Artificial Intelligence (AI) and Endogenous Productivity: Evidence from South Korean Firms *

Ziyao Wang† Jae Wook Jung ‡

January 30, 2026

Click here for the most recent version

Abstract

This paper studies how Artificial Intelligence (AI) adoption affects firm-level productivity in South Korea. Using a 2017–2023 panel covering all major market sectors, we combine direct survey measures of AI adoption with a structural model of endogenous productivity to address selection into adoption. We find AI adoption raises productivity by 5% on average, with substantial heterogeneity by time since adoption that is consistent with a delayed J-curve, by sector with the largest gains in Information and Communication Technology (ICT) and losses in Manufacturing, and by application area with gains concentrated in product/service development and sales and marketing rather than production processes. Using the estimated production parameters, we further examine whether AI adoption is associated with changes in firms' market power.

Keywords: Artificial Intelligence (AI), Technology Adoption, Endogenous Productivity, Production Estimation, Market Power, South Korean Firms
JEL Codes: D24, L23, O31, O33

*We are grateful to Santiago Caicedo, Jianfei Cao, James Dana, Joe Emmens, Jordi Jaumandreu, John Kwoka, Matthias Mertens, Imke Reimers, and participants of the Northeastern University Economics Workshop, MIT Futuretech Seminar, for their insightful suggestions. Any remaining errors are our sole responsibility.

†Ph.D. Candidate in Economics, Northeastern University. Email: wang.ziyao1@northeastern.edu
‡Ph.D. Candidate in Economics, Northeastern University. Email: jung.jaew@northeastern.edu

1 Introduction

Artificial Intelligence (AI) is widely viewed as a General-Purpose Technology (GPT) with the potential to reshape firm productivity across a broad range of industries. But its effects are unlikely to be immediate or uniform across firms. Productivity gains may come with delay, and the magnitude of the impact may vary across sectors and across the application areas in which AI is deployed within the firm. However, rigorous firm-level evidence remains limited: partly because available datasets cover only a short recent window of years, and firms do not adopt AI at random: adoption decisions are correlated with underlying productivity.

This paper asks whether AI adoption raises firm productivity, and how the effects vary across time, sectors, and application areas. To study these questions, we estimate the effect of AI adoption on firm endogenous productivity in South Korea from 2017 to 2023. We find that AI adoption increases firms' productivity by about 5% on average, with substantial heterogeneity over time, across sectors, and across AI application areas.

We use the Korean Statistical Office's Survey of Business Activities (SBA), an annual panel of incorporated firms covering all major market sectors from 2017 to 2023. The survey provides a particularly informative setting for studying AI as a general-purpose technology. First, it directly records organizational AI use over a seven-year horizon rather than relying on indirect proxies. Second, conditional on adoption, it records the primary area of application-product/service development, production processes, sales/marketing, or organizational management—allowing us to link productivity effects to how AI is deployed inside the firm.

Our empirical challenge is the self-selection of AI adoption: the AI adoption decision is correlated with firm productivity. To address this concern, we model productivity as a latent firm-specific state that evolves dynamically and allow AI adoption to shift the productivity transition, motivated by the estimating endogenous productivity literature (Doraszelski and Jaumandreu, 2013; De Loecker, 2013). The framework accounts for the endogeneity of AI

adoption decisions and allows us to recover both production parameters, productivity, and the dynamic role of AI in shaping productivity.

Although the empirical analysis focuses on South Korea, the strategy is general. It applies an endogenous productivity framework to AI as a GPT and can be used in other settings where firm-level adoption measures and standard input–output data are available.

We find that AI adoption raises firms' endogenous productivity by about 5% on average. The gains are not immediate: effects are small in the first year after adoption and become economically meaningful only after firms have had time to reorganize around the technology, consistent with adjustment costs and learning that generate a productivity J-curve. There is significant heterogeneity around this average effect. Productivity gains are large and precisely estimated in ICT (Information and Communication Technology), the sector with the highest adoption intensity, while effects are statistically weak in Trade and Producer–Consumer Services. In Manufacturing, we find a modest negative effect, consistent with transitional costs when AI is deployed in production-related tasks. Heterogeneity by application area helps clarify these patterns: productivity improvements are concentrated in product/service development and sales/marketing, whereas AI used primarily in production processes has effects that are economically small.

Although we interpret deflated revenue mainly as quantity, residual within-industry price variation may remain and may respond to AI adoption; markups provide a direct way to capture this pricing-power channel. We examine whether AI adoption is associated with higher markups. We find no evidence that markups rise around adoption once we condition on productivity, which alleviates concerns that our productivity estimates are driven by a pricing-power channel.

We further distinguish the adoption or use of AI from the development of AI by augmenting the productivity process with an AI-development indicator. The estimated AI development effect is negative but statistically insignificant. This is consistent with development being an investment activity whose returns may not be realized in the short run absent

successful commercialization.

This paper contributes to two related literatures at the intersection of AI adoption and structural productivity estimation. First, it adds firm-level evidence on AI and productivity using a direct use-based measure of organizational AI deployment. We document systematic heterogeneity across sectors, adoption vintages, and application areas. This complements a rapidly growing empirical literature that studies AI and productivity through intangible capital and productivity dynamics (Brynjolfsson et al., 2021; Aghion et al., 2024; Acemoglu, 2024), task reallocation and inequality (Acemoglu and Restrepo, 2019, 2022), robotics and automation (Graetz and Michaels, 2018; Brynjolfsson et al., 2023), and experimental/field evidence on generative AI at work (Noy and Zhang, 2023; Peng et al., 2023; Brynjolfsson et al., 2025). Related work at the firm-level often relies on capability proxies such as worker resumes (Babina et al., 2024); our setting is closer to "actual use" because it records deployment.

Our results are closely related to Czarnitzki et al. (2023), McElheran et al. (2025), and Park and Shin (2025). Czarnitzki et al. (2023) study German firms, treat AI as an input, and use an IV strategy. In contrast, we model AI adoption as an endogenous firm choice that shifts the productivity law of motion, which yields a structural interpretation of how AI affects productivity dynamics. Using U.S. Census data and an IV design, McElheran et al. (2025) document a J-curve in manufacturing, with short-run losses giving way to longer-run gains; we find a similar delayed pattern for South Korea, and we also document a negative effect in manufacturing. Park and Shin (2025) use the same survey data but adopt a difference-in-differences design and find no evidence that AI raises labor productivity, measured as value added per worker. Our approach differs from the existing literature in two respects. First, we recover firm productivity from a production-function framework rather than relying on labor-productivity proxies (e.g., value added per worker) or profit-based measures that can be influenced by market power. Second, we embed AI adoption in the productivity transition to address self-selection into adoption by conditioning on lagged

productivity.

Second, the paper connects to the endogenous productivity estimation literature by allowing AI (GPT) adoption to enter the productivity law of motion and estimating its effect within that framework. The approach builds on production function identification strategies (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015). It follows a tradition that studies how firm choices shift productivity dynamics, such as R&D (Doraszelski and Jaumandreu, 2013), exporting (Aw et al., 2011; De Loecker, 2013), mergers (Braguinsky et al., 2015), and privatization (Chen et al., 2021). By placing AI in the Markov transition for productivity rather than as a contemporaneous input, the framework directly addresses self-selection into adoption driven by prior productivity. We are also related to estimating firm market power (markup) literature, such as De Loecker and Warzynski (2012) and Raval (2023).

The remainder of the paper is organized as follows. Section 2 describes the data and AI measures. Section 3 presents the production function, productivity dynamics, and estimation. Section 4 reports baseline and heterogeneous effects. Section 5 investigates AI development effects, the AI-markup relationship, and whether productivity drives selection into adoption. Section 6 provides a robustness check on the effect of AI by an alternative estimation strategy. Section 7 concludes.

2 Background: AI Adoption in South Korea

This section describes the data and key variables used in the analysis. We first introduce the Survey of Business Activities and the construction of our 2017-2023 firm-level panel. We then introduce our AI measures, covering adoption status, application area, and development indicators, and document how these patterns vary across sectors. Finally, we present the production-function variables and cost shares that motivate our estimation approach.

2.1 Survey of Business Activities

The analysis uses the annual Survey of Business Activities administered by the Korean Statistical Office. The survey covers firms in all industries with at least 50 regular employees and at least 300 million Korean Won in paid-in capital (approximately USD 260,000 at the 2020 exchange rate). These criteria remain constant throughout the survey period. Industry codes follow the Korean Standard Industrial Classification (KSIC), comparable to the International Standard Industrial Classification (ISIC) revision 4.

Our baseline panel spans 2017–2023, aligning pre-adoption covariates with the digital-technology questions (available from 2017 onward). We restrict to incorporated firms with non-missing core accounts and harmonize industry codes over time. Multi-activity firms are assigned to their primary industry (sales from secondary lines are consolidated to the main industry), consistent with the classification of the survey.

2.2 AI Related Variables

The Survey of Business Activities is designed to provide a comprehensive picture of firms' activities, covering performance, diversification and group affiliation, performance-management systems, and shifts in strategy and industry structure (Statistics Korea, 2022). Since 2005 it has collected rich firm characteristics. In 2017 it added a digital-technology module to track diffusion of new technologies. Firms report whether they have adopted any of nine technologies: Artificial Intelligence (AI), Robotics, Internet of Things (IoT), Cloud Computing, Big Data, mobile technologies and services (including 5G), Blockchain, 3D printing, and augmented/virtual reality (AR/VR). They also report which of four business functions employ these technologies: product/service development, production processes, sales/marketing strategy, and organizational management.

