

The Macroeconomics of Generative AI: Price Dynamics and Firm Entry

Jose Carreño*

University of Oregon

October 2026

First version: March 2026

Abstract

We study the effects of generative AI on U.S. price dynamics and firm entry. A dynamic continuous-treatment event study using predetermined sectoral exposure shows little evidence of systematic pre-existing differential price trends. From 2023 onward, however, more exposed sectors exhibit progressively lower relative producer prices. An interquartile exposure difference is associated with 4.2 log points lower cumulative price growth by 2025. Controlling for supply-chain stress and sectoral conditions preserves this pattern. More exposed sectors also exhibit stronger growth in business applications. A parsimonious model links these responses to efficiency gains, with welfare benefits from lower prices and greater variety.

JEL Classification: E31, L26, O33

Keywords: Generative AI, Producer Prices, Price Dynamics, Firm Entry, General-Purpose Technology

*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, generative AI tools have become increasingly available for business use. Figure 1 presents data from the U.S. Census Bureau’s Business Trends and Outlook Survey (BTOS), showing an upward trend in the share of firms reporting current AI use alongside a consistently higher share expecting to use AI over the next six months. With enterprise adoption still far from complete, this ongoing diffusion 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 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 complementary margins of adjustment in product markets: price dynamics and firm entry.




Figure 1: U.S. Firm AI Adoption: Current Use and Six-Month Expectations

Notes: The figure reports survey estimates every two weeks of the share of U.S. firms using artificial intelligence in producing goods or services during the previous two weeks, alongside the share expecting to use AI over 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](#)). By reducing the cost of activities such as information processing, documentation, software development, customer communication, marketing, and coordination, generative AI may reduce marginal production costs in sectors where exposed tasks account for a larger share of activity. 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 same capabilities may also reduce the fixed costs of starting and operating a business, encouraging entry in more exposed sectors. Lower prices and stronger entry can therefore arise as joint consequences of a common improvement in efficiency, without requiring either response to cause the other.

We first examine price dynamics by combining sectoral AI exposure measures constructed from occupational task evaluations and pre-release employment and wage data with an annual panel of producer price indices. We estimate a dynamic continuous-treatment event study that traces relative price trajectories and assesses pre-period trends, while controlling 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 little evidence of systematic differential trends before 2023. From 2023 onward, more exposed sectors experience lower relative producer-price growth, and the relative price gap widens over time. An interquartile difference in AI exposure is associated with approximately 4.2 log points lower cumulative producer-price growth between 2019 and 2025. Additional controls for sectoral supply and demand conditions and import competition leave the estimated trajectory largely unchanged. A complementary cross-sectional inflation-gap exercise yields a similar economic magnitude. These findings are consistent with AI-related efficiency gains gradually passing through to producer prices as the technology diffuses.

We complement the price analysis with evidence on firm entry, using granular U.S. Census Bureau data on business applications as a proxy. A dynamic event study shows little evidence of systematic pre-existing trends in the entry-growth differential associated with AI exposure. The 2024 coefficient is positive and statistically significant. For an interquartile difference in AI exposure, the annual Davis–Haltiwanger–Schuh (DHS) growth differential in business applications is approximately 6.2 percentage points higher in 2024 than in 2022. This result is robust to excluding individual industries and broader industry families, as well as to using an alternative reference year. Complementary evidence at a broader industry level extends through 2025 and shows that the positive relationship also appears among applications with a high propensity to become employer businesses.

To connect these findings and clarify their broader economic implications, we develop a parsimonious static model of monopolistic competition with heterogeneous firms, building on [Melitz \(2003\)](#). The model provides an organizing framework with closed-form comparative statics. In the model, generative AI improves cognitive labor efficiency, reducing marginal production costs and the fixed overhead costs associated with operating and starting a busi-

ness. Lower marginal costs reduce firm-level prices, while lower fixed operating requirements support a larger equilibrium mass of firms. The expansion in product variety further lowers the aggregate CES price index through the standard love-for-variety channel. Within this benchmark, which abstracts from labor-market displacement, transition costs, and distributional effects, lower prices and greater variety also raise real welfare.

The empirical 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 price and entry 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 effects of AI on productivity (Acemoglu and Restrepo, 2019; Eloundou et al., 2024; Brynjolfsson et al., 2025b; Acemoglu, 2025; Autor, 2024; Brynjolfsson et al., 2025a). We complement this literature by studying how exposure to generative AI is associated with changes in producer prices and firm entry. This product-market perspective highlights channels through which the economic effects of generative AI may extend beyond labor-market adjustment.

Second, 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.

Third, we contribute to the literature on firm dynamics and technological adoption (Hopenhayn, 1992; Decker et al., 2014). This literature emphasizes how fixed operating costs and entry-related requirements shape business formation and market structure. Recent empirical work shows that AI adoption is associated with growth and product innovation among incumbent firms (Babina et al., 2024). More directly related to the entry margin, Cai et al. (2025) document stronger increases in new-firm registrations following the release of ChatGPT in Chinese locations with pre-existing AI patenting activity. We complement this evidence by studying how business-application growth in the United States varies with pre-determined sectoral exposure to generative AI. Our exposure measure captures the potential for generative AI to reduce the labor requirements of occupational tasks, aggregated using industries’ pre-release employment and wage data. A dynamic event-study design traces changes in the entry-growth differential and assesses whether such differences were already present before the technology’s diffusion. We find stronger growth in business applications in more exposed sectors, with complementary evidence for both High-Propensity and Low-Propensity Applications. By studying this entry margin alongside producer-price dynamics, we connect both responses within a broader analysis of the product-market implications of generative AI.

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 design, results, and robustness checks for price dynamics. Section 4 studies firm entry. Section 5 develops the theoretical framework. Section 6 concludes. An in-paper appendix provides the model derivations. A separate Supplemental Appendix presents additional data details, empirical results, and robustness checks.

2 Data and Measurement

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 producer price indices and administrative data on business applications to study sectoral price dynamics and firm entry. Supplemental Appendix A provides detailed construction procedures and supplementary descriptive statistics.

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 whether large language models (LLMs), used directly or through complementary software, could reduce the time required to complete occupational tasks by at least 50 percent while maintaining equivalent quality. We focus on their β measure, which incorporates both direct LLM capabilities and exposure through complementary applications. To reduce reliance on any single annotation procedure, we define our occupational exposure score as the simple average of the human-annotated and GPT-4-generated evaluations provided in their dataset.

We construct the industry-level index at the 3-digit NAICS level using the May 2022 data from the Bureau of Labor Statistics (BLS) Occupational Employment and Wage Statistics (OEWS) program. Following the approach used 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:

$$AI_Exposure_j = \sum_{o \in j} \left(\frac{w_{o,j} E_{o,j}}{\sum_{o \in j} w_{o,j} E_{o,j}} \right) Exposure_Score_o, \quad (1)$$

where $E_{o,j}$ and $w_{o,j}$ denote employment and mean annual wages for occupation o in industry j , and $Exposure_Score_o$ is the occupational generative AI exposure score. Weighting by the wage bill gives greater weight to occupations accounting for a larger share of each industry’s labor costs.

Employment and wage weights are fixed at their May 2022 values, before the public release of ChatGPT in late 2022. This prevents subsequent changes in occupational composition or wages from altering the exposure index. The resulting measure captures potential technological exposure based on each industry’s pre-release occupational structure.

For the complementary analyses based on the standard sectoral Business Formation

Statistics, we also construct the exposure index at the 2-digit NAICS level using the same procedure.¹

2.2 Producer Price Indices

We use Producer Price Indices (PPI) from the BLS, obtained through the Federal Reserve Economic Data (FRED) database. The PPI measures changes in the selling prices received by domestic producers for their output, making it well suited to studying price responses to sectoral cost changes.

We construct a monthly panel at the 3-digit NAICS level to balance cross-industry variation against potential measurement noise in more disaggregated price 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. For the analysis, we average the monthly indices within each year to focus on medium-run price dynamics. Supplemental Appendix A provides additional details on the construction of the PPI panel.

