The Macroeconomics of Generative AI: Firm Entry and Price Dynamics

Jose Carreño*
University of Oregon

June 2026

First version: March 2026

Abstract

Much of the early economic literature on generative artificial intelligence has focused on its labor-market consequences. This paper studies a complementary margin of adjustment: the product-market effects of generative AI on firm entry and price dynamics. It constructs a granular panel combining predetermined measures of generative AI exposure with administrative data on U.S. business applications and sectoral Producer Price Indices. Using dynamic continuous-treatment event-study designs, the paper documents two margins of product-market adjustment. First, more exposed industries experience stronger growth in firm entry after the commercialization of generative AI. Moving across the interquartile range of AI exposure is associated with approximately 6.2 percentage points higher business-application growth by 2024. Second, more exposed sectors experience lower relative price growth after 2023, in specifications that control for differential exposure to pandemic-era supply-chain stress. This price difference becomes more pronounced over time, consistent with gradual diffusion and pass-through of AI-related efficiency gains to producer prices. Moving across the interquartile range of AI exposure is associated with approximately 4.2 percentage points lower cumulative excess inflation by 2025. The paper then develops a static model of monopolistic competition with heterogeneous firms to organize the empirical findings. Within the benchmark economy, these product-market forces generate a consumer welfare dividend: lower prices raise purchasing power, while a larger mass of firms expands the set of available varieties.

JEL Classification: E31, L26, O33

Keywords: Generative AI, Firm Entry, Price Dynamics, General-Purpose Technology, Consumer Welfare

*I am grateful to Haritz Garro, Huiyu Li, and Ben Pugsley for helpful comments and discussions. All errors are my own. Email: josecarreno.econ@gmail.com.

1 Introduction

Generative artificial intelligence may become one of the defining general-purpose technologies of the coming decades. Since its widespread commercialization in late 2022, the use of generative AI tools by firms has expanded rapidly. Data from the U.S. Census Bureau's Business Trends and Outlook Survey (BTOS) show a steady increase in the share of firms reporting current AI use, while near-term expected use remains persistently above current adoption, pointing to continued expansion in enterprise adoption intent (Figure 1). At the same time, aggregate enterprise adoption remains far from complete. This combination of rapid diffusion and still-limited penetration creates a useful empirical window for studying the early economic effects of generative AI and identifying emerging macroeconomic trends before they fully mature. Much of the early public debate and economic research has focused, naturally, on labor-market exposure, worker displacement, and changes in the demand for cognitive skills (e.g., Autor, 2024; Brynjolfsson et al., 2025a). These are central issues. This paper studies a complementary margin of adjustment: the effect of generative AI on product markets.

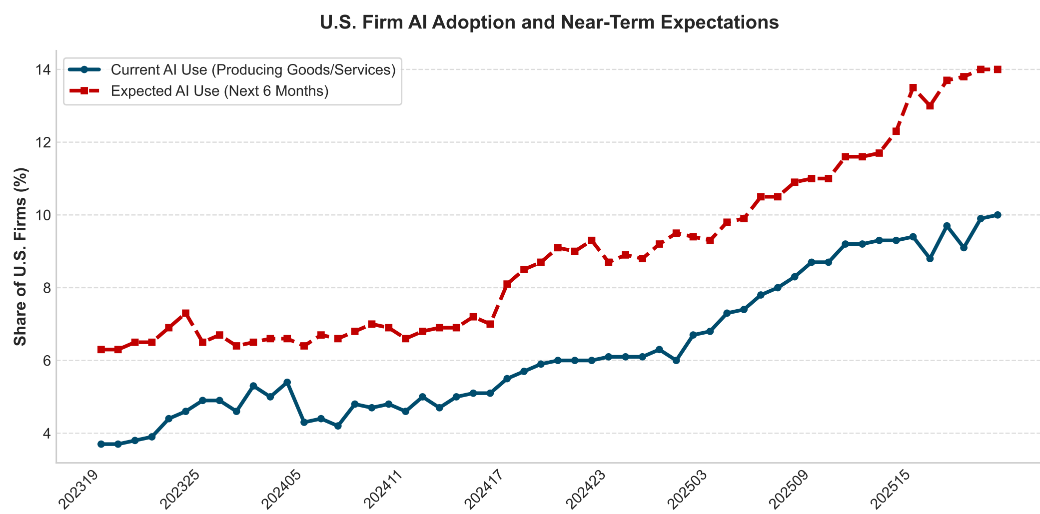

Figure 1: U.S. Firm AI Adoption: Current Use and 6-Month Projections

Notes: The figure plots the bi-weekly share of U.S. firms reporting current use of artificial intelligence in producing goods or services, alongside the share of firms expecting to use artificial intelligence within the next six months. Data are from the U.S. Census Bureau's Business Trends and Outlook Survey (BTOS).

The product-market effects of generative AI follow from its ability to automate or augment cognitive, administrative, and organizational tasks (Eloundou et al., 2024). These tasks matter not only for workers and incumbent firms, but also for potential entrants and for the cost structure of production. By reducing the cost of activities such as information processing, documentation, software development, customer communication, marketing, and coordination, generative AI may lower entry-related overhead and reduce marginal production costs in sectors where exposed tasks account for a larger share of activity. This suggests two product-market margins through which the technology may operate. First, lower overhead may reduce barriers to entry and support stronger firm formation. Second, lower marginal costs may put downward pressure on relative prices, with the strength and timing of pass-through depending on adoption, task reorganization, and market conditions. The two outcomes need not be mechanically linked in the short run. Rather, they can be viewed as complementary responses to the same underlying technological force: a reduction in the cost of performing exposed cognitive and administrative tasks.

First, we study firm entry. We construct a granular industry-level panel that combines predetermined measures of generative AI exposure with administrative data on U.S. business applications at the 3-digit NAICS level. We measure entry growth using Davis-Haltiwanger-Schuh growth rates, which make growth comparable across sectors with different baseline levels of applications. In descriptive cross-sectional regressions, more exposed sectors experience faster growth in business applications during the initial post-commercialization period. We then estimate a dynamic continuous-treatment event-study design. The pre-period coefficients show no systematic evidence that more exposed sectors were already on faster entry-growth paths before the diffusion of generative AI. After 2022, the pattern changes: by 2024, the coefficient becomes positive, statistically significant, and economically meaningful. Moving from the 25th to the 75th percentile of AI exposure is associated with approximately 6.2 percentage points higher DHS growth in business applications relative to 2022. Complementary evidence from the 2-digit Business Formation Statistics extends the analysis through 2025 and shows that the relationship appears for both High-Propensity and Low-Propensity applications, suggesting that the entry response is not confined to applications with a low ex-ante probability of becoming employer businesses.

Second, we study price dynamics. If generative AI lowers the cost of exposed cognitive and administrative tasks, sectors with greater exposure should experience lower relative price growth as adoption progresses and cost reductions pass through to producer prices. To test this hypothesis, we construct a granular panel of sectoral Producer Price Indices and estimate a dynamic continuous-treatment event-study design. This design traces the timing of relative price changes, tests for differential pre-existing trends, and includes controls for differential exposure to pandemic-era supply-chain stress using predetermined input-output measures interacted with the Global Supply Chain Pressure Index. The estimates show no systematic evidence that more and less exposed sectors were on different price trajectories before the diffusion of generative AI. After 2023, however, more exposed sectors experience lower relative price growth, and this relative price difference becomes more pronounced over time. This dynamic pattern is consistent with a diffusion process in which AI-related efficiency gains gradually accumulate and pass through to sectoral prices. A complementary cross-sectional inflation-gap exercise yields a closely aligned magnitude: moving from the 25th to the 75th percentile of AI exposure is associated with approximately 4.2 percentage points lower cumulative excess inflation by 2025.

Third, we develop a simple theoretical framework to organize the two empirical findings. It takes the form of a static model of monopolistic competition with heterogeneous firms, building on Melitz, 2003, in which generative AI is represented as a reduced-form improvement in cognitive labor efficiency. The model is designed as an organizing framework rather than a quantitative account of the transition path observed in the data. Its purpose is to show, in a parsimonious setting with closed-form comparative statics, how an improvement in cognitive labor efficiency can simultaneously increase the equilibrium mass of firms and lower the aggregate price index. In the model, this improvement lowers marginal production costs and reduces fixed operating and entry-related overhead. Lower overhead supports a larger equilibrium mass of firms, while lower marginal costs reduce firm-level prices. In a CES environment, the larger mass of available varieties further lowers the aggregate price index through the standard love-for-variety channel. Within the benchmark economy, these product-market forces generate a consumer welfare dividend: lower prices raise purchasing power, while a larger mass of firms expands the set of available varieties.

These findings should be interpreted in light of the broader timeline of technological adoption. As indicated by current enterprise usage rates (Figure 1), generative AI remains in an early stage of macroeconomic diffusion. This matters because general-purpose technologies—from electrification to the internet—have historically required extended periods of organizational restructuring, complementary investment, and intangible capital accumulation before their aggregate effects were fully realized (David, 1990; Mokyr, 1990; Brynjolfsson et al., 2021). Generative AI may diffuse more rapidly than earlier technologies, but its economic effects are still likely to unfold over time as firms experiment with the technology, reorganize tasks, and integrate AI-enabled tools into production and administrative processes. The empirical estimates in this paper therefore provide an early evaluation of an ongoing adjustment rather than a measure of the full long-run effect. The product-market responses documented here may capture only the initial phase of this process. As adoption deepens, the entry and price margins studied here may become increasingly important.

This paper contributes to three main strands of literature. First, we contribute to the rapidly growing literature on the economic effects of artificial intelligence. Much of this work has focused on labor-market exposure, task displacement, changes in the demand for cognitive skills, and the productivity effects of AI tools within firms (Acemoglu and Restrepo, 2018; Eloundou et al., 2024; Brynjolfsson et al., 2025b; Acemoglu, 2025; Autor, 2024; Brynjolfsson et al., 2025a). We complement this literature by studying the product-market side of the macroeconomic adjustment to generative AI. Rather than focusing only on workers, occupations, or incumbent-firm productivity, we examine how exposure to generative AI is associated with changes in firm entry and price dynamics. This perspective highlights a channel through which the economic effects of generative AI may extend beyond labor-market adjustment. In the data, more exposed sectors experience stronger entry growth and lower relative price growth after the commercialization of generative AI. In the model, these product-market forces generate a consumer welfare dividend: lower prices raise purchasing power, while a larger mass of firms expands the set of available varieties. The paper therefore connects the labor- and task-based view of AI to a complementary product-market margin through which the technology may raise consumer welfare.

Second, we contribute to the literature on firm dynamics and technological adoption (Hopenhayn, 1992; Decker et al., 2014). Theoretical models of firm dynamics emphasize that reductions in fixed operating costs and entry-related costs can affect business formation, market structure, and aggregate dynamism. However, empirical evaluations of previous general-purpose technologies—such as the ICT revolution—often found that aggregate effects became visible only gradually, as firms reorganized production, made complementary investments, and accumulated new intangible capital, a pattern closely related to the Solow productivity paradox and the productivity J-curve (Brynjolfsson, 1993; Brynjolfsson et al., 2021). Against this background, generative AI provides a useful setting for studying whether product-market responses can be detected at an earlier stage of diffusion. We study one such margin: the cost of cognitive, administrative, and organizational tasks faced by potential entrants. Recent empirical work shows that AI adoption is associated with growth and product innovation among incumbent firms (Babina et al., 2024). We complement this evidence by studying *de novo* firm formation. Using granular industry-level variation in predetermined AI exposure, we show that more exposed sectors experience stronger growth in business applications after the commercialization of generative AI. Complementary evidence from the standard Business Formation Statistics indicates that this relationship appears for both High-Propensity and Low-Propensity Applications, suggesting that the entry response is not confined to applications with a low ex-ante probability of becoming employer businesses.

Third, we contribute to the macroeconomic literature on technological change and price dynamics. Earlier work has linked information technology, automation, and other productivity-enhancing innovations to changes in costs, productivity, and prices (Oliner et al., 2007; Graetz and Michaels, 2018). Evidence on the price effects of generative AI, however, remains scarce. This is partly because the technology is still at an early stage of enterprise diffusion, and partly because its commercialization occurred shortly after the pandemic-era volatility in supply chains and prices. We address these challenges by studying relative price dynamics across sectors with different predetermined exposure to generative AI in a dynamic continuous-treatment event-study design. The design traces the timing of relative price changes, tests for differential pre-existing trends, and controls for differential exposure to pandemic-era supply-chain stress using predetermined input-output measures interacted with the Global Supply Chain Pressure Index. The results provide early evidence that more exposed sectors experience lower relative price growth after generative AI begins to diffuse. This price divergence becomes more pronounced over time, a pattern consistent with gradual diffusion and pass-through rather than a one-time price adjustment at the moment of commercialization.

The remainder of the paper is organized as follows. Section 2 describes the data construction and the measurement of generative AI exposure. Section 3 presents the empirical analysis of firm entry. Section 4 studies price dynamics. Section 5 develops the theoretical framework. Section 6 concludes.

2 Data and Measurement

This section describes the data sources and measurement procedures used to construct our industry-level panels. We combine occupational task evaluations with granular employment and wage data to construct a predetermined sector-level index of generative AI exposure. We then pair this index with administrative data on business formations and producer price indices to construct our macroeconomic outcomes. To maintain the focus on the core economic mechanisms and empirical results, we relegate detailed data construction procedures and supplementary descriptive statistics to Appendix A.

2.1 Measuring Generative AI Exposure

To quantify sectoral exposure to generative AI, we build on the occupational task-level framework developed by Eloundou et al., 2024. Their methodology assesses the extent to which access to a Large Language Model (LLM) can reduce the time required for a human worker to perform specific occupational tasks by at least 50 percent. We focus on their β measure, which accounts not only for direct LLM capabilities but also for the broader software and tooling ecosystem built on top of these models. To reduce reliance on any single annotation procedure, we define our baseline occupational exposure score as the simple average of the human-annotated and machine-generated (GPT-4) evaluations provided in their dataset.

Because our macroeconomic outcomes—firm entry and producer prices—are observed at the sectoral level, we map occupational exposures into industry-level exposure indices. Following the standard approach in the literature on task-biased technological change and automation (Autor, 2013; Acemoglu and Restrepo, 2020), we define the AI exposure of industry j as the wage-bill-weighted average exposure of its constituent occupations:

$$AI\_Exposure_j = \sum_{o \in j} \left(\frac{w_o E_{o,j}}{\sum_{o \in j} w_o E_{o,j}} \right) Exposure\_Score_o, \tag{1}$$

where $E_{o,j}$ is employment of occupation o in industry j, w_o is the mean annual wage of occupation o, and $Exposure\_Score_o$ is the occupational generative AI exposure measure. Weighting by the wage bill, rather than raw employment, gives greater weight to occupations that account for a larger share of labor costs within each industry's production structure.

We construct this index using data from the Bureau of Labor Statistics (BLS) Occupational Employment and Wage Statistics (OEWS) program. We use the May 2022 OEWS release, which predates the public release of ChatGPT in late 2022. Fixing occupational employment and wage weights before the commercialization of generative AI ensures that the sectoral exposure measure is predetermined with respect to the post-release outcomes studied below and is not affected by subsequent changes in occupational composition.

We construct the exposure index at both the 3-digit and 2-digit NAICS levels. The 3-digit index is used in the main analyses of firm entry and price dynamics, where the more granular industry panel provides richer cross-sectional variation. The 2-digit index is used for complementary analyses based on the standard sectoral Business Formation Statistics series, including the decomposition of applications into High-Propensity and Low-Propensity categories.[1]

[1] For the 2-digit NAICS aggregation, we follow the U.S. Census Bureau standard practice of grouping Manufacturing (31–33), Retail Trade (44–45), and Transportation and Warehousing (48–49). See Appendix A for more details.

2.2 Firm Entry: Business Applications

To measure firm entry, we use the Business Formation Statistics (BFS) compiled by the U.S. Census Bureau. The BFS are based on applications for Employer Identification Numbers (EINs) submitted to the Internal Revenue Service and provide a high-frequency measure of early-stage business formation. Prior work shows that these applications are informative about subsequent business formation and job creation (Bayard et al., 2018).