AI Adoption Indicator Statistics Korea (2022) define AI as the technology underpinning programs capable of human-like learning, reasoning, perception, and understanding of

natural language, etc. Following the survey's definition, we classify a firm as an AI user if it reports employing the technology in the four functions. Firms reported AI adoption based on official organizational usage, such as corporate-level adoption or encouragement, or formal utilization accompanied by cost payments, rather than informal, unrecognized personal use.[1]

This direct, use-based measure improves on common proxies in prior work, such as AI-related patents (Da Silva Marioni et al., 2024) or job postings (Babina et al., 2024). These proxies capture development activity or hiring intent rather than actual deployment in the production process. If a firm's use of AI is measured based on the possession of AI-related patents or employee resumes, there is a strong correlation with whether the company is developing AI technologies. However, this does not necessarily indicate whether these technologies are being utilized in the production process (Park and Shin, 2025).

In the estimation, we define an indicator $AI_{it} = 1$ if a firm i in year t reports using AI for business purposes, and 0 otherwise. This variable will be used in the baseline estimation.

Areas of AI Application Indicator The survey records the area where a firm applies AI: product or service development, production processes, sales/marketing strategy, or organizational management. This yields a mutually exclusive indicator at the firm and year level. We exploit this indicator to estimate heterogeneous semi-elasticities by application.

AI Developing Indicator The survey also enables identification of whether a firm is developing AI tools. We separate AI development from usage because development may differ from adoption in that it requires investment and labor inputs rather than boosting productivity. This variable allows us to examine how AI development affects productivity and its subsequent impact on firm revenue.

[1]Based on the response from Statistics Korea (2022), the question asks about AI in actual management activities. Thus, both company-level adoption of external AI solutions and employees' use of tools such as ChatGPT count as "AI use" when such usage is officially adopted or encouraged by the firm or involves paid access. Reported adoption therefore reflects official organizational usage rather than purely private or informal use.

2.3 AI Adoption and Application

As shown in Table 1, overall AI adoption increases markedly from 1.2% in 2017 to 5.3% in 2023, though the pace of adoption varies considerably across sectors. To facilitate analysis, we reclassify the original 12 industries into four economically coherent sectors: Manufacturing (comprising Manufacturing and Construction), ICT ((Information & Communication and Professional Science & Technology), Trade (Wholesale & Retail and Transportation & Warehouse), and Producer-Consumer Services. This reclassification stabilizes adoption rates in industries with few firms and improves comparability in subsequent production function estimates.

Among the four sectors, ICT leads adoption by a substantial margin, rising from 6.4% to 21.2% over the period. Manufacturing shows a more gradual uptake, increasing from 0.8% to 3.8%, while Trade follows a similar trajectory, climbing from 1.3% to 4.0%. The Producer-Consumer Services aggregate also grows from 0.9% to 4.2%, although this figure masks substantial variation within the sector. Specifically, the electricity and gas industries adopt at relatively higher rates, while the arts and sports industries remain near zero throughout the period.

In our data, product development is the most commonly reported application, followed by production and sales/marketing, while organizational management is comparatively rare. Figure 1 visualizes the application mix by sector and year (2017–2023), documenting both cross-sector differences and the shift toward broader use over time. Across sectors, product and service development is the most common AI application, while Manufacturing stands out because firms most often deploy AI in production processes. By contrast, relatively few firms report using AI primarily in organizational management or in sales and marketing.

Table 1: AI Adoption Rate by Sector and Industry in 2017 and 2023 (4-Sector Split)

	2017		2023	
Sector / Industry	**AI-adopter share**	**Firms**	**AI-adopter share**	**Firms**
Manufacturing (aggregate)	0.8%	6,095	3.8%	6,784
Manufacturing	0.8%	5,640	3.8%	6,253
Construction	0.4%	455	3.8%	531
ICT (aggregate)	6.4%	566	21.2%	749
Information and communication	7.8%	434	24.8%	544
Professional science and technology	1.5%	132	11.7%	205
Trade (aggregate)	1.3%	522	4.0%	960
Wholesale and Retail	1.3%	474	3.7%	882
Transportation and warehouse	2.1%	48	6.4%	78
Producer–Consumer Services (aggregate)	0.9%	572	4.2%	706
Electricity and gas	5.4%	37	18.0%	50
Accommodation and food service	0.6%	174	3.7%	190
Real estate	0.0%	30	2.4%	42
Business facility management	0.0%	89	6.1%	115
Arts and sports	0.0%	128	0.0%	145
Other services	1.8%	114	3.7%	164
Total	1.2%	7,755	5.3%	9,199

Notes: Data source: Survey of Business Activities.

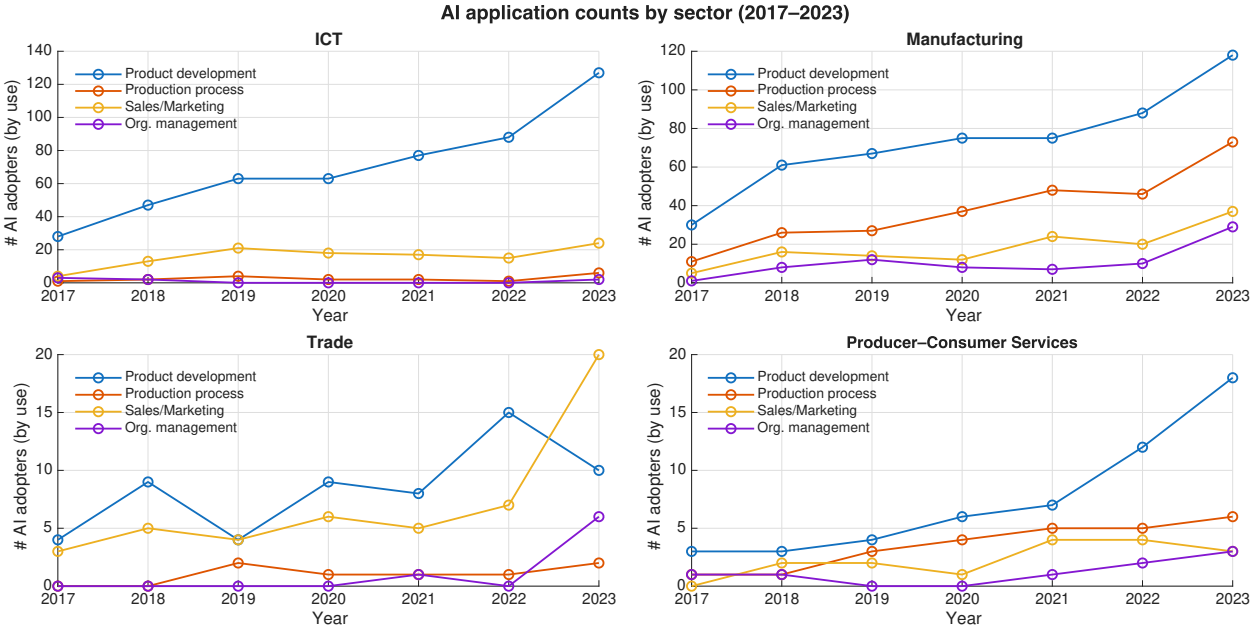


Figure 1: AI Application Count by Sector (2017-2023)

2.4 Production Variables

Table 2 presents summary statistics for key production variables. The number of firms increases from 7,755 to 9,199 over the sample period, with the count remaining relatively stable at approximately 8,100 firms from 2018 to 2021. Average firm revenue fluctuated around 220-250 billion KRW throughout the period, peaking at 248 billion KRW in 2022 before declining to 226 billion KRW in 2023. The capital stock (tangible and intangible assets) shows slight variations ranging from 278 billion KRW in 2018 to 308 billion KRW in 2021. Labor input, measured in number of employees, remained relatively stable around 330-360 workers per firm. Materials costs, the largest component of firm expenditures, ranged from approximately 124 billion KRW in 2019 to 143 billion KRW in 2022. All monetary values have been deflated to 2020 as the base year, with revenue and materials adjusted using the producer price index and capital deflated using the Gross Fixed Capital Formation Deflator.

Table 2: Summary Statistics by Year: Means

Year	No. of firms	Revenue	Capital	Labor	Materials
2017	7,755	241,949.52	281,488.26	361.24	130,573.39
2018	8,125	236,301.05	277,923.40	351.68	129,756.17
2019	8,138	232,893.55	284,157.99	355.53	127,950.90
2020	8,056	228,867.45	293,158.08	362.17	124,360.57
2021	8,145	243,745.43	308,911.92	360.65	135,893.40
2022	8,510	248,778.87	307,644.75	353.98	142,839.42
2023	9,199	226,614.76	295,434.97	338.55	129,652.39

Notes: For each year, it shows the number of firms in the sample and the means for four key variables: revenue, capital, number of employees, and material costs. All monetary values are expressed in millions of Korean won (KRW 1M) and deflated to 2020 as the base year. Revenue and materials have been deflated using the producer price index, while capital has been deflated using the Total Fixed Capital Formation Deflator. Data source: Survey of Business Activities.

