiple administrative and survey sources on firms in Turkey between 2021 and 2024. The data merge five complementary components. First, a nationally representative survey on information and communications technology and AI usage reports whether firms adopt AI and identifies the specific technologies employed. Second, matched employer–employee records provide detailed information on employment and occupations, allowing us to track adjustments in ICT and non-ICT labor. Third, firm registry and financial accounts provide information on location, sector, and performance. Fourth, customs data allow us to measure firms' export activity. Finally, we incorporate district-level measures of digital infrastructure, including broadband download speed, latency, and proximity to data centers. This combined dataset allows us to characterize AI diffusion across firms, technologies, and regions and to study how infrastructure shapes both adoption and firm performance.

The emerging-economy context is particularly informative for identifying the mechanisms linking digital infrastructure to AI adoption. In settings where connectivity constraints remain binding, variation in broadband quality generates observable shifts in firms' adoption decisions that are harder to detect in more digitally saturated economies. This makes the Turkish case well suited to studying the conditions under which firms enter the AI frontier.

We organize the analysis around a simple conceptual framework in which firms choose whether and how to adopt AI technologies. The model highlights two key margins of heterogeneity. First, firms differ in productivity and organizational capacity, which affects their ability to absorb fixed adoption costs. Second, AI technologies differ in their dependence on digital infrastructure. In particular, software-intensive AI applications rely heavily on broadband connectivity and cloud access, whereas hardware-intensive applications such as autonomous robots depend more on on-site capital integration. Improvements in digital infrastructure therefore reduce adoption thresholds most strongly for software-intensive technologies and for firms that are close to the margin of adoption.

Our empirical strategy exploits exogenous variation generated by the staged expansion of Turkey's national natural gas pipeline network operated by the state-owned BOTAŞ Petroleum Pipeline Corporation. As documented by Demir et al. (2024), fiber-optic cables are frequently deployed alongside pipeline corridors because joint trenching substantially reduces installation costs and regulatory barriers. Crucially, the routing and timing of pipeline projects are determined by energy-distribution priorities rather than local digital demand. As a result, the rollout of pipeline-linked fiber infrastructure generates plausibly exogenous spatial and temporal variation in broadband connectivity. We exploit this variation in two complementary ways. First, we estimate a staggered difference-in-differences design that measures how connectivity improvements affect AI adoption. Second, we use pipeline-induced broadband expansion as an instrument for AI adoption to estimate the causal effect of adoption on firm outcomes.

The analysis delivers three main results. First, we document large disparities in AI adoption across firms and regions. Adoption is strongly concentrated among large firms and in locations with faster broadband and closer proximity to data centers. The correlation is particularly strong for software-intensive AI technologies, consistent with the hypothesis that digital connectivity is a key enabling input for these applications.

Second, exploiting the staggered expansion of the BOTAŞ-linked fiber network, we show that improvements in broadband connectivity causally increase AI adoption. The effects are largest for software-intensive technologies and for SMEs, which rely more heavily on external cloud infrastructure and are therefore more sensitive to local connectivity conditions.

Third, we use infrastructure-induced variation in AI adoption to estimate its causal effects on firm outcomes. Instrumental-variables estimates indicate that AI adoption raises labor productivity and export intensity, increases the demand for ICT workers, and reduces non-ICT employment. These adjustments suggest that AI adoption is accompanied by gradual organizational change rather than purely technological substitution.

By linking detailed microdata with quasi-experimental variation in digital infrastructure, this paper contributes to three related literatures. First, it extends the literature on the economic effects of broadband networks to the domain of AI adoption. A large body of work shows that broadband expansion affects productivity, employment, and innovation (Czernich et al., 2011; Akerman et al., 2015; Hjort and Poulsen, 2019). More broadly, digital infrastructure shapes firms' ability to access and use data-intensive technologies (Goldfarb and Tucker, 2019; D'Andrea and Limodio, 2024). Closest to our empirical strategy, DeStefano et al. (2025) use variation in fiber broadband speeds as an instrument for cloud computing adoption and show that cloud technologies raise firm performance. We build on this insight by moving from cloud infrastructure to the technologies that run on top of it, showing that improvements in connectivity causally affect the diffusion of frontier technologies such as AI, particularly for applications that rely on cloud computing and large-scale data transmission.

Second, the paper contributes to the literature on the diffusion of general-purpose technologies and firm heterogeneity. Previous research highlights the importance of complementary investments in skills and organizational capital for the adoption of new technologies (Bresnahan et al., 2002; Bloom et al., 2012). Our results emphasize the role of digital infrastructure as an additional constraint that shapes the timing and location of adoption, especially in environments where connectivity remains uneven. In this sense, the paper complements recent evidence documenting disparities in AI adoption across firms, sectors, and countries (Calvino and Fontanelli, 2023; McElheran et al., 2024; Bick et al., 2026).

Finally, the paper relates to the emerging literature on the economic effects of AI and automation. Recent work documents the effects of AI adoption on firm performance, innovation, and worker productivity, using both firm-level data and micro evidence on task-level outcomes (Babina et al., 2024; Brynjolfsson et al., 2025). Experimental evidence shows that AI can generate substantial productivity gains, particularly for information-processing tasks and for workers previously constrained by access to expertise (Noy and Zhang, 2023). Additional firm-level evidence points in a similar direction, documenting productivity gains without short-run employment declines (Aldasoro et al., 2026). This literature complements earlier work on the labor-market effects of automation technologies (Acemoglu and Restrepo, 2020). While most of this work focuses on the effects of AI adoption, relatively little is known about its determinants. Our analysis highlights local digital infrastructure as a key factor shaping which firms adopt AI and how they adjust their production structure. Consistent with this mechanism, AI adoption is closely linked to cloud computing rather than on-premise IT capital, underscoring the importance of infrastructure-sensitive technologies (McElheran et al., 2025).

The remainder of the paper proceeds as follows. Section 2 describes the data sources and presents descriptive evidence on AI diffusion across firms and regions. Section 3 introduces a simple model of AI adoption that frames the empirical analysis. Section 4 documents reduced-form relationships between digital infrastructure and AI adoption. Section 5 presents causal evidence on the effects of broadband expansion on AI adoption using a staggered difference-in-differences design. Section 6 estimates the causal effects of infrastructure-induced AI adoption on firm performance. Section 7 concludes.

2 Data and Descriptive Statistics

The analysis combines administrative firm-level data with the ICT Usage Survey conducted by the Turkish Statistical Institute (TÜİK), which provides detailed information on digital technology adoption, including AI. All datasets are accessed through TÜİK under secure data protocols. The survey spans four annual waves (2021–2024) and follows a stratified design by two-digit NACE sector and firm size (10–49 and 50–249 employees), with firms of 250 or more employees surveyed with certainty. The sampling frame covers firms with at least 10 employees in the main non-agricultural sectors, and each firm is assigned a sampling weight equal to the inverse of its selection probability within sector–size cells. Reflecting mandatory participation under Turkish Statistical Law (No. 5429), response rates exceed 97% in all waves, alleviating concerns about non-response bias. While the baseline analysis does not use sampling weights, robustness checks show that weighted estimates are virtually identical.[1]

The survey is merged with administrative data covering the universe of registered firms, yielding a firm-level panel for Turkey over 2021–2024 that links information on AI adoption, employment, financial accounts, trade activity, and local digital infrastructure. The panel is unbalanced: firms in the 10–49 and 50–249 size classes are sampled independently across years within sector–size strata, whereas large firms (250+ employees) are surveyed with certainty in each wave. The sample comprises 62,033 firm-year observations, with approximately 14,000 to 16,000 firms per year.[2]

This section describes the construction of the main variables and documents descriptive

[1] Additional details on the survey design and data are provided in Appendix A.

[2] Appendix B, Table B.1, reports the number of firm-year observations by survey year and firm size class. The sample size is stable across waves, reflecting the repeated cross-sectional nature of the stratified survey design, with large firms consistently represented due to census sampling.

patterns that motivate the empirical strategy. In particular, we highlight systematic heterogeneity in AI adoption across firm size, geography, and infrastructure quality, consistent with the sorting mechanism developed in Section 3.

2.1 AI Adoption

Our main outcome variable is a binary indicator equal to one if a firm reports using at least one AI technology in a given year. The survey collects detailed information on seven distinct AI technologies, which allows us to analyze both overall AI adoption and technology-specific diffusion patterns.

Because the central mechanism of the paper concerns how broadband quality affects AI adoption through cloud access and data transmission, we classify technologies primarily according to their infrastructure intensity. This distinction aligns directly with the conceptual framework and the empirical identification strategy. We define *software-intensive AI* as applications that rely primarily on computational processing and high-throughput data transmission: text mining, speech recognition, image recognition, natural language processing (NLP), deep learning, and robotic process automation (RPA). These technologies are typically cloud-dependent and therefore highly sensitive to broadband quality and data-center proximity. In contrast, we define *hardware-intensive AI* as applications that require physical machinery and on-site capital integration, namely autonomous robots and drones. Because these technologies rely more heavily on local equipment and less on continuous cloud access, their adoption is expected to be less sensitive to broadband performance. This distinction is central to our empirical strategy: software-intensive AI corresponds most closely to the cloud-dependent mode in the conceptual framework, while hardware-intensive AI provides a natural negative-control category in the event-study and instrumental-variable analyses.

In addition to infrastructure intensity, we also report results using a secondary classification based on functional capability. We distinguish between *Generative AI* and *Predictive AI*. Generative AI includes technologies designed to generate content or complex representations, comprising NLP and deep learning. Predictive AI encompasses applications focused on recognition, optimization, and task automation, including text mining, speech recognition, image recognition, RPA, and autonomous robots or drones. While this functional split is informative for descriptive purposes and for documenting heterogeneity within software-intensive AI, it is not the primary dimension for identification. Table 1 summarizes both classifications.

Table 1: Classification of AI Technologies by Functionality and Infrastructure

AI Category	Main Feature	Included Technologies
Generative AI	Content creation and complex pattern generation	Natural Language Processing (NLP), Deep learning
Predictive AI	Analysis, recognition, and process automation	Text mining, Speech recognition, Image recognition, Software robotics (RPA), Autonomous robots or drones
Hardware-Intensive AI	Physical implementation and movement	Autonomous robots or drones
Software-Intensive AI	Computational processing and data transmission	Text mining, Speech recognition, Image recognition, Natural Language Processing (NLP), Deep learning, Software robotics (RPA)

Notes: The software- versus hardware-intensive classification reflects differential reliance on broadband connectivity and cloud-based computation. This distinction is central to the empirical identification strategy, as software-intensive AI is more sensitive to digital infrastructure conditions.

The software-intensive category therefore corresponds most closely to the cloud-dependent adoption mode described in Section 3, while hardware-intensive applications approximate technologies that rely more heavily on local capital integration and are less directly constrained by broadband quality.

2.2 Digital Infrastructure

To measure digital infrastructure, we construct annual province-level indicators of broadband quality and computational access. Download speed and latency are derived from administrative data of the national ICT authority and Ookla's Open Data platform, aggregated from geolocated tiles to province-year means. These measures capture the performance of fixed broadband connections.

To account for spatial access to cloud infrastructure, we compute the distance from each firm to the nearest operational data center and construct province-level measures of data center capacity.

Together, these variables capture the local digital environment in which firms operate. Importantly, they vary both across provinces and over time, providing the basis for the difference-in-differences and instrumental-variable strategies implemented in subsequent sections.

2.3 Summary Statistics

Table 2 reports descriptive statistics separately for small and medium-sized enterprises (SMEs) and large firms (250+ employees). The SME subsample comprises 38,792 firm-year observations, while the large-firm subsample includes 23,241 observations.

Table 2: Summary Statistics by Firm Size

	SMEs (<250 employees)					**Large Firms (250+ employees)**				
Variable	**Mean**	**SD**	**Min**	**Max**	**N**	**Mean**	**SD**	**Min**	**Max**	**N**
Panel A. AI Adoption Groups										
AI adoption (any)	6.30	24.40	0	100	38,792	13.98	34.60	0	100	23,241
Generative AI	4.10	19.90	0	100	38,792	8.30	27.70	0	100	23,241
Predictive AI	5.80	23.30	0	100	38,792	13.10	33.70	0	100	23,241
Software-Intensive AI	5.60	23.00	0	100	38,792	12.20	32.70	0	100	23,241
Hardware-Intensive AI	3.60	18.60	0	100	38,792	8.70	28.20	0	100	23,241
Panel B. Individual AI Technologies										
Text mining	2.60	16.00	0	100	38,792	4.30	20.40	0	100	23,241
Speech recognition	2.10	14.30	0	100	38,792	3.40	18.00	0	100	23,241
Natural language processing	2.30	14.90	0	100	38,792	3.40	18.20	0	100	23,241
Image recognition	3.10	17.40	0	100	38,792	7.50	26.30	0	100	23,241
Deep learning	3.30	18.00	0	100	38,792	7.60	26.40	0	100	23,241
Software robotics (RPA)	3.30	17.80	0	100	38,792	9.30	29.00	0	100	23,241
Autonomous systems	1.60	12.70	0	100	38,792	3.90	19.40	0	100	23,241
Panel C. Digital Infrastructure and Employment										
Download speed (Mbps)	48.13	10.20	26	72	38,792	48.95	10.38	26	72	23,241
Latency (ms)	32.35	7.13	24	69	38,792	31.59	6.52	24	69	23,241
Distance to data center (km)	39.53	82.63	0	513.57	38,792	27.98	68.70	0	513.57	23,241
Export intensity	24.60	31.80	0	100	13,908	28.10	30.90	0	100	12,627

Notes: AI variables are reported in percentages (binary indicators multiplied by 100). Export intensity is exports as a share of total revenue in percentage points. Sample sizes differ for export intensity due to missing revenue or export information in some firm-years.

AI adoption exhibits a pronounced firm-size gradient. Among SMEs, only 6.30% report using at least one AI technology, compared with 13.98% of large firms. Adoption of software-intensive AI exceeds hardware-intensive adoption in both groups, consistent with the model's emphasis on infrastructure sensitivity.

The ICT Usage Survey follows the harmonized Eurostat questionnaire, ensuring direct cross-country comparability of AI adoption in terms of both technology definitions and sampling frame. This comparability allows us to benchmark Turkey's level of AI adoption relative to European economies. Over the period 2021–2024, the share of firms using at least one AI technology in Turkey is approximately 9%, placing it in the lower quartile of the European distribution and below the EU average (around 13.5% in 2024). The gap is particularly pronounced for advanced AI applications, such as machine learning and natural language processing, which are more prevalent in high-income EU countries. Beyond aggregate adoption rates, the harmonized design also enables comparison of the composition of AI use across technologies and firm size classes. In both Turkey and EU countries, AI adoption is strongly increasing in firm size, but the gradient is steeper in Turkey, reflecting lower adoption rates among small and medium-sized enterprises. These cross-country patterns provide useful context for interpreting our results and support the view that Turkey represents an emerging economy where AI diffusion remains relatively limited.[3]

2.4 Stylized Facts

Three stylized facts motivate the empirical strategy. First, AI adoption increases sharply with firm size (Figure 1), consistent with a sorting mechanism in which more productive firms adopt more fixed-cost-intensive technologies.

Second, AI adoption is spatially concentrated in metropolitan regions with dense broadband networks (Figure 2). Regions such as Istanbul, Ankara, and Izmir exhibit both high fiber intensity and high AI adoption, suggesting a strong infrastructure gradient.

Third, AI adoption is positively associated with broadband performance measures. Figure 3 shows that areas with faster download speeds and lower latency host significantly more

[3] Further details on international comparability can be found in Appendix A.4, while additional summary statistics are reported in Appendix B, Table B.2.

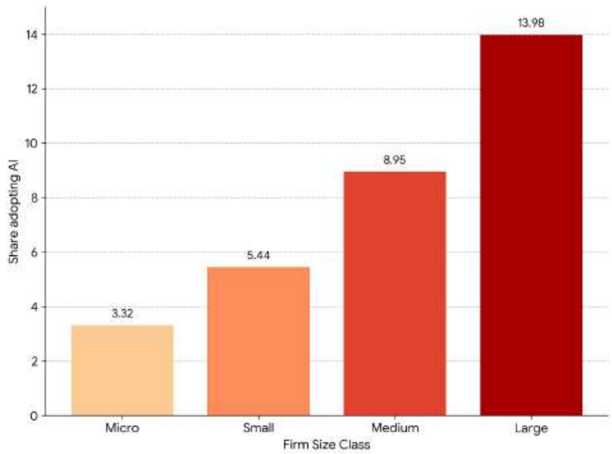

Figure 1: AI adoption by firm size class.

Notes: Share of firms reporting the use of at least one AI technology, by size class based on number of employees. Pooled sample, 2021–2024.

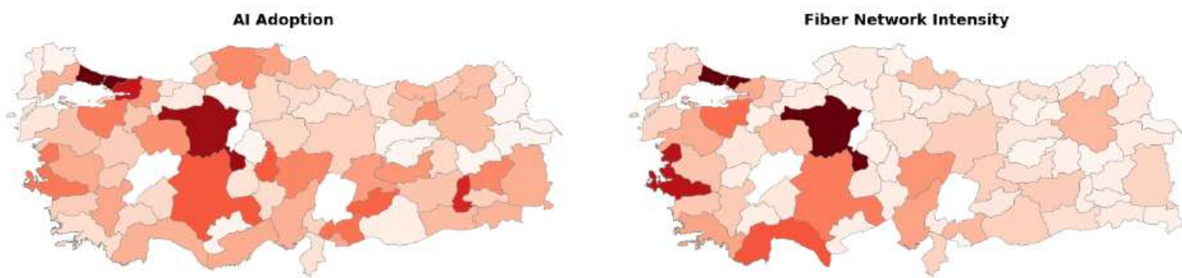

Figure 2: Spatial Clustering of AI Adoption and Fiber Intensity.

Notes: Left panel: province-level AI adoption rate (share of firms reporting any AI use in the ICT survey). Right panel: fiber intensity (total fiber-optic network length). Darker shading indicates higher values. Pooled sample, 2021–2024.

AI adopters. This relationship persists within industries and firm-size classes, indicating that infrastructure quality, rather than mere sector composition matters.

Taken together, these descriptive patterns suggest that digital infrastructure acts as a key determinant of AI diffusion. The remainder of the paper exploits exogenous variation in broadband expansion to establish a causal relationship between infrastructure quality and firm-level AI adoption.

3 A Simple Model of AI Adoption and Infrastructure Sensitivity

This section presents a stylized partial-equilibrium model of firm-level AI adoption that organizes the empirical analysis. The model is not intended to fully characterize AI production

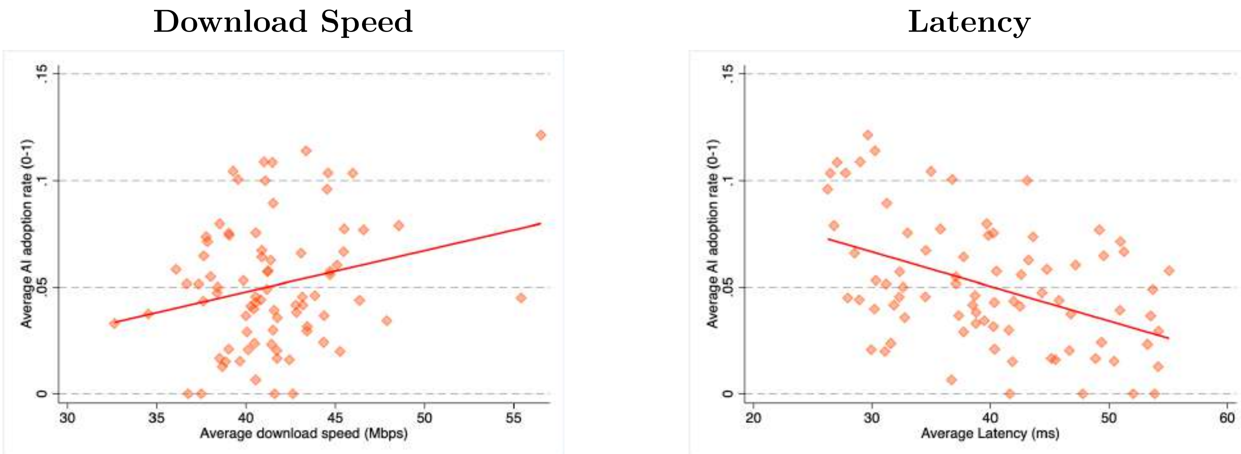

Figure 3: AI adoption and local internet quality.

Notes: Each panel plots the province-level AI adoption rate against local broadband quality. Left: average download speed (Mbps). Right: average latency (ms). Pooled sample, 2021–2024.

technologies or to be structurally estimated. Instead, it highlights how digital infrastructure shifts adoption thresholds and generates heterogeneous responses across firms, technologies, and locations.