We use two BFS data products. The standard sectoral BFS series provides monthly business applications at the 2-digit NAICS level and includes the Census Bureau's classification of applications by their ex-ante propensity to become employer businesses. This series is useful for studying the composition of firm entry. The Census Bureau's Weekly BFS by Industry product provides more granular application counts at the 3-digit NAICS level. This series offers richer cross-sectional variation across industries and serves as the basis for our main firm-entry analysis.

For the 2-digit analysis, we use the standard sectoral BFS series. This dataset provides three measures of business applications. Total Business Applications (BA) measure the aggregate number of EIN applications within each sector. High-Propensity Applications (HPA) identify applications with characteristics historically associated with a high probability of becoming employer businesses, such as corporate status, a planned date for first payroll, or industries in which non-employer firms are rare. We define Low-Propensity Applications (LPA) as the residual difference between total applications and high-propensity applications, $LPA_{j,t} = BA_{j,t} - HPA_{j,t}$. This residual captures applications less likely to transition into employer businesses, including independent contractors, zero-employee startups, and other small-scale entrepreneurial activity. The 2-digit NAICS panel spans 2011–2025 and covers 19 private sectors.

For the main firm-entry analysis, we use the Census Bureau's Weekly BFS by Industry product, which reports application counts at the 3-digit NAICS level. This substantially more granular series allows us to exploit within-sector variation that is not visible in the standard 2-digit panel. The Census Bureau assigns 3-digit NAICS codes using an automated industry-coding program based on the business names and descriptions provided on EIN applications. We aggregate the underlying weekly application counts into annual totals to match the frequency of the rest of the analysis and to focus on medium-run entry dynamics rather than high-frequency fluctuations.

The additional granularity of the 3-digit series comes with classification noise. Some applications cannot be assigned to a specific 3-digit subsector and are instead reported as

"Not Further Defined" (NFD) or left unclassified within their broader parent sector. To limit the influence of unresolved classifications, we apply the data-quality filter described in Appendix A, which excludes 3-digit subsectors belonging to parent sectors with persistently high NFD shares. The resulting 3-digit NAICS panel spans 2011–2024 and covers 63 private industries.

2.3 Producer Price Indices

To study sectoral price dynamics, we use Producer Price Indices (PPI) from the Bureau of Labor Statistics (BLS). The PPI measures the average change over time in the selling prices received by domestic producers for their output. This makes it well suited for studying price responses to sectoral cost changes, since it captures producer prices closer to the point of production than consumer price measures such as the CPI, which also reflect retail margins, imported final goods, indirect taxes, and consumer expenditure weights.

To align the price data with our occupational exposure measure, we construct a monthly panel of producer prices at the 3-digit NAICS level. This level of aggregation provides substantial cross-sectional variation across industries while avoiding some of the measurement noise associated with highly disaggregated product-level price series.

We obtain industry-level PPI series from the Federal Reserve Economic Data (FRED) database, which republishes BLS PPI series. The BLS provides 3-digit NAICS price aggregates for most sectors in our analysis, but some sectors are only available at finer levels of disaggregation. To maximize sectoral coverage while applying a systematic rule, we reconstruct missing 3-digit aggregates using the unweighted average growth rate of their available sub-sectoral components. This procedure yields a final panel of 65 distinct 3-digit NAICS sectors spanning January 2011 through April 2026. Appendix A provides additional details on the construction of the PPI panel.

2.4 Supply-Chain Stress and Exposure Measures

Because the dynamic price analysis spans the COVID-19 period, we also construct controls for differential exposure to supply-chain stress. The time-series component of these controls is the Global Supply Chain Pressure Index (GSCPI) produced by the Federal Reserve Bank of New York. The GSCPI summarizes disruptions in global supply chains using information from transportation costs and manufacturing supply indicators, and is normalized in standard-deviation units relative to its historical distribution. We aggregate the monthly index to the annual level using its calendar-year mean, matching the frequency of our sector-level price panel.

To capture heterogeneous sectoral exposure to these aggregate disruptions, we construct two predetermined input-output measures using the 2017 detailed BEA Supply-Use framework and the 2017 BEA Import Matrix. The first measure captures the share of gross output spent on goods-producing and transportation/warehousing inputs. This variable identifies sectors whose production processes rely more heavily on physical inputs and logistics services. The second measure captures imported-input intensity, defined as imported intermediate inputs divided by total intermediate input use. This variable identifies sectors more exposed to disruptions in global rather than purely domestic supply chains.

We use the 2017 BEA detailed tables because they are the latest detailed benchmark input-output tables available before both the pandemic shock and the commercialization of generative AI. The resulting exposure measures are fixed at their 2017 values throughout the analysis. Appendix A provides construction details and summary statistics.

2.5 Sectoral Characteristics

We also construct an industry-level measure of labor intensity to account for baseline structural differences across sectors that may influence firm entry and pricing behavior.

Labor share is defined as total employee compensation divided by industry value added. We obtain both series from the BEA "Components of Value Added by Industry" accounts. To keep this measure predetermined relative to the commercialization of generative AI and avoid the unusual movements in factor shares during the COVID-19 period, we compute labor share as an average over 2016–2019:

$$Labor\_Share_j = \frac{\sum_{t=2016}^{2019} Compensation_{j,t}}{\sum_{t=2016}^{2019} Value\_Added_{j,t}}. \tag{2}$$

We construct this measure at both the 2-digit and 3-digit NAICS levels, aggregating the underlying BEA industry codes to match the sector definitions used in each empirical exercise.

3 Firm Entry

We first examine whether generative AI exposure is associated with changes in firm entry. A central implication of generative AI is that it can automate or augment cognitive, administrative, and organizational tasks (Eloundou et al., 2024). These tasks are important not only for incumbent firms, but also for potential entrants: writing code, producing marketing material, managing customer interactions, preparing business documents, analyzing information, and coordinating operations. If generative AI lowers the cost of these activities, it may reduce barriers to entry in sectors where exposed tasks account for a larger share of production.

Our entry hypothesis follows from this task-based channel. Sectors with greater exposure to generative AI should experience stronger entry responses after the technology begins to diffuse, because potential entrepreneurs and young firms in those sectors can substitute AI tools for some tasks that would otherwise require labor, specialized expertise, or external services. The effect need not be immediate. New firms may respond only after a period of experimentation, task reorganization, and identification of business opportunities. We therefore ask whether sectors with greater generative AI exposure exhibit stronger firm-entry growth after the commercialization of the technology, and whether such differences were already present before.

To evaluate this hypothesis, we use the 3-digit NAICS Business Formation Statistics panel described in Section 2.2. We measure firm entry using business applications and focus on DHS growth rates, which allow us to compare entry growth across sectors with different baseline levels of applications. The 3-digit panel provides substantially richer cross-sectional variation than the standard 2-digit sectoral series, making it well suited for studying differential entry dynamics across levels of AI exposure.

We study the relationship between AI exposure and firm entry through two complementary exercises. The first provides descriptive cross-sectional evidence, relating AI exposure to post-commercialization growth in business applications across 3-digit sectors. This exercise summarizes the raw entry-growth pattern in the granular data. The second implements a dynamic continuous-treatment event-study design. This design uses the full annual panel to trace the timing of relative entry growth and to examine whether sectors with different exposure levels followed similar trajectories before the commercialization of generative AI. Together, the two exercises provide evidence on the magnitude, timing, and pre-period dynamics of the firm-entry response.

3.1 Descriptive Evidence

We begin with a simple descriptive relationship between generative AI exposure and subsequent firm-entry growth in the 3-digit NAICS panel. The purpose of this exercise is not to provide the main identification test, but to summarize the cross-sectional pattern that motivates the dynamic analysis below. We measure entry growth using the Davis-Haltiwanger-Schuh growth rate (Davis et al., 1996), defined as the change in business applications between two years divided by their average. This measure is symmetric around zero and is commonly used in firm-dynamics settings to compare growth rates across units with different baseline levels.

Figure 2 plots the DHS growth rate in business applications from 2023 to 2024 against the sectoral generative AI exposure index. The figure shows a positive relationship: sectors with greater exposure to generative AI experienced faster growth in business applications over this period. The fitted line is unconditional, and the scatter is meant to provide a transparent visual summary of the raw cross-sectoral relationship in the granular entry data.

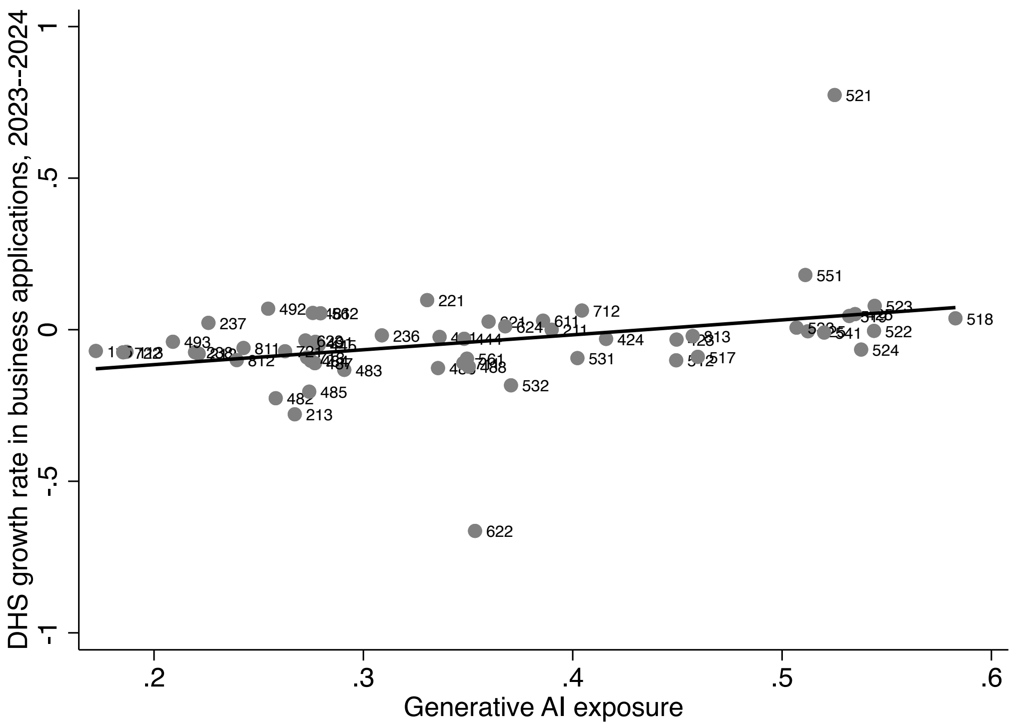

Figure 2: Generative AI Exposure and Firm Entry Growth

Notes: The figure plots the DHS growth rate in business applications from 2023 to 2024 against the generative AI occupational exposure index across private 3-digit NAICS sectors. The DHS growth rate is defined as the change in applications divided by the average of applications in 2023 and 2024. The sample applies the classification-quality filter described in Appendix A. The fitted line is the unconditional linear relationship and does not include the labor-share control.

Table 1 reports corresponding cross-sectional regressions. Column (1) uses DHS growth from 2023 to 2024, matching the window shown in the figure. Column (2) uses the broader

2022–2024 window. Both specifications include pre-pandemic labor share as a control, since AI-related reductions in cognitive and administrative tasks may have different implications for entry depending on the sector's baseline labor intensity. The coefficient on generative AI exposure is positive and statistically significant in both columns. In the 2023–2024 specification, the coefficient is 0.501. Evaluated over the interquartile range of the 3-digit NAICS AI exposure index, from 0.273 to 0.402 (Table A.1), this estimate implies that moving from the 25th to the 75th percentile of exposure is associated with a 6.5 percentage point higher DHS growth rate in business applications. The corresponding estimate for 2022–2024 implies a similar 5.8 percentage point difference.

Table 1: Generative AI Exposure and Firm Entry Growth

	(1) Growth 2023–2024	(2) Growth 2022–2024
Generative AI Exposure (β)	0.501**	0.450**
	(0.198)	(0.191)
Labor Share (2016–2019)	0.021	0.210*
	(0.103)	(0.111)
Observations	55	55
R^2	0.118	0.145

Notes: The dependent variable is the DHS growth rate in business applications. Column (1) uses growth from 2023 to 2024, while Column (2) uses growth from 2022 to 2024. The DHS growth rate is defined as the change in applications divided by the average of applications in the initial and final years. The sample consists of 55 private 3-digit NAICS sectors that pass the classification-quality filter described in Appendix A and have matched AI exposure and labor-share data. The full filtered 3-digit entry panel contains 63 sectors; the smaller regression sample reflects limited availability of matched labor-share data. All specifications include pre-pandemic Labor Share as a control. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

These descriptive results suggest that sectors more exposed to generative AI experienced faster entry growth during the initial post-commercialization period. They do not, by themselves, establish the timing of the divergence or rule out differential pre-existing trends. The next subsection therefore turns to a dynamic event-study design, which uses the annual panel to examine whether the relationship emerges after the commercialization of generative AI and whether high- and low-exposure sectors followed similar trajectories before that point.

3.2 Dynamic Event Study

We next estimate a dynamic event-study specification to examine the timing of the relationship between generative AI exposure and firm-entry growth. The descriptive evidence above shows that more exposed sectors experienced faster growth in business applications after the commercialization of generative AI. The event-study design asks whether this relationship emerges after 2022 and whether more and less exposed sectors displayed systematically different entry-growth patterns before that point.

We estimate the following two-way fixed effects specification:

$$DHS\_Growth_{j,t} = \alpha_j + \tau_t + \sum_{k=2011, k\neq 2022}^{2024} \beta_k \left(AI\_Exposure_j \times \mathbb{I}\{t = k\}\right) + \varepsilon_{j,t}. \tag{3}$$

The dependent variable, $DHS\_Growth_{j,t}$, is the Davis-Haltiwanger-Schuh growth rate in business applications for sector j between year $t-1$ and year t. Sector fixed effects, α_j, absorb time-invariant differences in average entry growth across industries. Year fixed effects, τ_t, absorb aggregate shocks to business formation common to all sectors. The coefficients of interest are the dynamic exposure coefficients, β_k, which measure the year-specific difference in entry growth associated with a one-unit higher generative AI exposure index, relative to the omitted reference year 2022. The identifying assumption is that, absent the diffusion of generative AI, sectors with different exposure levels would not have experienced systematically different changes in business applications after the reference year. The pre-2022 coefficients provide a diagnostic for this assumption. Standard errors are clustered at the 3-digit NAICS sector level.

We use 2022 as the reference year because it is the last year before the main post-commercialization adjustment in business applications can plausibly be observed. ChatGPT was released publicly at the end of November 2022, and business formation responses require time as firms and potential entrepreneurs experiment with new tools, reorganize tasks, and initiate new ventures.

Figure 3 plots the estimated dynamic exposure coefficients. The pre-2022 coefficients fluctuate around zero and do not display a systematic upward pattern. Some estimates are positive and others are negative, and all are statistically indistinguishable from zero. This pattern suggests that sectors with higher AI exposure were not already experiencing persistently faster entry growth before the commercialization of generative AI.