2.3 Firm Entry: Business Applications

We use business applications from the U.S. Census Bureau’s Business Formation Statistics (BFS) as a proxy for firm entry. These data are based on applications for Employer Identification Numbers (EINs) submitted to the Internal Revenue Service and are informative about subsequent business formation and job creation (Bayard et al., 2018). Our main firm-entry analysis uses the Weekly BFS by Industry product, which reports application counts at the 3-digit NAICS level and provides richer cross-industry variation than the standard 2-digit series. We aggregate weekly counts into annual totals.

To limit classification noise, we exclude 3-digit subsectors belonging to parent sectors with persistently high shares of applications classified as “Not Further Defined” (NFD), following the filter described in Supplemental Appendix A. The resulting panel spans 2011–2024 and covers 63 private industries.

¹For 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 Supplemental Appendix A for more details.

For complementary evidence on the composition of applications, we use the standard monthly BFS series at the 2-digit NAICS level. This panel covers 19 private sectors over 2011–2025, extending the analysis by one year. Total Business Applications (BA) comprise EIN applications retained after the Census Bureau’s standard filters. High-Propensity Applications (HPA) identify applications with characteristics associated with a higher probability of becoming employer businesses. We define Low-Propensity Applications (LPA) as the residual difference, $LPA_{j,t} = BA_{j,t} - HPA_{j,t}$. Supplemental Appendix A provides further construction details for both panels.

2.4 Supply-Chain Stress and Sectoral Controls

Given that the price analysis spans the COVID-19 period, we 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 use its calendar-year mean to match the annual frequency of the price analysis.

To capture heterogeneous sectoral exposure to these aggregate disruptions, we construct two predetermined input-output measures using the 2017 detailed Supply-Use framework and Import Matrix from the Bureau of Economic Analysis (BEA). The first measures purchases of goods-producing and transportation/warehousing inputs relative to gross output. The second measures imported intermediate inputs as a share of total intermediate input use. These measures capture dependence on physical inputs and logistics, and on foreign intermediate inputs, respectively. We use 2017, the latest detailed benchmark year preceding both the pandemic and the commercialization of generative AI, and hold both exposure measures fixed throughout the analysis. Supplemental Appendix A provides construction details and summary statistics.

For the exercises in Subsection 3.3, we also use a goods-versus-services indicator based on the BLS industry classification and three additional sectoral measures constructed from the 2017 BEA tables: direct energy-input exposure, orientation toward private fixed investment, and exposure to import competition in industries’ product markets. Supplemental Appendix D documents their construction.

Finally, we construct an industry-level measure of labor intensity to capture baseline structural differences across sectors. 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 divide total compensation over 2016–2019 by total value added over the same period:

$$Labor_Share_j = \frac{\sum_{t=2016}^{2019} Compensation_{j,t}}{\sum_{t=2016}^{2019} Value_Added_{j,t}}. \quad (2)$$

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 Price Dynamics

This section examines whether sectors with greater exposure to generative AI experience lower relative producer-price growth as the technology diffuses. By automating or augmenting exposed cognitive tasks (Eloundou et al., 2024), generative AI may improve productivity and reduce labor requirements and administrative overhead. To the extent that these gains lower marginal production costs, they should exert downward pressure on producer prices, with the strength and timing of pass-through depending on the pace of adoption, task reorganization, and market conditions, including competition.

To evaluate this hypothesis, we estimate a dynamic continuous-treatment event study using the annual sectoral PPI panel described in Section 2. This approach traces when relative price differences emerge and how they evolve, while controlling for differential exposure to pandemic-era supply-chain stress. A complementary cross-sectional exercise in Supplemental Appendix C summarizes cumulative price divergence by 2025 relative to projections based on each sector’s pre-pandemic trend, providing an endpoint comparison alongside the dynamic analysis.

3.1 Empirical Design

We use a continuous-treatment event-study design to examine the dynamic relationship between generative AI exposure and sectoral prices. This framework allows us to assess whether more and less AI-exposed sectors followed similar price trajectories before the post-2022 diffusion period and to trace the subsequent evolution of relative prices. We also include controls for differential exposure to pandemic-era supply-chain stress, a key potential

confounder given industries' differing reliance on goods, logistics, and imported inputs.

We estimate the following two-way fixed effects specification:

$$\ln(PPI_{j,t}) = \alpha_j + \tau_t + \sum_{\substack{k=2012 \\ k \neq b}}^T \beta_k (AI_Exposure_j \times \mathbb{I}\{t = k\}) + \mathbf{SC}'_{j,t} \boldsymbol{\lambda} + \varepsilon_{j,t}. \quad (3)$$

Here, $\ln(PPI_{j,t})$ denotes the log producer price index for sector j in year t , and T denotes the final year of the estimation sample. Sector fixed effects, α_j , absorb time-invariant unobserved heterogeneity across industries, including persistent differences in production technologies and structural factor intensities. Year fixed effects, τ_t , absorb annual price movements common to all sectors, including those associated with 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 difference in log price changes from the omitted reference year b to year k associated with a one-unit higher exposure to generative AI. The identifying variation therefore comes from comparing price changes across sectors with different predetermined AI exposure, after accounting for permanent sectoral differences, common annual shocks, and differential exposure to supply-chain stress.

The vector $\mathbf{SC}_{j,t}$ controls for differential exposure to supply-chain stress. Year fixed effects absorb aggregate price movements common across sectors, but they do not absorb heterogeneous responses to a common shock arising from differences in 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). If these input dependencies are correlated with AI exposure, the event-study coefficients could partly reflect pandemic-related cost pressures.

We construct these controls by interacting the annual GSCPI with the two predetermined sectoral exposure measures from the 2017 BEA tables described in Section 2.4: dependence on goods-producing and transportation/warehousing inputs, and imported-input intensity. Supplemental Appendix A provides construction details.

The baseline specification allows supply-chain pressure to affect sectoral prices flexibly in two ways. First, for each exposure measure, we include interactions with both the contemporaneous GSCPI and its one-year lag, allowing supply-chain disruptions to pass through to producer prices with delay. Second, we allow the coefficients on both terms to differ between normal and high-stress years, using the current year's stress classification. 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.3). 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.² 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 disruption 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 Supplemental Appendix B.

Identification relies on a conditional parallel-trends assumption: absent the diffusion of generative AI, price changes relative to the reference year would not have varied systematically with predetermined AI exposure after accounting for sector and year fixed effects and the supply-chain controls. 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.

3.2 Dynamic Price Responses

Figure 2 presents the dynamic exposure coefficients from the baseline event-study specification. Each point measures the differential log price change associated with a one-unit higher AI exposure, relative to 2019, the omitted reference year. The figure reports both 90 percent and 95 percent confidence intervals. The 2026 coefficient uses January–April data 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, take both positive and negative values, and are statistically indistinguishable from zero. This pattern is consistent with similar pre-pandemic price trajectories across sectors with different AI exposure.

²In 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.




Figure 2: Generative AI Exposure and Sectoral Price Dynamics

Notes: The figure plots the dynamic exposure coefficients β_k from Equation 3. Each coefficient measures the difference in log price changes from 2019 to year k associated with a one-unit higher generative AI exposure index. The omitted reference-year coefficient 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 pointwise 90% and 95% confidence intervals, respectively. Standard errors are clustered at the NAICS industry level. [†]The 2026 coefficient uses January–April data and should be interpreted as suggestive. Full estimation results are reported in Supplemental Appendix B.

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 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 business-adoption patterns documented in Figure 1 show that reported AI use increased but remained limited over the period shown. Translating generative AI capabilities into lower producer prices can involve firm-level experimentation, integration into business processes, task reorganization, cost reductions, and eventual pass-through to producer prices. Under this mechanism, price effects need not be large immediately after the public release of the technology. As adoption broadens and firms learn to use the technology more effectively, efficiency gains may accumulate and be progressively reflected in sectoral producer prices. The widening relative price gap in Figure 2 is consistent with this gradual transmission.

This dynamic profile also helps discipline alternative interpretations. A central concern is that the estimates reflect residual pandemic-era supply-chain forces. Aggregate supply-chain pressures were concentrated in 2020–2022 and had largely normalized by 2023, as shown in Table A.3. Yet the pandemic-window point estimates fluctuate rather than displaying a sustained decline, and the relative price gap becomes progressively more pronounced thereafter. This timing weakens an explanation driven solely by contemporaneous bottleneck pressures, which would more naturally generate a divergence concentrated in the high-stress period.