Table 3 presents the annual average cost shares for capital, labor, and materials in the total input cost from 2017 to 2023. Materials constitute the largest cost component, accounting for approximately 60-63% of total costs throughout the period. Labor costs represent

the second largest share, averaging around 30% with a gradual increase until 2020 when they peaked at 32.7%, followed by a slight decline in subsequent years. Capital costs represent the smallest share, typically ranging between 5-8% of total costs, though with notable variation including a low of 5.0% in 2021 and a peak of 10.9% in 2023. The relative stability of labor cost shares over the sample period provides limited evidence of biased technological change. So we focus on how AI adoption affects Hicks-neutral productivity.

Table 3: Average Annual Input Cost Shares, 2017–2023

Year	Capital	Labor	Materials
2017	0.066	0.300	0.634
2018	0.071	0.302	0.627
2019	0.076	0.309	0.616
2020	0.058	0.327	0.615
2021	0.050	0.321	0.628
2022	0.078	0.299	0.622
2023	0.109	0.296	0.595
Mean	0.073	0.308	0.620

Notes: This table reports annual average cost shares for capital, labor, and materials across the full sample. Cost shares are calculated at the firm-year level and then averaged by year. Data source: Survey of Business Activities.

3 Empirical Framework

This section introduces the empirical framework used to quantify the effect of AI adoption on firm productivity. We begin with a gross output Cobb-Douglas production function and model unobserved productivity as a persistent state that evolves according to a Markov process. AI adoption is introduced as a shifter of the productivity law of motion rather than as a contemporaneous input. Identification follows a two-step control-function approach.

3.1 Production Function

We use a gross output specification because manufacturing firms, for which intermediate inputs constitute a substantial share of output, comprise more than half of our Korean sample.[2] We model firm i's gross output in year t as

$$y_{it} = \beta_0 + \beta_K k_{it} + \beta_L \ell_{it} + \beta_M m_{it} + \omega_{it} + \varepsilon_{it}, \tag{1}$$

where y_{it} is log deflated revenue which can be a proxy of physical output quantity,[3] and $k_{it}, \ell_{it}, m_{it}$ denote logs of deflated capital, employee count, and deflated materials, respectively. The coefficients β_K, β_L, and β_M are the output elasticities with respect to capital, labor, and materials. Unobserved productivity is captured by ω_{it}, and the disturbance ε_{it} represents measurement error that cannot be anticipated when input decisions are made, which are identically distributed across firms and over time with zero mean.

We do not include AI as a contemporaneous input in Equation (1). Embedding "AI adoption" directly in the production function as input would impose a common deterministic shift and, under Cobb–Douglas, a unit elasticity of substitution with standard inputs, as discussed in De Loecker (2013). Instead, AI is allowed to shift the evolution of productivity. Another practical reason is measurement. Treating AI as a separate input is more appropriate when firm-year expenditures or investment in AI are observed, but such data are rarely available. With only an adoption indicator, we cannot map AI into a meaningful flow (or price–quantity pair) or define an elasticity, so we treat AI as shifting the productivity law of motion rather than as a contemporaneous input.

[2]Several papers on estimating endogenous productivity change use this gross output production function setting, like Doraszelski and Jaumandreu (2013); De Loecker (2013); Braguinsky et al. (2015). Recent work of Chen et al. (2021) extend it to a non-parametric production function and estimate the privatization effect on productivity using framework of Gandhi et al. (2020).

[3]That is, deflating revenue by an industry–year price index removes most common price movements and allows us to read changes in y_{it} mainly as changes in quantities produced. At the same time, the deflator is not firm-specific, so y_{it} may still embed residual within-industry price variation and product-quality differences.

3.2 Productivity Dynamics

Following the literature on endogenous productivity estimation (Doraszelski and Jaumandreu, 2013; De Loecker, 2013; Braguinsky et al., 2015), we model firm productivity dynamics as a first-order Markov process in which AI adoption shifts the law of motion:

$$\omega_{it} = g\big(\omega_{i,t-1}, AI_{i,t-1}\big) + \xi_{it}, \tag{2}$$

$$\text{Baseline (linear):} \qquad \omega_{it} = \rho\,\omega_{i,t-1} + \gamma\,AI_{i,t-1} + \xi_{it}, \tag{3}$$

where $AI_{i,t-1}$ indicates whether firm i has adopted AI by $t-1$, ρ captures persistence, γ is the productivity effect of adoption, and ξ_{it} is the innovation to productivity. To study heterogeneity, productivity process in Section 4 allows the AI effect to vary by interacting the adoption indicator with calendar year, year since adoption, sector, and application area.

One could also augment $g(\cdot)$ with other firm choices as state variables such as exporting or R&D. In our setting covering 2017–2023 in South Korea, however, within-firm transitions in export status are not usual and the survey does not report R&D expenditures, so we do not include these variables.[4] The framework also allows $g(\cdot)$ to be nonlinear.[5] Here we focus on identifying the AI effect directly, so we keep the AI term linear. In principle, we use a parsimonious parameterization of the productivity process due to the low variation of AI adopters in our dataset.

It is important to stress the timing of decisions in this context, as the AI adoption decision is endogenous. At the start of period t, the predetermined state variables $(\omega_{i,t-1}, k_{it}, AI_{i,t-1})$ are given, where k_{it} is determined by prior investment. The innovation ξ_{it} is then real-

[4] In the past literature, it is common that the productivity process only include one state variable which is the mainly focus, see Doraszelski and Jaumandreu (2013) embed R&D inside $g(\cdot)$, De Loecker (2013) embed an exporting indicator, and Braguinsky et al. (2015) embed a merger indicator.

[5] In Aw et al. (2011) and Braguinsky et al. (2015), $g(\cdot)$ takes a nonlinear form of ω_{t-1} with a cubic polynomial, while determinants of productivity such as export status and mergers enter linearly. In our setting, we model the AI effect explicitly and estimate it directly, so we specify AI as entering linearly. A fully nonlinear productivity process would require approximating the AI effect rather than estimating it directly.

ized and the firm learns ω_{it}. Conditional on this information, the firm chooses variable inputs (ℓ_{it}, m_{it}), production occurs, and the transitory shock ε_{it} is realized. This timing, predetermined capital and adoption status, with flexible inputs chosen after observing productivity—is standard in endogenous productivity estimation (Doraszelski and Jaumandreu, 2013; De Loecker, 2013; Braguinsky et al., 2015; Chen et al., 2021) and also supports the orthogonality condition $\mathbb{E}[\xi_{it} \mid \omega_{i,t-1}, AI_{i,t-1}] = 0$.

We let AI enter with a one-period lag for two reasons.[6] First, in South Korea AI is most commonly applied to product and service development, so improvements may take time to translate into output. Second, adoption can occur within the year, implying that contemporaneous outcomes may reflect only partial exposure, with the full effect realized in subsequent periods. With this timing assumption, AI adoption in $t-1$ is uncorrelated with productivity innovation shock in t.[7] We also provide alternative timing assumptions as a robustness check, see Subsection 6.1.

To interpret the AI effect on productivity, we identify the AI effect by conditioning on last period's productivity. In particular, ω_{it-1} fixes each firm's starting point, so comparing adopters and non-adopters at the same ω_{it-1} contrasts firms that begin the period equally productive. We impose a timing/orthogonality restriction that AI is chosen before the innovation shock. Under this assumption, adopters and non-adopters with the same ω_{it-1} face the same expected innovation in period t; thus any difference in realized productivity ω_{it} is attributed to γ, the effect of AI on the productivity dynamics.

3.3 Identification and Estimation

Identification follows Levinsohn and Petrin (2003), which is common in the endogenous productivity estimation literature; see Doraszelski and Jaumandreu (2013), De Loecker (2013),

[6] Research in endogenous productivity estimation has assumptions that R&D, exporting, and merger have a one-period lag for productivity gain (Doraszelski and Jaumandreu, 2013; De Loecker, 2013; Braguinsky et al., 2015). Chen et al. (2021) exploit a lag structure in the timing of recorded ownership changes.

[7] An alternative is to assume AI adoption decision is made at the beginning of t and is uncorrelated with ξ_{it}, which would allow AI_{it} to enter the law of motion directly. We adopt the lag specification to match the institutional timing and to avoid partial-exposure concerns.

and Braguinsky et al. (2015).[8] We combine a control-function approach that proxies for unobserved productivity in the production function with the productivity law of motion in Equation (3). We assume the material (intermediate) input demand function $m_{it} = f_t\big(\omega_{it}, k_{it}\big)$ is strictly increasing in productivity, $\partial m_{it}/\partial \omega_{it} > 0$. This monotonicity ensures the existence of an inverse $\omega_{it} = h_t\big(m_{it}, k_{it}\big)$, which allows us to express the unobserved productivity as a function of observables. We also assume that the innovation shock is mean independent of all inputs, $E[\varepsilon_{it} \mid \omega_{it}, k_{it}, \ell_{it}, m_{it}] = 0$.

Given the assumptions above, we estimate the production function parameters in two stages. In the first stage, we can write Equation (1) as

$$y_{it} \;=\; \beta_0 \;+\; \beta_L \ell_{it} \;+\; \Phi_t\big(m_{it}, k_{it}\big) \;+\; \varepsilon_{it}, \tag{4}$$

where $\Phi_t(\cdot) \equiv \beta_K k_{it} + \beta_M m_{it} + h_t(m_{it}, k_{it})$ absorbs both the capital and materials coefficients along with productivity. We approximate $\Phi_t(\cdot)$ flexibly using year specific polynomials in (m, k).[9] This stage identifies $\widehat{\beta}_0$ and $\widehat{\beta}_L$ and delivers estimates $\widehat{\Phi}_{it}$. The coefficients β_K and β_M are not separately identified here as they are absorbed in $\Phi_t(\cdot)$.