Motivated by evidence that firms deploy AI through different technological architectures from fully cloud-based services to hybrid systems combining cloud access with local computing (Kergroach and Héritier, 2025; Brynjolfsson et al., 2025), we consider two AI modes that differ in fixed costs, productivity gains, and sensitivity to infrastructure. In the empirical analysis, the cloud-dependent mode corresponds to software-intensive applications, while the hybrid mode captures higher fixed-cost implementations that rely less on continuous connectivity. This distinction is best interpreted as a reduced-form characterization of infrastructure sensitivity rather than a literal technological taxonomy.

3.1 Environment and Technology Choice

The economy consists of a continuum of monopolistically competitive firms producing differentiated varieties under CES preferences with elasticity of substitution $\sigma > 1$. Upon entry, each firm draws productivity φ from a distribution $G(\varphi)$ with support on $(0, \infty)$ and takes aggregate demand as given.

Firms operate in locations indexed by d, characterized by broadband quality $B_d > 0$, where higher values reflect faster and more reliable connectivity. In an extension below, locations also differ in distance to the nearest data center, denoted $D_d \geq 0$. Each firm chooses one of three mutually exclusive modes: no AI ($m = 0$), cloud-dependent AI ($m = S$), or hybrid AI ($m = H$). Mode m is characterized by a fixed cost $F_m(B_d)$ and a productivity multiplier $A_m(B_d)$. Fixed costs are

$$F_0 = 0, \qquad F_S(B_d) = f_S B_d^{-\kappa_S}, \qquad F_H(B_d) = f_H B_d^{-\kappa_H}, \tag{1}$$

with $f_H > f_S > 0$ and $\kappa_H \geq \kappa_S > 0$. Hybrid AI entails higher upfront costs but at least as strong sensitivity to broadband improvements. Productivity multipliers are

$$A_0(B_d) = B_d^{\eta_0}, \qquad A_S(B_d) = a_S B_d^{\eta_S}, \qquad A_H(B_d) = a_H B_d^{\eta_H}, \tag{2}$$

with $a_H > a_S > 1$ and $\eta_H \geq \eta_S \geq \eta_0 \geq 0$. Hybrid AI therefore delivers larger productivity gains and is more complementary to broadband quality. Effective productivity is $\tilde{\varphi}_m = \varphi A_m(B_d)$.

3.2 Profits and Sorting

Under CES demand, operating profits take the form

$$\pi^{var}(\tilde{\varphi}_m) = K \tilde{\varphi}_m^{\sigma-1}, \tag{3}$$

where $K > 0$ is a demand shifter. Total profits are

$$\Pi_m(\varphi; B_d) = K \left(\varphi A_m(B_d)\right)^{\sigma-1} - F_m(B_d). \tag{4}$$

Firms select the mode that maximizes profits subject to non-negativity. The adoption decision is characterized by two productivity thresholds. The first, $\varphi_S(B_d)$, equates profits between no AI and cloud-dependent AI; the second, $\varphi_H(B_d)$, equates cloud-dependent and hybrid AI. These satisfy

$$K \varphi_S^{\sigma-1} \left(A_S(B_d)^{\sigma-1} - A_0(B_d)^{\sigma-1}\right) = F_S(B_d), \tag{5}$$

$$K \varphi_H^{\sigma-1} \left(A_H(B_d)^{\sigma-1} - A_S(B_d)^{\sigma-1}\right) = F_H(B_d) - F_S(B_d). \tag{6}$$

Under the maintained parameter restrictions, $0 < \varphi_S(B_d) < \varphi_H(B_d)$. Firms with $\varphi < \varphi_S(B_d)$ do not adopt AI, those with $\varphi \in [\varphi_S(B_d), \varphi_H(B_d))$ adopt cloud-dependent AI, and those with $\varphi \geq \varphi_H(B_d)$ adopt hybrid AI. More productive firms therefore select into more fixed-cost-intensive but more productive technologies.

Improvements in broadband reduce both thresholds. Since $F_m(B_d)$ decreases and $A_m(B_d)$ increases with B_d, we have

$$\frac{\partial \varphi_S}{\partial B_d} < 0, \qquad \frac{\partial \varphi_H}{\partial B_d} < 0.$$

Thus, better infrastructure induces entry into AI adoption and upgrading across modes. Firms near $\varphi_S(B_d)$ are the marginal adopters most responsive to infrastructure changes; these are the firms identified in the instrumental-variables estimates.

3.3 Spatial Frictions and Data Centers

To introduce spatial heterogeneity in cloud access, suppose firms face a friction

$$\tau(D_d, B_d) = \exp(\psi D_d - \omega \ln B_d), \quad \psi, \omega > 0, \tag{7}$$

Accordingly, distance to data centers increases frictions, while broadband mitigates them.

Productivity becomes

$$A_S(B_d, D_d) = a_S B_d^{\eta_S} \tau(D_d, B_d)^{-\mu_S}, \tag{8}$$

$$A_H(B_d, D_d) = a_H B_d^{\eta_H} \tau(D_d, B_d)^{-\mu_H}, \tag{9}$$

with $\mu_S > \mu_H > 0$. Cloud-dependent AI is therefore more exposed to spatial frictions than hybrid AI, implying that distance to data centers disproportionately reduces the attractiveness of cloud-based adoption. This asymmetry motivates the empirical use of hardware-intensive AI applications as negative controls: because they rely less on continuous cloud access, their adoption should be less sensitive to broadband improvements.

3.4 Testable Implications

The model yields several predictions that guide the empirical analysis. First, broadband expansion increases AI adoption at the extensive margin by lowering adoption thresholds.

Second, responses are strongest among marginal firms near these thresholds. Third, infrastructure improvements disproportionately affect cloud-dependent, data-intensive AI technologies, while hardware-intensive applications respond weakly, if at all. Fourth, distance to data centers reduces adoption primarily for cloud-dependent AI. Finally, industries with larger AI-related productivity gains exhibit stronger responses.

These implications provide a framework for interpreting the empirical results. Event-study estimates capture reduced-form shifts in adoption induced by broadband expansion, while instrumental-variables estimates identify causal effects for marginal adopters induced by these shifts. Heterogeneity across technologies, firm characteristics, and spatial frictions maps directly to the comparative statics of the model.

4 Reduced-Form Evidence

Guided by the conceptual framework in Section 3, this section documents reduced-form relationships between local digital infrastructure and firm-level AI adoption in Turkey over 2021–2024. The goal is to assess whether the cross-sectional and within-period patterns in the data align with the model's comparative statics, in particular the idea that improvements in broadband quality disproportionately facilitate the diffusion of infrastructure-sensitive (cloud-dependent, software-intensive) AI technologies and that spatial frictions matter more for smaller firms. These results are descriptive in nature; Sections 5 and 6 subsequently exploit quasi-experimental variation to move from correlation to causal inference.

4.1 Identification Logic and Empirical Variation

The variation exploited in this section comes from differences in broadband quality across districts and years, combined with within-industry adoption heterogeneity across firms. We focus on three dimensions of digital infrastructure: download speed, latency, and proximity to data centers. The model predicts that these factors should matter most for AI applications that rely on cloud access and large-scale data transmission. This motivates the emphasis on infrastructure intensity as the key dimension of heterogeneity: software-intensive AI should respond more strongly to broadband quality than hardware-intensive AI, which relies more on on-premise capital integration.

4.2 Variables and Measurement

Each firm i in district d and year t is observed within a stratified survey design in which firms are sampled by size category. Sampling weights w_i ensure representativeness of the underlying firm population. We incorporate these weights by estimating survey-weighted probit models that account for the stratification structure, while clustering standard errors at the district level throughout.

To facilitate comparability across regressors, we standardize all continuous exposure variables to have mean zero and standard deviation one. These include download speed (Mbps), latency (ms), firm size (log employment), and distance to the nearest data center. Coefficients are reported as average marginal effects (AMEs) and can therefore be interpreted as percentage point changes in AI adoption associated with a one-standard-deviation increase in each regressor.

Because download speed and latency are mechanically correlated, we also construct an orthogonalized latency measure. Specifically, we regress latency on download speed and use the residualized component

$$z_{\text{lt\_orth}} = \frac{\text{Latency}_{idt} - \widehat{\text{Latency}}_{idt}}{\sigma_{\text{Latency}}}, \qquad \widehat{\text{Latency}}_{idt} = \alpha + \rho\,\text{DownloadSpeed}_{idt},$$

which captures variation in network delay that is independent of bandwidth capacity.

4.3 Estimating Equations

We begin with the baseline specification relating AI adoption to broadband quality:

$$\text{AI}_{idt} = \alpha + \beta_1 z_{\text{dl},idt} + \beta_2 z_{\text{lt},idt} + \beta_3 z_{\text{L},idt} + \lambda_{st} + \varepsilon_{idt}, \tag{10}$$

where AI_{idt} is an indicator equal to one if firm i adopts at least one AI technology, z_{dl} and z_{lt} denote standardized download speed and latency, z_{L} is standardized firm size, and λ_{st} are four-digit NACE Rev. 2 industry–year fixed effects.

We then extend the baseline model to capture spatial frictions and technology heterogeneity:

$$\text{AIType}_{idt} = \alpha + \beta_1 z_{\text{dl},idt} + \beta_2 z_{\text{lt},idt} + \beta_3 z_{\text{L},idt} + \beta_4 z_{\text{dist\_dc},idt} + \lambda_{st} + \varepsilon_{idt}, \tag{11}$$

where AIType_{idt} denotes adoption of a given AI category or a specific AI technology. In both equations, we estimate survey-weighted probit models and report AMEs.

4.4 Broadband Quality and AI Adoption

Table 3 reports baseline estimates of equation (10), separately for SMEs and large firms. Across specifications, download speed is a robust and economically meaningful correlate of AI adoption. For SMEs, a one-standard-deviation increase in download speed is associated with an increase of roughly 0.8–0.9 percentage points in the probability of adopting AI. Latency enters with the expected negative sign but is not precisely estimated, and the orthogonalized latency measure yields similar conclusions. To distinguish persistent differ-

Table 3: Connectivity and AI Adoption: Download Speed and Latency by Firm Size

	SMEs			**Large Firms**		
	(1) Baseline	(2) Orth. Latency	(3) Perm./Trans.	(1) Baseline	(2) Orth. Latency	(3) Perm./Trans.
z(Download)	0.008***	0.009***		0.026	0.027*	
	(0.002)	(0.002)		(0.014)	(0.013)	
z(Latency)	-0.002			-0.006		
	(0.002)			(0.009)		
z(Orth. Latency)		-0.002			-0.006	
		(0.002)			(0.008)	
$z(\text{Download}_{perm})$			0.007***			0.011
			(0.002)			(0.009)
$z(\text{Download}_{trans})$			0.001			0.066**
			(0.004)			(0.024)
$z(\text{Latency}_{perm})$			-0.002			-0.008
			(0.002)			(0.007)
$z(\text{Latency}_{trans})$			-0.000			-0.002
			(0.001)			(0.005)
$z(L)$	0.016***	0.016***	0.016***	0.026***	0.026***	0.026***
	(0.001)	(0.001)	(0.001)	(0.007)	(0.007)	(0.007)
Observations	38,216	38,216	38,216	18,009	18,009	18,009

Notes: Entries report average marginal effects (AMEs) in percentage points for a one standard deviation increase in each regressor. All models are estimated using survey-weighted Probit regressions with industry-year fixed effects at the four-digit NACE Rev. 2 level. Sampling design accounts for firm-level probability weights and stratification by firm size. Standard errors are clustered at the district level. Firm size is measured by standardized number of employees, $z(L)$. * p<0.05, ** p<0.01, *** p<0.001.

ences in digital capacity from short-term fluctuations, Table 3 also reports results using a permanent–transitory decomposition. For SMEs, the correlation is concentrated in the permanent component of download speed, consistent with the idea that sustained digital capacity matters for adoption decisions with fixed-cost components. For large firms, the correlation is less precisely estimated in the baseline and orthogonal-latency columns, while the decomposition suggests that short-run upgrades may coincide with adoption changes, though the estimates are noisier due to the thinner distribution of large firms across four-digit industries.

4.5 Spatial Frictions: Proximity to Data Centers

Table 4 augments equation (10) with distance to the nearest data center as a proxy for spatial frictions in cloud access. For SMEs, distance is a robust negative correlate of adoption: a one-standard-deviation increase in distance reduces the likelihood of AI adoption by roughly 0.5 percentage points. For large firms, the coefficient remains negative but is not statistically significant, consistent with the idea that larger firms can partly substitute away from local infrastructure constraints through private networks or on-premise computing.

Table 4: Connectivity and AI Adoption: Data Centers and Geographic Mechanisms

	SMEs			**Large Firms**		
	(1) Baseline	(2) Orth. Latency	(3) Perm./Trans.	(1) Baseline	(2) Orth. Latency	(3) Perm./Trans.
z(Download)	0.007***	0.007***		0.026	0.026	
	(0.002)	(0.002)		(0.014)	(0.014)	
z(Latency)	-0.000			-0.005		
	(0.002)			(0.008)		
z(Orth. Latency)		-0.000			-0.005	
		(0.002)			(0.008)	
$z(\text{Download}_{perm})$			0.006**			0.009
			(0.002)			(0.009)
$z(\text{Download}_{trans})$			0.002			0.066**
			(0.004)			(0.024)
z(Dist-DC)	-0.005**	-0.005**	-0.005**	-0.004	-0.006	-0.004
	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)	(0.004)
$z(L)$	0.016***	0.016***	0.016***	0.026***	0.026***	0.026***
	(0.001)	(0.001)	(0.001)	(0.007)	(0.007)	(0.007)
Observations	38,216	38,216	38,216	18,009	18,009	18,009

Notes: Average marginal effects (AMEs) in percentage points for a one standard deviation (SD) change in each variable. All models are estimated with survey-weighted Probit regressions with industry-year fixed effects. Sampling design accounts for stratification by firm size. z(Dist-DC) represents the standardized distance to the nearest data center. * p<0.05, ** p<0.01, *** p<0.001.

4.6 Heterogeneity by AI Category and Technology

To test the model's key implication that infrastructure quality matters most for infrastructure-sensitive AI, we estimate equation (11) separately for AI categories and individual technologies. Tables 5 and 6 report results for SMEs and large firms, respectively. Across both firm-size groups, the most consistent pattern is that broadband speed is most strongly associated with software-intensive AI adoption, while hardware-intensive AI exhibits weak or statistically insignificant responses. This contrast anticipates the negative-control logic used later in the causal sections.

The functional split between generative and predictive AI is informative but secondary to infrastructure intensity. In the SME sample, software-intensive and generative categories show the strongest associations with download speed, whereas hardware-intensive AI shows no clear dependence on speed. Disaggregating further, technologies such as text mining, NLP, deep learning, and RPA tend to display stronger speed correlations than autonomous systems, consistent with the interpretation that cloud-dependent applications drive the infrastructure gradient.

Spatial frictions also matter differentially by firm size. For SMEs, distance to data centers remains a robust correlate for several software-intensive technologies, particularly deep learning and RPA. For large firms, distance is generally insignificant across categories, consistent with their ability to mitigate geographic constraints.

4.7 Interpretation and Limitations

Overall, the reduced-form evidence points to a robust infrastructure gradient in AI diffusion. Download speed is consistently associated with higher adoption, particularly for SMEs and for software-intensive technologies. Distance to data centers is negatively associated with adoption for SMEs, suggesting that spatial frictions in cloud access constrain smaller firms. These patterns are consistent with the conceptual framework, in which broadband quality reduces effective adoption frictions for cloud-dependent AI and shifts the adoption threshold for marginal firms.

Table 5: Heterogeneity and Diffusion in AI Adoption: SMEs

Panel A: Category-Level Heterogeneity

	(1) Generative AI	(2) Predictive AI	(3) Hardware-Intensive	(4) Software-Intensive
z(Download)	0.006***	0.006**	0.001	0.008***
	(0.002)	(0.002)	(0.002)	(0.002)
z(Latency)	-0.002	0.000	0.002	-0.002
	(0.002)	(0.002)	(0.001)	(0.002)
z(Dist-DC)	-0.004*	-0.004*	-0.004**	-0.005*
	(0.002)	(0.002)	(0.002)	(0.002)
$z(L)$	0.009***	0.015***	0.010***	0.013***
	(0.001)	(0.001)	(0.001)	(0.001)
Observations	37,071	38,133	36,997	37,888

Panel B: Technology-Specific Diffusion

	(1) Text Mining	(2) NLP	(3) Deep Learning	(4) RPA	(5) Image	(6) Speech	(7) Autonomous
z(Download)	0.007***	0.006***	0.004*	0.004*	0.003	0.003	-0.000
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
z(Latency)	-0.000	-0.001	-0.001	-0.000	0.001	-0.001	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
z(Dist-DC)	-0.001	-0.001	-0.005**	-0.005**	-0.003	-0.001	-0.002
	(0.002)	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)
$z(L)$	0.005***	0.004***	0.009***	0.010***	0.009***	0.004***	0.004***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	34,045	34,192	36,245	36,210	36,535	33,228	34,090

Notes: Average marginal effects (AMEs) in percentage points per one–standard deviation (SD) change. Standard errors in parentheses. * p<0.05, ** p<0.01, *** p<0.001. Panel A reports heterogeneity across broader AI categories (Generative vs. Predictive, Hardware- vs. Software-Intensive), while Panel B disaggregates to specific technologies. Both panels include survey weights, four-digit NACE Rev. 2 industry-year fixed effects, and controls for firm size ($z(L)$).

Table 6: Heterogeneity and Diffusion in AI Adoption: Large Firms

Panel A: Category-Level Heterogeneity

	(1) Generative AI	(2) Predictive AI	(3) Hardware-Intensive	(4) Software-Intensive
z(Download)	0.029*	0.026	0.014	0.026
	(0.013)	(0.014)	(0.013)	(0.014)
z(Latency)	-0.002	-0.005	-0.001	-0.005
	(0.007)	(0.008)	(0.007)	(0.008)
z(Dist-DC)	0.001	-0.004	-0.003	-0.004
	(0.004)	(0.005)	(0.005)	(0.005)
$z(L)$	0.013**	0.024***	0.013**	0.021**
	(0.005)	(0.007)	(0.004)	(0.007)
Observations	17,674	17,989	17,839	17,940

Panel B: Technology-Specific Diffusion

	(1) Text Mining	(2) NLP	(3) Deep Learning	(4) RPA	(5) Image	(6) Speech	(7) Autonomous
z(Download)	0.026*	0.007*	0.029*	0.026	0.014	0.006*	-0.005
	(0.010)	(0.003)	(0.013)	(0.013)	(0.013)	(0.002)	(0.003)
z(Latency)	0.006	-0.004	-0.002	-0.004	-0.001	0.000	-0.003
	(0.006)	(0.003)	(0.007)	(0.008)	(0.007)	(0.002)	(0.002)
z(Dist-DC)	-0.000	0.004	-0.002	-0.006	-0.003	0.001	-0.004
	(0.003)	(0.003)	(0.005)	(0.005)	(0.004)	(0.002)	(0.003)
$z(L)$	0.006**	0.004**	0.012**	0.017**	0.011**	0.004**	0.007***
	(0.002)	(0.001)	(0.004)	(0.005)	(0.004)	(0.001)	(0.002)
Observations	16,737	16,752	17,540	17,647	17,705	16,263	16,080

Notes: Average marginal effects (AMEs) in percentage points per one–standard deviation (SD) change. Standard errors in parentheses. * $p<0.05$, ** $p<0.01$, *** $p<0.001$. Panel A reports heterogeneity across broader AI categories for large firms, while Panel B disaggregates to specific technologies. Both panels include survey weights, four-digit NACE Rev. 2 industry-year fixed effects, and controls for firm size ($z(L)$).

At the same time, these estimates should be interpreted as correlations rather than causal effects. Broadband quality and proximity to data centers may be correlated with unobserved local economic conditions, firm networks, or complementary investments in digital capital. The next section addresses these concerns by exploiting quasi-exogenous variation in broadband expansion induced by the staggered rollout of Turkey's natural gas pipeline network to identify causal effects of infrastructure on AI adoption.

5 Causal Effects on Adoption

The reduced-form evidence in Section 4 shows that AI adoption is systematically correlated with local digital infrastructure, especially for software-intensive technologies and for smaller firms. Those correlations, however, may still reflect omitted local characteristics. Districts with more skilled labor, denser business networks, or stronger local demand for digital services may both attract better infrastructure and exhibit higher AI adoption. This section addresses that concern by exploiting plausibly exogenous variation in broadband availability generated by the staggered expansion of Turkey's natural gas pipeline network.

The identification strategy uses the timing of BOTAŞ-linked fiber rollout to estimate the causal effect of broadband expansion on firm-level AI adoption. Consistent with the conceptual framework, the analysis focuses on whether exogenous improvements in connectivity shift adoption toward cloud-dependent, software-intensive applications and whether those effects are largest for firms facing the strongest infrastructure constraints.