More generally, an omitted factor would need to account for the widening post-2022 price differential across sectors with different AI exposure, beyond the common annual shocks and differential supply-chain exposure captured by the baseline controls. The relatively flat pre-pandemic profile also weighs against an explanation based on the continuation of pre-existing differential price trends. Delayed pandemic-related adjustments or new sectoral shocks correlated with AI exposure remain possible. Subsection 3.3 therefore examines whether the estimated trajectory is sensitive to additional controls for goods-versus-services movements, energy dependence, investment-oriented demand, and competition from imported products, allowing these dimensions to have distinct annual price effects from 2020 onward.

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. The complementary cross-sectional estimates in Table C.1 of Supplemental Appendix C yield a similar economic magnitude: the same exposure difference is associated with a 4.2 percentage point lower price gap in 2025 relative to projections based on each sector’s 2015–2019 trend. The accompanying scatterplot (Figure C.1) displays the unadjusted relationship between AI exposure and this price gap.

Supplemental Appendix B provides full estimation results and complementary checks for the baseline specification. Table B.1 reports the coefficients underlying Figure 2. Figure B.1 re-estimates the specification after excluding the partial 2026 observation; the estimated trajectory through 2025 is virtually unchanged. Figure B.2 and Table B.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 slightly noisier estimates, but preserves a similar dynamic 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.

3.3 Alternative Explanations and Robustness

The post-2022 price divergence occurs during a period of substantial sectoral adjustment. We therefore examine whether the estimated relationship is sensitive to additional controls for sectoral supply and demand conditions and exposure to import competition. Table 1 summarizes these comparisons on the same sample of 65 industries and 921 observations. Each estimate expresses the price differential relative to 2019 in log points, using the interquartile difference in AI exposure of 0.129 reported above. The columns distinguish 2025 from the partial observation for 2026. Rows (2)–(5) add each control separately to the baseline; row (6) includes all four blocks jointly. Figure 3 compares the full dynamic profiles for the four individual controls with the baseline. Across these comparisons, the post-2022 downward trajectory remains close to the baseline. Supplemental Appendix D describes the specifications and reports the complete set of profiles, including the joint specification.




Figure 3: Dynamic Price Responses with Additional Sectoral Controls

Notes: Panels (a)–(d) correspond to rows (2)–(5) of Table 1. Blue points show estimates with the indicated control added to the baseline; gray points repeat the baseline. Each additional control is interacted with every year from 2020 onward. Points report log-point price differentials relative to 2019 for an interquartile AI-exposure difference of 0.129, $100 \times 0.129 \times \hat{\beta}_t$. Inner thick and outer thin bars show pointwise 90% and 95% confidence intervals, respectively, with standard errors clustered by industry. All specifications retain the baseline supply-chain controls and use the same 65 industries and 921 observations. Slight horizontal offsets separate the intervals. Shading marks 2020–2022; the dotted line separates the post-2022 period. 2026† averages January–April and is suggestive.

Table 1: Price Differentials Across Alternative Specifications

Specification	2025	2026 [†]
(1) Baseline	−4.23** (2.11)	−5.85** (2.77)
(2) Baseline + goods/services	−5.18** (2.07)	−6.18** (2.70)
(3) Baseline + energy inputs	−4.62** (2.18)	−5.72** (2.81)
(4) Baseline + investment orientation	−4.25* (2.13)	−5.81** (2.82)
(5) Baseline + imported-product competition	−4.50** (2.07)	−5.93** (2.79)
(6) Baseline + all four controls	−4.19* (2.17)	−5.00* (2.86)

Notes: Entries are $100 \times 0.129 \times \hat{\beta}_t$: log-point price differentials relative to 2019 for the interquartile AI-exposure difference of 0.129 reported in Table A.1. Industry-clustered standard errors are in parentheses; inference uses 64 degrees of freedom. All rows use the same 65 industries and 921 observations and retain the baseline GSCPI controls. Rows (2)–(5) separately add the indicated control interacted with each year from 2020 onward. Row (6) includes all four blocks jointly. [†]2026 averages January–April; 2025 is a complete year. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Goods and services. Pandemic-related changes in spending and supply conditions may have shifted the relative prices of goods and services persistently. We add a goods-industry indicator interacted with every year from 2020 onward. These terms absorb a distinct annual path for goods relative to services, including permanent shifts and continuing adjustments. The estimated price differentials become −5.18 in 2025 and −6.18 in 2026, compared with −4.23 and −5.85 in the baseline. The full dynamic profile remains similar to the baseline, with small pre-pandemic coefficients and a progressively widening negative price differential after 2022. The persistence of the negative pattern weakens an explanation based on movements common to these broad groups. The specification estimates a common AI-exposure coefficient conditional on those movements; it does not establish equal AI-related responses across the two groups.

Energy costs. Energy-price shocks could contribute to the observed price divergence if industries’ energy dependence is correlated with AI exposure. We separately allow prices to evolve differently according to industries’ direct purchases of energy products relative to gross output in 2017. Interacting this exposure with every year from 2020 onward accommodates persistent as well as transitory movements associated with energy dependence. The estimated differentials are −4.62 and −5.72 log points in 2025 and 2026, respectively, and the dynamic trajectory remains close to the baseline. This comparison addresses direct

energy-input exposure, including feedstocks; it does not capture every indirect energy cost or the determination of energy producers' selling prices.

Investment-oriented demand. Weaker investment demand, including that associated with monetary tightening, could depress relative prices in industries supplying investment goods and services. We interact each year from 2020 onward with the share of 2017 domestic production allocated directly to private fixed investment. The dynamic profile remains close to the baseline, with estimated differentials of -4.25 and -5.81 log points in 2025 and 2026, respectively. This stability weakens an explanation based on differential price movements associated with prior investment orientation. The measure includes equipment, nonresidential structures, residential investment, and intellectual property; it is not an estimated interest-rate sensitivity.

Competition from imported products. Changes in the prices or availability of competing imports could contribute to the observed divergence in domestic producer prices. Imported-input exposure concerns what an industry buys. To address competition in what it sells, we separately interact each year from 2020 onward with the 2017 import share in the domestic markets for its products, weighted by its production mix. The estimated differentials are -4.50 and -5.93 log points in 2025 and 2026, respectively, and the full dynamic profile remains close to the baseline. Thus, the association is also stable to differential price movements along this distinct trade dimension, conditional on the baseline GSCPI controls.

Including all four controls jointly leaves the estimated price differentials close to the baseline, at -4.19 and -5.00 log points in 2025 and 2026, respectively, with standard errors also close to their baseline values. The dynamic profile also remains close to the baseline. Taken together, these comparisons weaken explanations based on differential price movements associated with the measured sectoral characteristics. They do not rule out unobserved sectoral changes correlated with AI exposure.

Finally, Supplemental Appendix E examines coverage and influence using the baseline specification through 2025. Requiring complete monthly coverage throughout 2012–2025 leaves the full estimated trajectory close to that of the reference sample over the same period. Table E.1 reports the corresponding 2025 comparison. The estimated 2025 differential relative to 2019 also remains negative and comparable in magnitude to the reference estimate when each industry's entire history is omitted in turn, and when broader industry families are omitted (Table E.2). These checks address the composition of the observed panel and the influence of individual sectors; they complement the common-sample comparisons in Table 1.

4 Firm Entry

We next examine whether generative AI exposure is associated with changes in firm entry. Generative AI 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. By lowering the cost of these activities, generative AI may reduce barriers to entry in sectors where exposed tasks account for a larger share of business activities. AI-related efficiency gains may thus affect both prices and entry, without requiring either response to cause the other.

This channel predicts stronger entry responses in more exposed sectors as generative AI diffuses. Potential entrepreneurs and young firms in these sectors can use AI tools for tasks that would otherwise require labor, specialized expertise, or external services. The response need not be immediate: translating these capabilities into new business activity may require experimentation, task reorganization, and the identification of business opportunities.