The post-2022 pattern is different. The 2023 coefficient remains statistically insignificant, consistent with a gradual adjustment process rather than an immediate entry response. By 2024, however, the coefficient becomes positive, precisely estimated, and statistically significant at the 1 percent level. This timing is consistent with a diffusion mechanism in which generative AI affects entry after a period of experimentation and task reorganization,

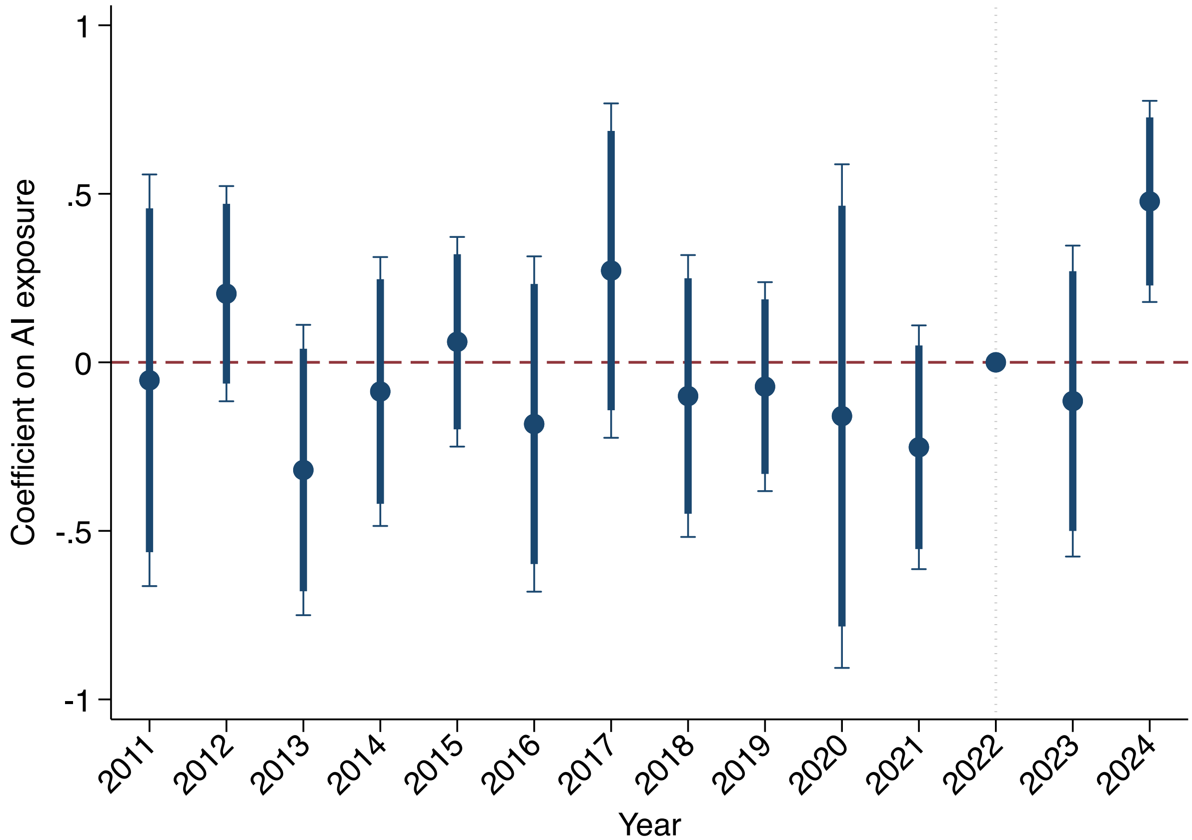

Figure 3: Dynamic Effect of Generative AI Exposure on Firm Entry

Notes: The figure plots the dynamic exposure coefficients β_k from a continuous-treatment event-study regression of DHS growth in business applications on interactions between year indicators and the generative AI occupational exposure index. Each coefficient measures the year-specific difference in firm-entry growth associated with a one-unit higher generative AI exposure index, relative to the omitted reference year 2022, which is normalized to zero. The specification includes 3-digit NAICS fixed effects and year fixed effects. The inner thick and outer thin vertical lines represent 90% and 95% confidence intervals, respectively. Standard errors are clustered at the 3-digit NAICS sector level. Full estimation results are reported in Appendix C.

rather than producing an immediate discrete break at the time of public release.

The magnitude of the 2024 coefficient is economically meaningful. The point estimate is 0.477. Evaluated over the interquartile range of the 3-digit NAICS AI exposure index, from 0.273 to 0.402 (Table A.1), this implies that moving from the 25th to the 75th percentile of exposure is associated with approximately 6.2 percentage points higher DHS growth in business applications relative to 2022. This magnitude is closely aligned with the descriptive cross-sectional estimates in Table 1.

Taken together, the event-study evidence indicates that the positive relationship between generative AI exposure and firm-entry growth is not driven by a smooth pre-existing trend in more exposed sectors. Instead, the differential entry response becomes visible after the commercialization of generative AI: the 2024 estimate is positive, statistically significant, and economically meaningful. While the design cannot rule out all sector-specific shocks correlated with AI exposure, the timing of the estimates narrows the most direct alternative interpretation: more exposed sectors were not already on a systematically higher entry-growth path before the post-2022 diffusion period. Appendix C reports the full set of coefficients underlying Figure 3.

The 3-digit panel provides the cross-sectional variation needed for the dynamic event-study design, but it does not contain the Census Bureau's high- versus low-propensity classification of applications. Appendix B therefore reports complementary evidence using the standard 2-digit NAICS BFS sectoral series. Although the 2-digit panel contains only 19 broad private sectors and is therefore less suitable for dynamic inference, it has two advantages: it extends the analysis through 2025, and it separates applications into High-Propensity and Low-Propensity categories. The results are consistent with the main 3-digit evidence. In the standard 2-digit sectoral series, sectors with greater AI exposure also display stronger post-2023 growth in business applications through 2025. Importantly, this relationship appears for both High-Propensity and Low-Propensity Applications, suggesting that the entry response is not confined to applications with a low ex-ante probability of becoming employer businesses.

4 Price Dynamics

The previous section showed that sectors more exposed to generative AI experienced a stronger increase in firm entry after the commercialization of the technology. We now turn to a second, distinct product-market margin: sectoral price dynamics. The two outcomes need not be linked mechanically in the short run. Rather, both can be viewed as complementary responses to the same underlying technological force: a reduction in the cost of performing cognitive, administrative, and organizational tasks.

Our pricing hypothesis follows directly from this cost-based channel. To the extent that generative AI automates or augments exposed cognitive tasks (Eloundou et al., 2024), adoption should reduce labor dependencies, administrative overhead, and marginal production costs in more exposed sectors. Under standard pricing models, lower marginal costs should exert downward pressure on producer prices, with the strength and timing of pass-through depending on the speed of adoption, the extent of task reorganization, and market conditions. We therefore ask whether sectors with greater exposure to generative AI exhibit lower relative producer-price growth after the technology begins to diffuse.

To evaluate this hypothesis, we use a granular panel of sectoral Producer Price Indices (PPI) from the Bureau of Labor Statistics, aggregated to annual frequency to focus on medium-run price dynamics rather than high-frequency fluctuations. We study the relationship between AI exposure and sectoral prices through two complementary empirical strategies. The first measures cumulative inflation gaps relative to pre-pandemic price trends. The second uses a dynamic continuous-treatment event-study design.

The cross-sectional exercise provides a simple endpoint comparison. It asks whether sectors with greater AI exposure exhibit lower cumulative excess inflation by 2025, after removing their own pre-pandemic price trends. This approach is useful because it focuses on price divergence after the acute phase of the pandemic-era supply-chain disruption had largely subsided. At the same time, it necessarily summarizes the adjustment in a single endpoint and relies on estimated counterfactual price paths.

The event-study design provides the main dynamic test. It uses the full annual panel to trace the timing of relative price changes, examine whether high- and low-exposure sectors followed similar pre-period trajectories, and explicitly control for differential exposure to pandemic-era supply-chain stress. Taken together, the two exercises provide complementary evidence on the emergence, timing, and magnitude of relative price divergence across levels of generative-AI exposure.

4.1 Cross-Sectoral Inflation Gaps

To evaluate the disinflationary hypothesis, we begin with a cross-sectional exercise that measures cumulative price divergence across sectors relative to their pre-pandemic trajectories. Specifically, we construct a sector-level inflation gap that removes secular, industry-specific price trends and asks whether sectors more exposed to generative AI exhibit systematically lower excess inflation by 2025.

We estimate the baseline inflation trend for each sector j over the pre-pandemic period (2015–2019). Given the high degree of persistence characteristic of price indices, we model the log-level of the Producer Price Index (PPI) as a function of time:

$$\ln(PPI_{j,t}) = \alpha_j + \delta_j t + \varepsilon_{j,t} \tag{4}$$

Using the estimated parameters, we project the counterfactual price trajectory implied by each sector's pre-pandemic trend. Specifically, the expected price index in 2025 is given by

$\tilde{PPI}_{j,2025} = \exp(\hat{\alpha}_j + \hat{\delta}_j t_{2025})$. We then define our primary dependent variable, the *Inflation Gap*, as the percentage deviation of the actual observed prices in 2025 from this baseline counterfactual:

$$Inflation\_Gap_j = \frac{PPI_{j,2025} - \tilde{PPI}_{j,2025}}{\tilde{PPI}_{j,2025}} \tag{5}$$

This metric captures the cumulative excess inflation experienced by each sector relative to its own pre-pandemic trajectory. If generative AI acts as a deflationary force by reducing marginal costs, sectors with higher occupational exposure should exhibit systematically smaller—or even negative—inflation gaps relative to less exposed sectors.

Evaluating this cumulative gap at the end of our sample period in 2025 serves a dual methodological purpose. First, the commercialization and enterprise adoption of generative AI models is a progressive, cumulative process. Measuring the gap in 2025 captures the maximum extent of technological diffusion available in current data. Second, evaluating price divergence several years after the acute phase of the pandemic mitigates the influence of transitory supply-chain disruptions on the cross-sectional comparison. By 2024 and 2025, the severe bottlenecks and demand-supply imbalances that characterized the immediate post-pandemic period had largely normalized.[2] This timing does not eliminate all pandemic-related concerns, but it makes the cross-sectional gap a useful measure of persistent price divergence after the most severe supply-chain disruptions had subsided. The dynamic analysis in Section 4.2 addresses this issue more directly by controlling for sectoral exposure to pandemic-era supply-chain stress in an event-study framework.

To account for fundamental differences in production factor intensities, we retain the sector's historical labor share as a structural control. The microeconomic rationale is straightforward: in highly capital-intensive sectors, AI-driven labor savings represent a relatively small fraction of total marginal costs, limiting the scope for downward price pass-through. This control helps ensure that the estimated relationship is not driven by broad differences in price dynamics between capital- and labor-intensive sectors during the post-pandemic recovery. We formalize the cross-sectional relationship through the following specification:

$$Inflation\_Gap_j = \gamma_0 + \gamma_1 AI\_Exposure_j + \theta Labor\_Share_j + v_j \tag{6}$$

[2] This macroscopic normalization is empirically well documented. The Federal Reserve Bank of New York's Global Supply Chain Pressure Index (GSCPI)—which peaked at an unprecedented 4.3 standard deviations above its historical mean in late 2021—had reverted to its historical average by February 2023 (Federal Reserve Bank of New York, 2023). Corroborating this timeline, the Federal Open Market Committee assessed in late 2023 that the healing of global supply chains and labor-supply imbalances was largely complete (Federal Open Market Committee, 2023). Consequently, the 2024–2025 period provides a cleaner empirical window in which transitory supply bottlenecks are less likely to dominate relative price dynamics.

where γ_1 is the coefficient of interest. Under the disinflationary hypothesis, $\gamma_1 < 0$. To assess robustness to the functional form and horizon of the baseline projection, we estimate Equation (6) using the baseline 2015–2019 trend, a longer 2011–2019 linear trend, and a more flexible quadratic pre-trend that accommodates potential non-linearities in pre-pandemic price growth.

Table 2 reports the cross-sectional estimates of Equation (6). Across all specifications, the coefficient on AI exposure is negative and statistically significant, consistent with lower cumulative excess inflation in more AI-exposed sectors.

Column (1) presents our preferred specification, which defines the counterfactual projection based on the most recent pre-pandemic trend (2015–2019). The point estimate of -0.326 is significant at the 5% level. To interpret this magnitude, we evaluate the effect across the interquartile range of our 3-digit NAICS sample (see Table A.1). Moving from a sector at the 25th percentile of technological exposure ($\beta = 0.273$) to one at the 75th percentile ($\beta = 0.402$) is associated with a 4.2 percentage point reduction in the cumulative excess inflation gap.

This pattern is robust to alternative ways of constructing the counterfactual price trajectory. Column (2) extends the linear pre-trend to the 2011–2019 period, while Column (3) employs a flexible quadratic trend to allow for non-linear historical price dynamics. In both alternative specifications, the point estimates on AI exposure become larger in magnitude (-0.445 and -0.520, respectively) and remain statistically significant. The baseline estimate is therefore not driven by the particular choice of a short 2015–2019 linear projection.

Table 2: Cross-Sectoral Estimates: Generative AI and Inflation Gaps

	(1) Short Trend	(2) Long Trend	(3) Quadratic Trend
Generative AI Exposure (β)	-0.326**	-0.445**	-0.520**
	(0.148)	(0.219)	(0.233)
Labor Share (2016-2019)	0.126	-0.513*	0.252
	(0.092)	(0.286)	(0.161)
Observations	57	57	57
R^2	0.128	0.156	0.149

Notes: The dependent variable is the *Excess Inflation Gap*, defined as the percentage deviation of the observed 2025 Producer Price Index from its counterfactual projection based on historical pre-pandemic trends. Column (1) projects from the 2015–2019 linear trend; Column (2) from the 2011–2019 linear trend; and Column (3) from the 2011–2019 quadratic trend. The estimation sample comprises 57 NAICS 3-digit industries due to the limited availability of historical labor share data for a small subset of 6 sectors. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

These cross-sectional results provide suggestive evidence of cumulative relative price divergence by 2025. At the same time, the gap approach necessarily relies on estimated counterfactual price trajectories and summarizes the adjustment in a single endpoint. It also does not directly trace how the relationship evolves over the pandemic and post-pandemic period. We therefore turn to a dynamic event-study design, which examines whether highly exposed sectors followed parallel pre-period trajectories, traces the timing of the subsequent price divergence, and explicitly addresses differential exposure to pandemic-era supply-chain stress.

4.2 Dynamic Event Study

We next implement a continuous-treatment event-study design—a continuous difference-in-differences framework—to examine the dynamic relationship between generative AI exposure and sectoral prices. This design serves three purposes. First, it allows us to assess pre-period price dynamics: whether more and less AI-exposed sectors followed similar trajectories before the commercialization of generative AI. Second, it traces the full dynamic profile of relative price differences, rather than summarizing the adjustment in a single endpoint. Third, it allows us to address a key confounder in this setting: supply-chain stress during the pandemic period, which affected sectors differentially depending on their exposure to goods, logistics, and imported inputs.

We estimate the following two-way fixed effects specification:

$$\ln(PPI_{j,t}) = \alpha_j + \tau_t + \sum_{k=2011,k\neq b}^{T} \beta_k \left(AI\_Exposure_j \times \mathbb{I}\{t=k\}\right) + \mathbf{SC}'_{j,t}\boldsymbol{\lambda} + \varepsilon_{j,t} \tag{7}$$

Here, $\ln(PPI_{j,t})$ denotes the log producer price index for sector j in year t. Sector fixed effects, α_j, absorb time-invariant unobserved heterogeneity across industries, including persistent differences in production technologies, average price levels, and structural factor intensities. Year fixed effects, τ_t, absorb aggregate shocks common to all sectors, such as economy-wide inflation, monetary policy, and macroeconomic demand conditions.

The coefficients of interest are the dynamic exposure coefficients, β_k. Each β_k measures the year-specific price differential associated with a one-unit higher exposure to generative AI, relative to the omitted reference year b. The identifying variation therefore comes from whether sectors with higher predetermined exposure to generative AI experience different price dynamics in a given year relative to less exposed sectors, after removing permanent sectoral differences and common annual shocks.

The vector $\mathbf{SC}_{j,t}$ controls for differential exposure to supply-chain stress. Year fixed effects absorb aggregate price shocks that are common across sectors, but they do not absorb the possibility that a common shock affects sectors differently depending on their input structure. This distinction is central in our setting because the pandemic-era supply-chain disruption raised costs more sharply in sectors that depended on physical goods, logistics services, and imported inputs (Young et al., 2021). Without accounting for this heterogeneity, the event-study coefficients could partly capture differential exposure to the COVID-era supply-chain shock rather than the dynamic relationship between generative AI exposure and sectoral prices.