5.1 Identification Strategy

Our identification strategy rests on the institutional fact that investment and expansion decisions for Turkey's natural gas network are driven by gas-distribution priorities, energy security, and engineering costs rather than by expected returns to broadband provision (Demir et al., 2024). Fiber is installed alongside pipeline routes as a complementary input for monitoring and control, implying that the timing and routing of the resulting fiber backbone are determined primarily by gas infrastructure planning rather than by local digital demand.

In telecommunications infrastructure, a large share of deployment cost reflects excavation rather than the cable itself. Because BOTAŞ already digs trenches for its natural gas pipelines, the marginal cost of laying fiber-optic conduit along these routes is low. In addition, Supervisory Control and Data Acquisition (SCADA) systems are required to monitor pressure, temperature, and flow in natural gas pipelines, so BOTAŞ has to install fiber for its own operational needs.[4] When BOTAŞ installs fiber, it typically lays high-capacity cable and leases unused strands ("dark fiber") to operators. Following the deregulation of broadband services in October 2011, the Ministry of Transport encouraged state-owned enterprises to grant private operators access to this infrastructure to avoid duplicative excavation. As a result, private internet providers gradually obtained access to optical fiber embedded in the BOTAŞ gas pipeline system.

The key implication is that the geographic routing and sequencing of BOTAŞ expansions are determined by geological constraints and energy-transport considerations rather than by local economic development strategies or anticipated digital demand. Districts connected earlier to the BOTAŞ-linked fiber backbone gained access to high-capacity broadband sooner, whereas other districts were connected only in later phases. This staggered rollout generates plausibly exogenous spatial and temporal variation in broadband availability.

We define treatment at the district-year level. A district becomes treated in the first year in which it experiences an expansion, upgrade, or capacity increase in the BOTAŞ-linked pipeline network that supports fiber deployment. Once treated, the district remains treated thereafter. Firms are indexed by i, districts by d, and years by t.

Our approach builds on the same institutional setting as Demir et al. (2024), but exploits a different source of variation. Whereas Demir et al. (2024) use distance to the pre-existing 2011 BOTAŞ network interacted with year dummies to instrument broadband connectivity, we exploit the district-level timing of BOTAŞ-linked upgrades in a staggered difference-in-differences framework. This allows us to estimate dynamic treatment effects within firms, test

[4] SCADA is a control-system framework that combines computing, networked communication, and graphical interfaces to oversee and regulate complex machinery and processes.

for pre-trends directly, and study persistence in AI adoption following broadband expansion.[5]

5.2 Estimating Equation

We estimate dynamic treatment effects using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). Let g_d denote the first year in which district d becomes treated. We compute cohort-time average treatment effects, $ATT(g_d, t)$, and aggregate them into an event-study profile indexed by relative event time $k = t - g_d$.

For exposition, the underlying event-study specification can be written as

$$AI_{idt} = \sum_{k \neq -1} \beta_k \mathbf{1}\{t - g_d = k\} + \alpha_i + \tau_t + \lambda_{jt} + \varepsilon_{idt}, \tag{12}$$

where AI_{idt} is an indicator for whether firm i adopts AI in year t, α_i are firm fixed effects, τ_t are year fixed effects, and λ_{jt} are one-digit NACE Rev. 2 industry-year fixed effects. The omitted period is $k = -1$, so the coefficients β_k trace the evolution of adoption relative to the year immediately before treatment. In the figures below, outcomes are reported in percentage points. Standard errors are clustered at the district level throughout.

The identifying assumption is a parallel-trends condition: absent BOTAŞ-enabled fiber expansion, early-treated and later-treated districts would have followed similar trends in AI adoption.[6] The institutional setting and the absence of differential pre-trends provide support for this assumption.

5.3 Main Effects on Adoption

Figure 4 presents the event-study estimates for overall AI adoption, separately for SMEs and large firms. The pre-treatment coefficients are small and statistically indistinguishable from zero, consistent with parallel trends and the absence of anticipatory adoption. Following treatment, AI adoption rises in connected districts relative to not-yet-treated districts. Two years after connection, the cumulative increase reaches about 10.2 percentage points for SMEs and 6.5 percentage points for large firms.

[5] Appendix B, Figures B.1–B.5, illustrates the geographic distribution of BOTAŞ pipeline expansion and AI adoption.

[6] Appendix C, Table C.5, reports pre-treatment covariate balance.

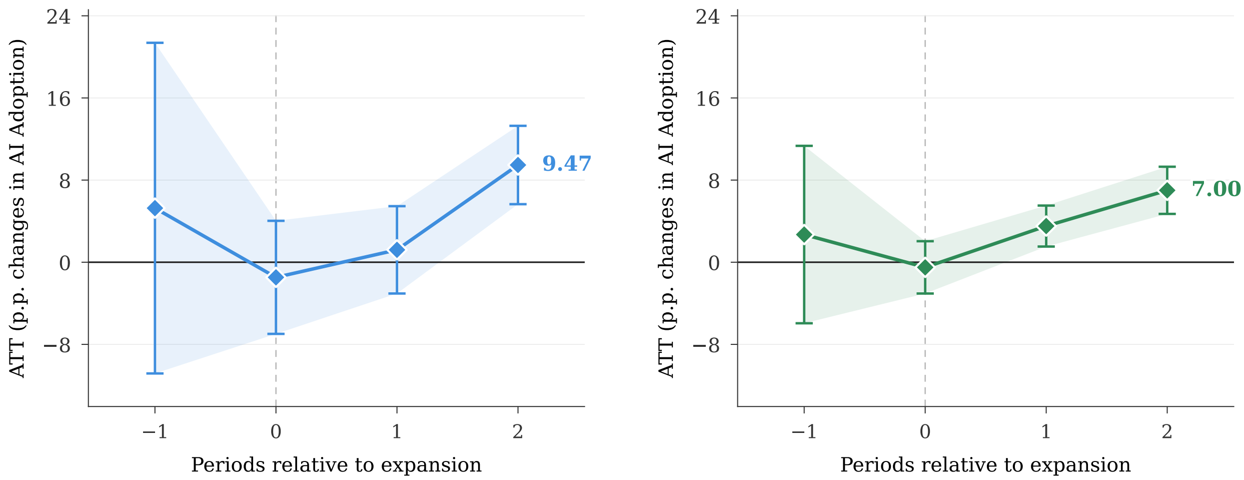

(a) SMEs (b) Large Firms

Figure 4: Effect of BOTAŞ Fiber Expansion on AI Adoption

Notes: This figure reports the dynamic average treatment effects on the treated (ATT) by event time, estimated using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). The outcome is an indicator equal to one if the firm reports using at least one AI technology. Period 0 denotes the first year in which the BOTAŞ-linked fiber expansion becomes active in the firm's district. The dashed vertical line marks the onset of treatment. Diamonds indicate point estimates; vertical bars show 95% confidence intervals based on standard errors clustered at the district level. Shaded bands depict the pointwise confidence region. Both panels share a common vertical axis. Panel (a) restricts the sample to SMEs (<250 employees); Panel (b) to large firms (≥250 employees).

These effects are economically large relative to baseline adoption rates. The stronger response among SMEs is also consistent with the conceptual framework: smaller firms are more likely to rely on cloud-based solutions and therefore benefit more from a reduction in connectivity constraints. By contrast, large firms can rely more heavily on internal computing capacity and appear to adjust more gradually.

5.4 Heterogeneity and Negative Controls

We next examine whether the treatment effect is concentrated in the technologies that the conceptual framework identifies as most infrastructure-sensitive. Figures 6 and 7 report the event-study profiles separately by AI technology. Several patterns emerge. First, pre-treatment coefficients remain close to zero across technologies, reinforcing the identifying assumption. Second, post-treatment responses are strongest for deep learning, text mining, NLP, and RPA. These applications are data-intensive and rely heavily on remote computation and storage, so improved broadband directly relaxes their adoption bottlenecks. This pattern is most evident for SMEs, where adoption often exhibits a threshold-like response once fiber connectivity arrives. For large firms, the response is more gradual but still concentrated in software-intensive technologies. These results are consistent with the reduced-form evidence in Section 4 and with the model's prediction that broadband mainly lowers the cost of cloud-dependent AI adoption.[7]

To further probe the mechanism linking broadband expansion to AI adoption, we exploit heterogeneity across AI technologies in their dependence on digital infrastructure. In particular, we distinguish between software-intensive AI applications that rely heavily on cloud computing and high-throughput data transmission, and hardware-intensive applications such as autonomous robots that depend primarily on on-site capital and local networks. If broadband expansion operates through improved digital connectivity, adoption responses should be concentrated among software-intensive technologies, while hardware-intensive applications should display weaker or no responses. Figure 5 provides direct evidence on this mechanism. The results reveal a clear technological asymmetry. Adoption of software-intensive AI increases sharply following broadband expansion, particularly among SMEs, whereas hardware-intensive technologies display little systematic response.[8] This pattern supports the interpretation that improvements in digital connectivity primarily relax constraints for cloud-based and data-intensive AI applications rather than for physical automation technologies.

A useful falsification test is provided by autonomous systems, which serve as a negative-control outcome. These technologies are capital- and energy-intensive and require complementary investments in physical machinery and organizational restructuring, but they do not rely heavily on continuous cloud access. If BOTAŞ-linked infrastructure expansion were affecting firms primarily through a broad local stimulus or through improved energy availability, one would expect autonomous systems to respond as well. Figure 8 shows that they do not: treatment effects are small, statistically insignificant, and display no systematic

[7] The underlying dynamic treatment-effect estimates for the technology-specific event studies are reported in Appendix C, Table C.2.

[8] Additional event-study evidence for broader AI groupings, including generative and predictive AI, is reported in Appendix C, Figure C.1.

Figure 5: Dynamic Effects of Broadband Expansion on AI Adoption: Software- vs. Hardware-Intensive AI

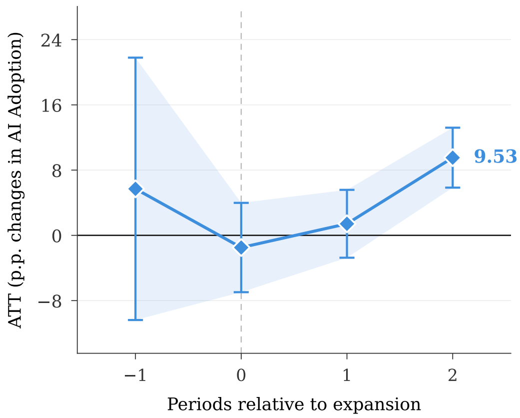

(a) Software-Intensive – SMEs

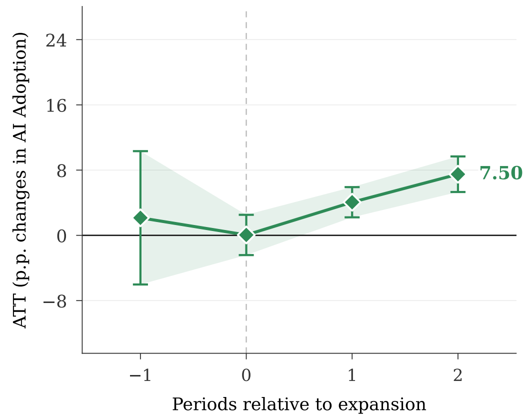

(b) Software-Intensive – Large Firms

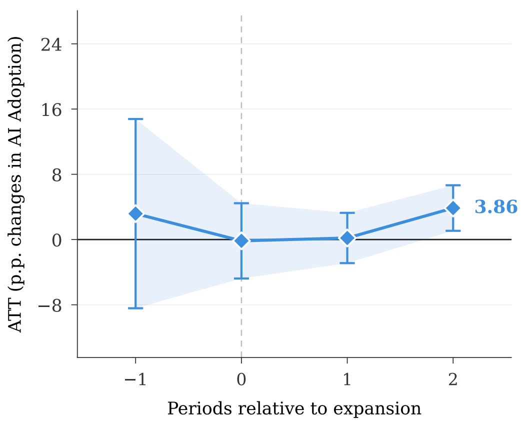

(c) Hardware-Intensive – SMEs

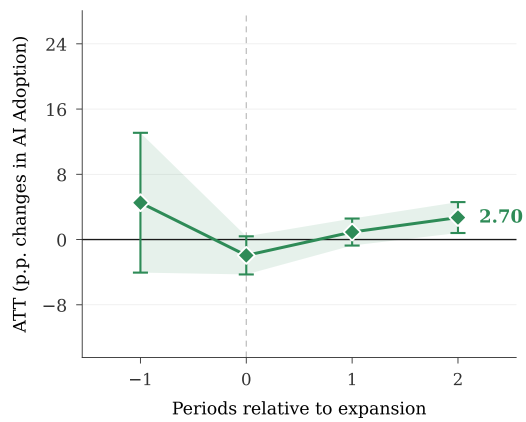

(d) Hardware-Intensive – Large Firms

Notes: Each panel reports the dynamic average treatment effects on the treated (ATT) by event time, estimated using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). The outcome is an indicator for whether the firm reports adopting at least one AI technology in the specified group. Software-intensive AI includes text mining, speech recognition, image recognition, NLP, deep learning, and robotic process automation. Hardware-intensive AI includes autonomous robots or drones. Period 0 corresponds to the first year in which the BOTAŞ-linked fiber expansion becomes active in a district. The dashed vertical line marks the onset of treatment. Diamonds indicate point estimates; vertical bars show 95% confidence intervals; shaded bands depict the pointwise confidence region. Standard errors are clustered at the district level. All panels share a common vertical axis. Left column: SMEs (<250 employees); right column: large firms (≥250 employees).

Figure 6: Dynamic Effects of Broadband Expansion on AI Adoption by Application Type (I)

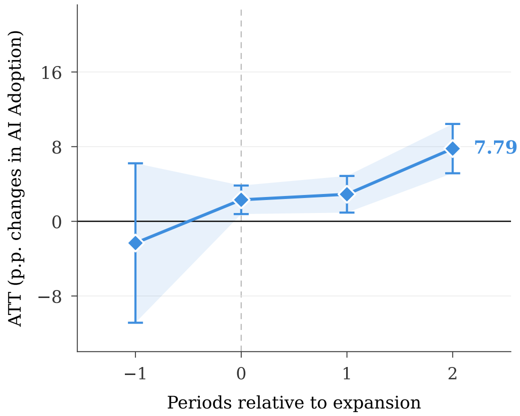

(a) Text Mining – SMEs

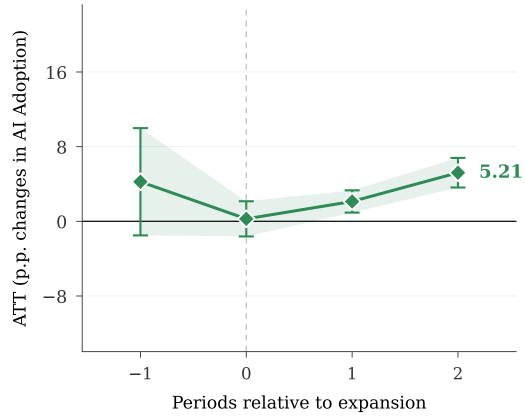

(b) Text Mining – Large Firms

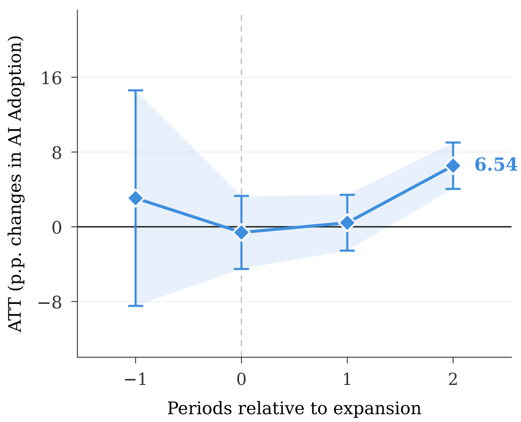

(c) NLP – SMEs

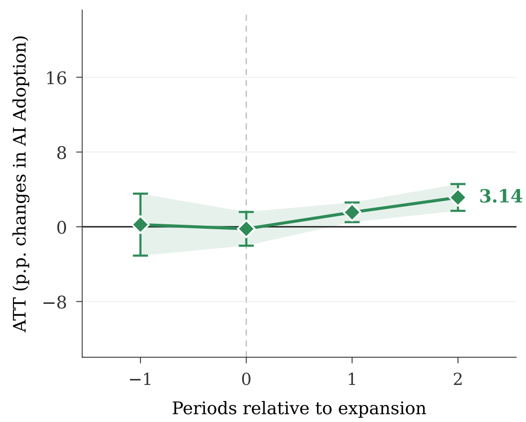

(d) NLP – Large Firms

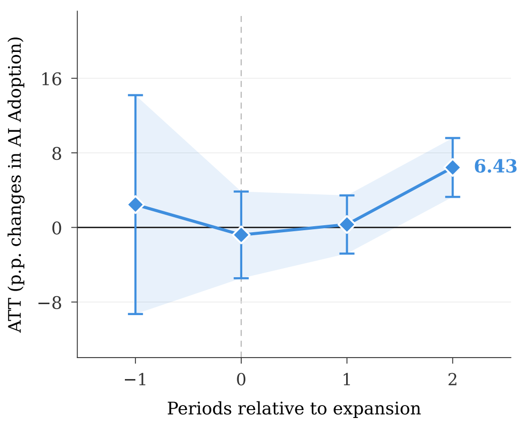

(e) Deep Learning – SMEs

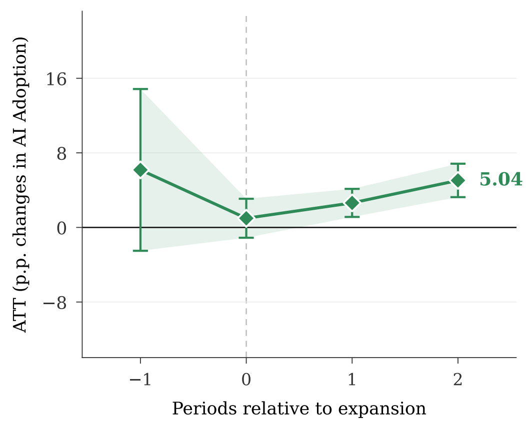

(f) Deep Learning – Large Firms

Notes: Each panel reports the dynamic average treatment effects on the treated (ATT) by event time, estimated using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). The outcome is an indicator for whether the firm reports adopting the specified AI technology. Period 0 corresponds to the first year in which the BOTAŞ-linked fiber expansion becomes active in a district. The dashed vertical line marks the onset of treatment. Diamonds indicate point estimates; vertical bars show 95% confidence intervals; shaded bands depict the pointwise confidence region. Standard errors are clustered at the district level. All panels share a common vertical axis. Left column: SMEs (<250 employees); right column: large firms (≥250 employees).