To evaluate this hypothesis, we use the 3-digit NAICS Business Formation Statistics panel described in Section 2.3. We use business applications as a proxy for firm entry and measure their growth using Davis–Haltiwanger–Schuh (DHS) growth rates, which scale changes by the average number of applications in the two years being compared. The 3-digit panel provides more detailed industry variation than the standard 2-digit series.

We estimate a dynamic continuous-treatment event study to trace differences in entry growth across sectors with different AI exposure and assess whether such differences were already present before the commercialization of generative AI. Supplemental Appendix F provides complementary cross-sectional evidence at the 3-digit and 2-digit NAICS levels, with the latter also distinguishing between high- and low-propensity business applications.

4.1 Empirical Design

We use a dynamic event-study design to examine whether the relationship between generative AI exposure and annual entry growth changes after 2022 and to assess entry-growth patterns

before that point. We estimate the following two-way fixed effects specification:

$$DHS_Growth_{j,t} = \alpha_j + \tau_t + \sum_{\substack{k=2011 \\ k \neq 2022}}^{2024} \beta_k (AI_Exposure_j \times \mathbb{I}\{t = k\}) + \varepsilon_{j,t}. \quad (4)$$

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 persistent differences in average entry growth across industries. Year fixed effects, τ_t , absorb annual movements in entry growth common to all sectors.

The coefficients of interest are the dynamic exposure coefficients, β_k . Each β_k measures the change, relative to 2022, in the annual entry-growth differential associated with a one-unit higher generative AI exposure index. The comparison therefore allows sectors to differ persistently in their average entry growth and asks whether growth in more exposed sectors rises or falls relative to that in less exposed sectors.

We use 2022 as the reference year to anchor the comparison immediately before the post-2022 diffusion period. ChatGPT was released publicly in late November 2022, and business formation responses may take time to emerge as potential entrepreneurs experiment with new tools, reorganize tasks, and initiate new ventures. Supplemental Appendix G reports the corresponding estimates using 2019 as an alternative reference year.

The identifying assumption is that, absent the diffusion of generative AI, changes in annual entry growth relative to 2022 would not have varied systematically with predetermined AI exposure, after accounting for sector and year fixed effects. The pre-2022 coefficients provide a diagnostic for this assumption. Standard errors are clustered at the 3-digit NAICS sector level.

4.2 Dynamic Entry Responses

Figure 4 plots the estimated dynamic exposure coefficients relative to 2022. The pre-2022 coefficients fluctuate around zero and are individually statistically indistinguishable from zero. This pattern provides little evidence of a systematic pre-existing trend in the entry-growth differential between more and less AI-exposed sectors.

The post-2022 pattern is different. The 2023 coefficient is imprecisely estimated and statistically indistinguishable from zero. By 2024, however, the coefficient is positive and statistically significant at the 1 percent level. This timing is consistent with an entry response




Figure 4: Generative AI Exposure and Entry Growth

Notes: The figure plots the dynamic exposure coefficients β_k from Equation 4. Each coefficient measures the change, relative to 2022, in the annual DHS-growth differential in business applications associated with a one-unit higher sectoral generative AI exposure index. The reference-year coefficient 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 Supplemental Appendix G.

that takes time to emerge as potential entrepreneurs experiment with AI tools, reorganize tasks, and identify business opportunities.

The magnitude of the 2024 estimate is economically meaningful. The coefficient of 0.477 implies that, for an interquartile difference in AI exposure—from 0.273 to 0.402 (Table A.1)—the annual DHS-growth differential in business applications is approximately 6.2 percentage points higher in 2024 than in 2022. This magnitude is closely aligned with the complementary cross-sectional estimates in Table F.1 in Supplemental Appendix F.1, which also presents the corresponding scatterplot (Figure F.1).

The estimates provide little evidence of a systematic pre-existing trend in the entry-growth differential associated with AI exposure, although sector-specific shocks correlated

with exposure could still contribute to the post-2022 pattern. Supplemental Appendix G reports the full coefficient estimates and examines sensitivity to sectoral influence and the choice of reference year. The 2024 differential remains positive and statistically significant at the 5 percent level when each industry or broader industry family is omitted in turn, as well as when 2019 is used as the reference year.

Supplemental Appendix F.2 reports complementary evidence using the standard 2-digit NAICS BFS series for 19 broad private sectors. These data extend the analysis through 2025 and distinguish between High-Propensity and Low-Propensity Applications, a classification unavailable in the 3-digit panel. In this panel, greater AI exposure is associated with stronger post-2023 growth in business applications relative to pre-pandemic trends, consistent with the main 3-digit evidence. Supplemental Appendix F.2.4 reports checks using alternative specifications and sectoral exclusions. The positive relationship appears for both High-Propensity and Low-Propensity Applications (Table F.3 in Supplemental Appendix F.2.3), suggesting that the entry response is not confined to applications with a low ex-ante probability of becoming employer businesses.

5 Theoretical Framework

The empirical analysis documents two related product-market patterns following the introduction of generative AI. More exposed sectors experience lower relative price growth and stronger firm-entry growth. Both patterns 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.

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 lower the aggregate price index and increase the equilibrium mass of firms.

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 stylized cost structure isolates 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 provide 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 a larger mass of firms and lower prices, while the empirical analysis examines the early responses along these two margins.

The comparative statics are straightforward. Lower fixed operating 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: a larger mass of firms and lower prices can arise 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 A 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, $W = 1$.³ 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:

$$U = \left(\int_{i \in \Omega} q_i^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}, \quad (5)$$

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 incurs 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

³The wage normalization expresses prices and costs in units of labor. The model abstracts from labor-market frictions and, with inelastic labor supply and market clearing, holds aggregate employment fixed at L . This allows us to focus on adjustment through firm entry and pricing.

cost. This yields net operating profits that are strictly increasing in productivity:

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

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

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,

$$G(A) = 1 - \left(\frac{A_{\min}}{A} \right)^k, \quad A \geq A_{\min}, \quad (7)$$

where $A_{\min} > 0$ and the shape parameter satisfies $k > \varepsilon - 1$. A firm operates only if its net operating profit is nonnegative, which defines a zero-cutoff-profit survival threshold A^* . We focus on an interior-selection equilibrium with $A^* > A_{\min}$.⁴ Free entry requires ex-ante expected operating profits, accounting for the possibility of not operating, to equal the sunk entry cost f_e .

Combining the zero-cutoff-profit and free-entry conditions with labor-market clearing yields the pre-AI equilibrium mass of operating firms:

$$N_{pre} = \frac{L}{\varepsilon f} \left(\frac{k - (\varepsilon - 1)}{k} \right). \quad (8)$$

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. We assume that this improvement reduces 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. Rather, it provides a parsimonious way to capture a common labor-saving force: the technology economizes on cognitive and administrative tasks that matter both

⁴This requires $f_e/f < (\varepsilon - 1)/[k - (\varepsilon - 1)]$, as shown in Appendix A.

for production and for the operation and creation of firms. By leaving the ratio between fixed operating and sunk entry costs unchanged, this assumption preserves the cutoff for underlying productivity draws, allowing us to isolate the effects on the mass of operating firms and the aggregate price index.

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

Proposition 1 (The Firm Entry Effect): *In the symmetric benchmark, a cognitive labor-efficiency gain $\gamma > 0$ 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

$$N_{post} = \frac{L}{\varepsilon \left(\frac{f}{1+\gamma} \right)} \left(\frac{k - (\varepsilon - 1)}{k} \right) = N_{pre}(1 + \gamma). \quad (9)$$

Thus, the equilibrium mass of operating firms rises in the same proportion as labor efficiency. Because fixed operating and sunk entry costs fall in the same proportion, their ratio remains unchanged, preserving the cutoff for underlying productivity draws. The increase in firm mass 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 comparative static concerns the equilibrium mass of firms and provides a mechanism for interpreting the stronger growth in business applications observed in more AI-exposed sectors in Section 4.

Proposition 2 (The Price Effect): *In the symmetric benchmark, a cognitive labor-efficiency gain $\gamma > 0$ lowers the aggregate CES price index relative to the wage through lower marginal costs and a larger mass of available varieties.*

Intuition. The post-adoption price index reflects two changes. First, CES mean 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

$$P_{post} = P_{pre}(1 + \gamma)^{-\frac{\varepsilon}{\varepsilon - 1}}. \quad (10)$$

Since $\varepsilon > 1$, the price index is strictly decreasing in γ . The expression combines two forces.