In the second stage, we identify the remaining parameters using the Markov structure of productivity. Define $\tilde{y}_{it} \equiv y_{it} - \widehat{\beta}_L \ell_{it} - \widehat{\beta}_0$ as output net of the first stage estimates. From Equation 3, current productivity depends on lagged productivity, which we can express as $\omega_{i,t} \;=\; \rho(\widehat{\Phi}_{i,t-1} - \beta_K k_{i,t-1} - \beta_M m_{i,t-1}) + \gamma\, AI_{i,t-1} + \xi_{it}$, using the first stage estimates. Substituting into the production function yields

$$\tilde{y}_{it} \;=\; \beta_K(k_{it} - \rho\, k_{i,t-1}) \;+\; \beta_M(m_{it} - \rho\, m_{i,t-1}) \;+\; \rho\, \widehat{\Phi}_{i,t-1} \;+\; \gamma\, AI_{i,t-1} \;+\; u_{it}, \tag{5}$$

[8] Ackerberg et al. (2015) show that in the LP first stage, labor is functionally dependent on the same variables that enter the nonparametric proxy term—indeed l_{it} is a time-varying function of (m_{it}, k_{it})—so l_{it} is perfectly collinear with the proxy and β_L is not identified. Their solution estimates all coefficients in the second stage, and is typically implemented with a value-added production function that excludes materials m. This choice is further supported by Bond and Söderbom (2005) and Gandhi et al. (2020), who show that m_{it-1} is not a valid instrument for identifying the coefficient on m_{it} in a gross output framework. As a robustness check, we also estimate an ACF-style identification in Subsection 6.2.

[9] Olley and Pakes (1996) and Levinsohn and Petrin (2003) allow Φ_t to vary by periods. In our estimation, we allow Φ_t to vary across every year and employ up to third order polynomials.

where $u_{it} \equiv \xi_{it} + \varepsilon_{it}$ is the composite error. Since capital is predetermined with respect to the productivity innovation and lagged materials are determined before ξ_{it} is realized, these variables are valid instruments. We estimate $\theta = (\beta_K, \beta_M, \rho, \gamma)^\top$ by two step generalized method of moments with instrument set $z_{it} = (k_{i,t}, k_{i,t-1}, m_{i,t-1}, AI_{i,t-1})$, which yields just-identification. Standard errors are clustered at the firm level to account for within-firm correlation. Given the estimates $\widehat{\beta}_0$, $\widehat{\beta}_K$, $\widehat{\beta}_M$, and $\widehat{\Phi}_{it}$, we recover productivity as $\widehat{\omega}_{it} = \widehat{\Phi}_{it} - \widehat{\beta}_K k_{it} - \widehat{\beta}_M m_{it} - \widehat{\beta}_0$.[10]

4 Results

This section presents our estimation results, proceeding from a pooled baseline to increasingly flexible specifications. We begin with a baseline that imposes a common AI effect across firms and years, then progressively relax this restriction to examine heterogeneity along four dimensions, namely calendar time, time since adoption, sector, and AI application area.

The baseline estimates indicate a statistically significant positive effect of AI adoption on productivity. Allowing the effect to vary by time since adoption reveals a delayed pattern consistent with a J-curve, where productivity gains are negligible in the first two years of adoption but become sizable and significant after two more years, suggesting that complementary organizational adjustments take time. Sector-level heterogeneity shows that AI adoption yields large productivity gains in ICT but negative short-run effects in manufacturing, consistent with transitional adoption costs. Finally, decomposing by application area indicates that product (service) development and sales/marketing applications generate the positive short-run gains, while production process applications show no detectable effect.

[10] Olley and Pakes (1996) approximate productivity as $\widehat{\omega}_{it} = y_{it} - \widehat{\beta}_K k_{it} - \widehat{\beta}_L \ell_{it} - \widehat{\beta}_M m_{it} - \widehat{\beta}_0$ because $\widehat{\Phi_t}$ cannot be constructed for plants with zero investment. Since our dataset does not have missing values in $\widehat{\Phi_t}$, we follow Levinsohn and Petrin (2003) and De Loecker (2013) in recovering productivity.

4.1 Baseline Results

We first estimate the production function and the productivity process with a homogeneous AI adoption effect across year and sector. Table 4 reports the two-step GMM estimates. In the first stage, the labor elasticity is precisely estimated at $\widehat{\beta}_L = 0.356$. The second stage yields capital and materials elasticities of $\widehat{\beta}_K = 0.240$ and $\widehat{\beta}_M = 0.543$, both economically sizable and statistically significant. The AI coefficient is positive and statistically significant at conventional levels. We estimate $\widehat{\gamma} = 0.049$ (s.e. 0.014), implying a short run semi-elasticity of output with respect to lagged AI adoption of approximately 4.9%.

The persistence parameter is $\widehat{\rho} = 1.000$ (s.e. 0.001), with a 95% confidence interval of approximately [0.998, 1.002]. Thus ρ is statistically indistinguishable from unity, implying a random walk productivity process in which innovations are permanent while expected changes are unpredictable conditional on past information. Under this near unit root, a one time adoption of AI permanently shifts the level of the firm's productivity path.[11] Given this near unit root, the "long-run effect" $\gamma/(1-\rho)$ becomes mechanically large and numerically unstable. We therefore emphasize the short run semi-elasticity and refrain from interpreting long run magnitudes.

Under a Cobb–Douglas interpretation, the implied returns to scale are $\widehat{\nu} = \widehat{\beta}_L + \widehat{\beta}_K + \widehat{\beta}_M = 1.139$, suggesting mildly increasing returns.

Finally, recovered productivity exhibits a mild upward drift over the sample window. Average productivity increased by approximately 2% over seven years, consistent with gradual improvements within the firm after netting out input contributions.

[11] This near unit root persistence is common in firm level productivity estimation. We explored several flexible specifications for $\phi_t(k, m)$ and obtained $\widehat{\rho}$ very close to 1 in each case. See also Wang (2025), who report $\widehat{\rho} \approx 1.01$ in a related setting of labor-augmenting productivity process. The near unit estimate is therefore a robust feature of the data rather than an artifact of a particular first stage specification.

Table 4: Two-step GMM Estimates

Parameter	Estimate	Std. Err.
β_K	0.240	0.059
β_L	0.356	0.007
β_M	0.543	0.071
ρ	1.000	0.001
γ	0.049	0.014

Notes: Two-step GMM estimates with firm-clustered standard errors. Stage-1 sample size: $N = 57{,}928$; Stage-2 sample size: $N = 46{,}316$. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

4.2 Heterogeneity Effect of AI adoption by Year

To allow the productivity contribution of AI to vary over time, we replace the single semi–elasticity parameter with a set of year–specific effects. The key modification is the productivity law of motion, where lagged AI adoption enters with coefficients $\{\gamma_t\}$ interacted with calendar years. This specification captures evolving payoffs to AI (e.g., diffusion, complementary investments, or macro shocks) while keeping β_K, β_M, and ρ common across years. We estimate the $\{\gamma_t\}$ jointly by GMM using the same instrument set as in the baseline, augmented with year–interacted AI terms.

$$\omega_{it} = \rho\, \omega_{i,t-1} \; + \; \sum_{t \in \mathcal{T}} \gamma_t \, AI_{i,t-1} \, \mathbf{1}\{\text{year} = t\} \; + \; \xi_{it}, \tag{6}$$

where $\mathbf{1}\{\cdot\}$ is an indicator for calendar year t, and ξ_{it} is the innovation in productivity.

Table 5 shows precise core elasticities (labor $\widehat{\beta}_L \approx 0.356$, capital $\widehat{\beta}_K \approx 0.240$, materials $\widehat{\beta}_M \approx 0.543$) and nearly unit-root productivity ($\widehat{\rho} \approx 1.000$). The year-specific AI semi-elasticities $\{\widehat{\gamma}_t\}_{t=2018}^{2023}$ are positive in every year, ranging from $\widehat{\gamma}_{2018} = 0.021$ to $\widehat{\gamma}_{2019} = 0.074$, with a second peak in 2023 ($\widehat{\gamma}_{2023} = 0.069$). The simple (unweighted) average across years is $\frac{1}{6}\sum_{t=2018}^{2023} \widehat{\gamma}_t = 0.046$, which is close to the pooled estimate $\widehat{\gamma} = 0.049$. Thus, allowing time heterogeneity reveals the timing of gains—strongest around 2019 and again in 2023—without

materially changing the overall average effect.

Table 5: Two-step GMM with Year-varying AI Effect γ_t

Parameter	Estimate	Std. Err.
Panel A: Core elasticities and persistence		
β_K	0.240	0.059
β_L	0.356	0.007
β_M	0.543	0.123
ρ	1.000	0.001
Panel B: AI semi-elasticity by year, γ_t		
γ_{2018}	0.021	0.049
γ_{2019}	0.074	0.029
γ_{2020}	0.024	0.030
γ_{2021}	0.042	0.037
γ_{2022}	0.044	0.034
γ_{2023}	0.069	0.033

Notes: Two-step GMM estimates with firm-clustered standard errors. Stage-1 sample size: $N =$ 57,928; Stage-2 sample size: $N =$ 46,316. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

4.3 Heterogeneity Effect of AI Adoption by Time Since Adoption

To examine whether there is evidence of a J-curve, we next allow the impact of AI on productivity to vary with the time elapsed since adoption. Instead of a single adoption dummy, we construct three mutually exclusive indicators based on years since adoption at $t-1$. The first indicator captures firms in their first year of AI adoption. The second indicator captures firms in their second year after adoption. The third indicator captures firms that have used AI for more than two years.