Figure 7: Dynamic Effects of Broadband Expansion on AI Adoption by Application Type (II)

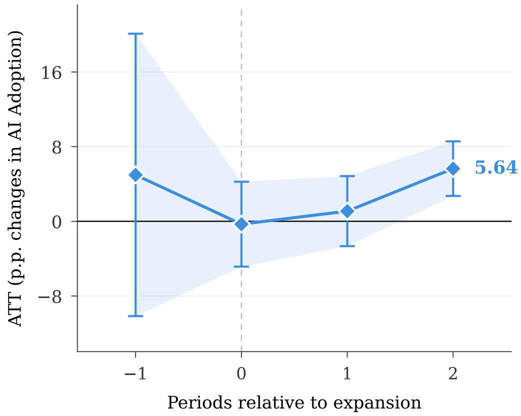

(a) RPA – SMEs

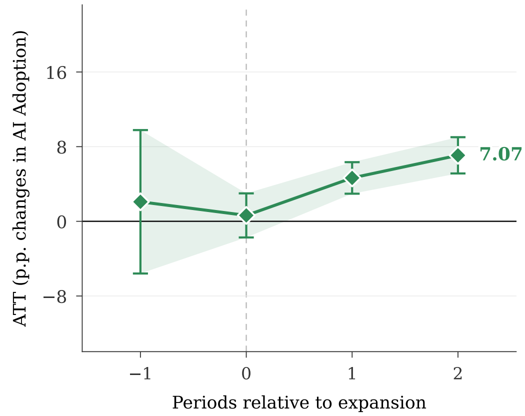

(b) RPA – Large Firms

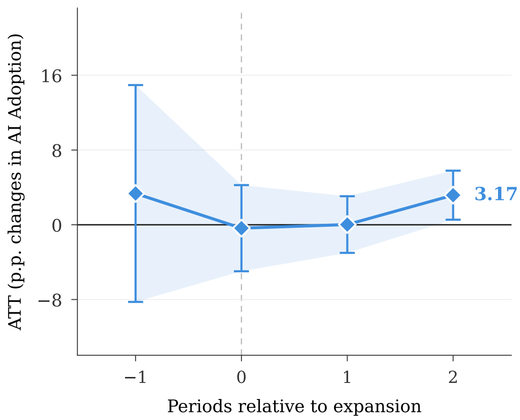

(c) Image Recognition – SMEs

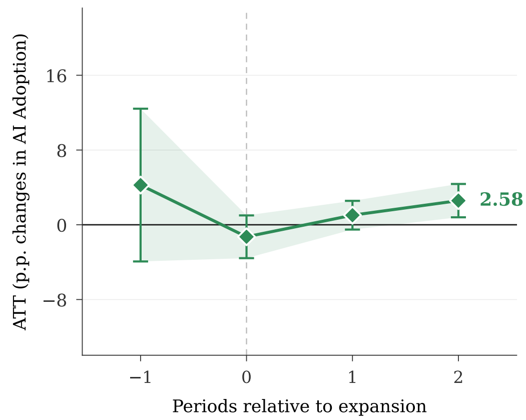

(d) Image Recognition – Large Firms

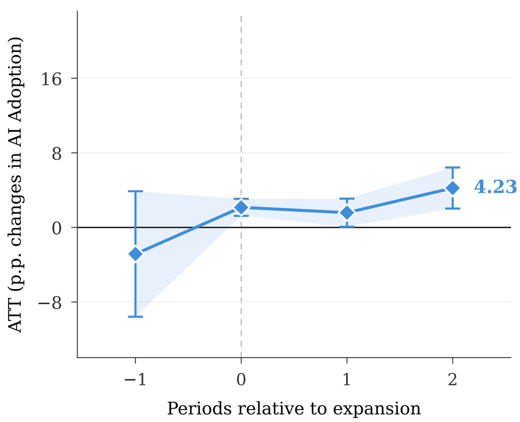

(e) Speech Recognition – SMEs

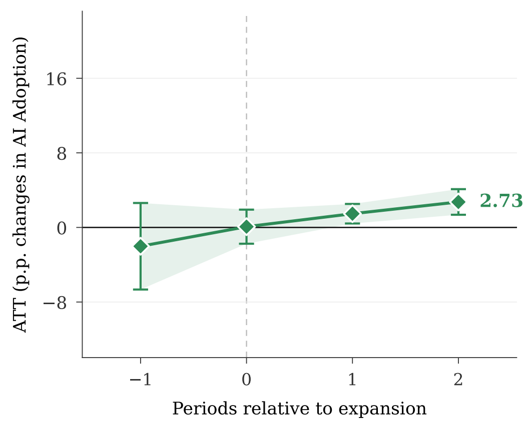

(f) Speech Recognition – Large Firms

Notes: Each panel reports the dynamic average treatment effects on the treated (ATT) by event time, estimated using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). The outcome is an indicator for whether the firm reports adopting the specified AI technology. Period 0 corresponds to the first year in which the BOTAŞ-linked fiber expansion becomes active in a district. The dashed vertical line marks the onset of treatment. Diamonds indicate point estimates; vertical bars show 95% confidence intervals; shaded bands depict the pointwise confidence region. Standard errors are clustered at the district level. All panels share a common vertical axis. Left column: SMEs (<250 employees); right column: large firms (≥250 employees).

post-treatment dynamics. This sharply contrasts with the sizable post-treatment effects for cloud- and data-intensive applications.

Figure 8: Autonomous Systems as a Negative Control

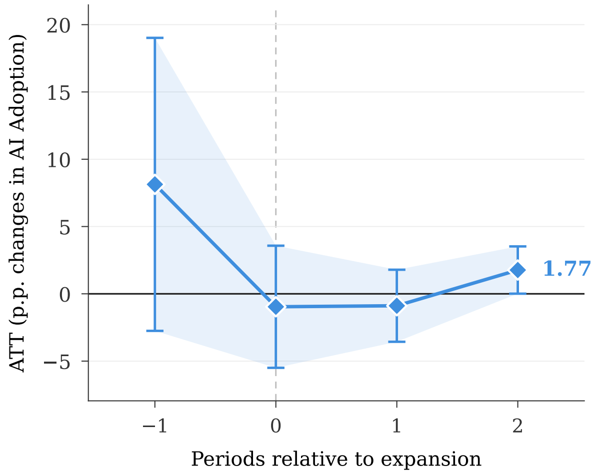

(a) SMEs

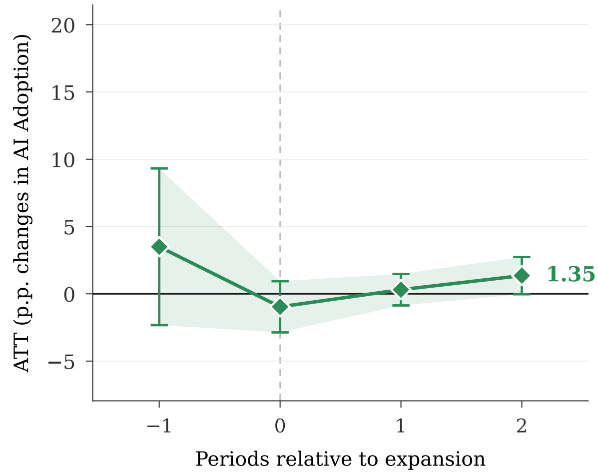

(b) Large Firms

Notes: This figure reports the dynamic average treatment effects on the treated (ATT) by event time, estimated using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). The outcome is an indicator for whether the firm reports adopting an autonomous system such as robots or drones. Period 0 corresponds to the first year in which the BOTAŞ-linked fiber expansion becomes active in a district. The dashed vertical line marks the onset of treatment. Diamonds indicate point estimates; vertical bars show 95% confidence intervals; shaded bands depict the pointwise confidence region. Standard errors are clustered at the district level. Both panels share a common vertical axis. Panel (a): SMEs (<250 employees); Panel (b): large firms (≥250 employees). The near-zero and statistically insignificant post-treatment coefficients are consistent with the hypothesis that broadband expansion affects software-based AI technologies but not hardware-dependent autonomous systems, supporting the exclusion restriction.

The negative-control evidence also bears directly on alternative channels. If the main effect of BOTAŞ expansion operated through natural-gas availability or general cost reductions, energy- and capital-intensive technologies such as robots should react at least as strongly as software-intensive AI. Instead, the opposite pattern emerges: the response is concentrated in cloud-oriented technologies, strongest for SMEs, and closely aligned with the infrastructure-sensitivity hierarchy documented in Section 4.

5.5 Mechanism: Software- versus Hardware-AI Exposure

While the results above strongly suggest an AI-specific digital mechanism, a subtler concern remains. Improved broadband may affect firm outcomes not only through AI adoption, but also through the adoption of other non-AI digital technologies such as enterprise resource planning systems, cloud storage, or e-commerce platforms. If so, the estimated treatment effects could reflect general digital upgrading rather than the deployment of AI itself.

To address this concern, we implement a triple-difference design that exploits heterogeneity in pre-determined industry-level exposure to different types of AI. The key insight is that general digital technologies should benefit firms across all industries, whereas AI-specific channels should generate differential responses depending on the type of AI an industry is predisposed to adopt. In particular, if broadband expansion operates through cloud-dependent AI, effects should be amplified in industries with higher baseline software-AI exposure, but not in industries with higher baseline hardware-AI exposure.

We construct baseline (2021) measures of industry-level exposure to software-intensive and hardware-intensive AI, defined as the share of firms in each four-digit NACE industry adopting each type of AI prior to the broadband expansion. These exposure measures are fixed over time and capture the structural propensity of an industry to adopt different AI technologies.[9]

We estimate

$$\begin{aligned} Y_{idt} = {} & \beta_1(\text{Pipeline}_{dt} \times \text{SWExposure}_j) + \beta_2(\text{Pipeline}_{dt} \times \text{HWExposure}_j) \\ & + X'_{idt}\delta + \alpha_j + \tau_t + \varepsilon_{idt}, \end{aligned} \tag{13}$$

where Y_{idt} denotes firm outcomes, α_j and τ_t are industry and year fixed effects, and X_{idt} includes log employment and firm-size controls. Standard errors are clustered at the district level.

Table 7 reports the results. Two findings emerge. First, the interaction with software-AI exposure significantly increases ICT employment in both samples and labor productivity among large firms.[10] Second, the interaction with hardware-AI exposure has either no effect or the opposite sign, particularly for employment composition. Among large firms, for example, software-AI exposure raises ICT employment while hardware-AI exposure lowers

[9] The corresponding first-stage estimates for adoption are reported in Appendix C, Table C.3.

[10] The ISCO-08 occupations used to define ICT workers are listed in Appendix C, Table C.1.

it and increases non-ICT employment. For SMEs, software-AI exposure is again associated with a positive ICT-employment response, whereas hardware-AI exposure does not generate comparable gains.

This contrast is highly informative. Broadband expansion shifts labor demand toward ICT-intensive tasks in industries predisposed to software-based AI, while industries oriented toward physical automation exhibit no comparable adjustment or even an offsetting response. That pattern is difficult to reconcile with explanations based on general digital upgrading or energy availability, both of which would predict similar effects across industries. Instead, the triple-difference evidence supports the interpretation that broadband expansion improves firm outcomes primarily through the adoption of cloud-dependent, software-intensive AI technologies.

Table 7: Triple-Difference: Firm Performance

	(1) Log Labor Prod.	(2) Export Int.	(3) Log ICT	(4) Log Non-ICT
Panel A: SMEs (<250 employees)				
Pipeline × SW Exposure	0.030	0.018	0.056***	-0.006
	(0.025)	(0.328)	(0.014)	(0.004)
Pipeline × HW Exposure	-0.026*	-0.146	-0.001	-0.000
	(0.010)	(0.379)	(0.007)	(0.001)
Pipeline	0.171***	6.295***	0.098***	-0.011***
	(0.049)	(0.843)	(0.012)	(0.001)
Observations	37,067	38,734	38,734	38,734
R^2	0.184	0.213	0.404	0.986
Panel B: Large Firms (≥250 employees)				
Pipeline × SW Exposure	0.062**	-0.825	0.056**	-0.005
	(0.019)	(0.595)	(0.016)	(0.007)
Pipeline × HW Exposure	-0.017	0.785	-0.046**	0.006***
	(0.015)	(0.557)	(0.014)	(0.002)
Pipeline	0.128***	5.812***	0.171***	-0.003
	(0.029)	(1.280)	(0.019)	(0.002)
Observations	22,930	23,218	23,218	23,218
R^2	0.379	0.432	0.457	0.986
Industry FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Size Controls	Yes	Yes	Yes	Yes

Notes: Each column reports OLS estimates of equation (13). SW Exposure and HW Exposure are standardized industry-level baseline (2021) adoption shares of software-intensive and hardware-intensive AI, respectively. All specifications include industry fixed effects, year fixed effects, log employment, and firm-size category controls. Standard errors are clustered at the district level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

5.6 Related Adjustment Margins and Interpretation

The adoption effects documented are accompanied by changes in firms' internal organization. Figure 9 reports estimates using the number of non-ICT and ICT workers as outcomes. Following broadband expansion, non-ICT employment declines while ICT employment initially falls and then rises. This pattern is consistent with a reorganization of production toward more digitally intensive tasks and with the idea that firms need time to restructure workflows and build complementary capabilities after the arrival of better connectivity.

Figure 9: Employment Composition

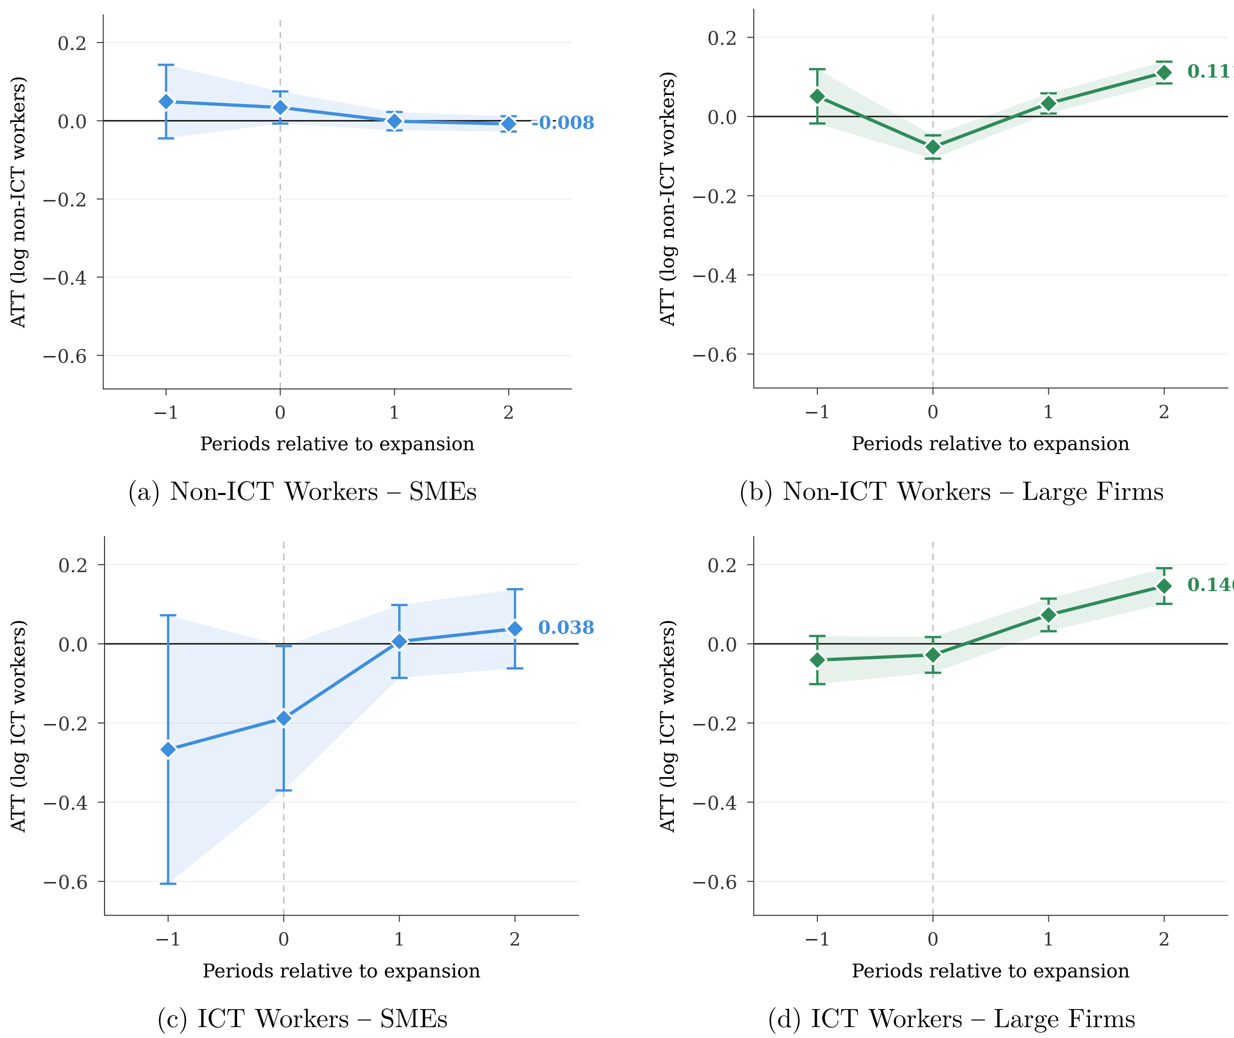

Notes: Each panel reports the dynamic average treatment effects on the treated (ATT) by event time, estimated using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). In panels (a)–(b) the outcome is the log number of non-ICT workers; in panels (c)–(d) the outcome is the log number of ICT workers. Period 0 corresponds to the first year in which the BOTAŞ-linked fiber expansion becomes active in a district. The dashed vertical line marks the onset of treatment. Diamonds indicate point estimates; vertical bars show 95% confidence intervals; shaded bands depict the pointwise confidence region. Standard errors are clustered at the district level. All panels share a common vertical axis. Left column: SMEs (<250 employees); right column: large firms (≥250 employees).

Figure 10 reports analogous event-study profiles for labor productivity and export intensity. These patterns are suggestive of subsequent performance gains, but they should not be interpreted as the main causal estimates of performance effects. Rather, they illustrate adjustment margins that accompany the adoption response and motivate the instrumental-variables analysis in Section 6, which isolates the causal impact of infrastructure-induced AI adoption on firm outcomes.

Figure 10: Labor Productivity and Export Intensity

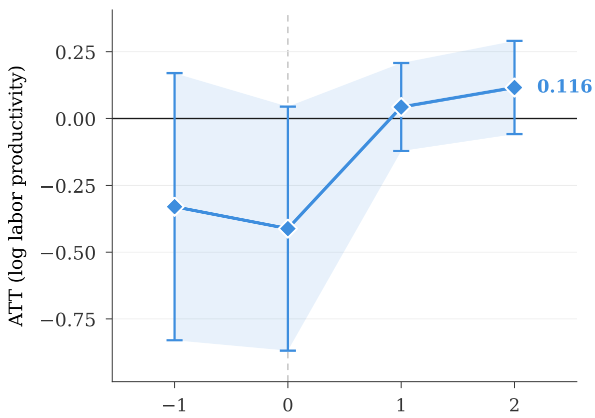

(a) Labor Productivity – SMEs

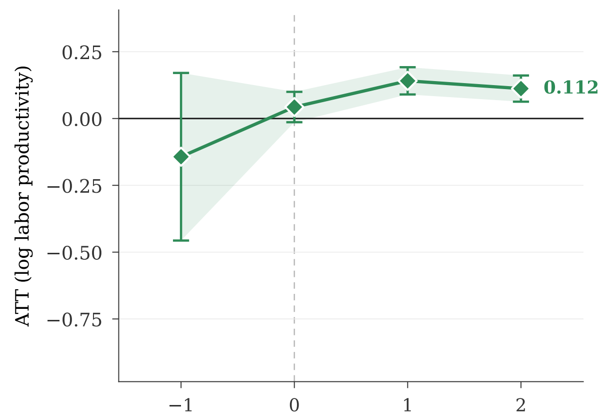

(b) Labor Productivity – Large Firms

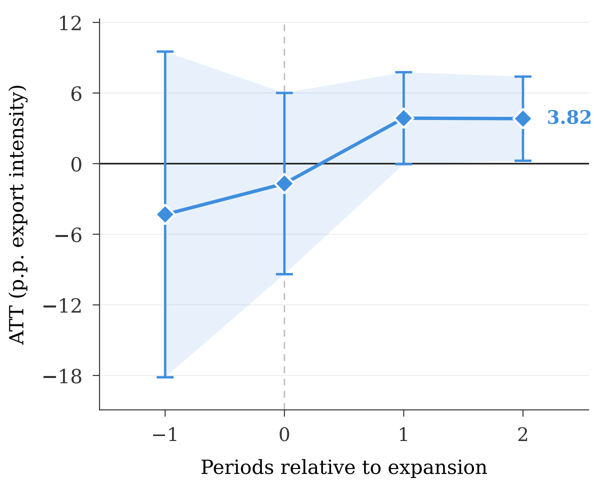

(c) Export Intensity – SMEs

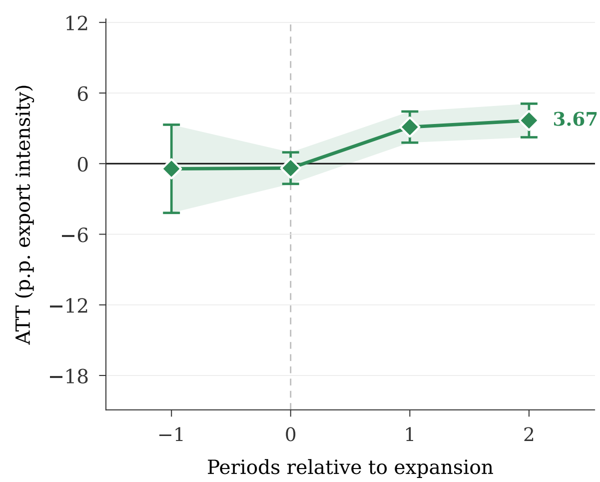

(d) Export Intensity – Large Firms

Notes: Each panel reports the dynamic average treatment effects on the treated (ATT) by event time, estimated using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). In panels (a)–(b) the outcome is log labor productivity (revenue per worker); in panels (c)–(d) the outcome is export intensity (exports as a share of total revenue, in percentage points). Period 0 corresponds to the first year in which the BOTAŞ-linked fiber expansion becomes active in a district. The dashed vertical line marks the onset of treatment. Diamonds indicate point estimates; vertical bars show 95% confidence intervals; shaded bands depict the pointwise confidence region. Standard errors are clustered at the district level. Within each outcome, panels share a common vertical axis. Left column: SMEs (<250 employees); right column: large firms (≥250 employees).