We therefore construct supply-chain controls with a shift-share structure. The time-series component is the New York Fed's Global Supply Chain Pressure Index (GSCPI), aggregated to the annual level. The sectoral exposure components are predetermined input-output measures constructed from the 2017 BEA detailed tables. The first measure captures the share of gross output spent on goods-producing and transportation/warehousing inputs. The second captures imported-input intensity, measured as imported intermediate inputs divided by total intermediate input use. These two exposures capture complementary dimensions of vulnerability to supply-chain disruptions: reliance on physical and logistics-intensive inputs, and reliance on foreign intermediate inputs. Details on the construction of these measures are provided in Appendix A.

The baseline specification allows supply-chain pressure to affect sectoral prices flexibly in two ways. First, we include both the contemporaneous GSCPI interaction and its one-year lag, allowing supply-chain disruptions to pass through to producer prices with delay. Second, we allow the pass-through coefficients to differ between normal years and high-stress years. We define high-stress years as years in which the annual mean GSCPI exceeds one standard deviation above its historical average; in our sample, this rule identifies 2020, 2021, and 2022 (see Table A.4). This flexible structure allows the pass-through of supply-chain pressure to differ between ordinary fluctuations in supply-chain conditions and the exceptional disruption of the pandemic period.

We use 2019 as the omitted reference year in the baseline specification. Although the public release of ChatGPT in late 2022 makes this year a natural conceptual benchmark, it is a less suitable normalization year for prices because it lies at the peak of the pandemic-era inflation and supply-chain episode.[3] Using 2019 instead anchors the event study in the last full pre-pandemic year and provides a stable reference point before both the COVID-19 dis-

[3] In 2022, the annual mean GSCPI stood 2.21 standard deviations above its historical average and reached a maximum monthly value of 3.76. Annual average CPI-U inflation was 8.0 percent in 2022, the highest value in our sample period.

ruption and the commercialization of generative AI. This normalization choice does not relax the identifying requirement: the pre-period coefficients remain informative about whether more and less AI-exposed sectors were on similar price trajectories before the diffusion of generative AI, including through the pandemic window. We report the corresponding event study normalized to 2022 in Appendix D.

Following standard practice in continuous-treatment event-study designs (Autor et al., 2006), the identifying assumption is that, absent the diffusion of generative AI, sectors with different exposure levels would not have experienced systematically different price dynamics after the reference period. The coefficients prior to the post-2022 diffusion period therefore provide a direct diagnostic for pre-trends. Standard errors are clustered at the NAICS industry level to allow for arbitrary within-sector autocorrelation over time.

Figure 4 presents the dynamic exposure coefficients from the baseline event-study specification. The coefficients are normalized relative to 2019, the omitted reference year, so each point measures the differential log price change associated with higher AI exposure relative to that benchmark. The figure reports both 90 percent and 95 percent confidence intervals. The 2026 coefficient is based on data available through April 2026 and should be read as suggestive.

The figure shows little evidence of systematic pre-trends in the pre-pandemic years. The coefficients from 2012 through 2018 are small in magnitude, alternate in sign, and are statistically indistinguishable from zero. This pattern suggests that, prior to the pandemic and the commercialization of generative AI, sectors with higher and lower AI exposure were not already following systematically different price trajectories.

The pandemic window is naturally a special period in this design. The coefficients for 2020–2022 become modestly negative, but they remain statistically insignificant and do not display a clear dynamic pattern: the estimate moves down in 2020, closer to zero in 2021, and down again in 2022. These are precisely the years in which differential exposure to supply-chain stress is most likely to matter, and the supply-chain controls are designed to absorb this source of heterogeneity. After accounting for these controls, the remaining pandemic-window estimates do not point to a clear or statistically reliable dynamic break by AI exposure before the post-2022 diffusion period.

After 2022, the estimates begin to display a distinct dynamic pattern. Starting in 2023, the coefficients move monotonically downward, becoming larger in magnitude and more clearly separated from zero over time. The 2023 estimate is negative but imprecise, consistent with the fact that generative AI adoption was still at an early stage. By 2024 and 2025, however, the estimates become progressively more negative and statistically distinguishable

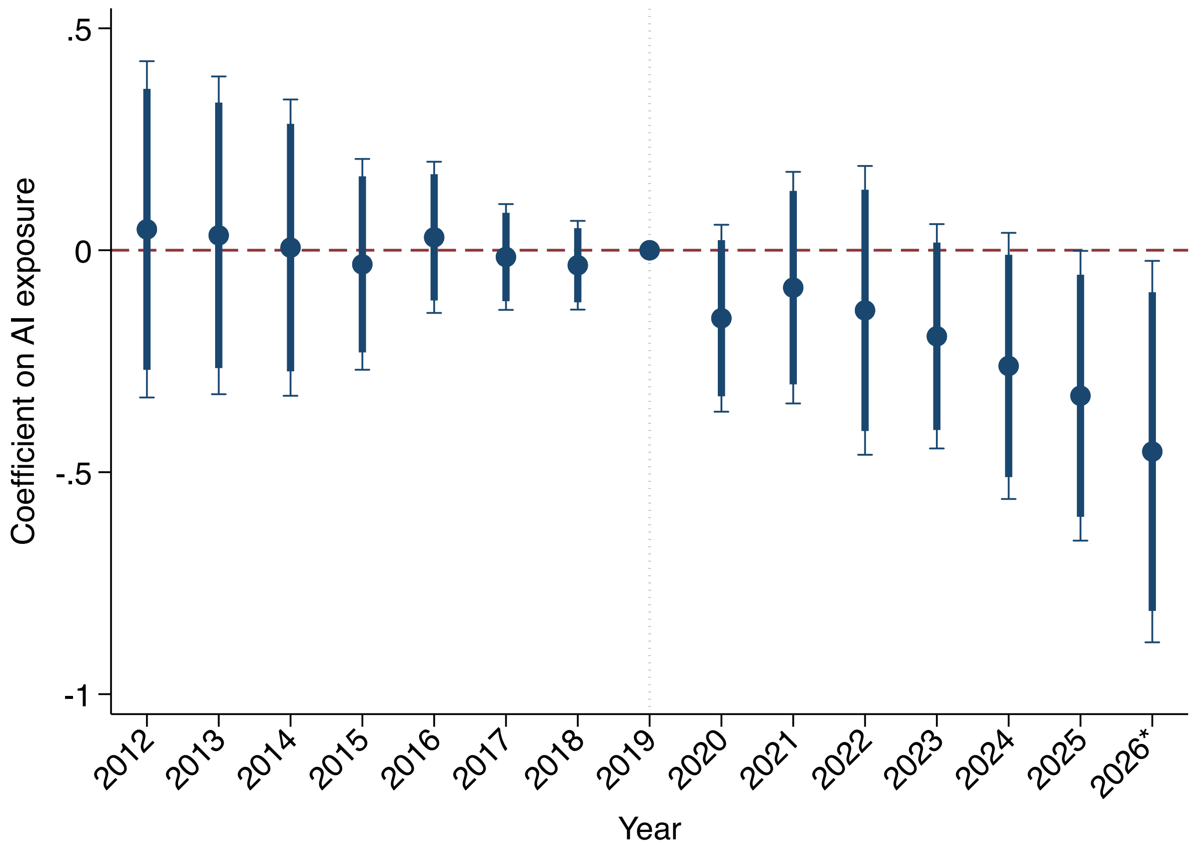

Figure 4: Dynamic Effect of Generative AI Exposure on Sectoral Prices

Notes: The figure plots the dynamic exposure coefficients β_k from Equation 7. Each coefficient measures the year-specific difference in log producer prices associated with a one-unit higher Generative AI exposure index, relative to the omitted reference year 2019, which is normalized to zero. The specification includes sector fixed effects, year fixed effects, and supply-chain controls based on GSCPI interactions with predetermined input-output exposure measures. The inner thick and outer thin vertical lines represent 90% and 95% confidence intervals, respectively. Standard errors are clustered at the NAICS industry level. The 2026 coefficient is based on data available through April 2026 and is shown as suggestive. Full estimation results are reported in Appendix D.

from zero. The partial-year 2026 coefficient continues this downward profile, although it should be read as suggestive. This gradual widening of the relative price gap is consistent with a diffusion process in which AI-related efficiency gains accumulate over time and are progressively reflected in sectoral producer prices.

This timing is consistent with a gradual diffusion mechanism rather than an immediate discrete price break. The adoption patterns documented in Figure 1 show that generative AI adoption remains limited and still expanding. Its effect on prices should therefore operate through firm-level experimentation, integration into business processes, task reorganization, cost reductions, and eventual pass-through to producer prices. Under this mechanism, one should not necessarily expect a large and precisely estimated effect immediately after the public release of generative AI. The more relevant empirical signature is a progressive widening of the relative price gap as diffusion deepens. The post-2022 trajectory in Figure 4 follows this pattern: the price differential emerges gradually, becomes more pronounced over time, and is strongest in the most recent years of the sample.

This dynamic profile also helps discipline alternative interpretations. The most immediate concern is that the estimates reflect residual pandemic-era supply-chain forces rather than AI-related price effects. This explanation is limited by the timing of the estimates. Supply-chain pressures were concentrated in 2020–2022 and had largely normalized by 2023, as shown in Table A.4, while the estimates do not display a clear or statistically reliable break during the high-stress window. Instead, the relative price gap widens precisely after the period in which supply-chain disruptions were most severe. A purely supply-chain-based explanation would therefore more naturally predict an early divergence during the bottleneck episode, followed by stabilization or reversal as those bottlenecks eased, rather than a progressively widening gap after 2023.

More general unobserved shocks face a similar timing restriction. To account for the estimates, such a shock would have to be differentially correlated with AI exposure, generate no systematic price divergence during the pre-pandemic decade, generate no clear break during the pandemic window, and then induce an increasingly negative relative-price trajectory precisely after the commercialization of generative AI. This does not eliminate all endogeneity concerns, as no macroeconomic empirical design can do so completely. But it substantially narrows the set of plausible alternatives: the pattern is difficult to reconcile with a transitory aggregate shock, a stable sectoral trend, or a one-time pandemic scarring effect, and is more naturally aligned with a gradual diffusion process.

The magnitude of the estimates is economically meaningful. The interquartile range of the 3-digit NAICS AI exposure index runs from 0.273 to 0.402, a difference of 0.129 (Table A.1). Applying this difference to the 2025 event-study coefficient implies that moving from the 25th to the 75th percentile of AI exposure is associated with approximately 4.2 log points lower producer prices relative to 2019. This magnitude is closely aligned with the cross-sectional gap exercise in Section 4.1, where the same interquartile movement in AI exposure is associated with a 4.2 percentage point reduction in cumulative excess inflation by 2025.

Appendix D reports several robustness checks for this baseline specification. Table D.1 provides the full set of coefficients underlying Figure 4. Figure D.1 re-estimates the baseline specification after excluding the partial 2026 observation; the results are virtually unchanged.

Figure D.2 and Table D.2 repeat the event study using 2022 as the omitted reference year. As expected given the exceptional inflation and supply-chain conditions in 2022, this normalization produces noisier estimates, but it preserves a similar post-2022 downward trajectory. Finally, the results are similar when the supply-chain controls are constructed using alternative annual aggregations of the GSCPI, including end-of-year and median values. Together, these checks indicate that the dynamic price pattern is not driven by the inclusion of the partial 2026 observation, the choice of normalization year, or the particular annual aggregation of supply-chain pressure.

5 Theoretical Framework

The empirical analysis documents two related product-market patterns following the introduction of generative AI. More exposed sectors experience stronger firm-entry growth and lower relative price growth. These outcomes need not be viewed as separate or contradictory forces. Both can arise from the same underlying mechanism: a technology that improves cognitive labor efficiency and reduces the labor requirements associated with administrative, organizational, and information-processing tasks.

This section develops a simple theoretical framework to organize and interpret this mechanism. We develop a static model of monopolistic competition with heterogeneous firms, building on Melitz, 2003. The model is designed as an organizing framework rather than a quantitative account of the transition path observed in the data. Its purpose is to show, in a parsimonious setting with closed-form comparative statics, how an improvement in cognitive labor efficiency can simultaneously increase the equilibrium mass of firms and lower the aggregate price index.

We model generative AI as a reduced-form improvement in cognitive labor efficiency. In the model, this improvement lowers marginal production costs and also reduces the fixed overhead costs associated with running or starting a business. This formulation captures the idea that generative AI can economize on tasks that enter both variable production costs and overhead costs, including information processing, documentation, coordination, customer communication, and other administrative activities. The exact cost structure is intentionally stylized: the goal is not to describe every channel of AI adoption, but to isolate a common labor-saving force that operates on both the production and entry margins.

The framework deliberately abstracts from several important dimensions. It does not model gradual technology diffusion, heterogeneous firm-level adoption, labor-market reallocation, or transition dynamics. The empirical estimates in Sections 3 and 4 should therefore be interpreted as evidence from an early stage of adjustment, while the model compares equilibria before and after the integration of the technology. This distinction is useful: the model clarifies the product-market forces that can jointly generate higher entry and lower prices, while the empirical analysis traces the early emergence of these forces in the data.

The comparative statics are straightforward. Lower fixed and entry-related overhead costs support a larger equilibrium mass of operating firms. Lower marginal costs reduce firm-level prices. In a CES environment, the expansion in product variety further lowers the aggregate price index through the standard love-for-variety channel. Thus, the model provides a conceptual bridge between the two empirical margins studied in the paper: generative AI can raise firm entry and reduce prices as joint consequences of a common cost-reducing technological force. Within the model, these same forces also raise welfare through lower prices and a larger set of available varieties. The main text presents the core intuition and equilibrium results; Appendix E provides the formal derivations.

5.1 Environment and Pre-AI Equilibrium

Consider an economy populated by a representative household that supplies L units of labor inelastically. Labor is the sole factor of production, and we set the nominal wage as the numéraire, $W = 1$.[4] The household allocates income across a continuum of available varieties Ω, whose endogenous total mass is N. Preferences take the standard Constant Elasticity of Substitution (CES) form:

$$U = \left(\int_{i \in \Omega} q_i^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{8}$$

where $\varepsilon > 1$ is the elasticity of substitution across varieties.

On the production side, there is an unbounded pool of prospective entrants. To produce quantity q_i, a firm requires a fixed operating cost f and a constant marginal cost $c_i = 1/A_i$, both measured in labor units. The term A_i denotes the firm's idiosyncratic labor productivity. Facing CES demand, a firm with productivity A_i sets a constant markup over marginal cost. This yields net operating profits that are strictly increasing in productivity:

$$\pi(A_i) = BA_i^{\varepsilon-1} - f, \tag{9}$$

where B is an endogenous aggregate demand shifter common to all firms.

[4] By normalizing the nominal wage and assuming inelastic labor supply, we abstract from nominal labor-market frictions and employment displacement dynamics. This abstraction allows us to focus on the transmission of the technological change through product markets, via firm entry and pricing.

Firms do not know their productivity prior to entry. To discover A_i, a prospective entrant must pay an upfront sunk entry cost f_e, capturing the labor requirements associated with initial setup, organization, and other entry-related overhead. Upon paying this cost, the firm draws productivity from a common Pareto distribution,

$$G(A) = 1 - \left(\frac{A_{\min}}{A}\right)^k, \tag{10}$$

with shape parameter $k > \varepsilon - 1$. A firm operates only if its net operating profit is nonnegative, which defines a zero-cutoff-profit survival threshold A^*. Free entry requires the ex-ante expected value of taking a productivity draw to equal the sunk entry cost f_e.

Solving the zero-cutoff-profit and free-entry conditions under the Pareto distribution yields the pre-AI equilibrium mass of operating firms:

$$N_{pre} = \frac{L}{\varepsilon f}\left(\frac{k - (\varepsilon - 1)}{k}\right). \tag{11}$$

5.2 The Macroeconomic Effects of Generative AI Adoption

We model the integration of generative AI as a reduced-form improvement in cognitive labor efficiency. Let $\gamma > 0$ denote the efficiency gain associated with the technology. Because labor is the sole factor of production in the model, we represent this improvement as a reduction in the labor requirements associated with both variable production activities and overhead tasks. Specifically, effective productivity for firm i rises from A_i to $A_i(1 + \gamma)$, lowering marginal cost. At the same time, fixed operating costs and entry-related overhead costs are scaled down to $f/(1 + \gamma)$ and $f_e/(1 + \gamma)$, respectively.