Keeping common input elasticities, we modify the productivity law of motion to

$$\omega_{it} \;=\; \rho\,\omega_{i,t-1} \;+\; \gamma_0 D^{(0)}_{i,t-1} \;+\; \gamma_1 D^{(1)}_{i,t-1} \;+\; \gamma_{2+} D^{(2+)}_{i,t-1} \;+\; \xi_{it},$$

where $D_{i,t-1}^{(0)}$ is an indicator for the first year of AI use, $D_{i,t-1}^{(1)}$ is an indicator for the second year after adoption, and $D_{i,t-1}^{(2+)}$ is an indicator for having used AI for at least two years. Non–adopters have all three indicators equal to zero.

The estimates in Table 6 reveal a clear delayed or J-curve effect of AI on productivity. The first year coefficient is close to zero and statistically insignificant, indicating that firms experience little immediate productivity gain when AI is first introduced. The second year coefficient $\widehat{\gamma}_1 = 0.045$ is positive but imprecisely estimated, consistent with modest short run improvements as firms begin to integrate AI into their operations. By contrast, the coefficient for firms that have used AI for more than two years is positive and precisely estimated at $\widehat{\gamma}_{2+} = 0.110$ (s.e. 0.026), implying that long run adopters enjoy productivity levels 11 percent higher than comparable non-adopters.

Overall, the pattern is consistent with learning and adjustment costs, consistent with a productivity J-curve. AI adoption requires time for complementary organizational changes, and sizable productivity gains materialize only after firms have had at least a couple of years to reorganize production around the new technology.

Table 6: Two-step GMM with Years-since-adoption AI Effects

Parameter	**Estimate**	**Std. Err.**
Panel A: Core elasticities and persistence		
β_K	0.241	0.036
β_L	0.356	0.007
β_M	0.541	0.071
ρ	1.000	0.001
Panel B: AI semi-elasticity by years since adoption		
γ_0 (first year)	0.010	0.020
γ_1 (second year)	0.045	0.030
γ_{2+} (2+ years)	0.110	0.026

Notes: Two-step GMM estimates with firm-clustered standard errors. Stage-1 sample size: N = 57,928; Stage-2 sample size: N = 46,316. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

4.4 Heterogeneity Effect of AI adoption by Sector

To allow the AI impact on productivity to differ across broad sectors, we keep common production function elasticities and modify the productivity law of motion to interact AI adoption with sector dummies. The specification becomes:

$$\omega_{it} \;=\; \rho, \omega_{i,t-1} \;+\; \sum_{g \in \{ICT, Mfg, Tr, PCS\}} \gamma_g, \mathrm{AI}^{g}_{i,t-1} \;+\; \xi_{it}.$$

where s_i denotes the firm's sector, $\mathrm{AI}_{i,t-1}$ is an indicator for AI adoption in $t-1$, and ρ is common across sectors. We merge Construction into Manufacturing and fold Professional Science & Technology into ICT, yielding four broad sectors. We estimate $\{\beta_K, \beta_M, \rho\}$ jointly and recover sector-specific semi-elasticities γ_g by interacting the lagged AI indicator with sector dummies in the second stage GMM.

Table 7 reports the results. The core elasticities remain stable with $\widehat{\beta}_L = 0.356$, $\widehat{\beta}K = 0.255$, and $\widehat{\beta}M = 0.509$. Persistence remains near unity at $\widehat{\rho} = 1.000$, so we interpret γ_g as short-run semi-elasticities. The slight movement in elasticities relative to the baseline specification reflects the tight correlation between sector membership and capital and materials intensity. Once we allow sector-specific AI effects, the model re-estimates the elasticities to satisfy all moment conditions.

The AI effect differs sharply across sectors. In ICT, adoption is associated with about a 21.2% increase in productivity. In Manufacturing, the short-run effect is about negative 5.0%. In Trade it is about negative 3.0%, and in Producer–Consumer Services it is about positive 7.6%, though these last two estimates are not statistically distinguishable from zero.

This pattern is consistent with technological complementarity. AI is closer to core production tasks in ICT, where digital infrastructure and data-intensive processes are central to firm operations. The negative short run effect in Manufacturing may reflect transitional adoption costs resembling a micro-level J-curve (Brynjolfsson et al., 2021), or input substitu-

tion that is not immediately productivity-enhancing.[12] This negative effect is also consistent with findings in McElheran et al. (2025), who document negative short run effects of AI adoption in U.S. manufacturing.

Table 7: Two-step GMM with Sector-specific AI Effect γ_g

Parameter	**Estimate**	**Std. Err.**
Panel A: Core elasticities and persistence		
β_K	0.255	0.035
β_L	0.356	0.007
β_M	0.509	0.068
ρ	1.000	0.001
Panel B: AI semi-elasticity by sector, γ_g		
γ (*ICT*)	0.212	0.031
γ (*Manufacturing*)	−0.050	0.013
γ (*Trade*)	−0.030	0.039
γ (*Producer–Consumer Services*)	0.076	0.072

Notes: Two-step GMM estimates with firm-clustered standard errors. Stage-1 sample size: $N = 57{,}928$; Stage-2 sample size: $N = 46{,}316$. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

4.5 Heterogeneity Effect of AI adoption by AI Application

We next examine whether the AI effect varies by how firms deploy the technology. We replace the single AI indicator with four mutually exclusive dummies that capture where AI is used within the firm. These application areas are product or service development, production processes, sales and marketing, and organizational management.

Keeping common input elasticities, we modify the productivity law of motion as

$$\omega_{it} \;=\; \rho\,\omega_{i,t-1} \;+\; \sum_{a\in\{\text{pdev, prod, sales, org}\}} \gamma_a\,\mathrm{AI}^{a}_{i,t-1} \;+\; \xi_{it},$$

where a indexes the four application areas. For each firm-year, $\mathrm{AI}^{a}_{i,t-1}$ is an indicator that

[12] One concern is that the negative effect in Manufacturing could be driven by imposing common input elasticities across all sectors. To address this, we separately estimate all parameters within Manufacturing alone and find that the negative AI effect remains statistically significant.

equals one if firm i has adopted AI by $t-1$ and primarily applies it in area a, and zero otherwise. Since firms report their primary application area, at most one application indicator equals one in a given period.

Table 8: Two-step GMM with Application-specific AI Effect

Parameter	**Estimate**	**Std. Err.**
Panel A: Core elasticities and persistence		
β_K	0.240	0.036
β_L	0.356	0.007
β_M	0.541	0.070
ρ	1.000	0.001
Panel B: AI semi-elasticity by application, γ_{app}		
γ (*Product (service) development*)	0.050	0.020
γ (*Production process*)	0.000	0.024
γ (*Sales / Marketing*)	0.061	0.027
γ (*Organizational management*)	0.179	0.067

Notes: Two-step GMM estimates with firm-clustered standard errors. Stage-1 sample size: N = 57,928; Stage-2 sample size: N = 46,316. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

Table 8 reports the results. The core elasticities remain stable with $\widehat{\beta}_L = 0.356$, $\widehat{\beta}_K = 0.240$, and $\widehat{\beta}_M = 0.541$. Persistence remains near unity at $\widehat{\rho} = 1.000$. The estimates reveal substantial heterogeneity across application areas. Organizational management shows the largest effect at $\widehat{\gamma}_{\text{org}} = 0.179$ (s.e. 0.067), followed by sales and marketing at $\widehat{\gamma}_{\text{sales}} = 0.061$ (s.e. 0.027). Product or service development shows a modest but statistically significant effect at $\widehat{\gamma}_{\text{pdev}} = 0.050$ (s.e. 0.020). Production processes show no significant effect with $\widehat{\gamma}_{\text{prod}} = 0.000$ (s.e. 0.024).

These findings help explain the negative manufacturing effect documented in Subsection 4.4. Most manufacturing adopters deploy AI in production process tasks, where near term productivity gains appear limited and implementation costs can temporarily outweigh efficiency improvements. The large organizational management estimate should be interpreted with caution given the small adopter sample for this application area, which makes the coefficient imprecise and potentially sensitive to specification choices.

We also estimate heterogeneous effects of AI adoption by sector and application area jointly in Appendix B. These estimates should be interpreted carefully. For some sector-application cells, AI adoption varies little over 2017 to 2023, leading to imprecise or unstable estimates.

5 Discussion

This section uses our estimated productivity measures to address three related questions. First, we examine whether AI adoption is related to market power. Second, we distinguish using AI from developing AI by augmenting the productivity process with an AI development indicator, which helps assess whether development activities translate into short-run output gains. Third, we document selection into AI adoption by relating new adoption to pre-adoption productivity (and size) across sectors.

5.1 AI Adoption and Market Power

Although we interpret deflated revenue primarily as a proxy for quantity, residual within-industry–year variation in prices and demand may remain and may respond to AI adoption. In that case, part of the estimated productivity gains could reflect higher markups rather than higher physical efficiency. Markups therefore provide a direct way to capture this pricing-power channel.