Taken together, the difference-in-differences evidence identifies a clear causal effect of broadband expansion on AI adoption. The effect is concentrated in software-intensive, cloud-reliant technologies, is strongest for SMEs, and is absent for hardware-intensive autonomous systems. The triple-difference estimates further show that these effects are amplified in industries predisposed to software-intensive AI and not in industries oriented toward hardware-intensive AI. These results are difficult to reconcile with explanations based solely on natural-gas availability or broad digital upgrading and instead support the interpretation that BOTAŞ-linked fiber expansion lowers the cost of adopting data-intensive AI.[11] In the next section, we use this exogenous variation to identify the causal effect of infrastructure-induced AI adoption on firm performance.

6 Causal Effects on Performance

Sections 4 and 5 establish two core results. First, local digital infrastructure is strongly associated with AI adoption, especially for software-intensive and cloud-reliant technologies. Second, plausibly exogenous improvements in broadband connectivity generated by BOTAŞ-linked fiber expansion causally increase AI adoption, with the largest responses observed among small and medium-sized enterprises. We now ask whether this infrastructure-induced adoption translates into improvements in firm performance.

To do so, we use BOTAŞ-driven broadband expansion as a source of exogenous variation in AI adoption and estimate the causal effect of adoption on firm outcomes. Consistent with the conceptual framework, the parameter of interest is the effect of adoption for marginal firms whose AI decisions are shifted by infrastructure improvements. The resulting estimates therefore have a Local Average Treatment Effect (LATE) interpretation: they capture the effect of AI adoption for firms that adopt because broadband constraints are relaxed.

[11] Event-study estimates for the pooled sample are reported in Appendix C, Figures C.2–C.3 and Table C.4, and confirm similar dynamic patterns in the aggregate data.

6.1 Identification Strategy

The identifying variation is the same as in Section 5. BOTAŞ pipeline expansions lower the cost of fiber deployment through rights-of-way and joint trenching, thereby generating plausibly exogenous variation in local broadband quality.

The exclusion restriction requires that, conditional on the controls and fixed effects included below, pipeline exposure affects firm performance only through the digital infrastructure channel and the induced adoption of AI. The evidence in Section 5, especially the absence of pre-trends, the concentration of treatment effects in software-intensive AI, the null effects for autonomous systems, and the triple-difference evidence by AI exposure, makes this interpretation plausible.

Our preferred interpretation is that BOTAŞ expansion improves download speeds, which in turn facilitates AI adoption and ultimately enhances firm performance. This mechanism is consistent with the model in Section 3: broadband lowers adoption thresholds, particularly for cloud-dependent technologies, and the IV estimates identify the effect of adoption for marginal firms induced to cross them.

6.2 Estimating Equations

We implement the IV strategy in three steps. First, we verify that BOTAŞ expansion strongly predicts district-level download speed:

$$z_{\mathrm{dl},dt} = \gamma_0 + \gamma_1 \mathrm{BOTAŞ}_{dt} + \gamma_2 \ln(\mathrm{Emp})_{idt} + \lambda_{st} + \theta_{\mathrm{size}} + \eta_{idt}, \tag{14}$$

where $z_{\mathrm{dl},dt}$ is standardized download speed in district d and year t, λ_{st} are four-digit NACE Rev. 2 industry-year fixed effects, and θ_{size} denotes firm-size category controls.

Second, we project firm-level AI adoption on the component of speed explained by pipeline-induced broadband variation:

$$AI_{idt} = \pi_0 + \pi_1 \widehat{z}_{\mathrm{dl},dt} + \pi_2 \ln(\mathrm{Emp})_{idt} + \lambda_{st} + \theta_{\mathrm{size}} + \nu_{idt}, \tag{15}$$

where $\widehat{z}_{\mathrm{dl},dt}$ is the fitted value from equation (14).

Third, we estimate the effect of infrastructure-induced adoption on outcomes:

$$Y_{idt} = \beta_0 + \beta_1 \widehat{AI}_{idt} + \beta_2 \ln(\text{Emp})_{idt} + \lambda_{st} + \theta_{\text{size}} + \varepsilon_{idt}, \tag{16}$$

where Y_{idt} denotes log labor productivity, export intensity, log ICT employment, or log non-ICT employment. Standard errors are clustered at the district level throughout. The coefficient β_1 identifies the LATE for firms whose AI adoption decisions respond to BOTAŞ-induced improvements in broadband quality.[12] In light of Sections 4 and 5, these compliers are most plausibly firms that are constrained by digital infrastructure, especially SMEs and firms adopting software-intensive AI.

6.3 Main Effects on Performance

Table 8 reports the main diagnostics and IV estimates. Panel A confirms relevance at both steps of the chain. Pipeline exposure raises local download speed substantially, and the resulting predicted speed strongly increases AI adoption. First-stage F-statistics exceed conventional thresholds, alleviating weak-instrument concerns in the baseline specification.

Panel B reports the second-stage estimates. Infrastructure-induced AI adoption increases labor productivity and export intensity, raises ICT employment, and slightly reduces non-ICT employment. These patterns are economically coherent and closely mirror the organizational adjustments documented in Section 5: firms that adopt AI in response to improved connectivity become more productive, more export oriented, and more ICT intensive.

In magnitude, the estimates imply sizeable performance effects.[13] Combining the pipeline-to-speed first stage, the speed-to-adoption first stage, and the second-stage productivity coefficient yields a non-trivial implied reduced-form productivity gain, consistent with the idea that broadband-enabled AI adoption represents a meaningful technological upgrade rather than a marginal digital adjustment.

[12] Reduced-form estimates linking pipeline exposure directly to firm outcomes are reported in Appendix D, Table D.2.

[13] Appendix D, Table D.1, compares alternative infrastructure channels, showing that the baseline results are strongest for download speed and data-center access, whereas latency is weakly identified.

Table 8: Broadband-Induced AI Adoption and Firm Performance

	(1) Log Labor Productivity	(2) Export Intensity	(3) Log ICT Workers	(4) Log Non-ICT Workers
Panel A: First-Stage Diagnostics				
Infrastructure stage: Speed ← Pipeline				
BOTAŞ Pipeline	1.209***	1.209***	1.209***	1.209***
	(0.147)	(0.147)	(0.147)	(0.147)
First-stage F-statistic	20.19	21.77	21.77	21.77
Adoption stage: AI ← Predicted Speed				
Predicted Speed ($\widehat{z}_{\text{dl}}$)	2.030***	2.030***	2.030***	2.030***
	(0.435)	(0.435)	(0.435)	(0.435)
Panel B: Second-Stage Estimates				
Broadband-Induced AI Adoption	0.060**	2.035***	0.046***	-0.002***
	(0.020)	(0.270)	(0.009)	(0.000)
Industry × Year FE	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes
Observations	60,072	62,033	62,033	62,033

Notes: This table reports the main IV results linking BOTAŞ-induced broadband improvements to firm performance through AI adoption. Panel A reports the infrastructure first stage (download speed on pipeline exposure) and the adoption first stage (AI adoption on predicted speed). Panel B reports second-stage estimates for firm outcomes. All specifications include four-digit NACE Rev. 2 industry-year fixed effects, firm-size controls, and log employment. Standard errors are clustered at the district level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

6.4 OLS Benchmarks and LATE Interpretation

To benchmark the IV estimates, Table 9 compares OLS and IV estimates of the effect of AI adoption on firm outcomes. Both specifications include four-digit NACE Rev. 2 industry-year fixed effects, firm-size controls, and log employment, with standard errors clustered at the district level. The OLS estimates capture the conditional correlation between AI adoption and firm performance, while the IV estimates identify the causal effect for compliers whose adoption is shifted by BOTAŞ-induced broadband expansion.

Three patterns emerge. First, the IV estimates consistently exceed the OLS estimates by a substantial margin for both overall AI adoption and software-intensive AI. For labor productivity, the OLS coefficient is close to zero, whereas the IV estimates are economically large for both any AI and software-intensive AI. This pattern is consistent with two complementary interpretations. Under classical measurement error in AI adoption, OLS is attenuated toward zero, as a binary self-reported indicator is a noisy proxy for underlying AI intensity. Under heterogeneous treatment effects, the IV identifies a LATE for compliers.

These are firms whose adoption is induced by infrastructure improvements and that may benefit more from AI than the average adopter because they were previously constrained by connectivity rather than by organizational capacity.

Second, OLS estimates for hardware-intensive AI are positive and similar in magnitude to those for software-intensive and overall AI, whereas the corresponding IV estimates are extremely imprecise. This divergence reflects the weak first stage for hardware-intensive AI documented below: because BOTAŞ expansion does not meaningfully shift hardware-AI adoption, the IV second-stage coefficients for that category are not informative. The OLS estimates instead capture a positive cross-sectional association, likely reflecting selection of larger and more productive firms into physical automation.

Third, the sign pattern is preserved across both estimators. AI adoption is associated with higher labor productivity, greater export intensity, more ICT employment, and less non-ICT employment. The qualitative story is therefore the same; the IV magnitudes are simply larger, which is exactly what one would expect if the instrument isolates firms with especially high marginal returns to adoption and if OLS suffers from attenuation bias.

Table 9: OLS versus IV Estimates of AI Adoption on Firm Performance

	Log Labor Prod.		Export Intensity		Log ICT Workers		Log Non-ICT Workers	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Panel A: Any AI Adoption								
AI Adoption	0.002***	0.060**	0.052***	2.035***	0.004***	0.046***	-0.000***	-0.002***
	(0.000)	(0.020)	(0.005)	(0.270)	(0.000)	(0.009)	(0.000)	(0.000)
Panel B: Software-Intensive AI								
AI Adoption	0.002***	0.054***	0.052***	1.843***	0.004***	0.042***	-0.000***	-0.002***
	(0.000)	(0.016)	(0.007)	(0.223)	(0.000)	(0.006)	(0.000)	(0.000)
Panel C: Hardware-Intensive AI								
AI Adoption	0.002***	0.306	0.047***	9.120*	0.003***	0.206	-0.000**	-0.011*
	(0.000)	(0.215)	(0.006)	(4.182)	(0.000)	(0.111)	(0.000)	(0.005)
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Size Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	60,072	60,072	62,033	62,033	62,033	62,033	62,033	62,033

Notes: This table compares OLS and IV estimates of the effect of AI adoption on firm outcomes. AI adoption is scaled to percentage points (0–100), so coefficients represent the effect of a one-percentage-point increase in adoption. OLS columns report coefficients from regressing each outcome on the AI adoption indicator with industry-year fixed effects, firm-size controls, and log employment. IV columns reproduce the second-stage estimates from Table 8 and from the category-specific IV specifications reported below. Standard errors are clustered at the district level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

This comparison clarifies the interpretation of the IV estimates.[14] The LATE recovered here pertains to firms whose AI adoption responds to broadband expansion, not to the average AI adopter. The large IV-to-OLS gap therefore suggests that the marginal adopters induced by improved connectivity experience above-average returns to AI adoption. This is especially plausible in the present setting, where the compliers are firms that appear organizationally ready to adopt cloud-dependent AI but are initially constrained by inadequate digital infrastructure.

6.5 Heterogeneity and Negative Controls

The central exclusion concern is that BOTAŞ expansion may affect firm outcomes through channels other than AI adoption, most notably through energy availability or other non-AI digital investments. The evidence from Section 5 already suggests otherwise: treatment effects are concentrated in software-intensive AI and absent for autonomous systems. We now bring the same logic to the IV estimates.

Table 10 reports IV estimates by broad AI category. The strongest and most precisely estimated effects arise for generative, predictive, and software-intensive AI, all of which display strong first stages. By contrast, hardware-intensive AI has a weak first stage and yields imprecise second-stage estimates. This pattern is exactly what the digital-channel interpretation predicts.[15] If BOTAŞ expansion mainly operated through a broad energy or local-demand channel, one would expect stronger responses for hardware-intensive applications. Instead, the causal effects concentrate in the AI categories that depend most on connectivity and remote computation.[16]

Table 11 sharpens this argument using autonomous systems as an explicit negative control. Pipeline exposure does not predict autonomous-system adoption, and the second-stage estimates are statistically insignificant and extremely imprecise. This is difficult to rec-

[14] Appendix D, Tables D.3 and D.4, reports the corresponding IV estimates separately for SMEs and large firms.

[15] Additional tests distinguishing digital from energy channels are reported in Appendix D, Tables D.7–D.9.

[16] Technology-specific IV estimates are reported in Appendix D, Table D.5.

Table 10: Heterogeneous Effects by AI Category

	(1) Generative AI	(2) Predictive AI	(3) Hardware -Intensive	(4) Software -Intensive
Panel A: First-Stage Diagnostics				
BOTAŞ Pipeline → AI Category	2.132***	2.265***	0.547	2.708***
	(0.283)	(0.492)	(0.302)	(0.447)
First-stage F-statistic	50.88	19.65	2.40	33.08
Panel B: Second-Stage – Labor Productivity				
Infrastructure-Induced AI Adoption	0.070***	0.065**	0.306	0.054***
	(0.018)	(0.022)	(0.215)	(0.016)
Panel C: Second-Stage – Export Intensity				
Infrastructure-Induced AI Adoption	2.341***	2.203***	9.120*	1.843***
	(0.353)	(0.304)	(4.182)	(0.223)
Panel D: Second-Stage – Log ICT Workers				
Infrastructure-Induced AI Adoption	0.053***	0.050***	0.206	0.042***
	(0.006)	(0.010)	(0.111)	(0.006)
Panel E: Second-Stage – Log Non-ICT Workers				
Infrastructure-Induced AI Adoption	-0.003***	-0.003***	-0.011*	-0.002***
	(0.001)	(0.000)	(0.005)	(0.000)
Industry × Year FE	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes
Observations	60,072	60,072	60,072	60,072

Notes: This table reports category-specific IV estimates. Panel A gives the first-stage relationship between BOTAŞ exposure and adoption of each AI category. Panels B–E report second-stage estimates for firm outcomes. Hardware-intensive AI exhibits a weak first stage and its second-stage estimates should therefore be interpreted cautiously. Standard errors are clustered at the district level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

oncile with a broad energy or local-growth mechanism and instead supports the exclusion restriction underlying the main IV estimates.

Table 11: Negative Control: Autonomous Systems

	(1) Log Labor Productivity	(2) Export Intensity	(3) ICT Workers	(4) Non-ICT Workers
Panel A: First-Stage – BOTAŞ → Autonomous Systems				
BOTAŞ Pipeline	0.193	0.193	0.193	0.193
	(0.198)	(0.198)	(0.198)	(0.198)
First-stage F-statistic	0.95	0.95	0.95	0.95
Panel B: Second-Stage – Outcomes				
Infrastructure-Induced Adoption	1.105	25.824	0.583	-0.030
	(1.724)	(25.532)	(0.593)	(0.029)
Industry × Year FE	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes
Observations	60,072	62,033	62,033	62,033

Notes: This table reports a falsification exercise using autonomous systems as the adoption margin of interest. The near-zero first stage indicates that BOTAŞ expansion does not shift adoption of this hardware-intensive technology. The absence of informative second-stage effects is consistent with the interpretation that the main IV results operate through broadband-sensitive AI rather than through broader local or energy channels. Standard errors are clustered at the district level.

6.6 Interpretation and Scope

Overall, the IV evidence delivers three main conclusions. First, BOTAŞ expansion is a strong predictor of local download speed and therefore provides relevant quasi-experimental variation in digital infrastructure. Second, broadband-induced AI adoption causally increases labor productivity and export intensity and shifts labor demand toward ICT-intensive occupations. Third, these gains are concentrated in AI categories that are most dependent on cloud access and data transmission, while hardware-intensive autonomous systems provide no comparable evidence of a causal channel.

These estimates should be interpreted as LATEs for firms whose adoption decisions are responsive to infrastructure improvements. They therefore speak most directly to marginal adopters located near the theoretical adoption thresholds in Section 3. They should not be interpreted as average treatment effects for all firms or all AI technologies. In particular, the large and imprecise coefficients for hardware-intensive AI reflect weak first stages rather than persuasive evidence of large causal effects.

A final issue is whether the baseline IV results depend on the weighting structure of the ICT survey. Because the survey uses a stratified design in which large firms are sampled with certainty and smaller firms receive probability weights, one might worry that unweighted IV estimates overrepresent medium-sized firms. Appendix D shows that the SME IV estimates are very similar with and without survey weights, both in magnitude and in significance (Table D.6). The main conclusions therefore do not depend on the survey-weighting choice.

A remaining concern is that exposure to the BOTAŞ network may contain a structural geographic component correlated with firm outcomes, even if the precise year of connection reflects quasi-random planning variation. Districts located along major transport corridors or closer to gas-exporting regions may be systematically more likely to receive pipeline infrastructure earlier. Following Borusyak and Hull (2023), we address this concern by constructing a recentered instrument that isolates the idiosyncratic component of the BOTAŞ expansion. Specifically, we estimate a district-level assignment model and define $\widetilde{D}_{dt} = D_{dt} - \mu_{dt}$, where μ_{dt} is the predicted treatment probability based on baseline characteristics. Re-estimating the IV chain using $\widetilde{D}_{dt}$ yields results that are very close to the baseline estimates, particularly for SMEs.[17] The export, ICT-employment, and non-ICT-employment effects are essentially unchanged, while the productivity effect is modestly attenuated but remains statistically significant. The autonomous-systems negative control also continues to show no meaningful first stage. These findings indicate that the main IV results are not driven by the structural geographic placement of the pipeline network but instead reflect the idiosyncratic timing variation exploited in our design.

Extended checks on alternative infrastructure channels, reduced-form magnitudes, firm-size splits, technology-specific estimates, and energy-versus-digital mechanisms are also reported in Appendix D. Those results reinforce the same conclusion: the dominant mechanism is a digital infrastructure shock that relaxes constraints on software-intensive AI adoption rather than a generalized energy or local-demand shock.

[17]See Appendix D, Tables D.10–D.12.

7 Conclusion

This paper studies how digital infrastructure shapes the diffusion and economic impact of artificial intelligence (AI) across firms. Using administrative and survey microdata on firms in Turkey between 2021 and 2024, we document large disparities in AI adoption across firms and regions. Adoption is concentrated among large firms and in digitally advanced locations with faster broadband and closer proximity to data centers. These patterns are particularly pronounced for software-intensive AI applications that rely on cloud computing and high-throughput data transmission.

To establish causal effects, we exploit the staggered expansion of Turkey's national natural gas pipeline network, along which fiber-optic cables are frequently deployed. Because pipeline routing and timing are driven by energy-distribution priorities rather than digital demand, this rollout generates plausibly exogenous variation in broadband connectivity. Difference-in-differences estimates show that improvements in connectivity significantly increase AI adoption, with the largest responses observed among small and medium-sized enterprises and for software-intensive AI technologies. We then use infrastructure-induced variation in adoption to estimate the causal effects of AI on firm outcomes. Instrumental-variables estimates indicate that AI adoption raises labor productivity and export intensity, increases demand for ICT workers, and reduces non-ICT employment. These patterns suggest that AI adoption is accompanied by gradual organizational adjustment, including the reallocation of labor toward more digitally intensive tasks.

Taken together, our results highlight digital infrastructure as a key determinant of both the diffusion and the economic impact of AI. Investments in broadband networks influence not only aggregate productivity but also which firms and regions participate in the AI frontier. In this sense, infrastructure policy can play an important role in shaping the inclusiveness of technological progress, particularly in middle-income economies where adoption gaps remain large. More broadly, the results echo lessons from earlier waves of general-purpose technologies. Just as organizational capabilities and skill complementarities conditioned the benefits of information technology diffusion, the gains from AI adoption appear to depend on firms' ability to combine improved connectivity with complementary human capital and organizational adjustment. The observed increases in ICT employment and the reallocation away from non-ICT labor suggest that firms rely on these complementary changes to realize the productivity benefits of AI. These patterns point to the importance of investments in skills and digital capabilities in translating infrastructure improvements into broader economic gains.

The setting of a rapidly digitalizing emerging economy allows us to identify mechanisms that are less visible in more advanced contexts. Where broadband constraints remain binding, infrastructure improvements directly shift firms' AI adoption decisions and subsequent performance. While the magnitude of these effects may differ across countries, the underlying channels (cloud-based computation, data access, and firm-level capability constraints) are likely to generalize beyond the Turkish case. Although our analysis focuses on Turkey, the mechanisms identified here are relevant to a broader set of environments in which access to high-quality digital infrastructure remains uneven across locations. Future work could examine how different layers of digital infrastructure (including mobile broadband, edge computing, and data-center networks) shape the geography of AI adoption across countries and industries.

Overall, our findings suggest that the geography of the AI revolution is not technologically predetermined but partly policy-contingent. Strengthening the digital backbone can accelerate AI diffusion, narrow spatial divides, and broaden the economic benefits of this emerging general-purpose technology.

References

Acemoglu, Daron and Pascual Restrepo, "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad, "The Skill Complementarity of Broadband Internet," *Quarterly Journal of Economics*, 2015, *130* (4), 1781–1824.

Aldasoro, Inaki, Leonardo Gambacorta, Csongor Pál, Debora Revoltella, Thomas Weiss, and Marcin Wolski, "AI Adoption, Productivity and Employment: Evidence from European Firms," Technical Report 21082, CEPR Discussion Paper 2026.