The term $(1 + \gamma)^{-1}$ captures the direct marginal-cost reduction. The term $(1 + \gamma)^{-1/(\epsilon-1)}$ 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 to show that a larger mass of firms and lower prices can be joint consequences of the same underlying cost-reducing force. The additional variety effect lowers the model’s CES price index, although its contribution is not separately identified by the empirical analysis. The direct marginal-cost channel provides a mechanism consistent with the relative price dynamics documented in Section 3, 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.* Free entry implies that aggregate profits net of entry costs are zero. With inelastic labor supply and the wage normalized to one, household income therefore equals L , and indirect utility is $U = L/P$. Since generative AI lowers the aggregate price index, real welfare rises. The welfare gain reflects the same two product-market forces highlighted above: lower firm-level 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 implications of generative AI. Much of the early economic discussion 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 complementary margins of adjustment in product markets: price dynamics and firm entry. We combine sectoral AI exposure measures constructed from occupational task evaluations and pre-release employment and wage data with producer price indices and business application data. We find that more exposed sectors experience lower relative producer-price growth and stronger growth in business applications following the commercialization of generative AI.

A dynamic event-study design documents a progressively widening relative price gap from 2023 onward, with little evidence of systematic differential trends in the pre-2022 pe-

riod. The specification controls for differential exposure to pandemic-era supply-chain stress. Additional controls for sectoral supply, demand conditions, and import competition leave the estimated trajectory largely unchanged, and a complementary cross-sectional inflation-gap exercise yields a similar economic magnitude. These findings are consistent with AI-related efficiency gains gradually passing through to producer prices.

Complementary evidence on firm entry shows stronger growth in business applications in more AI-exposed sectors. This result is robust to excluding individual industries and broader industry families. A parsimonious model of monopolistic competition with heterogeneous firms provides a common framework for interpreting the price and entry patterns. By reducing marginal production costs and fixed overhead requirements, improvements in cognitive labor efficiency can lower prices and support a larger equilibrium mass of firms. The model thus connects the two empirical margins as joint consequences of a common cost-reducing technological force. Within this benchmark, lower prices and greater product variety also raise real welfare.

The evidence should be interpreted as an early evaluation of an ongoing technological transition. Recent evidence on enterprise AI adoption suggests 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 price and entry margins studied here may become more important.

These findings have implications for both policy and future research. A fundamental value of early analysis lies in identifying emerging macroeconomic trends before they fully mature, helping to inform 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 price dynamics and firm entry. 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.

Appendix A: 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.

A.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 :

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

where

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

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

$$C(q_i) = f + \frac{q_i}{A_i}. \quad (\text{A.3})$$

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

$$\pi_i = p_i q_i - \frac{q_i}{A_i} - f. \quad (\text{A.4})$$

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

$$p(A_i) = \left(\frac{\varepsilon}{\varepsilon - 1} \right) \frac{1}{A_i}. \quad (\text{A.5})$$

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

$$\pi(A_i) = B A_i^{\varepsilon-1} - f, \quad (\text{A.6})$$

where

$$B \equiv \frac{R}{\varepsilon} \left(P \frac{\varepsilon - 1}{\varepsilon} \right)^{\varepsilon - 1} \quad (\text{A.7})$$

is an endogenous aggregate demand shifter common to all firms.

A.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:

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

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:

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

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:

$$\int_{A^*}^{\infty} \pi(A) dG(A) = f_e. \quad (\text{A.10})$$

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

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

A.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

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

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. We focus on an interior-selection equilibrium with $A^* > A_{\min}$, which requires $f_e/f < (\varepsilon - 1)/[k - (\varepsilon - 1)]$.

Substituting the Pareto density,

$$dG(A) = k A_{\min}^k A^{-k-1} dA, \quad (\text{A.13})$$

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

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

Rearranging gives the equilibrium survival threshold:

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

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 A.9 over the truncated Pareto distribution for active firms, $A \geq A^*$:

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

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

$$\bar{r} = \varepsilon(\bar{\pi} + f). \quad (\text{A.17})$$

Substituting Equation A.16 gives

$$\bar{r} = \varepsilon f \left(\frac{k}{k - (\varepsilon - 1)} \right). \quad (\text{A.18})$$

The equilibrium mass of operating firms is total expenditure divided by average firm

revenue, $N_{pre} = R/\bar{r}$. Using $R = L$, we obtain

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

A.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

$$c_{i,post} = \frac{1}{A_i(1 + \gamma)}. \quad (\text{A.20})$$

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

$$f_{post} = \frac{f}{1 + \gamma}, \quad f_{e,post} = \frac{f_e}{1 + \gamma}. \quad (\text{A.21})$$

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 A.15 under the new cost structure. Because fixed operating costs and sunk entry costs are scaled down symmetrically, their ratio remains unchanged:

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

It follows that the threshold for the underlying productivity draw is unchanged relative to the pre-AI benchmark:

$$A^{*,post} = A^{*,pre}. \quad (\text{A.23})$$

Substituting the lower fixed operating cost, $f_{post} = f/(1 + \gamma)$, into the equilibrium firm-mass

expression in Equation A.19 yields

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

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 Price Effect): The aggregate CES price index depends on the mass of active firms and their CES mean effective productivity, $\tilde{A}$. Before adoption, this mean is $\tilde{A}_{pre} \equiv [\mathbb{E}(A^{\varepsilon-1} \mid A \geq A^*)]^{1/(\varepsilon-1)}$. The price index is

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

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

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

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

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

Rearranging gives

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

Therefore,

$$P_{post} = P_{pre}(1 + \gamma)^{-\frac{\varepsilon}{\varepsilon-1}}. \quad (\text{A.29})$$

Since $\varepsilon > 1$ and $\gamma > 0$, the aggregate price index falls, $P_{post} < P_{pre}$. Equation A.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 $W = 1$ and $R = L$, indirect utility is $U = L/P$. Hence $P_{post} < P_{pre}$ implies higher real welfare in the benchmark economy. ■

References

- Acemoglu, D. (2025). “The Simple Macroeconomics of AI”. *Economic Policy* 40.121, pp. 13–58.
- Acemoglu, D. and P. Restrepo (2019). “Artificial Intelligence, Automation, and Work”. *The Economics of Artificial Intelligence: An Agenda*. Ed. by A. Agrawal, J. Gans, and A. Goldfarb. University of Chicago Press, 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.
- (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., 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.

- Cai, J. J., X. Gu, L. Sheng, M. Xia, L. Zhao, and W. Zhu (2025). “AI as “Co-founder”: GenAI for Entrepreneurship”. Working paper. URL: https://jh-cai.com/docs/AI_business_formation.pdf.
- 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.
- 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.
- 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*.

The Macroeconomics of Generative AI: Price Dynamics and Firm Entry

Supplemental Appendix

This supplement presents additional data details, empirical results, and robustness checks. Formal model derivations appear in the manuscript’s theoretical appendix.

Contents

A. Data Construction and Sources	S2
B. Price Dynamics	S9
C. Cross-Sectoral Inflation Gaps	S14
D. Alternative Explanations: Details	S18
E. Coverage and Sectoral Influence	S22
F. Firm Entry: Complementary Cross-Sectional Evidence	S26
G. Firm Entry: Dynamic Estimates and Robustness	S37

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.

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.

A.2. 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 :

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

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.2 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.3. 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.2: 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.3 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.3: 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:

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

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:

$$ImportShare_j = \frac{ImportedInputs_j}{TotalIntermediateInputs_j}. \quad (\text{A.3})$$

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.4 reports summary statistics for the two exposure measures used in the baseline event-study specification.

The construction of the additional sectoral controls used in Subsection 3.3 is documented in Appendix D.

Table A.4: 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.

A.4. 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 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.5. 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.5 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.5: 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.

B Price Dynamics

Table B.1 reports the full set of dynamic exposure coefficients underlying Figure 2. The table corresponds to the baseline event-study specification in Equation 3, 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 B.1 repeats the specification in Figure 2 after excluding the partial 2026 observation. The normalization remains 2019.