The symmetric scaling of marginal, fixed, and entry-related costs is a tractability assumption. It is not meant to imply that all real-world cost categories respond identically to generative AI, nor is the symmetry essential for the qualitative predictions of the model. Rather, it provides a parsimonious way to capture a common labor-saving force: the technology economizes on cognitive and administrative tasks that matter both for production and for the operation and creation of firms. This formulation keeps the selection margin transparent while isolating the product-market implications of the technology.

The propositions below summarize the main comparative statics; Appendix E provides the formal derivations.

Proposition 1 (The Firm Entry Effect): *A labor-saving cognitive technology, $\gamma > 0$, increases the equilibrium mass of operating firms.*

Intuition. In the pre-AI equilibrium, the mass of operating firms is decreasing in the fixed operating cost f. After AI integration, this cost falls to $f_{post} = f/(1+\gamma)$. Substituting this lower fixed operating cost into the equilibrium firm-mass expression yields

$$N_{post} = \frac{L}{\varepsilon\left(\frac{f}{1+\gamma}\right)}\left(\frac{k-(\varepsilon-1)}{k}\right) = N_{pre}(1+\gamma). \tag{12}$$

Thus, the equilibrium mass of operating firms increases proportionally with the cognitive labor-efficiency gain. The reduction in entry-related overhead costs preserves the ratio between fixed operating and sunk entry costs, leaving the productivity cutoff unchanged in this symmetric benchmark. The entry effect therefore operates through the lower fixed operating requirement: with fewer labor resources required to operate each firm, the economy can sustain a larger mass of active varieties. This result provides a simple mechanism consistent with the firm-entry pattern documented in Section 3.

Proposition 2 (The Disinflationary Effect): *A labor-saving cognitive technology, $\gamma > 0$, lowers the aggregate CES price index through lower marginal costs and a larger mass of available varieties.*

Intuition. The post-adoption price index reflects two changes. First, effective productivity rises from $\tilde{A}_{pre}$ to $\tilde{A}_{post} = \tilde{A}_{pre}(1+\gamma)$, reducing firm-level marginal costs and prices. Second, Proposition 1 implies that the mass of operating firms rises from N_{pre} to $N_{post} = N_{pre}(1+\gamma)$. Substituting these two terms into the aggregate CES price index gives

$$P_{post} = P_{pre}(1+\gamma)^{-\frac{\varepsilon}{\varepsilon-1}}. \tag{13}$$

Since $\varepsilon > 1$, the price index is strictly decreasing in γ. The expression separates two forces. The term $(1+\gamma)^{-1}$ captures the direct marginal-cost reduction. The term $(1+\gamma)^{-1/(\varepsilon-1)}$ captures the standard love-for-variety effect associated with the larger mass of firms.

This result should not be interpreted as implying that the early empirical price response is driven primarily by new-firm entry. The model abstracts from diffusion dynamics and compares equilibria before and after technology integration. Its purpose is instead to show that higher entry and lower prices can be joint consequences of the same underlying cost-reducing force. This mechanism is consistent with the relative price dynamics documented in Section 4, where more exposed sectors exhibit lower relative price growth as generative

AI begins to diffuse.

Corollary (The Welfare Dividend). *Generative AI integration raises real welfare in the benchmark economy.* With inelastic labor supply and the wage normalized to one, household nominal income is fixed at L, and indirect utility is proportional to real income, $U = L/P$. Since generative AI lowers the aggregate price index, real welfare rises. The welfare gain reflects the same two product-market forces highlighted above: lower prices from reduced marginal costs and a larger set of available varieties (the standard "love for variety" effect). This implication should be read within the scope of the model, which abstracts from labor-market displacement, transition costs, and distributional effects.

6 Conclusion

This paper studies the product-market effects of generative AI. Much of the early economic discussion of artificial intelligence has focused on labor-market exposure, worker displacement, and changes in the demand for cognitive skills. These are central issues. The evidence in this paper points to a complementary margin of adjustment: generative AI also affects how product markets evolve. We find that more exposed sectors adjust along two dimensions after the introduction of the technology: they experience stronger firm-entry growth and lower relative price growth.

We document these patterns using sector-level measures of generative AI exposure constructed from occupational task evaluations and pre-release employment and wage data. On the entry margin, a dynamic event-study design shows that more exposed 3-digit NAICS sectors experience stronger growth in business applications after 2022, with no systematic evidence of faster pre-existing entry growth. Complementary evidence from the standard 2-digit Business Formation Statistics shows that this relationship appears for both High-Propensity and Low-Propensity Applications.

On the price margin, a dynamic event-study design shows that more exposed sectors experience lower relative price growth after 2023. This relative price difference becomes more pronounced over time, consistent with a diffusion process in which AI-related efficiency gains gradually accumulate and pass through to sectoral prices. The estimates show no systematic evidence of differential pre-existing trends. The specification includes controls for differential exposure to pandemic-era supply-chain stress, using predetermined input-output exposure measures interacted with the Global Supply Chain Pressure Index. The same pattern is also visible in a cross-sectional inflation-gap exercise.

The theoretical framework provides a simple way to organize these findings. In a static model of monopolistic competition with heterogeneous firms, generative AI is represented as a reduced-form improvement in cognitive labor efficiency. This improvement lowers marginal production costs and reduces fixed operating and entry-related overhead. Lower overhead supports a larger equilibrium mass of firms, while lower marginal costs reduce firm-level prices. In a CES environment, the expansion in product variety further lowers the aggregate price index through the standard love-for-variety channel. Within the benchmark economy, these product-market forces generate a consumer welfare dividend: lower prices raise purchasing power, while a larger mass of firms expands the set of available varieties.

The evidence should be interpreted as an early evaluation of an ongoing technological transition. Current measures of enterprise AI adoption indicate that the U.S. economy remains in the early stages of integration. This is important because general-purpose technologies typically diffuse through a gradual process of organizational restructuring, complementary investment, and intangible capital accumulation. Even if generative AI diffuses more rapidly than earlier technologies, its aggregate effects are likely to unfold over time as firms experiment with the technology, reorganize tasks, and integrate AI-enabled tools into production and administrative processes. The product-market responses documented in this paper may therefore capture only the initial phase of a broader adjustment. As adoption deepens, the entry and price margins studied here may become more important.

These findings have implications for both policy and future research. A fundamental value of an early analysis lies in anticipating macro-trends before they fully mature, enabling the design of forward-looking policies and institutional frameworks that are better equipped to respond to the technological transition. For policymakers, the results suggest that generative AI should be tracked not only as a labor-market shock, but also as a force shaping business dynamism and inflation dynamics. For researchers, the findings point to the importance of studying AI as a broad economic technology whose effects operate across multiple markets. Tracing how these different forces co-evolve along the diffusion path of the technology will be essential for understanding the macroeconomic trajectory of the AI era.

References

Acemoglu, D. (2025). "The Simple Macroeconomics of AI". *Economic Policy* 40.121, pp. 13–58.

Acemoglu, D. and P. Restrepo (2018). "Artificial intelligence, automation, and work". *The economics of artificial intelligence: An agenda*, pp. 197–236.

– (2020). "Robots and jobs: Evidence from US labor markets". *Journal of Political Economy* 128.6, pp. 2188–2244.

Autor, D. H. (2013). "The labor market impacts of technological change: From unmeasured quality change to meaningful economic growth". *Annals of Economics and Statistics*, pp. 4–9.

Autor, D. H., J. J. Donohue III, and S. J. Schwab (2006). "The costs of wrongful-discharge laws". *The Review of Economics and Statistics* 88.2, pp. 211–231.

Autor, D. H. (2024). *Applying AI to Rebuild Middle Class Jobs*. Working Paper 32140. National Bureau of Economic Research.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). "Artificial Intelligence, Firm Growth, and Product Innovation". *Journal of Financial Economics* 151, p. 103745.

Bayard, K., E. Dinlersoz, T. Dunne, J. Haltiwanger, J. Miranda, and J. Stevens (2018). *Early-Stage Business Formation: An Analysis of Applications for Employer Identification Numbers*. Working Paper 24364. National Bureau of Economic Research.

Brynjolfsson, E. (1993). "The productivity paradox of information technology". *Communications of the ACM* 36.12, pp. 66–77.

Brynjolfsson, E., B. Chandar, and R. Chen (Nov. 2025a). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Working Paper. Stanford Digital Economy Lab.

Brynjolfsson, E., D. Li, and L. R. Raymond (2025b). "Generative AI at Work". *The Quarterly Journal of Economics* 140.2, pp. 889–942.

Brynjolfsson, E., D. Rock, and C. Syverson (2021). “The productivity J-curve: How intangibles complement general purpose technologies”. *American Economic Journal: Macroeconomics* 13.1, pp. 333–372.

David, P. A. (1990). “The dynamo and the computer: an historical perspective on the modern productivity paradox”. *The American Economic Review* 80.2, pp. 355–361.

Davis, S. J., J. C. Haltiwanger, and S. Schuh (1996). *Job Creation and Destruction*. Cambridge, MA: MIT Press.

Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2014). “The role of entrepreneurship in US job creation and economic dynamism”. *Journal of Economic Perspectives* 28.3, pp. 3–24.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). “GPTs are GPTs: Labor market impact potential of LLMs”. *Science,* 384.6702, pp. 1306–1308.

Federal Open Market Committee (2023). *Minutes of the Federal Open Market Committee, December 12–13, 2023*. Meeting Minutes. Board of Governors of the Federal Reserve System.

Federal Reserve Bank of New York (2023). *Global Supply Chain Pressure Index (GSCPI)*. `https://www.newyorkfed.org/research/policy/gscpi`. Data release and historical series.

Graetz, G. and G. Michaels (2018). “Robots at work”. *The Review of Economics and Statistics* 100.5, pp. 753–768.

Hopenhayn, H. A. (1992). “Entry, exit, and firm dynamics in long run equilibrium”. *Econometrica: Journal of the Econometric Society*, pp. 1127–1150.

Melitz, M. J. (2003). “The impact of trade on intra-industry reallocations and aggregate industry productivity”. *Econometrica* 71.6, pp. 1695–1725.

Mokyr, J. (1990). *The lever of riches: Technological creativity and economic progress*. Oxford University Press.

Oliner, S. D., D. E. Sichel, and K. J. Stiroh (2007). “Explaining a productive decade”. *Brookings Papers on Economic Activity* 2007.1, pp. 81–137.

Young, H. L., A. Monken, F. Haberkorn, and E. Van Leemput (Dec. 2021). “Effects of Supply Chain Bottlenecks on Prices Using Textual Analysis”. *FEDS Notes*.

Appendices

A Data Construction and Sources

A.1. Sectoral AI Exposure Construction

The construction of the sectoral AI exposure index relies on merging occupational scores from Eloundou et al., 2024 with the BLS OEWS May 2022 dataset. The original O*NET-SOC occupational codes from the OpenAI dataset (typically 10 digits, e.g., 11-1011.00) are truncated to their first 7 digits to perfectly match the detailed occupational classification used by the BLS. In cases where multiple 10-digit O*NET occupations collapse into a single 7-digit BLS occupation, we take the unweighted mean of the exposure scores before merging.

From the BLS OEWS dataset, we restrict the sample to detailed occupations to avoid double counting nested hierarchies. We drop any occupational cells where total employment or mean annual wages are suppressed by the BLS for confidentiality reasons. The wage bill for each occupation-industry cell is calculated as the product of total employment and the mean annual wage.

Table A.1 reports the summary statistics of the resulting AI exposure index at both the 2-digit and 3-digit NAICS aggregation levels. The mean exposure is approximately 0.35 across sectors, indicating that, on average, 35% of the sectoral wage bill is allocated to occupational tasks highly exposed to generative AI. The significant standard deviation (approximately 0.10) confirms the presence of substantial cross-sectional variation in AI susceptibility across the U.S. economy, providing the necessary identifying variation for our analysis.

A.2. Aggregation of Business Applications (2-Digit NAICS)

Data on aggregate and compositional Business Applications are obtained from the Federal Reserve Economic Data (FRED) database, which republishes the U.S. Census Bureau's Business Formation Statistics. We use the seasonally adjusted monthly series for Total Business Applications (`BABANAICSSAUS`) and High-Propensity Business Applications (`BAHBANAICSSAUS`).

The 2-digit NAICS data provide 19 broad private-sector series. Consistent with Census Bureau definitions, Manufacturing (NAICS 31–33), Retail Trade (NAICS 44–45), and

Table A.1: Summary Statistics: Sectoral AI Exposure Index

	NAICS 2-Digit	NAICS 3-Digit
Observations	20	85
Mean	0.362	0.350
Std. Dev.	0.107	0.104
Min	0.174	0.172
25th Percentile	0.296	0.273
Median	0.342	0.336
75th Percentile	0.423	0.402
Max	0.541	0.583

Notes: The table reports unweighted cross-sectional summary statistics for the constructed generative AI exposure index (β) at the 2-digit and 3-digit NAICS levels. The index is a wage-bill-weighted average of occupational exposure scores within each sector.

Transportation and Warehousing (NAICS 48–49) are reported as combined sectors. We exclude Public Administration, resulting in a balanced panel of 19 private sectors. Monthly flows are summed within each calendar year to construct annual business applications, $BA_{j,t} = \sum_{m=1}^{12} BA_{j,m,t}$. Low-Propensity Applications are constructed as the residual difference between Total Business Applications and High-Propensity Applications, $LPA_{j,t} = BA_{j,t} - HPA_{j,t}$.

A.3. Construction and Filtering of 3-Digit NAICS Business Applications

To obtain a more granular measure of firm entry by industry, we use the Census Bureau's "Special Projects: Weekly BFS by Industry" data product. This series reports weekly Business Applications by 3-digit NAICS subsector, substantially expanding the cross-sectional variation available relative to the standard 2-digit sectoral series. The Census Bureau assigns NAICS codes to applications using an automated industry-coding program based on the business names and descriptions provided on IRS Form SS-4. This procedure provides high-frequency industry detail but also introduces classification noise. Some applications cannot be assigned to a specific 3-digit subsector and are instead reported as "Not Further Defined" (NFD) within their broader parent sector.

We transform the raw weekly counts into an annual panel by summing application flows across weeks within each calendar year. To limit the influence of classification noise, we apply a cross-sectional quality filter based on long-run classification stability. For each 2-digit parent group, we compute the average NFD share over the 2011–2024 period. Table

A.2 reports these historical NFD shares by parent sector. We exclude 3-digit subsectors belonging to parent groups whose historical NFD share exceeds 10 percent, which primarily affects Manufacturing. This filter reduces reliance on parent sectors with persistently high rates of unresolved 3-digit classification.

We restrict the final 3-digit analytical panel to the 2011–2024 period. Although 2025 observations are available in the most recent preliminary release, that vintage exhibits an unusually large increase in unresolved industry classifications. The aggregate share of unclassified and NFD applications rises from a historical average of approximately 6 percent to more than 11 percent in 2025, with NFD shares exceeding 30 percent in some large parent sectors, including Retail Trade. Because the weekly NAICS product is revised as industry classifications are refined, these preliminary 2025 classifications may not yet have stabilized to their historical accuracy. We therefore exclude 2025 from the 3-digit firm-entry analysis and rely on the historically revised data through 2024. After applying the classification filter and merging with the occupational AI exposure index, the final unbalanced panel contains 63 private 3-digit NAICS sectors.