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

Bick, Alexander, Adam Blandin, David J. Deming, Nicola Fuchs-Schündeln, and Jonas Jessen, "Mind the Gap: AI Adoption in Europe and the U.S.," Technical Report 34995, National Bureau of Economic Research 2026.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen, "Americans Do IT Better? U.S. Multinationals and the Productivity Miracle," *American Economic Review*, 2012, *102* (1), 167–201.

Borusyak, Kirill and Peter Hull, "Non-Random Exposure to Exogenous Shocks," *Econometrica*, 2023, *91* (6), 2155–2185.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *Quarterly Journal of Economics*, 2002, *117* (1), 339–376.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond, "Generative AI at Work," *Quarterly Journal of Economics*, 2025, *140* (2), 889–942.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Calvino, Flavio and Luca Fontanelli, "A Portrait of AI Adopters across Countries: Firm Characteristics, Assets' Complementarities and Productivity," Science, Technology and Industry Working Papers 2023/02, OECD 2023.

Czernich, Nina, Oliver Falck, Tobias Kretschmer, and Ludger Woessmann, "Broadband Infrastructure and Economic Growth," *The Economic Journal*, 2011, *121* (552), 505–532.

Demir, Banu, Beata Javorcik, and Piyush Panigrahi, "Breaking Invisible Barriers: Does Fast Internet Improve Access to Input Markets?," Technical Report DP19827, Centre for Economic Policy Research (CEPR) 2024. Also published as CESifo Working Paper No. 11567.

DeStefano, Timothy, Richard Kneller, and Jonathan Timmis, "Cloud Computing and Firm Growth," *Review of Economics and Statistics*, 2025, *107* (6), 1638–1651.

D'Andrea, Angelo and Nicola Limodio, "High-speed internet, financial technology, and banking," *Management Science*, 2024, *70* (2), 773–798.

Goldfarb, Avi and Catherine Tucker, "Digital Economics," *Journal of Economic Literature*, 2019, *57* (1), 3–43.

Hjort, Jonas and Jonas Poulsen, "The Arrival of Fast Internet and Employment in Africa," *American Economic Review*, 2019, *109* (3), 1032–1079.

Kergroach, Sandrine and Julien Héritier, "Emerging Divides in the Transition to Artificial Intelligence," OECD Digital Economy Papers 386, OECD 2025.

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.

_ **, Zhi Yang, Zachary Kroff, and Erik Brynjolfsson**, "The Rise of Industrial AI in America: Microfoundations of the Productivity J-Curve(s)," Technical Report CES-WP-25-27, U.S. Census Bureau CES Working Paper 2025.

Noy, Shakked and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 2023, *381* (6654), 187–192.

Appendix A Data Appendix

This appendix provides additional details on the survey design, weighting scheme, data integration, and measurement of AI adoption. It also clarifies the relationship between the TÜİK ICT survey and the harmonized Eurostat framework.

A.1 Survey design, coverage, and weighting

The ICT Usage Survey conducted by TÜİK follows a stratified sampling design based on two-digit NACE Rev. 2 sectors and firm size classes. In particular, firms with 10–49 and 50–249 employees are sampled within sector–size cells, while firms with 250 or more employees are surveyed with certainty. Within the 10–49 and 50–249 size classes, selection is based on systematic random sampling. We use four annual waves from 2021 to 2024. The survey response rate is uniformly high in all waves: 98.7%, 99.1%, 97.5% and 99.1% in 2021, 2022, 2023 and 2024 respectively. This minimizes concerns about non-response bias.

The sampling frame covers firms with at least 10 employees in NACE Rev. 2 Sections C, D, E, F, G, H, I, J, K, L, M, N, and Group 95.1, corresponding to the main non-agricultural sectors of the economy. Each firm is assigned a sampling weight equal to the inverse of its selection probability within its sector–size cell. Operationally, the weight is constructed as the ratio of the population count to the sample count in the relevant cell. No additional post-stratification or calibration adjustments are applied. For firms with 250 or more employees, weights are equal to one by construction.

The baseline specifications do not use sampling weights. Instead, we include firm size fixed effects, which correspond to one of the stratification dimensions of the survey design. This choice is justified because the stratification variables (sector and size) are orthogonal to the geographic variation exploited by the broadband instrument. To assess sensitivity, we replicate the main IV specifications using sampling weights and find that the weighted estimates are quantitatively very similar to the baseline results. Event-study specifications are estimated without weights because the Callaway and Sant'Anna (2021) estimator does not accommodate probability weights.

A.2 Data integration and coverage

The survey is merged with administrative data covering the universe of registered firms, allowing us to jointly observe adoption outcomes and firm characteristics. The merged dataset links information on AI adoption, employment, financial accounts, trade activity, and local digital infrastructure conditions.

While the survey sampling frame includes only firms with at least 10 employees, a small number of firms in the merged dataset appear with fewer than 10 employees. This discrepancy arises because employment is measured using Social Security (SGK) records, which may differ from the business register used for survey sampling due to timing differences or workforce adjustments. These firms were part of the original 10+ sampling frame but subsequently appear below that threshold in the administrative data.

A.3 Measurement of AI adoption

AI adoption is measured using survey questions on the use of specific artificial intelligence technologies. Firms are asked whether they use each of the following technologies: text mining, speech recognition, natural language generation, image recognition, machine learning, AI-based process automation, and autonomous systems.

Our baseline measure of AI adoption is an indicator equal to one if the firm reports using at least one of these technologies. We also construct alternative measures based on the number and type of technologies used. In particular, the technology-specific indicators allow us to distinguish between software-intensive and hardware-intensive AI applications in the heterogeneity analysis.

A.4 International comparability

The ICT Usage Survey follows the Eurostat harmonized questionnaire, allowing direct comparison of AI adoption rates across countries. In particular, the definition and classification of AI technologies are aligned with those used in the Eurostat database *Artificial intelligence* *by size class of enterprise* (online data code: `isoc_eb_ai`), which reports firm-level adoption of AI by technology type, firm size class, and country for all EU member states. See https://ec.europa.eu/eurostat/cache/metadata/en/isoc_eb_esms.htm for additional details.

The harmonization covers both the list of AI technologies (including text mining, speech recognition, natural language generation, image recognition, machine learning, and AI-based automation systems) and the sampling frame, which restricts attention to firms with at least 10 employees in non-agricultural sectors. As a result, the Turkish data are directly comparable to the Eurostat aggregates without requiring further adjustments.

This comparability allows us to benchmark Turkey's level of AI adoption relative to European economies. Over the period 2021–2024, the share of firms using at least one AI technology in Turkey is approximately 9%, placing it in the lower quartile of the European distribution and below the EU average (around 13.5% in 2024). The gap is particularly pronounced for advanced AI applications, such as machine learning and natural language processing, which are more prevalent in high-income EU countries.

Beyond aggregate adoption rates, the harmonized design also enables comparison of the composition of AI use across technologies and firm size classes. In both Turkey and EU countries, AI adoption is strongly increasing in firm size, but the gradient is steeper in Turkey, reflecting lower adoption rates among small and medium-sized enterprises. These cross-country patterns provide useful context for interpreting our results and support the view that Turkey represents an emerging economy where AI diffusion remains relatively limited.

Appendix B Supplementary Descriptive Statistics

This appendix reports supplementary descriptive statistics for the main estimation sample and for key subsamples used in the paper. Table B.1 reports the number of firm-year observations by survey year and firm size class. The sample size is stable across waves, reflecting the repeated cross-sectional nature of the stratified survey design, with large firms consistently represented due to census sampling. Table B.2 reports AI adoption rates by EU firm-size class. Table B.3 summarizes the main firm-level outcomes, infrastructure measures, and treatment indicator for the full sample. Table B.4 reports adoption rates by AI category and specific AI technology. Table B.5 provides corresponding summary statistics for the low- and high-energy-intensity subsamples used in the mechanism analysis. Figures B.1–B.5 provide additional descriptive maps on the geographic distribution of treatment and AI adoption.

Table B.1: Number of Firms by Year and Size Class

	SMEs (<250)	Large (≥250)	Total
2021	8,767	5,320	14,087
2022	10,163	5,698	15,861
2023	9,933	6,009	15,942
2024	9,929	6,214	16,143
Total	38,792	23,241	62,033

Notes: The table reports the number of firm-year observations by size class. SMEs are firms with fewer than 250 employees; large firms have 250 or more employees.

Table B.2: AI Adoption by EU Firm Size Class

	Micro (<10)	Small (10–49)	Medium (50–249)	Large (≥250)
AI Adoption [%]	3.32	5.44	8.95	13.98

Notes: Share of firms reporting the use of at least one AI technology, by EU firm-size class based on number of employees. Pooled sample, 2021–2024.

Table B.3: Summary Statistics: Full Sample

	N	Mean	Std. Dev.	Min	Median	Max
AI Adoption [%]	62,033	9.171	28.862	0.000	0.000	100.000
Log Labor Productivity	60,072	9.576	1.351	0.031	9.561	16.454
Export Intensity [%]	62,033	23.605	36.009	0.000	0.000	97.771
Log ICT Workers	62,033	0.570	0.976	0.000	0.000	7.626
Log Non-ICT Workers	62,033	4.473	1.699	−0.539	4.382	11.381
Log Employment	62,033	4.498	1.691	0.000	4.407	11.381
Avg Download Speed [Mbps]	62,033	48.438	10.274	26.000	48.000	72.000
Average Latency [ms]	62,033	32.062	6.915	24.000	31.000	69.000
Distance to Data Center [km]	62,033	35.204	77.906	0.000	0.000	513.568
BOTAŞ Pipeline Treatment	62,033	0.320	0.466	0.000	0.000	1.000

Notes: Summary statistics for the full estimation sample. AI adoption is a binary indicator scaled to 0–100. Labor productivity is measured as log revenue per worker. Export intensity is the share of total revenue from exports. ICT and non-ICT workers are measured in logs. Download speed and latency are district-level averages. Distance to data center is a connectivity-based accessibility measure. BOTAŞ pipeline treatment indicates whether the district is connected in a given year.

Table B.4: AI Technology Adoption Rates

	N	Mean	Std. Dev.
Aggregate Measures			
Any AI Technology	62,033	0.092	0.289
By Functional Category			
Generative AI	62,033	0.057	0.232
Predictive AI	62,033	0.085	0.279
Hardware-Intensive AI	62,033	0.055	0.228
Software-Intensive AI	62,033	0.080	0.272
By Specific Technology			
Text Mining	62,033	0.033	0.178
Natural Language Processing	62,033	0.027	0.162
Deep Learning	62,033	0.049	0.216
Robotic Process Automation	62,033	0.055	0.229
Image Recognition	62,033	0.048	0.213
Speech Recognition	62,033	0.026	0.158
Autonomous Systems	62,033	0.025	0.156

Notes: Adoption rates for AI technologies across the full sample. All measures are binary indicators equal to one if the firm reports using the technology.

Table B.5: Summary Statistics by Energy Intensity

	Low Energy-Intensive			High Energy-Intensive		
	N	Mean	Std. Dev.	N	Mean	Std. Dev.
AI Adoption [%]	57,592	8.911	28.490	4,441	12.542	33.123
Log Labor Productivity	55,779	9.504	1.310	4,293	10.505	1.517
Export Intensity [%]	57,592	22.675	35.538	4,441	35.666	39.708
Log ICT Workers	57,592	0.571	0.987	4,441	0.554	0.820
Log Non-ICT Workers	57,592	4.457	1.704	4,441	4.677	1.627
Log Employment	57,592	4.484	1.695	4,441	4.682	1.628
Avg Download Speed [Mbps]	57,592	48.534	10.301	4,441	47.195	9.831
BOTAŞ Pipeline Treatment	57,592	0.324	0.468	4,441	0.271	0.444

Notes: Summary statistics separately for low- and high-energy-intensive industries. High-energy sectors are NACE 19, 20, 23, 24, and 35.

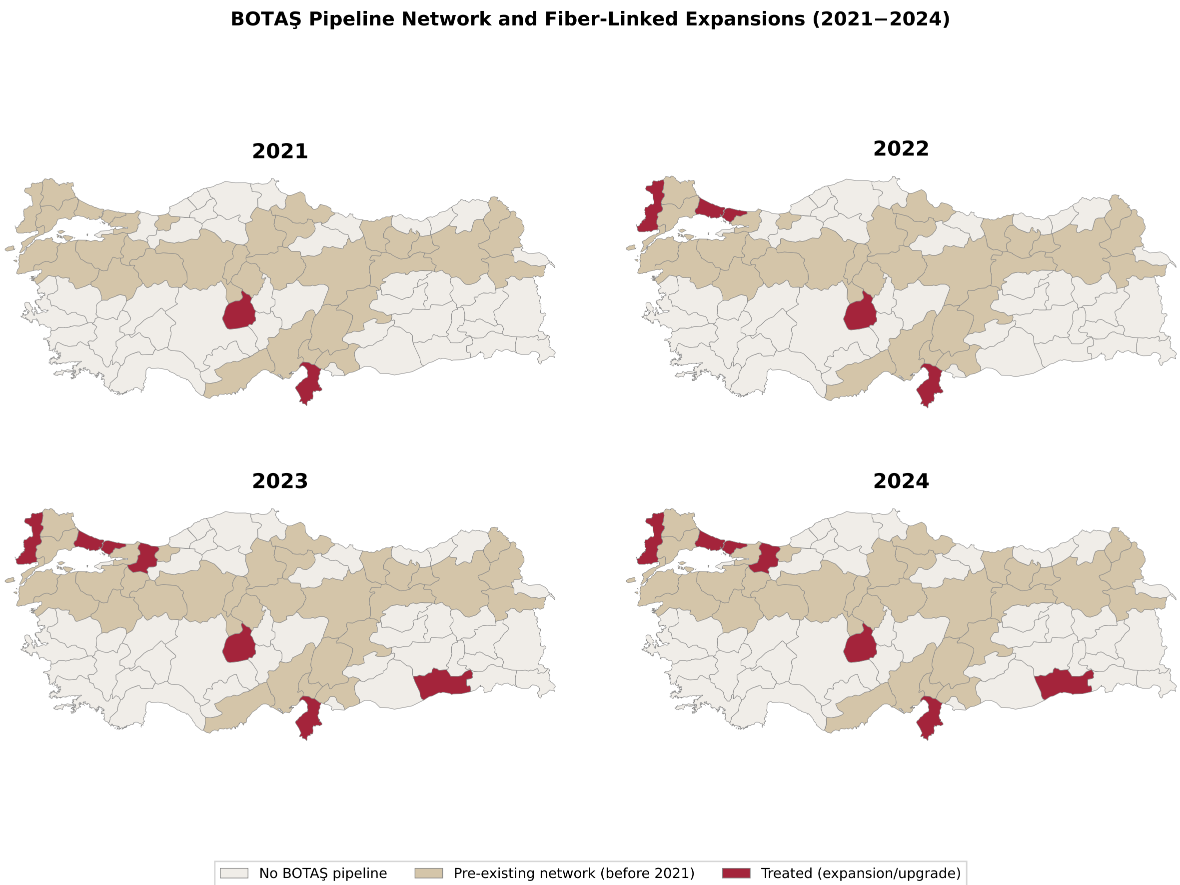

Figure B.1: BOTAŞ Pipeline Network and Fiber-Linked Expansions (2021–2024)

Notes: Beige provinces host pre-existing BOTAŞ trunk pipelines. Red provinces experience a fiber-linked pipeline expansion or upgrade during the sample period. No additional provinces are treated in 2024.

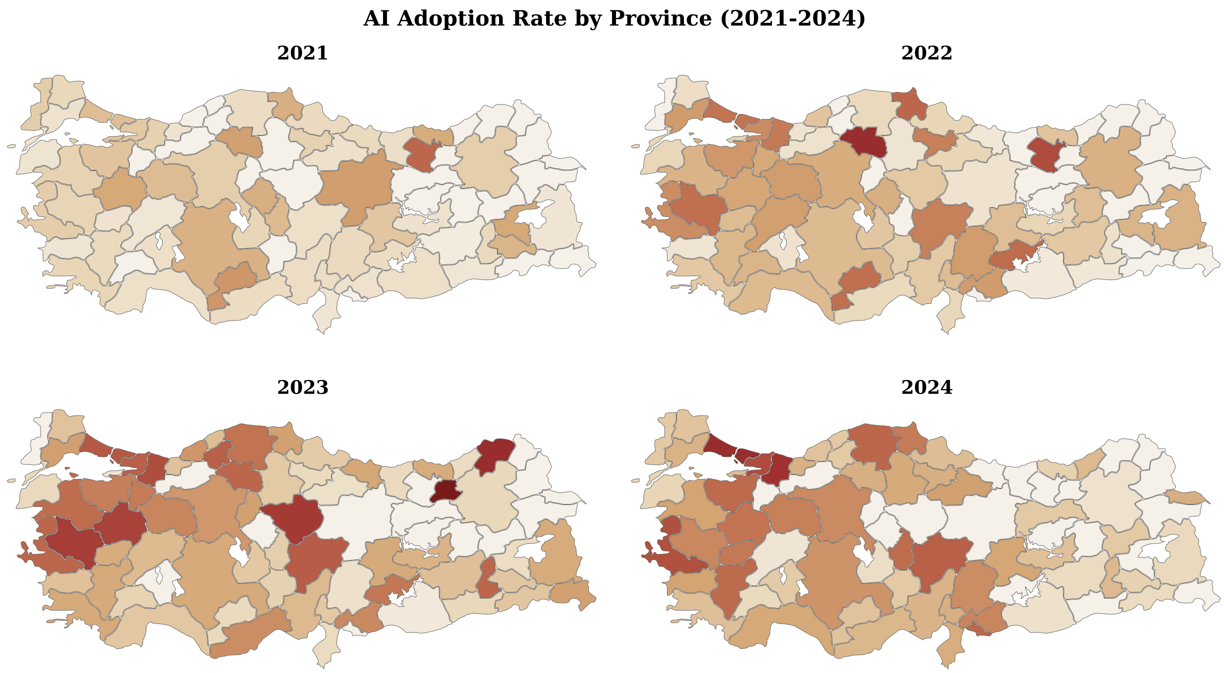

Figure B.2: AI Adoption Rate by Province (2021–2024)

Notes: Provinces are shaded by the share of firms reporting any AI use in the ICT survey. AI adoption increases from lighter to darker shades. The color scale is fixed across all four panels.

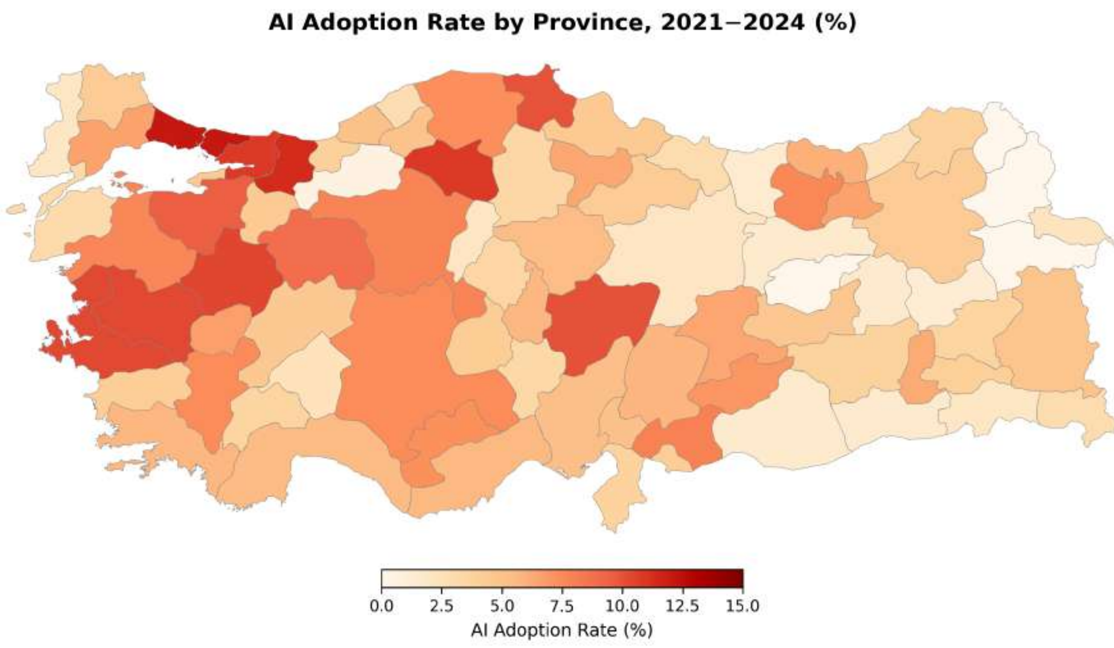

Figure B.3: AI Adoption Rate by Province, 2021–2024 (%)

Notes: Pooled mean across all firms.