As an additional normalization check, Figure B.2 and Table B.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 3.1.

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 B.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)
$\beta \times \mathbb{I}\{2026\}^\dagger$	-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 3. 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 B.1: Dynamic Effect of Generative AI Exposure on Sectoral Prices: Excluding 2026

Notes: The figure plots the dynamic exposure coefficients β_k from Equation 3, 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 B.2: Dynamic Effect of Generative AI Exposure on Sectoral Prices: 2022 Normalization

Notes: The figure plots the dynamic exposure coefficients β_k from Equation 3, 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 B.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)
$\beta \times \mathbb{I}\{2026\}^\dagger$	-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 3. 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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Cross-Sectoral Inflation Gaps

C.1 Predetermined Labor Share

The cross-sectional inflation-gap exercise controls for the predetermined industry-level labor-share measure described in Section 2.4 to account for baseline structural differences across sectors that may influence pricing behavior.

C.2 Construction and Estimates

To complement the dynamic event study in Section 3.1, this appendix presents 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:

$$\ln(PPI_{j,t}) = \alpha_j + \delta_j t + \varepsilon_{j,t} \quad (\text{C.1})$$

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 the dependent variable in this exercise, the *Inflation Gap*, as the percentage deviation of the actual observed prices in 2025 from this baseline counterfactual:

$$Inflation_Gap_j = \frac{PPI_{j,2025} - \tilde{PPI}_{j,2025}}{\tilde{PPI}_{j,2025}} \quad (\text{C.2})$$

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.

We evaluate this cumulative gap in 2025, the last complete calendar year in our price panel. First, this allows time for the commercialization and enterprise adoption of generative

AI to progress beyond the initial release of the technology. 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.¹ 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 3.1 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:

$$Inflation_Gap_j = \gamma_0 + \gamma_1 AI_Exposure_j + \theta Labor_Share_j + v_j \quad (C.3)$$

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 (C.3) 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.

Figure C.1 shows the unadjusted cross-sectional relationship between predetermined AI exposure and the 2025 inflation gap. The negative association indicates that more exposed sectors tend to exhibit smaller cumulative price increases relative to projections based on their own 2015–2019 trends.

Table C.1 reports the cross-sectional estimates of Equation (C.3). Across all specifications, the coefficient on AI exposure is negative and statistically significant, consistent with

¹This 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.




Figure C.1: Generative AI Exposure and Sectoral Inflation Gaps

Notes: The figure shows all 57 industries in column (1) of Table C.1, labeled by their 3-digit NAICS codes. The price gap is $100 \times (PPI_{j,2025}/\tilde{PPI}_{j,2025} - 1)$, where $\tilde{PPI}_{j,2025}$ projects each industry's 2015–2019 log-linear trend. The solid line is the ordinary least-squares fit using all 57 observations, without conditioning on labor share.

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 C.1: Cross-Sectoral Estimates: Generative AI and Inflation Gaps

	(1) Short Trend	(2) Long Trend	(3) Quadratic Trend
Generative AI Exposure (β)	-0.326^{**} (0.148)	-0.445^{**} (0.219)	-0.520^{**} (0.233)
Labor Share (2016-2019)	0.126 (0.092)	-0.513^* (0.286)	0.252 (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. The main event-study design in Section 3.1 instead 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.

D Alternative Explanations: Details

D.1 Common sample and specification

The comparisons in Table 1 use exactly the baseline estimation sample: 921 industry-year observations from 65 NAICS industries over 2012–2026. All specifications include industry and year fixed effects, annual interactions with AI exposure, and 2019 as the omitted year. The 2026 observation averages January–April, with all four months observed for each of the 63 industries present that year. It is not directly comparable to a full calendar-year average as a measure of annual inflation.

Let A_j denote AI exposure, $W_{j,t}$ the baseline GSCPI controls in Equation 3, and $Z_j = (G_j, E_j, I_j, C_j)'$ the four additional predetermined controls. The table specifications take the form

$$\begin{aligned} \ln PPI_{j,t} = & \alpha_j + \tau_t + \sum_{k \in \mathcal{T} \setminus \{2019\}} \beta_k A_j \mathbb{1}\{t = k\} \\ & + W'_{j,t} \gamma + \sum_{k=2020}^{2026} Z'_j \theta_k \mathbb{1}\{t = k\} + \varepsilon_{j,t}, \end{aligned} \quad (\text{D.1})$$

where $\mathcal{T} = \{2012, \dots, 2026\}$. Row (1) omits Z_j ; rows (2)–(5) include one component at a time; row (6) includes all four jointly. All rows retain the original GSCPI structure. The additional dimensions have separate coefficients from 2020 onward; they do not relax the historical supply-control structure before 2020.

Standard errors are clustered by industry. The common-sample comparisons use t critical values with 64 degrees of freedom. Reported magnitudes multiply β_k by 100×0.129 , preserving the exposure contrast used in the main text across all specifications. These are log-point price differentials relative to 2019, not changes in aggregate inflation. The confidence intervals are pointwise, not simultaneous bands.

D.2 Predetermined control measures

Goods and services. Following the BLS industry classification, goods industries comprise NAICS 11, 21, 23, and 31–33 (27 industries); the remaining 38 industries are classified as services, including utilities and trade. The indicator is an industry classification, not a classification of individual products or purchased inputs. The year interactions absorb the

difference in average paths between the groups and do not estimate separate AI-exposure coefficients for goods and services.

Energy inputs. From the 2017 BEA detailed Use table in the Supply–Use framework, we sum purchases of oil and gas extraction products (211000), coal (212100), electricity (221100), natural gas distribution (221200), and petroleum refining products (324110). For each mapped NAICS industry, the sum is divided by gross output at basic prices, using the same industry concordance as the existing input controls. The basket includes purchases used as feedstocks as well as energy used in production. It measures direct exposure and does not apply input-output multipliers to capture energy embodied in other inputs.

Private fixed investment. We use the 2017 BEA Make, Use, and Import tables before redefinitions at producers’ prices. For each product, we subtract the imported component from its allocation to private equipment, intellectual property, residential investment, and nonresidential structures (F02E00, F02N00, F02R00, and F02S00). We exclude the scrap and used-asset transfer accounts S00401 and S00402. Domestic product allocations are assigned to industries using their shares of product output in Make. The resulting industry amount is divided by gross output; NAICS aggregation sums numerators and denominators. In notation,

$$I_j = \frac{1}{g_j} \sum_{c \in \mathcal{C}} \frac{V_{jc}}{q_c} \sum_{f \in \mathcal{I}} (F_{cf} - M_{cf}), \quad (\text{D.2})$$

where V_{jc} is industry j ’s output of product c , q_c is domestic output of that product, g_j is industry gross output, and F and M denote final uses and their imported component. This direct allocation assumes that producers of a product participate in its uses according to their production shares. It includes own-account production and excludes inventories, government investment, capital-goods exports, household durable consumption, and indirect suppliers to investment.

Imported products. For each product c , let U_c^d be intermediate plus domestic final uses, excluding exports, changes in inventories, and the import-subtraction column. Let M_c^d be the corresponding uses in the Import matrix. Define the product’s import share as $p_c = M_c^d / U_c^d$. The industry measure weights these shares by its production mix:

$$C_j = \sum_c \frac{V_{jc}}{g_j} p_c. \quad (\text{D.3})$$

We aggregate industries to NAICS using gross output at producers’ prices. The measure uses BEA’s published import allocation and valuation at the domestic port, including the relevant duties and resident transportation and insurance services. Four special accounts do

not have comparable product-market shares; none is produced by the industries mapped to the estimation sample, so no sample observations are lost. Published rounding is retained without renormalizing the production shares. This measure concerns imported competition in domestic product markets; it does not measure foreign affiliates' local production, competition in export markets, or individual tariff and exchange-rate shocks.

The energy, investment, and import measures are fixed at 2017 values. They cover all 65 industry labels through the inherited concordance, which contains some broad mappings. In particular, NAICS 236 and 238 share a construction aggregate, and 449 and 459 share a trade aggregate. These measures therefore approximate the relevant sectoral exposures rather than observing every industry's exact transactions. Annual interactions allow movements proportional to the measured exposure, not arbitrary shocks to each industry.