We estimate markups following De Loecker and Warzynski (2012). As discussed in Raval (2023), we estimate markups using material costs rather than labor. Materials satisfy the static first-order condition required for identification, whereas labor is subject to adjustment costs such as hiring and firing frictions, which are particularly salient in South Korean labor markets.[13]

[13] Doraszelski and Jaumandreu (2019) note that the De Loecker and Warzynski (2012) approach can deliver inconsistent output-elasticity estimates when productivity is labor-augmenting and this channel is not accounted for. In our setting, however, the patterns in input cost shares suggest less evidence of labor-augmenting productivity over the sample period, which mitigates this concern.

We construct firm-level markups and markdowns using the estimated production-function elasticities and input cost shares. Let s_{it}^M denote the materials expenditure share in revenue: $s_{it}^M \equiv \frac{\text{MaterialCost}_{it}}{\text{Revenue}_{it}}$. Following De Loecker and Warzynski (2012), the markup for firm i in year t is computed as

$$\mu_{it} = \frac{\hat{\beta}_M}{s_{it}^M}, \tag{7}$$

where $\hat{\beta}_M$ is the estimated output elasticity of materials. Our production function estimates yield $\hat{\beta}_M = 0.543$ (from our baseline estimation in Table 4). The resulting markup distribution has a median of 1.148.[14]

AI Adoption and Markup Dynamics. Our panel also allows us to examine whether markup differs systematically within the group of adopters, and in particular whether there is a specific pattern around the timing of adoption—that is, before and after a firm begins (or stops) using AI.

Specifically, we estimate:

$$\ln \mu_{it} = \gamma_1 \text{Entry}_{it} + \gamma_2 \text{Exit}_{it} + \mathbf{b}_{it}'\boldsymbol{\sigma} + \alpha_i + \lambda_t + u_{it}, \tag{8}$$

where α_i and λ_t denote firm and year fixed effects, respectively. Here, $\text{Entry}_{it} = 1$ if firm i starts using AI in year t (and 0 otherwise), and $\text{Exit}_{it} = 1$ if firm i stops using AI in year t. The vector $\mathbf{b}_{it}$ collects the same firm controls as in the baseline regressions (e.g., log employment, log capital, and optionally estimated productivity), and standard errors are clustered at the firm level. The coefficients of interest are γ_1 and κ_1, which capture the change in markups/markdowns around AI entry, and γ_2 and κ_2, which capture the analogous change around AI exit.

Table 9 studies whether within-firm markups move around AI adoption events, using AI entry and exit indicators while absorbing permanent firm heterogeneity and common macro

[14]This ratio can be highly sensitive to extreme tails, in particular, very small s_{it}^M mechanically generates very large μ_{it}, so we emphasize the median of the markup distribution.

Table 9: Markup and AI Adoption Dynamics: Entry and Exit

	Markup (1)	Markup (2)
AI Entry	0.001	0.009
	(0.008)	(0.007)
AI Exit	0.012	−0.015
	(0.010)	(0.009)
Capital	0.293***	0.028***
	(0.012)	(0.010)
Employment	0.159***	0.327***
	(0.008)	(0.010)
Materials	−0.637***	
	(0.013)	
Productivity		1.360***
		(0.025)
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	46,316	46,316
R^2 (within)	0.6746	0.7009

Notes: Standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

shocks with firm and year fixed effects. Following the way De Loecker and Warzynski (2012) discuss entry-related markup changes in their exporting application, we view these regressions as descriptive: the coefficients summarize conditional comovement between markups and adoption transitions, rather than causal effects, since AI entry/exit may coincide with other unobserved adjustments (e.g., product upgrading, demand shifts, or organizational changes).

Columns (1) and (2) show that the estimated markup responses in the adoption year are small and statistically indistinguishable from zero. In column (1), conditioning on standard inputs (capital, employment, and materials) in addition to firm and year fixed effects, the estimated entry coefficient is essentially zero, and the exit coefficient is positive but imprecisely estimated. In column (2), replacing materials with estimated productivity (because of collinearity, we cannot keep both), the entry coefficient remains close to zero, and the exit coefficient turns negative, again without statistical significance at conventional levels. Taken together, these patterns suggest that, once we compare firms to themselves over time and control for standard inputs and/or productivity, there is little evidence of a discrete jump in markups exactly in the entry (or exit) year; any effect is modest relative to the within-firm variation explained by inputs and productivity.

Therefore, in the time series, we do not find that markups increase when firms adopt AI when we control for productivity. These results eliminate some of our concerns that our productivity estimates might be biased by the pricing channel product-market power.

5.2 Does Developing AI Boost Firm Productivity?

We distinguish between using AI and developing AI by augmenting the productivity process with separate indicators for each activity. This distinction is important because AI development represents an investment activity with upfront costs, whereas AI use reflects the deployment of existing technology in firm operations.

We modify the productivity law of motion to

$$\omega_{it} \;=\; \rho\,\omega_{i,t-1} \;+\; \gamma_{\text{adopt}}\,\text{AI}^{\text{adopt}}_{i,t-1} \;+\; \gamma_{\text{dev}}\,\text{AI}^{\text{dev}}_{i,t-1} \;+\; \xi_{it},$$

where $\text{AI}^{\text{adopt}}_{i,t-1}$ indicates whether firm i was using AI in period $t-1$ and $\text{AI}^{\text{dev}}_{i,t-1}$ indicates whether the firm was developing AI.

Table 10 reports the results. The core elasticities remain stable with $\widehat{\beta}_L = 0.356$, $\widehat{\beta}_K = 0.240$, and $\widehat{\beta}_M = 0.543$. Persistence remains near unity at $\widehat{\rho} = 1.000$. The AI use coefficient is positive at $\widehat{\gamma}_{\text{adopt}} = 0.067$ (s.e. 0.037), similar in magnitude to our baseline estimate. By contrast, the AI development coefficient is slightly negative and statistically insignificant at $\widehat{\gamma}_{\text{dev}} = -0.021$ (s.e. 0.039).

This pattern is consistent with the idea that AI development is an investment activity. Developing AI entails substantial costs including R&D expenditures, data infrastructure, and specialized labor. These investments may not translate into measurable output gains in the short run unless firms can successfully commercialize the resulting technology. Moreover, firms engaged in AI development may be diverting resources away from current production, leading to a temporary productivity drag. The positive effect of AI use, combined with the null effect of AI development, suggests that the output benefits of AI materialize primarily through deployment rather than through internal development activities during our sample period.

5.3 Does productivity affect AI adoption decision?

Figure 2 diagnoses selection into AI adoption. For each sector, we plot the probability of new adoption at t (among firms with no AI at $t-1$) across quartiles of prior productivity and quartiles of prior employment (both standardized within year). Patterns across quartiles reveal selection margins: monotone increasing ⇒ leader/scale selection; U-shaped ⇒ leaders and laggards adopt more; flat ⇒ little selection on that dimension. We use these visuals to

Table 10: Two-step GMM Estimates: AI Use and AI Development

Parameter	Estimate	Std. Err.
β_K	0.240	0.036
β_L	0.356	0.007
β_M	0.543	0.071
ρ	1.000	0.001
γ_{adopt}	0.067	0.037
γ_{dev}	-0.021	0.039

Notes: Two-step GMM estimates with firm-clustered standard errors. Stage-1 sample size: $N = 57{,}928$; Stage-2 sample size: $N = 46{,}316$. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

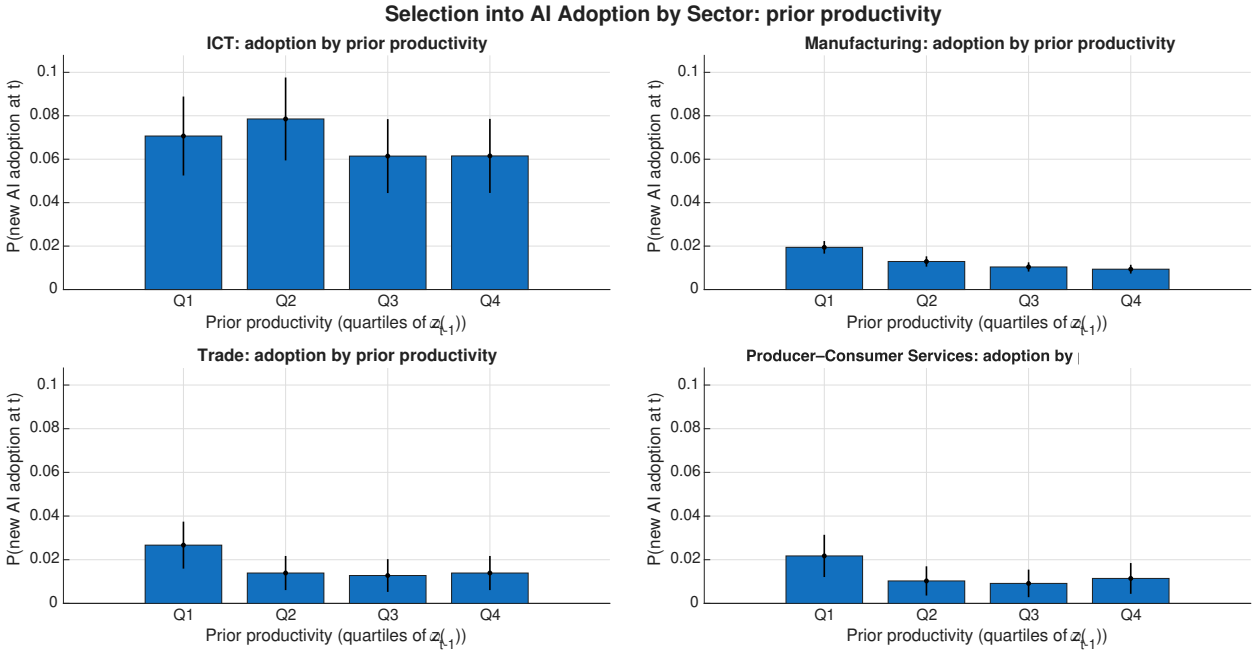


Figure 2: Selection into AI Adoption by Sector: Prior Productivity
Notes: For each sector, bars plot the probability of *new* adoption at t across **quartiles of prior productivity**, $z(\omega_{i,t-1})$, where ω_{it} is the productivity term from our Stage-1 common approximation (standardized within year).

motivate conditioning on pre-adoption productivity and size in our empirical design and to interpret sectoral heterogeneity in the estimated AI effects.