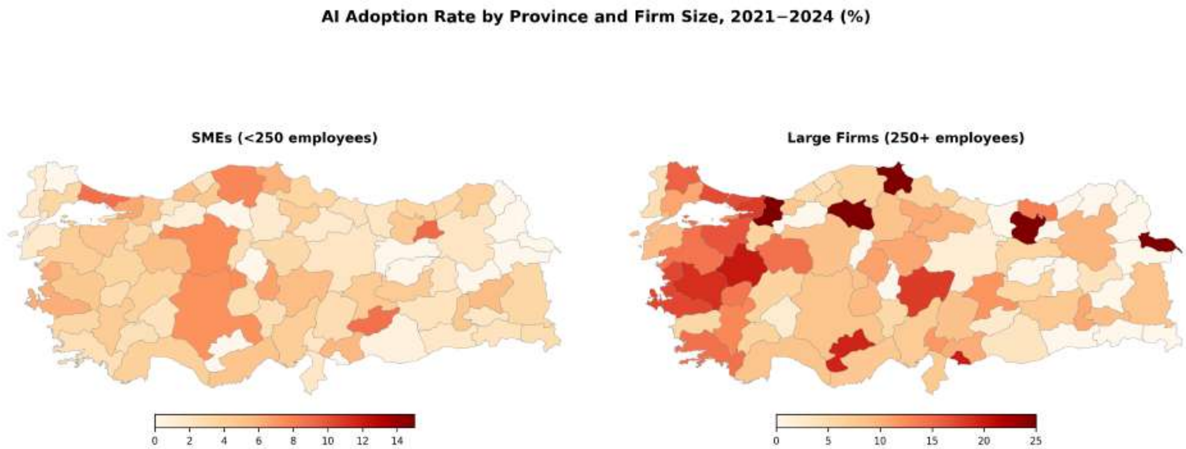

Figure B.4: AI Adoption Rate by Province and Firm Size, 2021–2024 (%)

Notes: Left: SMEs (<250, scale 0–15%). Right: Large (≥250, scale 0–25%).

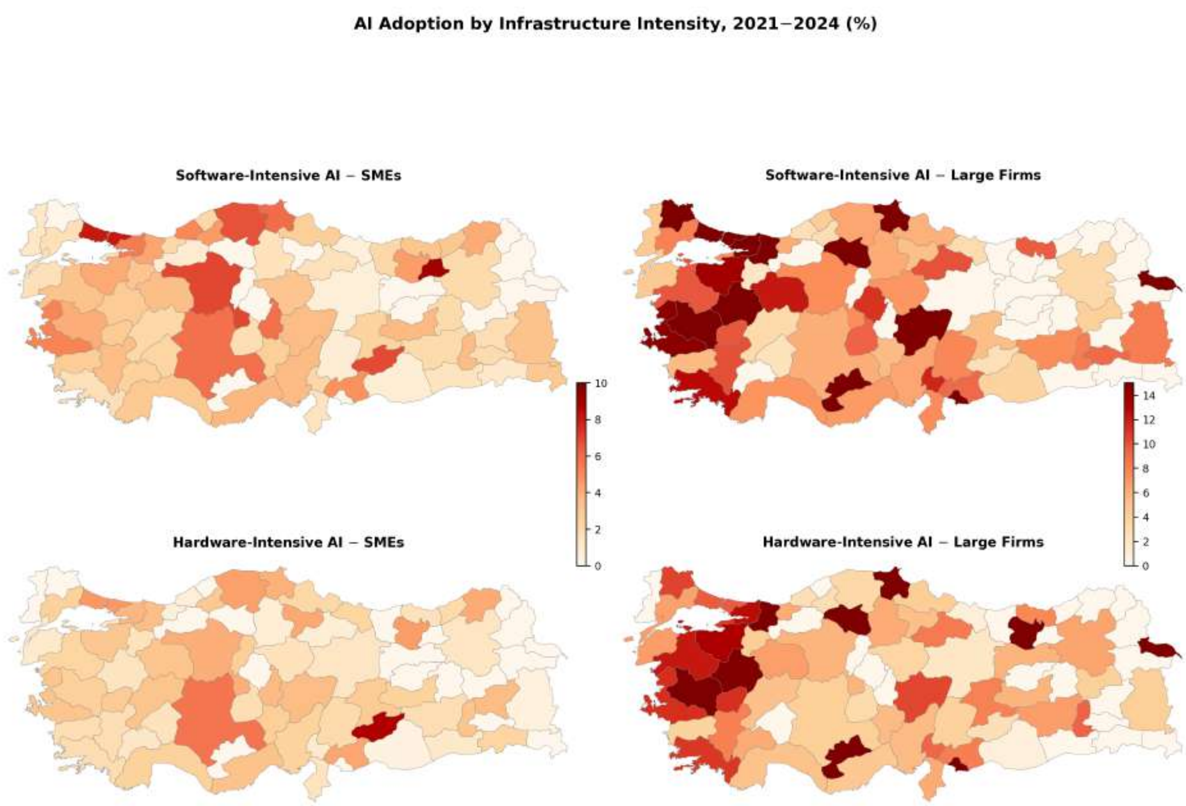

Figure B.5: AI Adoption by Infrastructure Intensity, 2021–2024 (%)

Notes: Top: software-intensive AI. Bottom: hardware-intensive AI. Left: SMEs. Right: large firms.

Appendix C Supplementary DiD Evidence

This appendix provides supplementary material for the difference-in-differences analysis in Section 5. Table C.1 defines the ISCO-08 occupations used to identify ICT workers in the employer–employee data. Figure C.1 reports event-study estimates for generative versus predictive AI. Table C.2 reports the underlying dynamic treatment-effect estimates for the technology-specific event studies, separately for SMEs and large firms. Table C.3 reports the first-stage validation for the triple-difference analysis. Figures C.2–C.3 and Table C.4 present pooled-sample event-study evidence. Finally, Table C.5 documents pre-treatment covariate balance.

Table C.1: ICT Workers by ISCO-08 Code

ISCO-08	**Title**
1330	ICT service managers
2511	Systems analysts
2512	Software developers
2513	Web and multimedia developers
2514	Applications programmers
2519	Software and applications developers and analysts
2521	Database designers and administrators
2522	Systems administrators
2523	Computer network professionals
2529	Database and network professionals
3511	ICT operations technicians
3513	Computer network and systems technicians
3514	Web technicians
3521	Broadcasting and audiovisual technicians
3522	Telecommunications engineering technicians

Notes: ISCO-08 occupational codes used to classify ICT workers in the employer–employee records.

Figure C.1: Dynamic Effects of Broadband Expansion on AI Adoption: Generative vs. Predictive AI

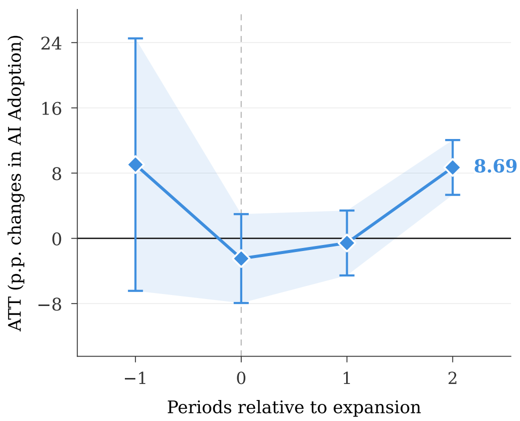

(a) Generative AI – SMEs

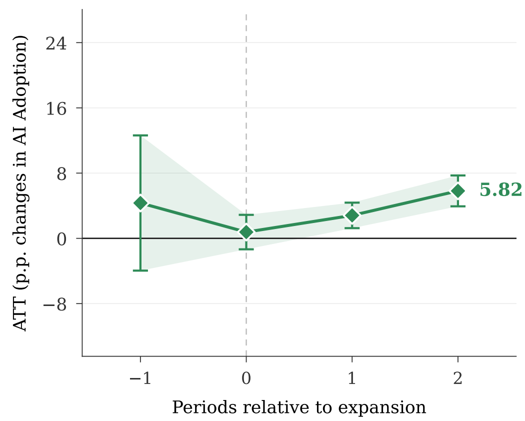

(b) Generative AI – Large Firms

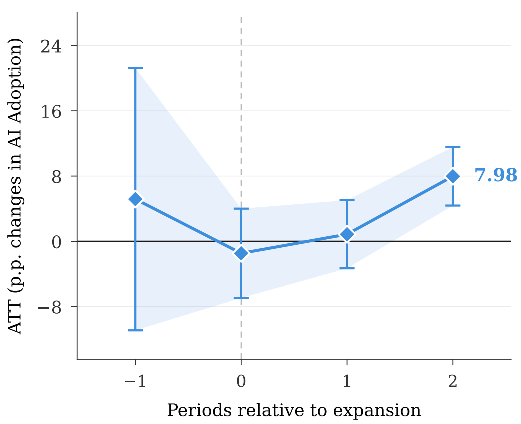

(c) Predictive AI – SMEs

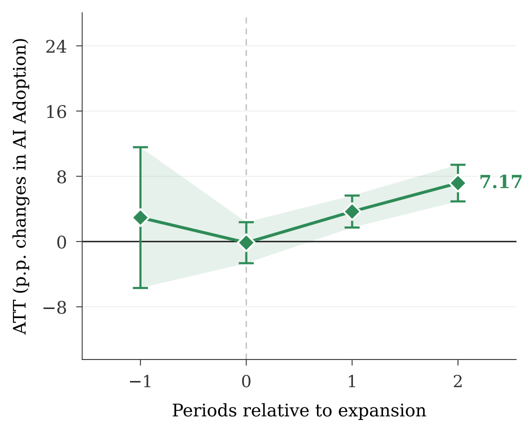

(d) Predictive AI – Large Firms

Notes: Each panel reports the dynamic average treatment effects on the treated (ATT) by event time, estimated using the staggered difference-in-differences procedure of Callaway and Sant'Anna (2021). The outcome is an indicator for whether the firm reports adopting at least one AI technology in the specified group. Generative AI comprises NLP and deep learning. Predictive AI encompasses text mining, speech recognition, image recognition, RPA, and autonomous robots or drones. Period 0 corresponds to the first year in which the BOTAŞ-linked fiber expansion becomes active in a district. The dashed vertical line marks the onset of treatment. Diamonds indicate point estimates; vertical bars show 95% confidence intervals; shaded bands depict the pointwise confidence region. Standard errors are clustered at the district level. All panels share a common vertical axis. Left column: SMEs (<250 employees); right column: large firms (≥250 employees).

Table C.2: Event Study Estimates: Impact of Fiber Expansion on Specific AI Technologies

Panel A: SMEs (<250 Employees)

	Text Mining	NLP	Deep Learning	Software Robotics	Image Rec.	Speech Rec.	Autonomous Robots
One year before event	−2.324	3.065	2.450	4.972	3.348	−2.849	8.133
	(4.358)	(5.885)	(5.989)	(7.719)	(5.920)	(3.435)	(5.554)
Event year (Expansion)	2.306**	−0.606	−0.801	−0.307	−0.363	2.154***	−0.962
	(0.780)	(1.995)	(2.373)	(2.321)	(2.353)	(0.466)	(2.315)
One year after event	2.899**	0.433	0.315	1.092	0.025	1.583*	−0.887
	(1.002)	(1.523)	(1.592)	(1.915)	(1.547)	(0.769)	(1.366)
Two years after event	7.791***	6.536***	6.432***	5.645***	3.168*	4.226***	1.767*
	(1.348)	(1.266)	(1.611)	(1.492)	(1.339)	(1.123)	(0.896)

Panel B: Large Firms (250+ Employees)

	Text Mining	NLP	Deep Learning	Software Robotics	Image Rec.	Speech Rec.	Autonomous Robots
One year before event	4.239	0.230	6.171	2.093	4.249	−2.019	3.496
	(2.932)	(1.692)	(4.427)	(3.920)	(4.170)	(2.368)	(2.970)
Event year (Expansion)	0.266	−0.218	0.984	0.634	−1.284	0.081	−0.965
	(0.962)	(0.921)	(1.067)	(1.205)	(1.167)	(0.933)	(0.970)
One year after event	2.128***	1.548**	2.631***	4.649***	1.025	1.472**	0.308
	(0.609)	(0.536)	(0.768)	(0.862)	(0.784)	(0.535)	(0.597)
Two years after event	5.209***	3.140***	5.043***	7.072***	2.584**	2.733***	1.354
	(0.809)	(0.732)	(0.915)	(0.991)	(0.910)	(0.701)	(0.707)

Notes: The table reports dynamic treatment effect estimates from the staggered difference-in-differences specification of Callaway and Sant'Anna (2021). The sample is split into SMEs (<250 employees) in Panel A and large firms (250+ employees) in Panel B. Dependent variables are binary indicators multiplied by 100 (percentage points). Standard errors clustered at the district level in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table C.3: Triple-Difference First Stage: AI Adoption

	(1) Software AI	(2) Hardware AI	(3) Any AI
Panel A: SMEs (<250 employees)			
Pipeline × SW Exposure	1.171**	0.144	1.018**
	(0.345)	(0.194)	(0.319)
Pipeline × HW Exposure	−0.315	0.060	−0.230
	(0.277)	(0.251)	(0.314)
Pipeline	1.698***	0.439	1.524***
	(0.371)	(0.254)	(0.394)
Observations	38,734	38,734	38,734
R^2	0.110	0.056	0.105
Panel B: Large Firms (≥250 employees)			
Pipeline × SW Exposure	3.015***	2.890***	3.315***
	(0.567)	(0.329)	(0.541)
Pipeline × HW Exposure	−0.716	−0.151	−0.958
	(0.436)	(0.528)	(0.487)
Pipeline	3.515**	0.474	3.217**
	(1.048)	(0.657)	(1.138)
Observations	23,218	23,218	23,218
R^2	0.196	0.157	0.196
Industry FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Size Controls	Yes	Yes	Yes

Notes: Each column regresses an AI adoption indicator (scaled to percentage points) on pipeline treatment interacted with pre-determined industry-level AI-type exposure (standardized 2021 baseline adoption shares), separately for SMEs and large firms. All specifications include industry fixed effects, year fixed effects, log employment, and firm-size category controls. Standard errors clustered at the district level in parentheses. * p<0.05, ** p<0.01, *** p<0.001.

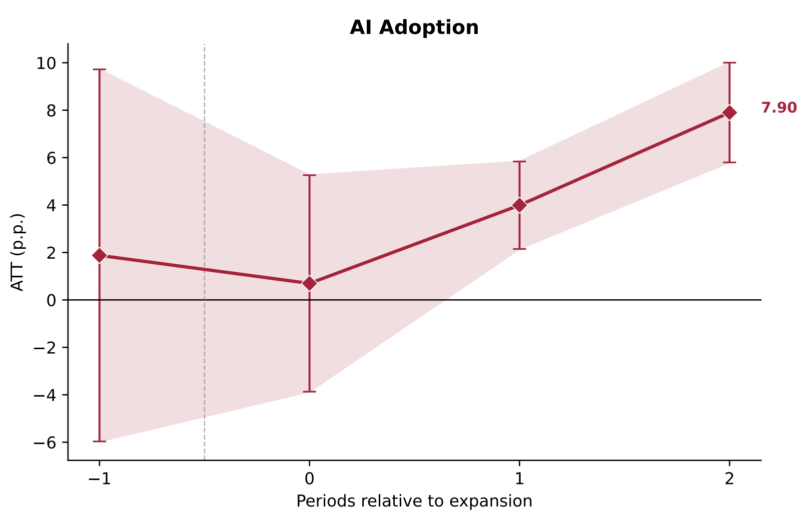

Figure C.2: Effect of BOTAŞ Fiber Expansion on AI Adoption — Pooled Sample

Notes: Callaway and Sant'Anna (2021) ATT estimates, pooled sample (SMEs + large firms). Outcome: AI adoption (any technology, percentage points). Period 0 is the first year of expansion. Standard errors are clustered at the district level.

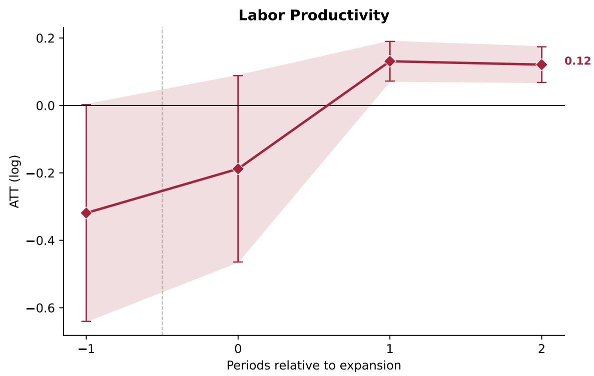

(a) Labor Productivity

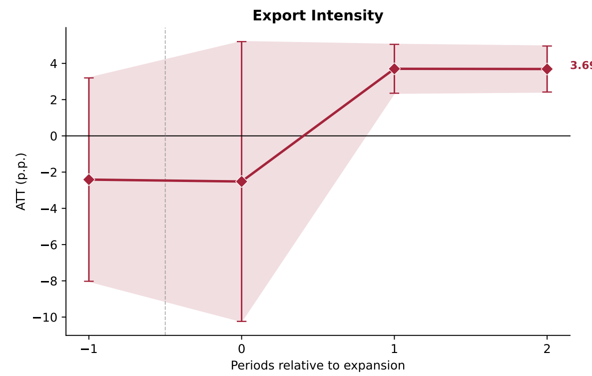

(b) Export Intensity

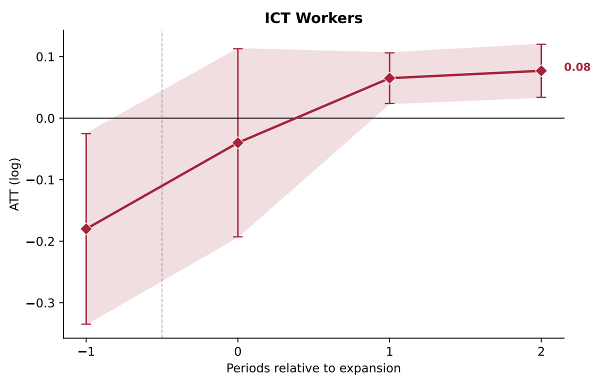

(c) ICT Workers

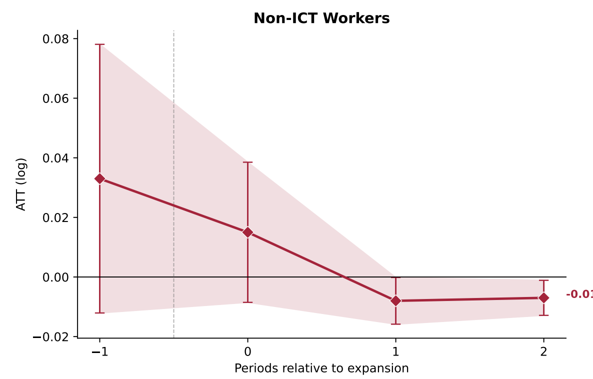

(d) Non-ICT Workers

Figure C.3: Dynamic Effects of Broadband Expansion on Firm Outcomes — Pooled Sample

Notes: Callaway and Sant'Anna (2021) ATT estimates, pooled sample. Panel (a) reports log labor productivity; panel (b) export intensity (percentage points); panel (c) log ICT workers; and panel (d) log non-ICT workers. Standard errors are clustered at the district level.

Table C.4: Event Study Estimates: Pooled Sample (SMEs + Large Firms)

	(1) AI Adoption (p.p.)	(2) Labor Prod. (log)	(3) Export Int. (p.p.)	(4) ICT Workers (log)	(5) Non-ICT Workers (log)
$k = -1$ (pre-trend)	1.877	−0.319	−2.413	−0.180*	0.033
	(4.002)	(0.164)	(2.863)	(0.079)	(0.023)
$k = 0$ (expansion year)	0.696	−0.188	−2.518	−0.040	0.015
	(2.329)	(0.141)	(3.939)	(0.078)	(0.012)
$k = +1$	3.992***	0.131***	3.701***	0.065**	−0.008*
	(0.941)	(0.030)	(0.688)	(0.021)	(0.004)
$k = +2$	7.900***	0.121***	3.689***	0.077***	−0.007**
	(1.073)	(0.027)	(0.648)	(0.022)	(0.003)

Notes: Callaway and Sant'Anna (2021) staggered difference-in-differences estimates for the pooled sample. Standard errors are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

Table C.5: Pre-Treatment Covariate Balance: Treated vs. Not-Yet-Treated (2021)

	SMEs (<250)		Large (≥250)	
	Treated	Control	Treated	Control
Log employment	3.43	3.23	6.30	6.20
Employment (levels)	48.1	36.6	847.5	697.9
AI adoption (%)	3.72	1.00	6.19	4.00
Software AI (%)	3.06	1.00	5.03	4.00
Hardware AI (%)	1.96	0.00	3.49	2.00
Log labor productivity	9.12	8.71	9.80	9.67
Export intensity (%)	14.4	19.7	27.3	22.0
Download speed (Mbps)	42.1	35.7	42.7	35.6
Latency (ms)	29.5	42.5	28.7	42.5

Notes: Baseline (2021) means for firms in early-treated versus not-yet-treated districts. Treated districts are those connected to the BOTAŞ-linked fiber network by 2021. Infrastructure variables differ by construction, since treatment directly affects broadband quality. The DRIPW estimator conditions on four-digit NACE Rev. 2 industry and log employment, adjusting for baseline differences through both propensity-score reweighting and outcome regression.