D.3 Dynamic profiles

Figure D.1 shows the complete profiles, with a common scale across specifications.




Figure D.1: Dynamic Price Differentials: Comparison with the Baseline

Notes: Panels correspond to rows of Table 1. Blue points show each specification; gray points repeat the baseline in panels (2)–(6). Slight horizontal offsets separate the intervals. Points are $100 \times 0.129 \times \hat{\beta}_t$ relative to 2019. Inner thick and outer thin bars show pointwise 90% and 95% industry-clustered confidence intervals. All models use 65 industries and 921 observations. Shading marks 2020–2022; the dotted line separates the post-2022 period. 2026[†] contains January–April and is suggestive.

E Coverage and Sectoral Influence

These checks use the baseline GSCPI specification through 2025, the last complete calendar year in the panel, with 2019 as the reference year. The reference sample has 858 observations and 65 industries. It differs from the sample including partial 2026 used in Subsection 3.3; all comparisons below are made within this shorter period. The fixed exposure contrast remains 0.129, including after sample exclusions.

E.1 Continuous coverage

Some industry PPI series do not span the full sample period. To assess sensitivity to this uneven coverage, we require all twelve monthly PPI observations in every year from 2012 through 2025. This balanced sample contains 59 industries and 826 observations, retaining all 17 reconstructed indices and 42 native indices. The estimated 2025 differential is -4.34 log points, compared with -4.27 in the shorter-period reference sample. Table E.1 reports both estimates and their confidence intervals. The similarity supports stability to this coverage restriction, while the restriction changes the set of observed industries.

Table E.1: Continuous Coverage and the 2025 Price Differential

Sample	Industries	Observations	2025 vs. 2019
Reference through 2025	65	858	-4.27 [$-8.44, -0.10$]
Continuous monthly coverage	59	826	-4.34 [$-8.62, -0.06$]

Notes: Both rows use the baseline GSCPI controls, 2019 as the reference, and data through 2025. Entries are log-point differentials for the fixed AI-exposure contrast of 0.129. Brackets report pointwise 95% confidence intervals with industry-clustered standard errors and, respectively, 64 and 58 degrees of freedom. The balanced sample requires twelve monthly observations in every year from 2012 to 2025.

E.2 Leaving out entire industry histories

We reestimate the same baseline after omitting each of the 65 industries in turn, and then after omitting each of 17 broader NAICS families. Each exclusion removes the unit’s entire observed history and reestimates all coefficients. Manufacturing (31–33), retail trade (44–45), and transportation and warehousing (48–49) are treated as combined families. Table E.2 reports every exclusion for the 2025 differential relative to 2019, with the estimation-specific

clustered confidence interval.

The point estimates remain negative and comparable in magnitude to the reference estimate across all 65 industry exclusions and all 17 family exclusions. The ranges are $[-5.60, -3.70]$ and $[-5.74, -3.59]$ log points, respectively. These are ranges across estimates, not confidence intervals. Differences across exclusions reflect changes in sample support and reestimation of the controls and cannot be interpreted as additive causal contributions of particular industries.

Table E.2: The 2025 Price Differential After Omitting Entire Industry Histories

Omitted NAICS	Industries	Estimate	95% confidence interval
<i>Individual industries</i>			
113	64	-5.32	$[-9.14, -1.51]$
211	64	-4.31	$[-8.51, -0.11]$
212	64	-3.70	$[-7.79, 0.39]$
213	64	-4.44	$[-8.64, -0.23]$
221	64	-4.28	$[-8.45, -0.10]$
236	64	-4.18	$[-8.36, -0.01]$
238	64	-4.08	$[-8.34, 0.18]$
311	64	-4.38	$[-8.69, -0.06]$
312	64	-4.19	$[-8.38, 0.01]$
313	64	-4.38	$[-8.62, -0.13]$
314	64	-4.44	$[-8.64, -0.25]$
315	64	-4.23	$[-8.40, -0.05]$
316	64	-4.30	$[-8.49, -0.12]$
321	64	-4.20	$[-8.47, 0.07]$
322	64	-4.29	$[-8.55, -0.04]$
323	64	-4.33	$[-8.51, -0.15]$
324	64	-4.21	$[-8.38, -0.03]$
325	64	-4.24	$[-8.42, -0.07]$
326	64	-4.26	$[-8.47, -0.05]$
327	64	-4.15	$[-8.35, 0.06]$
331	64	-4.13	$[-8.33, 0.08]$
332	64	-4.14	$[-8.32, 0.04]$
333	64	-4.33	$[-8.52, -0.15]$
334	64	-4.04	$[-8.42, 0.34]$
335	64	-4.48	$[-8.68, -0.27]$
336	64	-4.28	$[-8.45, -0.11]$
337	64	-4.23	$[-8.44, -0.03]$

Continued on next page

Table E.2 (continued)

Omitted NAICS	Industries	Estimate	95% confidence interval
339	64	-4.23	[-8.41, -0.04]
423	64	-4.82	[-8.98, -0.66]
424	64	-4.41	[-8.64, -0.19]
425	64	-4.27	[-8.44, -0.09]
441	64	-4.29	[-8.45, -0.12]
444	64	-4.36	[-8.53, -0.18]
445	64	-4.22	[-8.41, -0.03]
449	64	-4.21	[-8.38, -0.04]
455	64	-4.26	[-8.43, -0.08]
457	64	-4.24	[-8.41, -0.06]
458	64	-4.16	[-8.35, 0.02]
459	64	-4.05	[-8.21, 0.11]
481	64	-4.45	[-8.64, -0.26]
482	64	-4.43	[-8.64, -0.22]
483	64	-4.16	[-8.35, 0.02]
484	64	-4.28	[-8.48, -0.07]
486	64	-4.27	[-8.45, -0.10]
488	64	-4.25	[-8.42, -0.07]
491	64	-4.12	[-8.30, 0.06]
492	64	-4.03	[-8.22, 0.16]
493	64	-4.05	[-8.32, 0.22]
513	64	-4.40	[-8.66, -0.15]
517	64	-3.99	[-8.25, 0.27]
518	64	-3.77	[-8.34, 0.80]
522	64	-4.29	[-8.86, 0.27]
523	64	-5.60	[-9.19, -2.01]
524	64	-4.03	[-8.55, 0.50]
531	64	-4.18	[-8.38, 0.03]
532	64	-4.19	[-8.37, -0.01]
541	64	-4.21	[-8.69, 0.26]
561	64	-4.21	[-8.38, -0.04]
562	64	-4.22	[-8.41, -0.03]
621	64	-4.15	[-8.32, 0.02]
622	64	-4.22	[-8.39, -0.04]
623	64	-4.22	[-8.41, -0.02]
713	64	-4.35	[-8.55, -0.15]
721	64	-4.48	[-8.68, -0.29]
811	64	-3.87	[-8.01, 0.26]

Continued on next page

Table E.2 (continued)

Omitted NAICS	Industries	Estimate	95% confidence interval
<i>Industry families</i>			
11	64	-5.32	[-9.14, -1.51]
21	62	-3.91	[-8.06, 0.24]
22	64	-4.28	[-8.45, -0.10]
23	63	-3.99	[-8.25, 0.27]
31-33	44	-4.06	[-9.59, 1.46]
42	62	-4.99	[-9.22, -0.76]
44-45	57	-3.92	[-8.12, 0.28]
48-49	56	-3.87	[-8.32, 0.58]
51	62	-3.59	[-8.40, 1.22]
52	62	-5.74	[-10.10, -1.37]
53	63	-4.09	[-8.30, 0.12]
54	64	-4.21	[-8.69, 0.26]
56	63	-4.16	[-8.35, 0.03]
62	62	-4.05	[-8.25, 0.15]
71	64	-4.35	[-8.55, -0.15]
72	64	-4.48	[-8.68, -0.29]
81	64	-3.87	[-8.01, 0.26]

Notes: Every row reestimates the baseline GSCPI specification through 2025 after omitting the indicated unit's entire history. Estimates equal $100 \times 0.129 \times \hat{\beta}_{2025}$ relative to 2019. Confidence intervals use sector-clustered standard errors and the remaining number of industries minus one as degrees of freedom. The comparison does not use the common-sample specifications in Table 1 or include partial 2026.