Table A.2: Share of Unclassified (NFD) Applications by Parent Sector

	2011	2015	2019	2023	2024	Historical Avg.
31/32/33	9.11	9.66	11.21	13.84	11.18	11.11
42	5.41	5.94	7.66	12.96	11.60	8.37
51	7.74	6.35	7.81	8.25	9.03	8.14
44/45	9.21	6.77	7.06	7.72	6.33	7.40
48/49	5.23	5.29	5.35	6.66	6.02	5.39
21	2.59	2.05	4.41	6.53	7.41	4.91
52	3.42	3.85	4.23	3.94	2.26	3.76
71	3.19	2.37	2.98	4.32	3.61	3.16
62	3.16	2.85	2.98	3.93	2.79	3.02
23	2.64	2.31	2.50	3.31	2.98	2.65
72	1.09	0.33	0.35	7.52	7.28	2.50
53	1.84	1.41	1.04	6.56	5.69	2.31
81	0.03	0.07	1.44	3.64	3.68	1.05
56	0.03	0.06	0.18	0.13	0.18	0.12

Notes: This table reports the percentage of business applications categorized as "Not Further Defined" (NFD) relative to total applications within each broad NAICS parent group. The historical average is computed over the 2011–2024 period. To preserve space, only selected years are displayed. Parent sectors with a historical NFD share of zero, such as NAICS 11, 22, 54, 55, and 61, are omitted from the table.

A.4. Hierarchical Reconstruction of Sectoral Producer Prices

Of the 85 3-digit NAICS sectors present in our baseline AI exposure index, exact aggregate Producer Price Indices (PPI) are natively available in FRED for 48 sectors. For the remaining sectors, the BLS often publishes price indices only at finer levels of disaggregation (4-, 5-, or 6-digit NAICS) or restricts publication due to confidentiality and sampling limitations.

To prevent the systematic loss of these sectors, we implement a hierarchical reconstruction algorithm. For any missing 3-digit sector j, we identify the universe of its available sub-sectoral PPI series in the BLS master catalog. To avoid double-counting within nested hierarchies, we apply a strict survival filter: a sub-sectoral code is retained if and only if no broader parent code is available in the dataset.

Let Ω_j denote the set of structurally independent sub-sectors identified for the missing 3-digit sector j. We reconstruct the aggregate price index by computing the unweighted average of the monthly growth rates of the components in Ω_j:

$$\frac{\Delta P_{j,t}}{P_{j,t-1}} = \frac{1}{|\Omega_j|} \sum_{k \in \Omega_j} \frac{\Delta P_{k,t}}{P_{k,t-1}} \tag{A.1}$$

The reconstructed price index level $P_{j,t}$ is then recovered via cumulative multiplication. We favor an unweighted average over employment or value-added weighting schemes at this highly disaggregated level because dynamic revenue weights are rarely available at a monthly frequency for 6-digit NAICS sub-sectors, and static weights risk introducing mechanical measurement error into high-frequency price indices.

This procedure successfully recovers 17 additional 3-digit NAICS sectors, bringing our final analytical sample to 65 sectors. Table A.3 documents the exact composition of the reconstructed aggregates. The remaining 20 sectors (primarily non-market services, education, and specific public-adjacent sectors) are dropped from the panel due to a complete absence of reliable BLS producer price data.

A.5. Supply-Chain Stress and Sectoral Exposure Measures

Our dynamic price specification controls for differential exposure to supply-chain stress. The construction combines a time-series measure of aggregate supply-chain pressure with predetermined sectoral measures of exposure to goods, logistics, and imported inputs.

We measure aggregate supply-chain pressure using the Federal Reserve Bank of New

Table A.3: Composition of Reconstructed 3-Digit NAICS Price Indices

3-Digit NAICS	**Constituent Sub-Sector Codes Used for Aggregation**
113	1133
236	236211, 236221, 236222, 236223, 236224, 236400, 236500
238	23811X, 23816X, 2381MR, 23821X, 23822X
455	4551, 4552
458	45811
459	45940
518	51821
522	5221
531	53112, 53113, 5312, 53131, 531320
532	5321, 532412
541	5411, 541211, 5413, 5416, 54181
561	5613, 56151, 561612, 56172
562	5621
621	6211, 6212, 6215, 6216, 62199
623	6231, 62321
713	7131, 71391, 71394
811	8113

Notes: This table lists the 17 3-digit NAICS sectors for which aggregate producer price indices were reconstructed. The right column reports the BLS PCU sub-sector codes (Ω_j) that survived the hierarchical filter and were used to compute the monthly growth rates for the reconstructed 3-digit aggregate. Codes ending in "X" or "MR" follow standard BLS PCU nomenclature for specific industry groupings.

York's Global Supply Chain Pressure Index (GSCPI). The index is observed monthly and normalized in standard-deviation units relative to its historical distribution. For the annual sector-level panel, we aggregate the monthly index to the calendar-year level using the annual mean. In robustness exercises, we also use the December value of the index to capture end-of-year supply-chain conditions.

Table A.4 reports the behavior of the GSCPI across different windows. The index remains close to its historical average during the pre-pandemic period, rises sharply during 2020–2022, and then falls back toward normal values in 2023–2025.

Table A.4: Global Supply Chain Pressure Index Across Regimes

Period	Number of years	Annual mean	December value
2011–2019	9	-0.149	-0.087
2020–2022	3	2.314	2.505
2023	1	-0.615	-0.166
2024	1	-0.288	-0.270
2025	1	-0.003	0.553

Notes: The table reports summary statistics for the Federal Reserve Bank of New York's Global Supply Chain Pressure Index (GSCPI).

We construct sectoral exposure measures using the 2017 detailed BEA Supply-Use framework and the 2017 BEA Import Matrix. We use 2017 because detailed input-output tables are available only for Economic Census years, and 2017 is the latest detailed benchmark year that predates both the COVID-19 shock and the commercialization of generative AI.

Let $Input_{i,j}$ denote purchases of input category i by industry j, and let GO_j denote gross output. Our first exposure measure captures each sector's dependence on goods-producing and transportation inputs:

$$GoodsTransportShare_j = \frac{\sum_{i \in \mathcal{G} \cup \mathcal{L}} Input_{i,j}}{GO_j}, \tag{A.2}$$

where $\mathcal{G}$ includes goods-producing inputs—agriculture, mining, construction, and manufacturing—and $\mathcal{L}$ includes transportation and warehousing inputs. This measure captures the extent to which a sector's production process depends on physical inputs and logistics services that are plausibly exposed to supply-chain disruptions.

Our second exposure measure captures imported-input intensity. Using the BEA Import Matrix, we compute:

$$ImportShare_j = \frac{ImportedInputs_j}{TotalIntermediateInputs_j}. \tag{A.3}$$

This variable measures the share of a sector's intermediate input bundle that is sourced from imports, and therefore captures exposure to disruptions in global rather than purely domestic supply chains.

Because some BEA detailed industries are more granular than the 3-digit NAICS sectors in our price panel, while others do not map one-to-one into our target classification, we first construct the exposure measures at the detailed BEA industry level and then aggregate them to the 3-digit NAICS level using gross-output weights. The resulting exposure measures are fixed at their 2017 values throughout the analysis, ensuring that they are predetermined relative to both the pandemic shock and the post-2022 diffusion of generative AI. Table A.5 reports summary statistics for the two exposure measures used in the baseline event-study specification.

Table A.5: Summary Statistics: Supply-Chain Exposure Measures

	Mean	SD	P25	Median	P75	Min	Max
GoodsTransportShare	0.260	0.188	0.101	0.209	0.401	0.008	0.748
ImportShare	0.090	0.067	0.036	0.068	0.121	0.002	0.341

Notes: The table reports unweighted cross-sectional summary statistics across the 65 NAICS 3-digit industries in the producer-price panel. $GoodsTransportShare_j$ is the share of gross output spent on goods-producing and transportation/warehousing inputs, constructed from the 2017 detailed BEA Supply-Use framework. $ImportShare_j$ is imported intermediate inputs divided by total intermediate input use, constructed from the 2017 BEA Import Matrix. Both measures are fixed at their 2017 values.

B Firm Entry: Evidence from 2-Digit NAICS

This appendix reports complementary firm-entry evidence using the standard 2-digit NAICS Business Formation Statistics. The main firm-entry analysis in Section 3 uses the more granular 3-digit NAICS panel, which provides richer cross-sectional variation and supports a dynamic event-study design. The 2-digit series is nevertheless useful for two reasons: it corresponds to the official sectoral BFS product and it provides a decomposition of applications into High-Propensity and Low-Propensity categories. We therefore use it to assess whether the entry patterns documented in the main text are also visible in the broader official sectoral series and whether they extend across different types of business applications.

B.1 Empirical Strategy

This appendix analyzes whether the relationship between generative AI exposure and firm entry is also visible in the standard 2-digit NAICS Business Formation Statistics. The exercise is complementary to the main 3-digit analysis because the 2-digit panel contains only 19 private sectors, but it has the advantage of using the official sectoral BFS series and allowing us to distinguish between High-Propensity and Low-Propensity Applications.

The empirical challenge is that sectors differ in their baseline entry dynamics and that the COVID-19 pandemic generated a large level shift in business applications. A simple comparison of entry before and after the commercialization of generative AI would therefore mix three objects: sector-specific pre-pandemic trends, the pandemic-era increase in applications, and any subsequent differential change associated with AI exposure.

To summarize post-2023 entry dynamics in the 2-digit panel, we construct a measure of *excess firm entry.* We first estimate a log-linear trend for each sector j over the pre-pandemic period:

$$\ln(BA_{j,t}) = \alpha_j + \delta_j t + \varepsilon_{j,t}, \tag{B.1}$$

using either 2015–2019 or 2011–2019 as the estimation window. We then use the estimated trend growth rate, $\hat{\delta}_j$, to construct a counterfactual path after 2023. Rather than projecting forward from 2019, we anchor the counterfactual path at the observed 2023 level of business applications, $BA_{j,2023}$. This normalization absorbs the large pandemic-era level shift in applications and focuses the exercise on whether sectors grew faster or slower than their own historical trend after 2023.

We use 2023 as the normalization year because it is the first full year after the public release of ChatGPT and after the acute phase of the pandemic-era disruption. This choice is conservative in the sense that any entry response already present in 2023 is absorbed into the baseline. The resulting estimates should therefore be interpreted as measuring differential entry growth after 2023, rather than the full cumulative entry response since the initial release of generative AI tools.

Using this framework, the counterfactual expected entry level in 2025 is

$$\tilde{BA}_{j,2025} = BA_{j,2023} \times \exp(2\hat{\delta}_j). \tag{B.2}$$

We define the normalized *Excess Entry Gap* as the difference between observed and counterfactual applications in 2025, scaled by the 2023 baseline:

$$Excess\_BA_j = \frac{BA_{j,2025} - \tilde{BA}_{j,2025}}{BA_{j,2023}}. \tag{B.3}$$

This outcome measures whether applications in sector j grew more or less than predicted by its own pre-pandemic trend, after normalizing for the post-pandemic level of entry.

We include the sector's pre-pandemic labor share as a structural control. The motivation is that AI-related reductions in cognitive and administrative costs may matter more in sectors where labor inputs account for a larger share of value added, while physical capital requirements may remain more important barriers to entry in capital-intensive sectors.

We estimate the following cross-sectional regression across the 19 private 2-digit NAICS sectors:

$$Excess\_BA_j = \beta_0 + \beta_1 AI\_Exposure_j + \gamma Labor\_Share_j + u_j, \tag{B.4}$$

where $AI\_Exposure_j$ is the predetermined occupational exposure index. The model is estimated using heteroskedasticity-robust standard errors. We complement the baseline estimates with several robustness checks, including alternative pre-trend windows, quadratic trend specifications, direct growth-rate regressions controlling for pre-pandemic growth, and leave-one-out sensitivity exercises.

This 2-digit exercise should be interpreted as complementary evidence rather than as the main source of identification. The specification accounts for sector-specific pre-pandemic growth rates and the pandemic-era level shift in applications, but the small number of broad sectors limits the scope for dynamic tests and for absorbing all possible sector-specific shocks. The exercise therefore relies on the assumption that, after conditioning on historical growth and labor intensity, remaining post-2023 deviations in entry are not driven by unobserved shocks systematically correlated with AI exposure. Its value is primarily descriptive and compositional: it asks whether the entry patterns documented in the more granular 3-digit analysis are also visible in the standard official BFS sectoral series, and whether they appear across both high- and low-propensity applications.

B.2 Baseline Results

We begin with the cross-sectional relationship between generative AI exposure and excess firm entry in the 2-digit BFS panel. Figure B.1 plots the normalized excess entry gap in 2025 against the sectoral AI exposure index, using the 2015–2019 trend to construct the counterfactual. The relationship is positive: sectors with higher AI exposure tend to exhibit larger post-2023 deviations from their own historical entry trends. Because the outcome is normalized relative to the observed 2023 level, the figure should be interpreted as a comparison of relative entry growth across sectors rather than as a measure of the aggregate level of business formation.

Table B.1 reports the corresponding cross-sectional regressions. Column (1) uses the 2015–2019 pre-pandemic trend to construct the excess entry gap, while Column (2) uses the longer 2011–2019 trend. Both specifications include the sector's pre-pandemic labor share.

The estimated relationship between AI exposure and excess entry is positive in both specifications. In Column (1), the coefficient on generative AI exposure is 0.731 and is statistically significant at the 1 percent level. Interpreted over the interquartile range of 2-digit sectoral AI exposure, from 0.296 to 0.423, this estimate implies that moving from the 25th to the 75th percentile of exposure is associated with a 9.3 percentage point increase in excess firm entry relative to the 2023 baseline. The estimate is similar when the longer 2011–2019 pre-trend is used. The labor-share coefficient is also positive in both specifications, consistent with the idea that AI-related reductions in cognitive and administrative costs may

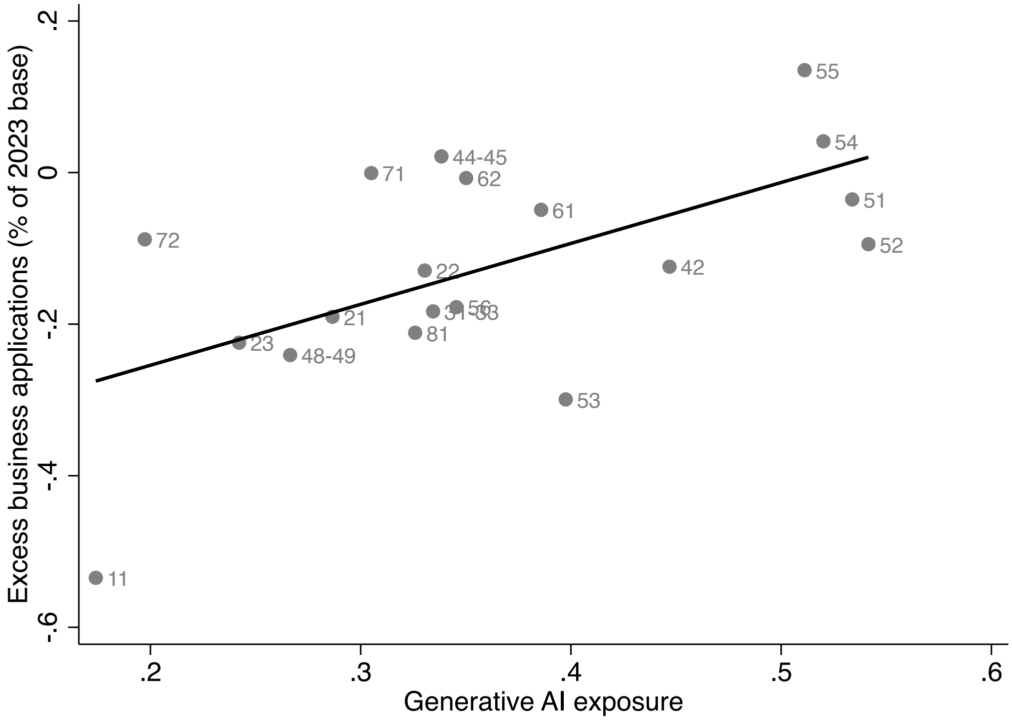

Figure B.1: Generative AI Exposure and Excess Firm Entry

Notes: The figure presents the cross-sectional scatter plot and linear fit between the generative AI occupational exposure index and the normalized *Excess Entry Gap* for 19 private 2-digit NAICS sectors. The excess gap is defined as the difference between observed 2025 business applications and the counterfactual projection based on the 2015–2019 trend, normalized by the 2023 level of business applications.

be more relevant in sectors where labor accounts for a larger share of value added.