Appendix D Supplementary IV Evidence

This appendix provides additional evidence supporting the instrumental-variables analysis in Section 6. It reports results on alternative infrastructure channels, reduced-form estimates, firm-size heterogeneity, technology-specific effects, robustness to survey weights, and tests distinguishing digital from energy channels. It also assesses robustness to non-random network exposure using the recentering procedure of Borusyak and Hull (2023).

Alternative infrastructure channels. Table D.1 compares alternative infrastructure margins, focusing on download speed, latency, and data-center proximity. The baseline results are strongest for download speed and data-center access, whereas latency displays a weak first stage and correspondingly imprecise second-stage estimates.

Table D.1: Alternative Infrastructure Channels: Download Speed, Latency, and Data Center Proximity

	(1) Download Speed (Baseline)	(2) Latency (Alternative)	(3) DC Distance (Alternative)
Panel A: First-Stage – Infrastructure			
BOTAŞ Pipeline	1.209***	−0.199	−1.042***
	(0.147)	(0.262)	(0.178)
First-stage F-statistic	20.19	2.17	19.66
Panel B: First-Stage – AI Adoption			
Predicted Infrastructure	2.030***	−12.318***	−2.354***
	(0.435)	(2.640)	(0.505)
Panel C: Second-Stage – Labor Productivity			
Broadband-Induced AI Adoption	0.060**	0.296	0.083***
	(0.020)	(0.204)	(0.018)
Industry × Year FE	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes
Observations	60,072	60,072	60,072

Notes: Re-estimates of the IV chain using three alternative infrastructure measures. Latency exhibits a weak first stage, whereas download speed and data-center distance display strong first stages. Standard errors are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

Reduced-form evidence. Table D.2 reports reduced-form regressions of pipeline exposure on firm outcomes. These estimates provide a benchmark for the IV magnitudes and confirm that pipeline exposure is associated with higher productivity, export intensity, and

ICT employment, alongside modest declines in non-ICT employment.

Table D.2: Reduced Form: Direct Effects of Pipeline Infrastructure on Firm Outcomes

	(1) Log Labor Productivity	(2) Export Intensity	(3) Log ICT Workers	(4) Log Non-ICT Workers
BOTAŞ Pipeline	0.143***	4.991***	0.113***	−0.006***
	(0.037)	(0.826)	(0.010)	(0.001)
Industry × Year FE	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes
Observations	60,072	62,033	62,033	62,033
Mean Dependent Variable	9.576	23.605	0.570	4.473

Notes: Reduced-form OLS estimates of pipeline exposure on firm outcomes. Standard errors are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

Firm-size heterogeneity. Tables D.3 and D.4 report IV estimates separately for SMEs and large firms. Consistent with the main text, the performance effects are stronger for SMEs, whereas large-firm estimates are generally smaller and less precisely estimated, reflecting the weaker first stage in that subsample.

Table D.3: LATE Estimates by Firm Size: SMEs versus Large Firms

	SMEs (<250 employees)		Large Firms (≥250 employees)	
	(1) Log Labor Productivity	(2) Export Intensity	(3) Log Labor Productivity	(4) Export Intensity
Panel A: First-Stage Diagnostics				
Infrastructure (Speed ← Pipeline):				
BOTAŞ Pipeline	1.190***	1.190***	1.217***	1.217***
	(0.159)	(0.159)	(0.127)	(0.127)
Adoption (AI ← Predicted Speed):				
Predicted Speed ($\hat{z}_{\mathrm{dl}}$)	1.332***	1.332***	2.553**	2.553**
	(0.323)	(0.323)	(0.959)	(0.959)
First-stage F-statistic	16.37	16.97	6.59	7.09
Panel B: Second-Stage Estimates				
Broadband-Induced AI Adoption	0.106***	3.510***	0.037*	1.717**
	(0.029)	(0.789)	(0.018)	(0.565)
Industry × Year FE	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes
Observations	37,119	38,792	22,953	23,241

Notes: Two-stage least squares estimates separately for SMEs and large firms. All specifications include four-digit NACE Rev. 2 industry-year fixed effects, size controls, and log employment. Standard errors are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

Technology-specific IV estimates. Table D.5 reports technology-specific IV estimates. As in the main text, software-intensive and cloud-dependent technologies exhibit strong first

Table D.4: Employment Effects by Firm Size

	SMEs (<250 employees)		Large Firms (≥250 employees)	
	(1) Log ICT Workers	(2) Log Non-ICT Workers	(3) Log ICT Workers	(4) Log Non-ICT Workers
Panel A: First-Stage Diagnostics				
First-stage F-statistic	16.97	16.97	7.09	7.09
Panel B: Second-Stage Estimates				
Broadband-Induced AI Adoption	0.050***	−0.005***	0.043**	−0.001
	(0.011)	(0.001)	(0.015)	(0.001)
Industry × Year FE	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes
Observations	38,792	38,792	23,241	23,241

Notes: Two-stage least squares estimates for employment outcomes by firm size. Standard errors are clustered at the district level. * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

stages and economically meaningful second-stage effects, whereas image recognition is more weakly identified.

Robustness to survey weights. Table D.6 compares weighted and unweighted IV estimates for the SME subsample. The results are very similar across specifications, indicating that the main IV findings are not driven by the weighting structure of the ICT survey.

Digital versus energy channels. Tables D.7–D.9 examine whether the IV estimates are more consistent with a digital-infrastructure channel than with an energy-cost channel. Across specifications, the evidence remains more supportive of the digital interpretation: the strongest and most stable effects are concentrated in software-intensive AI and in lower-energy settings where broadband constraints are likely to bind more directly.

Robustness to non-random network exposure. The baseline IV design exploits district-level variation in the timing of BOTAŞ-linked broadband expansion. While the institutional setting supports the view that this timing is plausibly exogenous, a remaining concern is that districts with greater expected exposure to the BOTAŞ network may differ systematically from other districts in ways that matter for firm outcomes. For example, districts located along major transport corridors or closer to gas-exporting borders may be structurally more

Table D.5: Heterogeneous Effects by Specific AI Technology

	(1) Text Mining	(2) Natural Language	(3) Deep Learning	(4) Robotic Process	(5) Image Recognition	(6) Speech Recognition
Panel A: First-Stage Diagnostics						
BOTAŞ Pipeline → Technology	1.715***	1.107***	1.865***	2.285***	0.661**	0.897***
	(0.178)	(0.158)	(0.255)	(0.364)	(0.223)	(0.148)
First-stage F-statistic	87.48	46.08	46.16	34.42	7.78	32.71
Panel B: Second-Stage – Labor Productivity						
Infrastructure-Induced Adoption	0.084***	0.130***	0.081***	0.065**	0.235*	0.167***
	(0.021)	(0.037)	(0.022)	(0.022)	(0.101)	(0.048)
Panel C: Second-Stage – Export Intensity						
Infrastructure-Induced Adoption	2.911***	4.509***	2.676***	2.184***	7.551***	5.563***
	(0.484)	(0.786)	(0.415)	(0.301)	(2.093)	(0.918)
Panel D: Second-Stage – ICT Workers						
Infrastructure-Induced Adoption	0.066***	0.102***	0.060***	0.049***	0.170**	0.126***
	(0.007)	(0.016)	(0.008)	(0.007)	(0.057)	(0.021)
Panel E: Second-Stage – Non-ICT Workers						
Infrastructure-Induced Adoption	−0.003***	−0.005***	−0.003***	−0.003***	−0.009***	−0.006***
	(0.001)	(0.001)	(0.001)	(0.000)	(0.003)	(0.001)
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations (Productivity)	60,072	60,072	60,072	60,072	60,072	60,072
Observations (Other Outcomes)	62,033	62,033	62,033	62,033	62,033	62,033

Notes: Technology-specific IV estimates. Image recognition has a weaker first stage and should be interpreted more cautiously. Standard errors are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

Table D.6: IV Estimates With and Without Survey Weights: SMEs

	Log Labor Prod.		Export Intensity		Log ICT Workers		Log Non-ICT Workers	
	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd
Panel A: Second-Stage Estimates								
Broadband-Induced AI	0.106***	0.115***	3.510***	3.758***	0.050***	0.044***	−0.005***	−0.005***
	(0.029)	(0.033)	(0.789)	(0.904)	(0.011)	(0.010)	(0.001)	(0.001)
Panel B: First-Stage Diagnostics								
BOTAŞ → Speed	1.190***	1.186***	1.190***	1.186***	1.190***	1.186***	1.190***	1.186***
	(0.159)	(0.165)	(0.159)	(0.165)	(0.159)	(0.165)	(0.159)	(0.165)
Speed → AI	1.332***	1.201***	1.332***	1.201***	1.332***	1.201***	1.332***	1.201***
	(0.323)	(0.312)	(0.323)	(0.312)	(0.323)	(0.312)	(0.323)	(0.312)
First-stage F	16.37	14.52	16.97	14.82	16.97	14.82	16.97	14.82
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Size Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Survey Weights	No	Yes	No	Yes	No	Yes	No	Yes
Observations	37,119	37,119	38,792	38,792	38,792	38,792	38,792	38,792

Notes: This table compares unweighted (baseline) and survey-weighted IV estimates for the SME subsample (<250 employees). Unweighted columns reproduce the results from Table D.3. Weighted columns incorporate the stratified ICT survey probability weights, which account for differential sampling rates across firm-size strata. Large firms (≥250 employees) are excluded because they are sampled with certainty ($w_i = 1$). Standard errors clustered at the district level in parentheses. * p<0.05, ** p<0.01, *** p<0.001.

Table D.7: AI Adoption Effects Interacted with Energy Intensity

	(1) Log Labor Productivity	(2) Export Intensity	(3) ICT Workers	(4) Non-ICT Workers
Broadband-Induced AI Adoption	0.061**	2.045***	0.047***	−0.002***
	(0.020)	(0.273)	(0.008)	(0.000)
AI × Energy-Intensive	−0.010	−0.191	−0.020***	0.001**
	(0.013)	(0.326)	(0.005)	(0.000)
High Energy-Intensive Industries	−0.150	−13.094	0.259	0.016*
	(0.490)	(7.213)	(0.190)	(0.007)
Industry × Year FE	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes
Observations	60,072	62,033	62,033	62,033

Notes: IV estimates allowing the effect of AI adoption to vary by energy intensity. Energy-intensive sectors are NACE 19, 20, 23, 24, and 35. Standard errors in parantheses are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

likely to receive pipeline infrastructure earlier. As Borusyak and Hull (2023) show, when treatment assignment contains both structural and idiosyncratic components, IV estimates may be biased if the structural component is correlated with potential outcomes.

To address this concern, we construct a recentered instrument that isolates the idiosyncratic component of the BOTAŞ expansion. We estimate a district-year assignment model in which treatment status depends on predetermined structural characteristics, including

Table D.8: IV Estimates by Energy Intensity Subsample

	Log Labor Productivity		Export Intensity		ICT Workers		Non-ICT Workers	
	High Energy	Low Energy	High Energy	Low Energy	High Energy	Low Energy	High Energy	Low Energy
Broadband-Induced AI	0.063*	0.059**	1.814**	2.042***	0.019*	0.048***	−0.001*	−0.002***
	(0.032)	(0.020)	(0.573)	(0.278)	(0.007)	(0.009)	(0.000)	(0.000)
First-Stage Coefficient	3.432	2.415	3.432	2.415	3.432	2.415	3.432	2.415
	(1.479)	(0.481)	(1.479)	(0.481)	(1.479)	(0.481)	(1.479)	(0.481)
First-Stage F-statistic	5.52	23.28	5.38	25.17	5.38	25.17	5.38	25.17
Industry × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Size Category Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,293	55,779	4,441	57,592	4,441	57,592	4,441	57,592

Notes: IV estimates separately for high- and low-energy industries. The high-energy subsample is small and its first-stage F-statistics fall below the conventional threshold, so those estimates should be interpreted cautiously. Standard errors in brackets are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

Table D.9: Technology-Specific Productivity Effects Across Energy Intensity Groups

	Text Mining	NLP	Deep Learning	RPA	Image Recog.	Speech Recog.	Autonomous Systems
Panel A: High Energy-Intensive Industries							
Infrastructure-Induced Adoption	0.110*	0.464	0.081*	0.086	0.144	0.136*	0.150
	(0.054)	(0.529)	(0.034)	(0.046)	(0.116)	(0.061)	(0.087)
First-Stage F-statistic	5.48	0.81	11.17	4.07	1.98	7.38	3.73
First-Stage Coefficient	2.181	0.557	2.792	2.669	1.642	1.814	1.419
	(0.824)	(0.535)	(0.790)	(1.252)	(1.073)	(0.574)	(0.748)
Observations	4,293	4,293	4,293	4,293	4,293	4,293	4,293
Panel B: Low Energy-Intensive Industries							
Infrastructure-Induced Adoption	0.081***	0.119***	0.080***	0.063**	0.244*	0.170***	2.829
	(0.021)	(0.035)	(0.021)	(0.021)	(0.103)	(0.051)	(10.411)
First-Stage F-statistic	108.58	59.47	44.69	44.36	8.79	39.08	0.07
First-Stage Coefficient	1.696	1.156	1.821	2.278	0.614	0.836	0.122
	(0.161)	(0.147)	(0.255)	(0.317)	(0.195)	(0.128)	(0.178)
Observations	55,779	55,779	55,779	55,779	55,779	55,779	55,779

Notes: IV estimates of the productivity effect of infrastructure-induced AI adoption separately by technology and energy intensity. The dependent variable is log labor productivity in all columns. Most first-stage F-statistics in the high-energy subsample fall below 10 and should be interpreted cautiously. Standard errors in brackets are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

distance to the nearest data center, baseline employment density, baseline number of firms, average download speed, and average latency, together with year effects capturing aggregate rollout trends. Let μ_{dt} denote the predicted treatment probability from this model. The recentered instrument is then defined as

$$\widetilde{D}_{dt} = D_{dt} - \mu_{dt}, \tag{17}$$

where D_{dt} is the observed treatment indicator. Positive values of $\widetilde{D}_{dt}$ indicate districts connected earlier than their structural characteristics would predict, while negative values indicate later-than-expected connection. By construction, $\widetilde{D}_{dt}$ is mean-zero conditional on the structural determinants, so identification relies on residual, as-good-as-random variation in treatment timing.

Table D.10 reports baseline and recentered IV estimates side by side for SMEs and large firms. The SME results are highly robust. Export intensity, ICT employment, and non-ICT employment effects are virtually unchanged, while the productivity coefficient declines modestly but remains statistically significant. The first-stage F-statistic falls, as expected, when identification relies on the smaller idiosyncratic share of the treatment variation. For large firms, the productivity effect becomes statistically insignificant, but the export-intensity and ICT-employment effects remain significant and economically meaningful. This pattern is consistent with the weaker first stage already present in the large-firm subsample.

Table D.10: Baseline versus Recentered IV Estimates by Firm Size

	SMEs (<250)		Large (≥250)	
	Baseline	Recentered	Baseline	Recentered
	Log Labor Productivity			
Broadband-Induced AI	0.106***	0.084*	0.037*	0.029
	(0.029)	(0.033)	(0.018)	(0.018)
First-stage F	16.37	8.23	6.59	6.05
	Export Intensity			
Broadband-Induced AI	3.510***	3.499***	1.717**	1.469**
	(0.789)	(0.922)	(0.565)	(0.529)
First-stage F	16.97	8.25	7.09	6.50
	Log ICT Workers			
Broadband-Induced AI	0.050***	0.048***	0.043**	0.037**
	(0.011)	(0.011)	(0.015)	(0.014)
First-stage F	16.97	8.25	7.09	6.50
	Log Non-ICT Workers			
Broadband-Induced AI	−0.005***	−0.005***	−0.001	−0.001
	(0.001)	(0.001)	(0.001)	(0.001)
First-stage F	16.97	8.25	7.09	6.50
Ind. × Year FE	Yes	Yes	Yes	Yes
Size Controls	Yes	Yes	Yes	Yes
Observations	37,119/38,792	37,119/38,792	22,953/23,241	22,953/23,241

Notes: Baseline columns reproduce the estimates from Tables D.3 and D.4. Recentered columns replace the observed BOTAŞ treatment indicator with the recentered instrument $\widetilde{D}_{dt} = D_{dt} - \mu_{dt}$, where μ_{dt} is the predicted treatment probability from a district-year assignment model following Borusyak and Hull (2023). The sequential IV chain is $\widetilde{D}_{dt}$ → download speed → AI adoption → firm outcomes. Observation counts differ across outcomes because value added is missing for some firm-years in the productivity regressions. Standard errors are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

Table D.11 reports the corresponding falsification exercise using autonomous systems. The first-stage F-statistics remain far below conventional thresholds, confirming that neither the baseline nor the recentered instrument predicts hardware-intensive AI adoption. The second-stage coefficients remain large, noisy, and statistically insignificant. This reinforces the exclusion restriction: even after removing the structural geographic component of pipeline exposure, the identifying variation does not operate through hardware-intensive automation.

Table D.11: Recentered IV: Autonomous Systems Negative Control

	(1) Log Labor Prod.	(2) Export Int.	(3) Log ICT	(4) Log Non-ICT
Panel A: Baseline				
Induced Adoption	1.105	25.824	0.583	−0.030
	(1.724)	(25.532)	(0.593)	(0.029)
First-stage F	0.95	0.95	0.95	0.95
Panel B: Recentered ($\widetilde{D}_{dt}$ as instrument)				
Induced Adoption	0.627	20.135	0.449	−0.024
	(0.810)	(16.966)	(0.390)	(0.020)
First-stage F	0.70	1.30	1.30	1.30
Ind. × Year FE	Yes	Yes	Yes	Yes
Size Controls	Yes	Yes	Yes	Yes
Observations	60,072	62,033	62,033	62,033

Notes: Panel A reproduces the baseline autonomous-systems falsification from Table 11. Panel B replaces the BOTAŞ treatment indicator with the recentered instrument $\widetilde{D}_{dt}$ following Borusyak and Hull (2023). The near-zero first stage confirms that neither the observed nor the recentered pipeline variation predicts hardware-intensive AI adoption. Standard errors are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.

Finally, Table D.12 reports the pooled-sample recentered IV estimates. Three of the four outcomes retain both sign and statistical significance under the recentered instrument, while the productivity coefficient is attenuated and becomes statistically insignificant. The first-stage F-statistic declines from the baseline but remains above conventional thresholds in the pooled sample. Overall, the recentered results indicate that the main findings are not driven by the structural geographic placement of the BOTAŞ network but instead by the idiosyncratic timing variation exploited in the design.

Table D.12: Baseline versus Recentered IV Estimates: Pooled Sample

	(1) Log Labor Prod.	(2) Export Int.	(3) Log ICT	(4) Log Non-ICT
Panel A: Baseline IV (observed BOTAŞ treatment)				
Broadband-Induced AI	0.060**	2.035***	0.046***	−0.002***
	(0.020)	(0.270)	(0.009)	(0.000)
First-stage F	20.19	21.77	21.77	21.77
Panel B: Recentered IV ($\widetilde{D}_{dt}$ as instrument)				
Broadband-Induced AI	0.045	1.894***	0.042***	−0.002***
	(0.023)	(0.273)	(0.009)	(0.000)
$\widetilde{D}_{dt} \rightarrow$ Speed	1.277	1.277	1.277	1.277
	(0.269)	(0.269)	(0.269)	(0.269)
Speed $\rightarrow$ AI	2.203	2.203	2.203	2.203
	(0.568)	(0.568)	(0.568)	(0.568)
First-stage F	14.28	15.07	15.07	15.07
Ind. $\times$ Year FE	Yes	Yes	Yes	Yes
Size Controls	Yes	Yes	Yes	Yes
Observations	60,072	62,033	62,033	62,033

Notes: Panel A reproduces the baseline IV estimates from Table 8. Panel B replaces the BOTAŞ treatment indicator with the recentered instrument $\widetilde{D}_{dt} = D_{dt} - \mu_{dt}$, where μ_{dt} is the predicted treatment probability from a district-year assignment model following Borusyak and Hull (2023). The sequential IV chain is $\widetilde{D}_{dt} \rightarrow$ download speed $\rightarrow$ AI adoption $\rightarrow$ firm outcomes. Standard errors are clustered at the district level. * p<0.05, ** p<0.01, *** p<0.001.