Using sector-specific quartiles of prior-year standardized productivity, we find clear selection patterns. We found that lower productivity firms are more likely to adopt AI in every sectors. This pattern is economically plausible and consistent with "catch-up" or performance-feedback mechanisms: when baseline efficiency is low, the marginal payoff to basic process/management improvements from a general-purpose technology is higher, so the adoption threshold is crossed earlier for laggards.

Therefore, AI adoption isn't random: firms decide to adopt based on things we don't fully observe—especially their underlying productivity. And our framework that embed AI into productivity process and soak up the selection on productivity level by conditioning on ω_{it-1} can address this selection.

6 Robustness Check

6.1 Timing Assumptions for AI Adoption

Our baseline specification treats AI adoption as affecting productivity with a delay, consistent with implementation frictions and learning. In particular, we model the productivity law of motion as

$$\omega_{it} = g(\omega_{i,t-1}, AI_{i,t-1}) + \xi_{it},$$

and assume the lagged adoption status is predetermined with respect to the current innovation shock (i.e., $AI_{i,t-1} \perp \xi_{it}$). This timing is standard in the endogenous productivity literature that models firm choices as shifting the productivity transition (e.g., Doraszelski and Jaumandreu (2013); Aw et al. (2011); De Loecker (2013); Braguinsky et al. (2015)).

As robustness, we consider alternative timing assumptions about when the adoption decision is made and whether it can react to within-year shocks. First, if adoption is planned

at the beginning of year t, then AI_{it} may be treated as predetermined within the year and can enter the transition directly (e.g., Chen et al. (2021)):

$$\omega_{it} = g(\omega_{i,t-1}, AI_{it}) + \xi_{it}, \quad AI_{it} \perp \xi_{it}.$$

Second, if adoption can respond to within-year innovations, then contemporaneous adoption is potentially endogenous, $\text{Cov}(AI_{it}, \xi_{it}) \neq 0$. In this case, we maintain the contemporaneous specification but instrument AI_{it} with its lag $AI_{i,t-1}$.

Across these timing assumptions, we re-estimate the productivity process and the implied AI effects to assess whether our main conclusions are sensitive to the precise adoption-timing interpretation.

Table 11: Robustness: Contemporaneous AI Effect (Predetermined AI_{it})

Parameter	**Estimate**	**Std. Err.**
β_K	0.240	0.036
β_L	0.356	0.007
β_M	0.542	0.071
ρ	1.000	0.001
γ	0.051	0.013

Notes: Two-step GMM estimates with firm-clustered standard errors. This specification allows a contemporaneous AI effect and treats AI_{it} as predetermined within year t. Stage-1 sample size: N = 57,928; Stage-2 sample size: N = 46,316. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

Table 11 and Table 12 report two alternative specifications that allow for a contemporaneous AI effect. When AI_{it} is treated as predetermined within year t, the estimated AI coefficient is positive and very close to our baseline estimation. When we instead allow AI_{it} to be endogenous and instrument it with $AI_{i,t-1}$, the estimated effect is larger. Across both specifications, the AI coefficient remains positive and statistically significant, suggesting that the main conclusion—AI adoption raises productivity—does not hinge on the baseline

delayed-effect timing assumption.

Table 12: Robustness: Contemporaneous AI Effect (Instrumenting AI_{it} with $AI_{i,t-1}$)

Parameter	Estimate	Std. Err.
β_K	0.241	0.036
β_L	0.356	0.007
β_M	0.538	0.070
ρ	1.001	0.001
γ	0.079	0.023

Notes: Two-step GMM estimates with firm-clustered standard errors. This specification allows a contemporaneous AI effect and treats AI_{it} as endogenous, instrumenting it with $AI_{i,t-1}$. Stage-1 sample size: $N = 57{,}928$; Stage-2 sample size: $N = 46{,}316$. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

6.2 ACF-style Identification

This subsection evaluates the robustness of our baseline LP-style estimates to an alternative identification strategy for production function elasticities. As discussed in Subsection 3.3, Ackerberg et al. (2015) argue that the first stage of the LP approach may not identify the labor coefficient because labor is a deterministic function of the same state variables that enter the nonparametric proxy, creating a collinearity problem. Their solution estimates all input coefficients in the second stage GMM.

We re-estimate the model using the ACF-style identification strategy. Table 13 reports the results. The estimated AI effect remains positive and statistically significant at $\widehat{\gamma} = 0.063$, very close to our baseline LP estimate of approximately 4.9%. This supports the robustness of our main finding that AI adoption has a significant positive effect on productivity.

The input elasticities shift somewhat under the ACF specification. The capital elasticity is 0.255, similar to our baseline. The materials elasticity increases to 0.636, while the labor

elasticity falls to = 0.114. This labor estimate is relatively low compared with the firm-level average labor cost share in our data, suggesting that the ACF specification may underestimate the contribution of labor in this gross output setting. Persistence remains near unity at $\widehat{\rho} = 1.000$.

Overall, the key result of a positive and statistically significant AI effect is robust to this alternative identification strategy, even though the input elasticities differ somewhat between the LP and ACF approaches.

Table 13: ACF-Style Two-step GMM Estimates

Parameter	**Estimate**	**Std. Err.**
β_K	0.255	0.047
β_L	0.114	0.050
β_M	0.636	0.090
ρ	1.000	0.001
γ	0.063	0.014

Notes: Two-step GMM estimates with firm-clustered standard errors. Stage-1 sample size: $N = 57{,}928$; Stage-2 sample size: $N = 46{,}316$. The model is just-identified, so the Hansen J-statistic is not informative. Data source: Survey of Business Activities.

7 Conclusion

This paper studies how AI adoption shapes firm-level productivity in South Korea using a 2017–2023 panel that directly measures organizational AI use and its primary application area. The key empirical challenge is selection: adoption is correlated with latent productivity. We address this by embedding lagged AI adoption in a Markov law of motion for productivity and estimating production parameters and productivity dynamics in a structural framework.

Four results emerge. First, in the pooled baseline, AI adoption raise productivity by about 5 percent. Second, the gains are delayed: effects are small in the first year and become economically meaningful only after firms have time to reorganize, consistent with

adjustment costs and learning that generate a productivity J-curve. Third, effects differ sharply across sectors, with large and precisely estimated gains in ICT, weak effects in Trade and Producer–Consumer Services, and modest negative short-run effects in Manufacturing. Fourth, application-area estimates show that gains are concentrated in product and service development and in sales and marketing, while AI used primarily in production processes delivers economically small effects on average.

These findings have two main implications. Substantively, AI's payoffs are not immediate and vary sharply by sector and by where the technology is deployed within the firm—patterns that align with the view that AI is a general-purpose technology whose returns depend on complementary organizational and intangible investments. Methodologically, modeling adoption as a shifter of the productivity law of motion provides a disciplined way to address selection into adoption and to interpret estimated effects as changes in endogenous productivity dynamics rather than contemporaneous correlations.

Several limitations point to natural extensions. Our AI measure is binary and does not capture intensity or spending; richer measures would sharpen heterogeneity and mechanism tests. The short sample window and near-unit persistence in productivity limit the interpretation of long-run magnitudes. Finally, productivity remains sensitive to demand conditions and pricing, motivating future work that jointly models markups and productivity or incorporates additional outcome measures beyond revenue.

Overall, the evidence indicates that AI adoption is associated with economically meaningful revenue gains on average, but that the timing and location of those gains—across sectors and application areas—are central to understanding when and for whom AI delivers productivity improvements.

An important direction for future research is to apply this framework to other countries as comparable data become available. In particular, U.S. Census and European firm-level datasets increasingly collect information on digital technologies, which would allow researchers to study whether the productivity dynamics associated with AI adoption differ

across institutional environments.

References

Acemoglu, D. (2024). The Simple Macroeconomics of AI. *Economic Policy*.

Acemoglu, D. and Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2):3–30.

Acemoglu, D. and Restrepo, P. (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, 90(6):1973–2016.

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification Properties of Recent Production Function Estimators. *Econometrica*, 83(6):2411–2451.

Aghion, P., Jones, B. F., and Jones, C. I. (2024). Artificial Intelligence and Economic Growth. *Econometrica*, 92(2):435–478.

Aw, B. Y., Roberts, M. J., and Xu, D. Y. (2011). R&D Investment, Exporting, and Productivity Dynamics. *American Economic Review*, 101(4):1312–1344.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Bond, S. and Söderbom, M. (2005). Adjustment costs and the identification of Cobb Douglas production functions. Technical report, IFS Working Papers.