The two-variable specification explains slightly more than half of the cross-sectional variation in excess entry dynamics, with an R^2 of 0.559 in the preferred specification and 0.531 when using the longer pre-trend. Given the small number of broad sectors, these results should be interpreted cautiously. The robustness checks in Section B.4 examine the sensitivity of the estimates to influential observations, alternative trend specifications, and direct growth-rate outcomes.

B.3 Compositional Heterogeneity: High vs. Low Propensity

The 2-digit BFS series allows us to examine whether the relationship between AI exposure and firm entry differs across types of business applications. This distinction is useful because High-Propensity Applications (HPA) are more likely to become employer businesses, while

Table B.1: Baseline Results: Generative AI Exposure and Firm Entry

	(1) Short Trend (2015–2019)	(2) Long Trend (2011–2019)
Generative AI Exposure (β)	0.731***	0.809***
	(0.243)	(0.254)
Labor Share (2016-2019)	0.326***	0.299**
	(0.089)	(0.106)
Observations	19	19
R^2	0.559	0.531

Notes: The dependent variable is the normalized *Excess Entry Gap* in 2025, computed using the 2015–2019 pre-trend in Column (1) and the 2011–2019 pre-trend in Column (2). Generative AI Exposure and Labor Share are measured as predetermined cross-sectional characteristics at the 2-digit NAICS level. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Low-Propensity Applications (LPA) capture applications less likely to transition into employer status, including independent contractors and zero-employee startups. We therefore repeat the baseline cross-sectional specification separately for the HPA and LPA components.

Table B.2: Heterogeneity Analysis: High vs. Low Propensity Firm Entry

	High-Propensity (HPA)		Low-Propensity (LPA)	
	(1) Short Trend	(2) Long Trend	(3) Short Trend	(4) Long Trend
Generative AI Exposure (β)	0.639***	0.670**	0.875***	0.954***
	(0.197)	(0.249)	(0.292)	(0.289)
Labor Share (2016-2019)	0.288**	0.266	0.317***	0.298**
	(0.103)	(0.153)	(0.102)	(0.105)
Observations	19	19	18	18
R^2	0.554	0.403	0.502	0.513

Notes: The dependent variable is the normalized *Excess Entry Gap* for High-Propensity Applications (Columns 1-2) and Low-Propensity Applications (Columns 3-4). The sample size for the LPA specifications is 18 due to a Census Bureau methodological feature that automatically categorizes all applications in the Accommodation and Food Services sector (NAICS 72) as High-Propensity. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.2 shows that the positive relationship between AI exposure and excess entry appears in both components of business applications. The coefficient on AI exposure is positive and statistically significant for High-Propensity Applications, indicating that the pattern is not confined to applications with a low ex-ante probability of becoming employer businesses.

The coefficient is also positive and statistically significant for Low-Propensity Applications, consistent with a response among independent contractors, zero-employee startups, and other small-scale entrepreneurial activity.

The LPA point estimates are larger than the HPA estimates, although the differences are not statistically distinguishable at conventional levels.[5] We therefore interpret the compositional evidence as suggesting a broad-based increase in applications across both margins rather than a response concentrated entirely in one type of business formation.

B.4 Robustness Checks

We report two sets of robustness checks for the 2-digit cross-sectional results. These exercises are intended to assess whether the positive relationship between AI exposure and excess entry is sensitive to influential sectors or to the construction of the excess-entry outcome.

First, we examine sensitivity to sample composition. Because the 2-digit BFS panel contains only 19 broad private sectors, individual observations could have a disproportionate effect on the estimates. Table B.3 reports leave-one-out re-estimations of the baseline specification, excluding one sector at a time. The coefficient on AI exposure remains remarkably stable and statistically significant at conventional levels. The table also reports a small set of dual-exclusion exercises that jointly remove pairs of knowledge-intensive sectors. The estimates remain stable in these specifications as well.

Second, we examine sensitivity to the construction of the dependent variable. Table B.4 reports specifications that replace the baseline log-linear excess-entry measure with alternative outcomes. Column (1) constructs the excess-entry gap using a quadratic pre-pandemic trend. Columns (2) and (3) replace the gap measure with the direct log growth in applications from 2023 to 2025, controlling for the sector's historical compound annual growth rate. Columns (4) and (5) use the log difference between average applications in 2022–2023 and average applications in 2024–2025. Across these alternatives, the coefficient on AI exposure remains large, positive, and statistically significant. These results suggest that the 2-digit relationship is not driven by a single influential sector or by the specific normalization used to construct the baseline excess-entry gap.

[5] Standard Wald tests fail to reject equality between the HPA and LPA coefficients.

Table B.3: Robustness: Leave-One-Out and Cluster Exclusion

Excluded Sector(s)	**AI Exposure Coef.** (β_1)	**p-value**
Panel A: Leave-One-Out (Single Sector)		
11 (Agriculture)	0.512**	0.013
21 (Mining)	0.753***	0.008
22 (Utilities)	0.739***	0.007
23 (Construction)	0.704**	0.020
31-33 (Manufacturing)	0.728**	0.010
42 (Wholesale Trade)	0.747***	0.009
44-45 (Retail Trade)	0.748***	0.007
48-49 (Transportation)	0.704**	0.017
51 (Information)	0.703**	0.020
52 (Finance)	0.834***	0.004
53 (Real Estate)	0.753***	0.008
54 (Professional Services)	0.729**	0.016
55 (Management)	0.698**	0.021
56 (Admin & Support)	0.717**	0.010
61 (Educational Services)	0.733***	0.009
62 (Health Care)	0.737***	0.009
71 (Arts & Entertainment)	0.777***	0.006
72 (Accommodation/Food)	0.857***	0.003
81 (Other Services)	0.705**	0.012
Panel B: Dual Exclusion (Knowledge-Intensive Clusters)		
51 & 54	0.694**	0.041
52 & 55	0.811**	0.013
53 & 56	0.736***	0.010

Notes: The table reports the point estimate and robust p-value for the generative AI exposure coefficient from re-estimations of the baseline 2-digit specification. Panel A excludes one NAICS sector at a time. Panel B jointly excludes selected pairs of knowledge-intensive sectors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.4: Robustness: Alternative Trends and Direct Growth Rates

	Quad. Gap	Growth (2023-2025)		Growth Avg. (22/23-24/25)	
	(1) Quad. Trend	(2) Short CAGR	(3) Long CAGR	(4) Short CAGR	(5) Long CAGR
Generative AI Exposure (β)	0.557**	0.918**	0.956**	0.590***	0.610***
	(0.221)	(0.349)	(0.347)	(0.169)	(0.175)
Observations	19	19	19	19	19
R^2	0.417	0.601	0.578	0.585	0.565

Notes: Column (1) estimates the excess-entry gap using a quadratic pre-pandemic trend. Columns (2)–(3) estimate the direct log difference in business applications from 2023 to 2025. Columns (4)–(5) estimate the log difference between average applications in 2022–2023 and average applications in 2024–2025. All specifications control for Labor Share and the corresponding pre-2020 historical growth measure. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Firm Entry: Dynamic Event-Study Estimates

Table C.1 reports the full set of dynamic exposure coefficients underlying Figure 3. The table corresponds to the baseline firm-entry event-study specification in Section 3.2, with 2022 as the omitted reference year and DHS growth in business applications as the dependent variable.

Table C.1: Dynamic Event Study: Firm Entry

	DHS growth in applications
$\beta \times \mathbb{I}\{2011\}$	-0.053
	(0.305)
$\beta \times \mathbb{I}\{2012\}$	0.204
	(0.160)
$\beta \times \mathbb{I}\{2013\}$	-0.319
	(0.215)
$\beta \times \mathbb{I}\{2014\}$	-0.087
	(0.200)
$\beta \times \mathbb{I}\{2015\}$	0.061
	(0.156)
$\beta \times \mathbb{I}\{2016\}$	-0.183
	(0.249)
$\beta \times \mathbb{I}\{2017\}$	0.272
	(0.248)
$\beta \times \mathbb{I}\{2018\}$	-0.100
	(0.209)
$\beta \times \mathbb{I}\{2019\}$	-0.072
	(0.155)
$\beta \times \mathbb{I}\{2020\}$	-0.160
	(0.374)
$\beta \times \mathbb{I}\{2021\}$	-0.252
	(0.181)
$\beta \times \mathbb{I}\{2022\}$ (Base)	0.000
	(.)
$\beta \times \mathbb{I}\{2023\}$	-0.115
	(0.231)
$\beta \times \mathbb{I}\{2024\}$	0.477^{***}
	(0.149)
Sector FE	Yes
Year FE	Yes
Reference year	2022
Number of sectors	63
Observations	882
Within R^2	0.252

Notes: The table reports the dynamic exposure coefficients β_k underlying Figure 3. The dependent variable is the DHS growth rate in business applications. The omitted reference year is 2022. The specification includes 3-digit NAICS fixed effects and year fixed effects. Standard errors, clustered at the 3-digit NAICS sector level, are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Price Dynamics

Table D.1 reports the full set of dynamic exposure coefficients underlying Figure 4. The table corresponds to the baseline event-study specification in Equation 7, with 2019 as the omitted reference year and the sample extended through the latest available 2026 producer-price observations.

As a conservative version of the baseline event study, Figure D.1 repeats the specification in Figure 4 after excluding the partial 2026 observation. The normalization remains 2019.

As an additional normalization check, Figure D.2 and Table D.2 repeat the baseline event-study specification using 2022 as the omitted reference year. This normalization anchors the coefficients around the public release of ChatGPT, but should be interpreted with the caveat that 2022 lies in the middle of the pandemic-era inflation and supply-chain episode discussed in Section 4.2.

As an additional robustness check, we also construct the supply-chain controls using alternative annual aggregations of the GSCPI, including the end-of-year and median values. The resulting event-study profiles are very similar to those reported here, and we omit them for brevity.

Table D.1: Dynamic Event Study: Baseline Specification

	$\ln(PPI)$
$\beta \times \mathbb{I}\{2012\}$	0.047
	(0.190)
$\beta \times \mathbb{I}\{2013\}$	0.034
	(0.179)
$\beta \times \mathbb{I}\{2014\}$	0.006
	(0.167)
$\beta \times \mathbb{I}\{2015\}$	-0.032
	(0.119)
$\beta \times \mathbb{I}\{2016\}$	0.029
	(0.085)
$\beta \times \mathbb{I}\{2017\}$	-0.015
	(0.060)
$\beta \times \mathbb{I}\{2018\}$	-0.034
	(0.050)
$\beta \times \mathbb{I}\{2019\}$ (Base)	0.000
	(.)
$\beta \times \mathbb{I}\{2020\}$	-0.153
	(0.105)
$\beta \times \mathbb{I}\{2021\}$	-0.084
	(0.131)
$\beta \times \mathbb{I}\{2022\}$	-0.135
	(0.163)
$\beta \times \mathbb{I}\{2023\}$	-0.194
	(0.126)
$\beta \times \mathbb{I}\{2024\}$	-0.261^{*}
	(0.150)
$\beta \times \mathbb{I}\{2025\}$	-0.328^{**}
	(0.163)
$\beta \times \mathbb{I}\{2026\}$	-0.454^{**}
	(0.215)
Sector FE	Yes
Year FE	Yes
Supply-chain controls	Yes
Reference year	2019
Number of sectors	65
Observations	921
Within R^2	0.747

Notes: The table reports the dynamic exposure coefficients β_k from Equation 7. The dependent variable is the log Producer Price Index. The omitted reference year is 2019. The specification includes sector fixed effects, year fixed effects, and supply-chain controls based on contemporaneous and lagged GSCPI interactions with predetermined input-output exposure measures. Standard errors, clustered at the NAICS industry level, are reported in parentheses. The 2026 coefficient is based on data available through April 2026 and should be interpreted as suggestive. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

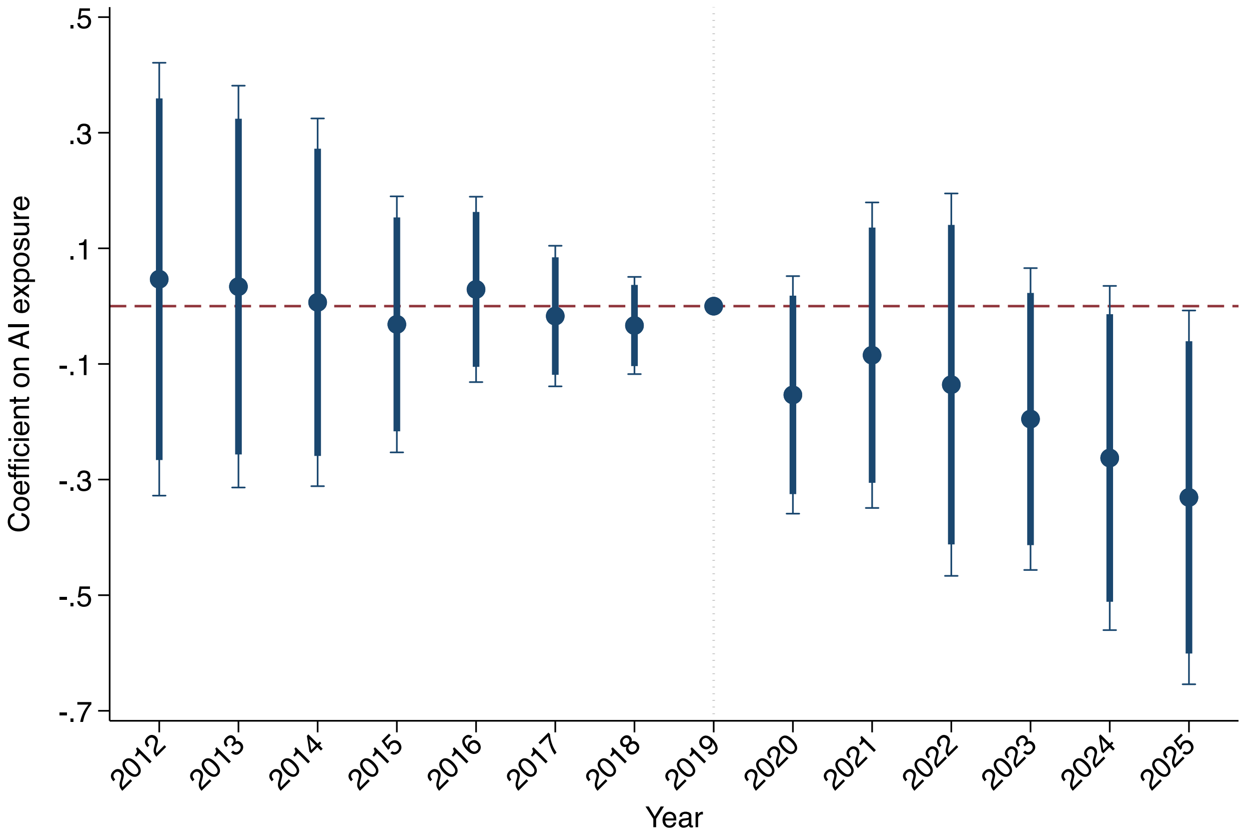

Figure D.1: Dynamic Effect of Generative AI Exposure on Sectoral Prices: Excluding 2026

Notes: The figure plots the dynamic exposure coefficients β_k from Equation 7, restricting the sample to exclude 2026. Each coefficient measures the year-specific difference in log producer prices associated with a one-unit higher Generative AI exposure index, relative to the omitted reference year 2019, which is normalized to zero. The specification includes sector fixed effects, year fixed effects, and supply-chain controls based on GSCPI interactions with predetermined input-output exposure measures. The inner thick and outer thin vertical lines represent 90% and 95% confidence intervals, respectively. Standard errors are clustered at the NAICS industry level.