Braguinsky, S., Ohyama, A., Okazaki, T., and Syverson, C. (2015). Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry. *American Economic Review*, 105(7):2086–2119.

Brynjolfsson, E., Buffington, C., Goldschlag, N., Li, J. F., Miranda, J., and Seamans, R. (2023). Robot Hubs: The Skewed Distribution of Robots in US Manufacturing. *AEA Papers and Proceedings*, 113:215–218.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2):889–942.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.

Chen, Y., Igami, M., Sawada, M., and Xiao, M. (2021). Privatization and Productivity in China. *The RAND Journal of Economics*, 52(4):884–916.

Czarnitzki, D., Fernández, G. P., and Rammer, C. (2023). Artificial Intelligence and Firm-Level Productivity. *Journal of Economic Behavior & Organization*, 211:188–205.

Da Silva Marioni, L., Rincon-Aznar, A., and Venturini, F. (2024). Productivity Performance, Distance to Frontier and AI Innovation: Firm-Level Evidence from Europe. *Journal of Economic Behavior & Organization*, 228:106762.

De Loecker, J. (2013). Detecting Learning by Exporting. *American Economic Journal: Microeconomics*, 5(3):1–21.

De Loecker, J. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.

Doraszelski, U. and Jaumandreu, J. (2013). R&D and Productivity: Estimating Endogenous Productivity. *Review of Economic Studies*, 80(4):1338–1383.

Doraszelski, U. and Jaumandreu, J. (2019). Using Cost Minimization to Estimate Markups. Working paper, CEPR Discussion Paper No. DP14114.

Gandhi, A., Navarro, S., and Rivers, D. A. (2020). On the Identification of Gross Output

Production Functions. *Journal of Political Economy*, 128(8):2973–3016.

Graetz, G. and Michaels, G. (2018). Robots at Work. *Review of Economics and Statistics*, 100(5):753–768.

Levinsohn, J. and Petrin, A. (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies*, 70(2):317–341.

McElheran, K., Yang, M.-J., Kroff, Z., and Brynjolfsson, E. (2025). The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve (s). Technical report.

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

Olley, G. S. and Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, 64(6):1263–1297.

Park, D. and Shin, K. (2025). Implications of Artificial Intelligence and Robots for Employment and Labor Productivity: Firm-Level Evidence from the Republic of Korea. Technical Report 769, Asian Development Bank, Economics Working Paper Series.

Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The Impact of AI on Developer Productivity: Evidence from Github Copilot. Technical report, arXiv.

Raval, D. (2023). Testing the production approach to markup estimation. *Review of Economic Studies*, 90(5):2592–2611.

Statistics Korea (2022). The relationship between management style and corporate performance using the manufacturing industry survey. Technical report, KOSTAT.

Wang, Z. (2025). Non-neutral Technological Change in Chinese Manufacturing. *Available at SSRN 5176447*.

Tables

Figures

Appendices

A Data Appendix

A.1 AI adoption rate

Table A.1: New Technology Adoption Rate by Year

Year	Overall	AI	Robotics	IoT	3D Printer	5G	Cloud	Big data	Blockchain	AR/VR
2017	8.2%	1.2%	1.0%	2.7%	1.3%	3.1%	2.7%	2.7%	0.6%	0.7%
2018	11.8%	2.4%	1.8%	4.3%	1.8%	2.9%	4.4%	3.8%	0.7%	1.1%
2019	12.9%	2.8%	2.2%	4.3%	2.2%	2.3%	5.1%	4.1%	0.5%	1.2%
2020	12.6%	3.0%	2.3%	3.9%	2.0%	2.3%	5.4%	4.1%	0.7%	1.3%
2021	13.4%	3.5%	2.4%	4.2%	1.9%	2.4%	5.7%	4.5%	0.7%	1.4%
2022	13.1%	3.7%	2.4%	4.1%	1.6%	2.1%	5.7%	4.2%	0.5%	1.2%
2023	17.3%	5.3%	3.2%	5.5%	2.3%	3.1%	8.4%	5.8%	0.7%	1.3%

Notes: This table reports the share of firms adopting each digital technology by year from 2017 to 2023. Overall adoption indicates whether the firm uses at least one of the nine technologies. Data source: Survey of Business Activities.

Table A.1 presents the adoption rates of various digital technologies among Korean manufacturing firms from 2017-2023. Overall technology adoption has shown a substantial increase over the period, rising from 8.2% in 2017 to 18.1% in 2023, with particularly notable acceleration in 2023. AI adoption, the focus of our study, demonstrates steady growth from 1.2% in 2017 to 5.6% in 2023, making it one of the faster-growing technologies alongside cloud and big data, which increased from 2.7% to 8.8% and from 2.7% to 6.2%, respectively, over the same period. IoT (Internet of Things) adoption also more than doubled, increasing from 2.7% to 5.8%. The adoption patterns reveal that firms are increasingly embracing complementary digital technologies, with cloud infrastructure and big data showing particularly strong uptake as foundational technologies that support other digital innovations. In contrast, emerging technologies like blockchain and AR/VR maintain relatively low adoption rates of 0.8% and 1.4% respectively by 2023.

To better understand how firms utilize AI once adopted, we examine the specific applications and development approaches. AI development includes self-development or outsourcing (domestic or abroad). AI usage encompasses product development, production processes, sales/marketing, and organizational management. Table A.1 shows AI usage patterns in 2023 across different applications and the distribution of where AI is being implemented. Among the 516 AI-adopting firms, product (service) development is the most common application, accounting for 56.4% of AI usage. Production processes and sales/marketing strategy are the next most frequent use cases at 18.0% and 17.4%, respectively. AI usage in organizational management (8.3%) is the least common. Notably, the 32 firms that develop AI internally are not included in any of the usage categories, as they create rather than apply AI technologies.

B Heterogeneity Effect of AI adoption by Sector × Application

To allow the impact of AI on productivity to vary jointly by sector and by application area, we extend the productivity law of motion to

$$\omega_{it} \;=\; \rho\,\omega_{i,t-1} \;+\; \sum_{g\in\{4\text{ sectors}\}} \;\sum_{a\in\{\text{pdev, prod, sales, org}\}} \gamma_{g,a}\,\mathrm{AI}^{g,a}_{i,t-1} \;+\; \xi_{it},$$

where s_i denotes the (time-invariant) sector of firm i, and $\mathrm{AI}^{a}_{i,t-1}$ is an indicator that equals one if firm i has adopted AI by period $t-1$ and applies it in area a in sector g. The coefficients $\gamma_{g,a}$ capture sector- and application-specific semi-elasticities of productivity with respect to lagged AI adoption.

Across sectors and applications, the estimated AI semi-elasticities display sharp heterogeneity. In manufacturing, AI in product development and sales appears weakly negative, suggesting that manufacturing firms may be more transitional adoption costs. By contrast, ICT firms exhibit large positive effects of AI in product development and sales/marketing, consistent with these firms being closer to the technology frontier and better able to translate AI tools into new products and customer acquisition. In the trade sector, AI in production/logistics and organizational functions is strongly negative, while sales-oriented AI is mildly positive, pointing to potential disruption and adjustment costs when AI is deployed in back-office and operational tasks. Finally, producer–consumer services show positive but imprecisely estimated effects of AI in product development, with other applications close to zero.

Note that we should be cautious when we interpret the large estimated effect in AI application like sales and organization management because of the small sample of adopters. Given relatively low adoption rates and small cell sizes—especially for some sector × application combinations—many of these coefficients will be imprecise and highly sensitive to specification. Overall, these results indicate that "where" AI is used inside the firm matters at least as much as "whether" it is adopted, and that the most productive applications differ systematically across sectors.

Table B.2: Two-step GMM with Sector- and Application-specific AI Effects

Parameter	**Estimate**	**Std. Err.**
Panel A: Core elasticities and persistence		
β_K	0.255	0.035
β_L	0.356	0.007
β_M	0.508	0.068
ρ	1.000	0.001
Panel B: AI semi-elasticity by sector × application, $\gamma_{s,a}$		
Manufacturing		
$\gamma_{\text{Mfg, pdev}}$	−0.107	0.018
$\gamma_{\text{Mfg, prod}}$	0.007	0.022
$\gamma_{\text{Mfg, sales}}$	−0.068	0.028
$\gamma_{\text{Mfg, org}}$	0.235	0.079
ICT		
$\gamma_{\text{ICT, pdev}}$	0.220	0.037
$\gamma_{\text{ICT, prod}}$	0.070	0.207
$\gamma_{\text{ICT, sales}}$	0.206	0.050
$\gamma_{\text{ICT, org}}$	0.079	0.103
Trade		
$\gamma_{\text{Trade, pdev}}$	−0.065	0.054
$\gamma_{\text{Trade, prod}}$	−0.292	0.096
$\gamma_{\text{Trade, sales}}$	0.076	0.057
$\gamma_{\text{Trade, org}}$	−0.364	0.003
Producer–Consumer Services		
$\gamma_{\text{PCS, pdev}}$	0.189	0.123
$\gamma_{\text{PCS, prod}}$	−0.029	0.112
$\gamma_{\text{PCS, sales}}$	−0.005	0.131
$\gamma_{\text{PCS, org}}$	−0.119	0.058

Notes: Two-step GMM estimates with firm-clustered standard errors. Stage-1 sample size: $N = 57{,}928$; Stage-2 sample size: $N = 46{,}316$. The model is just-identified, so the Hansen J-statistic is not informative.