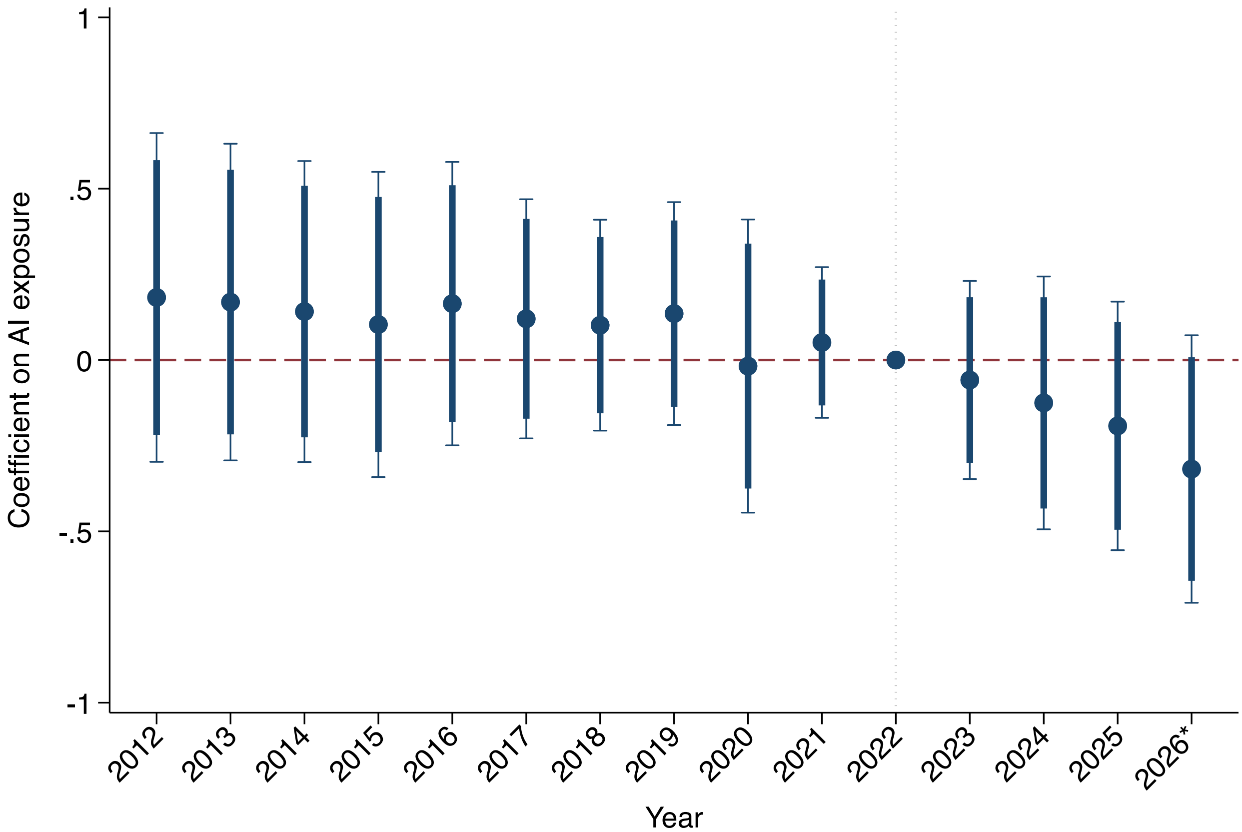

Figure D.2: Dynamic Effect of Generative AI Exposure on Sectoral Prices: 2022 Normalization

Notes: The figure plots the dynamic exposure coefficients β_k from Equation 7, using 2022 as the omitted reference year. Each coefficient measures the year-specific difference in log producer prices associated with a one-unit higher Generative AI exposure index, relative to 2022, which is normalized to zero. The specification includes sector fixed effects, year fixed effects, and supply-chain controls based on GSCPI interactions with predetermined input-output exposure measures. The inner thick and outer thin vertical lines represent 90% and 95% confidence intervals, respectively. Standard errors are clustered at the NAICS industry level. The 2026 coefficient is based on data available through April 2026 and is shown as suggestive.

Table D.2: Dynamic Event Study: 2022 Normalization

	$\ln(PPI)$
$\beta \times \mathbb{I}\{2012\}$	0.183
	(0.240)
$\beta \times \mathbb{I}\{2013\}$	0.169
	(0.231)
$\beta \times \mathbb{I}\{2014\}$	0.141
	(0.220)
$\beta \times \mathbb{I}\{2015\}$	0.104
	(0.223)
$\beta \times \mathbb{I}\{2016\}$	0.165
	(0.207)
$\beta \times \mathbb{I}\{2017\}$	0.120
	(0.175)
$\beta \times \mathbb{I}\{2018\}$	0.102
	(0.154)
$\beta \times \mathbb{I}\{2019\}$	0.135
	(0.163)
$\beta \times \mathbb{I}\{2020\}$	-0.018
	(0.214)
$\beta \times \mathbb{I}\{2021\}$	0.051
	(0.110)
$\beta \times \mathbb{I}\{2022\}$ (Base)	0.000
	(.)
$\beta \times \mathbb{I}\{2023\}$	-0.058
	(0.145)
$\beta \times \mathbb{I}\{2024\}$	-0.125
	(0.185)
$\beta \times \mathbb{I}\{2025\}$	-0.192
	(0.182)
$\beta \times \mathbb{I}\{2026\}$	-0.318
	(0.195)
Sector FE	Yes
Year FE	Yes
Supply-chain controls	Yes
Reference year	2022
Number of sectors	65
Observations	921
Within R^2	0.747

Notes: The table reports the dynamic exposure coefficients β_k from Equation 7. The dependent variable is the log Producer Price Index. The omitted reference year is 2022. The specification includes sector fixed effects, year fixed effects, and supply-chain controls based on contemporaneous and lagged GSCPI interactions with predetermined input-output exposure measures. Standard errors, clustered at the NAICS industry level, are reported in parentheses. The 2026 coefficient is based on data available through April 2026 and should be interpreted as suggestive. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Theoretical Appendix: Model Derivations

This appendix provides the formal derivations for the theoretical framework presented in Section 5. We first solve the firm's pricing problem, then characterize the zero-cutoff-profit and free-entry conditions under a Pareto productivity distribution, and finally derive the comparative statics associated with generative AI integration.

E.1 Firm Pricing and Relative Profits

The representative household allocates aggregate expenditure R across a continuum of varieties Ω. Solving the standard CES utility maximization problem yields the demand function for variety i:

$$q_i = \frac{R}{P}\left(\frac{p_i}{P}\right)^{-\varepsilon}, \tag{E.1}$$

where

$$P = \left(\int_{i\in\Omega} p_i^{1-\varepsilon}di\right)^{\frac{1}{1-\varepsilon}} \tag{E.2}$$

is the aggregate price index, and $\varepsilon > 1$ is the elasticity of substitution across varieties.

Firms produce using labor, which is supplied inelastically by households. Normalizing the wage rate to unity, $W = 1$, the cost function for a firm with idiosyncratic productivity A_i is

$$C(q_i) = f + \frac{q_i}{A_i}. \tag{E.3}$$

Facing CES demand, the firm chooses p_i to maximize net operating profits,

$$\pi_i = p_i q_i - \frac{q_i}{A_i} - f. \tag{E.4}$$

The first-order condition yields the standard constant-markup pricing rule:

$$p(A_i) = \left(\frac{\varepsilon}{\varepsilon-1}\right)\frac{1}{A_i}. \tag{E.5}$$

Substituting the pricing rule and the demand function into the profit equation, firm profits can be written as a function of the productivity draw:

$$\pi(A_i) = BA_i^{\varepsilon-1} - f, \tag{E.6}$$

where

$$B \equiv \frac{R}{\varepsilon}\left(P\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1} \tag{E.7}$$

is an endogenous aggregate demand shifter common to all firms.

E.2 The Zero-Cutoff Profit and Free-Entry Conditions

Firms operate only if their net operating profits are nonnegative. This defines the zero-cutoff-profit (ZCP) condition and the survival threshold A^*, corresponding to the productivity of the marginal active firm:

$$\pi(A^*) = 0 \quad \Longrightarrow \quad B(A^*)^{\varepsilon-1} = f. \tag{E.8}$$

Using the ZCP condition to substitute out the aggregate demand shifter B, the profit of any surviving firm, $A_i \geq A^*$, can be expressed relative to the marginal firm:

$$\pi(A_i) = f\left[\left(\frac{A_i}{A^*}\right)^{\varepsilon-1} - 1\right]. \tag{E.9}$$

Prior to entry, firms do not know their productivity. To take a productivity draw from the distribution $G(A)$, a prospective entrant must pay the sunk entry cost f_e. In a free-entry equilibrium, the ex-ante expected value of entry equals this sunk cost:

$$\int_{A^*}^{\infty} \pi(A)\, dG(A) = f_e. \tag{E.10}$$

Substituting the relative profit expression in Equation E.9 into the free-entry condition in Equation E.10 yields the equilibrium condition that determines the survival cutoff:

$$f\int_{A^*}^{\infty}\left[\left(\frac{A}{A^*}\right)^{\varepsilon-1} - 1\right] dG(A) = f_e. \tag{E.11}$$

E.3 Equilibrium under a Pareto Distribution

To obtain closed-form solutions for the equilibrium objects of interest, we parameterize ex-ante productivity draws using a Pareto distribution. The cumulative distribution function is

$$G(A) = 1 - \left(\frac{A_{\min}}{A}\right)^k \quad \text{for } A \geq A_{\min}, \tag{E.12}$$

where $A_{\min} > 0$ is the lower bound of the support and $k > \varepsilon - 1$ is the shape parameter governing the dispersion of productivities. A higher value of k corresponds to lower productivity dispersion.

Substituting the Pareto density,

$$dG(A) = kA_{\min}^k A^{-k-1} dA, \tag{E.13}$$

into the cutoff condition in Equation E.11 and evaluating the integral over $[A^*, \infty)$ yields

$$f\left[\frac{\varepsilon - 1}{k - (\varepsilon - 1)}\right]\left(\frac{A_{\min}}{A^*}\right)^k = f_e. \tag{E.14}$$

Rearranging gives the equilibrium survival threshold:

$$A^* = A_{\min}\left[\frac{f}{f_e}\left(\frac{\varepsilon - 1}{k - (\varepsilon - 1)}\right)\right]^{\frac{1}{k}}. \tag{E.15}$$

We next derive the equilibrium mass of operating firms. Free entry implies that aggregate profits net of entry costs are zero, so aggregate household expenditure equals total labor income. With the wage normalized to one and labor supply fixed at L, this gives $R = L$.

The average net operating profit of surviving firms, $\bar{\pi}$, is obtained by integrating the relative profit function in Equation E.9 over the truncated Pareto distribution for active firms, $A \geq A^*$:

$$\bar{\pi} = f\left[\frac{k}{k - (\varepsilon - 1)} - 1\right] = f\left(\frac{\varepsilon - 1}{k - (\varepsilon - 1)}\right). \tag{E.16}$$

From the firm's pricing problem, revenue is proportional to variable profits. Average firm revenue is therefore

$$\bar{r} = \varepsilon(\bar{\pi} + f). \tag{E.17}$$

Substituting Equation E.16 gives

$$\bar{r} = \varepsilon f\left(\frac{k}{k - (\varepsilon - 1)}\right). \tag{E.18}$$

The equilibrium mass of operating firms is total expenditure divided by average firm revenue, $N_{pre} = R/\bar{r}$. Using $R = L$, we obtain

$$N_{pre} = \frac{L}{\varepsilon f}\left(\frac{k-(\varepsilon-1)}{k}\right). \tag{E.19}$$

E.4 Generative AI Adoption

We model the integration of generative AI as a reduced-form improvement in cognitive labor efficiency. Let $\gamma > 0$ denote the efficiency gain associated with the technology. Because labor is the sole factor of production in the model, this improvement reduces the labor requirements associated with both variable production activities and overhead tasks.

For any firm i, effective productivity rises from A_i to $A_i(1+\gamma)$, reducing marginal cost to

$$c_{i,post} = \frac{1}{A_i(1+\gamma)}. \tag{E.20}$$

At the same time, fixed operating costs and entry-related overhead costs are scaled down to

$$f_{post} = \frac{f}{1+\gamma}, \qquad f_{e,post} = \frac{f_e}{1+\gamma}. \tag{E.21}$$

As discussed in the main text, the symmetric scaling is a tractability assumption that keeps the selection margin transparent while capturing a common labor-saving force operating on both the production and entry margins.

Proof of Proposition 1 (The Firm Entry Effect): To determine the post-adoption survival threshold, we evaluate the equilibrium cutoff in Equation E.15 under the new cost structure. Because fixed operating costs and sunk entry costs are scaled down symmetrically, their ratio remains unchanged:

$$\frac{f_{post}}{f_{e,post}} = \frac{f/(1+\gamma)}{f_e/(1+\gamma)} = \frac{f}{f_e}. \tag{E.22}$$

It follows that the survival threshold is unchanged relative to the pre-AI benchmark:

$$A^{*,post} = A^{*,pre}. \tag{E.23}$$

Substituting the lower fixed operating cost, $f_{post} = f/(1+\gamma)$, into the equilibrium firm-mass expression in Equation E.19 yields

$$N_{post} = \frac{L}{\varepsilon\left(\frac{f}{1+\gamma}\right)}\left(\frac{k-(\varepsilon-1)}{k}\right) = N_{pre}(1+\gamma). \tag{E.24}$$

Since $\gamma > 0$, $N_{post} > N_{pre}$. The reduction in fixed operating requirements allows the economy to sustain a larger equilibrium mass of operating firms. ■

Proof of Proposition 2 (The Disinflationary Effect): The aggregate CES price index can be written as a function of the equilibrium firm mass and the price charged by a firm with weighted average productivity $\tilde{A}$:

$$P = N^{-\frac{1}{\varepsilon-1}} p(\tilde{A}) = N^{-\frac{1}{\varepsilon-1}}\left(\frac{\varepsilon}{\varepsilon-1}\right)\frac{1}{\tilde{A}}. \tag{E.25}$$

Because the survival threshold is unchanged, the underlying distribution of surviving productivity draws is unchanged. However, effective productivity for every surviving firm is scaled by $(1+\gamma)$, so

$$\tilde{A}_{post} = \tilde{A}_{pre}(1+\gamma). \tag{E.26}$$

Using $N_{post} = N_{pre}(1+\gamma)$ and $\tilde{A}_{post} = \tilde{A}_{pre}(1+\gamma)$ in Equation E.25, we obtain

$$P_{post} = [N_{pre}(1+\gamma)]^{-\frac{1}{\varepsilon-1}}\left(\frac{\varepsilon}{\varepsilon-1}\right)\frac{1}{\tilde{A}_{pre}(1+\gamma)}. \tag{E.27}$$

Rearranging gives

$$P_{post} = \underbrace{\left[N_{pre}^{-\frac{1}{\varepsilon-1}}\left(\frac{\varepsilon}{\varepsilon-1}\right)\frac{1}{\tilde{A}_{pre}}\right]}_{} P_{pre}(1+\gamma)^{-\frac{1}{\varepsilon-1}}(1+\gamma)^{-1}. \tag{E.28}$$

Therefore,

$$P_{post} = P_{pre}(1+\gamma)^{-\frac{\varepsilon}{\varepsilon-1}}. \tag{E.29}$$

Since $\varepsilon > 1$ and $\gamma > 0$, the aggregate price index falls, $P_{post} < P_{pre}$. Equation E.28 separates the two forces behind this result: the term $(1+\gamma)^{-1}$ captures the direct marginal-cost reduction, while the term $(1+\gamma)^{-1/(\varepsilon-1)}$ captures the love-for-variety effect associated with the larger mass of operating firms.

Finally, with the wage normalized to one and labor supply fixed at L, nominal household income is $R = L$. Indirect utility is therefore proportional to real income, $U = L/P$. Since $P_{post} < P_{pre}$, real welfare rises in the benchmark economy. ■