F Firm Entry: Complementary Cross-Sectional Evidence

This appendix collects complementary cross-sectional evidence on business applications. Appendix F.1 reports the 3-digit scatterplot and growth regressions. Appendix F.2 presents the 2-digit evidence, its High-Propensity and Low-Propensity decomposition, and robustness checks. The full event-study estimates and their robustness checks are reported separately in Appendix G.

F.1 Evidence from 3-Digit NAICS

We report descriptive evidence on the relationship between generative AI exposure and subsequent growth in business applications in the 3-digit NAICS panel. This exercise summarizes the cross-sectional pattern that complements the dynamic analysis in Section 4.1. We measure growth in applications 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 F.1 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 summarizes the raw cross-sectoral relationship.

Table F.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 DHS growth rate in business applications that is 6.5 percentage points higher. The corresponding estimate for 2022–2024 implies a similar 5.8 percentage point difference.




Figure F.1: Generative AI Exposure and Firm Entry Growth

Notes: The figure plots the DHS growth rate in business applications from 2023 to 2024 against the sectoral generative AI 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.

These descriptive results suggest that sectors more exposed to generative AI experienced faster growth in business applications 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 dynamic event study in Section 4.1 therefore 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.

Table F.1: Generative AI Exposure and Firm Entry Growth

	(1) Growth 2023–2024	(2) Growth 2022–2024
Generative AI Exposure (β)	0.501** (0.198)	0.450** (0.191)
Labor Share (2016–2019)	0.021 (0.103)	0.210* (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. Both specifications include pre-pandemic labor share as a control. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F.2 Evidence from 2-Digit NAICS

This subsection reports complementary firm-entry evidence using the standard 2-digit NAICS Business Formation Statistics. The main firm-entry analysis in Section 4 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.

F.2.1 Empirical Strategy

The analysis covers 19 private 2-digit NAICS sectors. Data construction is described in Appendix A.

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 combine three components: 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:

$$\ln(BA_{j,t}) = \alpha_j + \delta_j t + \varepsilon_{j,t}, \quad (\text{F.1})$$

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.

The counterfactual number of business applications in 2025 is

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

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

$$Excess_BA_j = \frac{BA_{j,2025} - \tilde{BA}_{j,2025}}{BA_{j,2023}}. \quad (\text{F.3})$$

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:

$$Excess_BA_j = \beta_0 + \beta_1 AI_Exposure_j + \gamma Labor_Share_j + u_j, \quad (\text{F.4})$$

where $AI_Exposure_j$ is the predetermined sectoral AI exposure index. We estimate the regression by OLS and report 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.

F.2.2 Baseline Results

We begin with the cross-sectional relationship between generative AI exposure and excess firm entry in the 2-digit BFS panel. Figure F.2 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.




Figure F.2: Generative AI Exposure and Excess Firm Entry

Notes: The figure presents the cross-sectional scatter plot and linear fit between the sectoral generative AI 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.

Table F.2 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 normalized excess-entry gap that is 9.3 percentage points higher, measured relative to the 2023 level of business applications. 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 be more relevant in sectors where labor accounts for a larger share of value added.

Table F.2: Baseline Results: Generative AI Exposure and Firm Entry

	(1) Short Trend (2015–2019)	(2) Long Trend (2011–2019)
Generative AI Exposure (β)	0.731*** (0.243)	0.809*** (0.254)
Labor Share (2016–2019)	0.326*** (0.089)	0.299** (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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 Appendix F.2.4 examine the sensitivity of the estimates to influential observations, alternative trend specifications, and direct growth-rate outcomes.

F.2.3 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 Low-Propensity Applications (LPA) capture applications less likely to transition into em-

ployer status, including independent contractors and zero-employee startups. We therefore repeat the baseline cross-sectional specification separately for the HPA and LPA components.

Table F.3: Heterogeneity Analysis: High vs. Low Propensity Firm Entry

	High-Propensity (HPA)		Low-Propensity (LPA)	
	(1)	(2)	(3)	(4)
	Short Trend	Long Trend	Short Trend	Long Trend
Generative AI Exposure (β)	0.639*** (0.197)	0.670** (0.249)	0.875*** (0.292)	0.954*** (0.289)
Labor Share (2016-2019)	0.288** (0.103)	0.266 (0.153)	0.317*** (0.102)	0.298** (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). Short and long trends refer to 2015–2019 and 2011–2019, respectively. Each outcome is normalized by the 2023 level of its own application category. All specifications include pre-pandemic labor share. 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 F.3 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.² 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.

F.2.4 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

²Standard Wald tests fail to reject equality between the HPA and LPA coefficients.

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 F.4 reports leave-one-out re-estimations of the baseline specification, excluding one sector at a time. The coefficient on AI exposure remains positive and statistically significant at the 5 percent level across all exclusions. 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 F.5 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 log growth in applications from 2023 to 2025, controlling for average annual log growth over 2015–2019 and 2011–2019, respectively. Columns (4) and (5) use the log difference between average applications in 2022–2023 and average applications in 2024–2025, with the same historical-growth controls. Across these alternatives, the coefficient on AI exposure remains 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.

Table F.4: Robustness: Leave-One-Out and Cluster Exclusion

Excluded Sector(s)	AI Exposure Coef. (β_1)	p-value
<i>Panel A: Leave-One-Out (Single Sector)</i>		
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
<i>Panel B: Dual Exclusion (Knowledge-Intensive Clusters)</i>		
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 F.5: Robustness: Alternative Trends and Direct Growth Rates

	Quad. Gap	Growth (2023-2025)		Growth Avg. (22/23-24/25)	
	(1)	(2)	(3)	(4)	(5)
	Quad. Trend	2015–2019	2011–2019	2015–2019	2011–2019
Generative AI Exposure (β)	0.557** (0.221)	0.918** (0.349)	0.956** (0.347)	0.590*** (0.169)	0.610*** (0.175)
Observations	19	19	19	19	19
R^2	0.417	0.601	0.578	0.585	0.565

Notes: Column (1) uses an excess-entry gap constructed from a quadratic log trend estimated over 2011–2019 and anchored at observed applications in 2023. Columns (2)–(3) use the log difference in business applications from 2023 to 2025. Columns (4)–(5) use the log difference between average applications in 2022–2023 and average applications in 2024–2025. All specifications control for pre-pandemic labor share. Columns (2) and (4) additionally control for average annual log growth over 2015–2019; Columns (3) and (5) use the corresponding measure for 2011–2019. Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

G Firm Entry: Dynamic Estimates and Robustness

G.1 Baseline estimates

Table G.1 reports the full set of dynamic exposure coefficients underlying Figure 4. The estimates use the baseline firm-entry event-study specification in Equation 4, with 2022 as the omitted reference year and DHS growth in business applications as the dependent variable.

Table G.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 4. 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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

G.2 Sectoral influence

We omit each of the 63 industries and each of the 18 broader NAICS families in turn, removing their entire histories and reestimating all coefficients. We retain 2022 as the reference year and the fixed exposure contrast of 0.129. Retail trade (44–45) and transportation and warehousing (48–49) are combined families.

Table G.2 summarizes the 2024 estimates. Every exclusion yields a positive estimate significant at the 5 percent level. Figure G.1 reports the estimates and confidence intervals for every family exclusion.

Table G.2: Sectoral Influence on the 2024 Entry-Growth Differential

Exclusions	Reestimations	Minimum	Maximum	$p < 0.05$
Individual industries	63	5.38	6.80	63/63
Industry families	18	5.58	7.24	18/18

Notes: Entries are $100 \times 0.129 \times \hat{\beta}_{2024}$: percentage-point differences in annual DHS growth in business applications relative to 2022. The full-sample estimate is 6.16, with a 95% confidence interval of [2.31, 10.01]. Minimum and maximum refer to the range across reestimations, not confidence bounds. Every regression includes industry and year fixed effects and uses 2011–2024. Inference uses industry-clustered standard errors and the remaining number of industries minus one as degrees of freedom.




Figure G.1: The 2024 Entry-Growth Differential After Omitting Each Industry Family

Notes: Points report the 2024 differential after excluding each family’s entire history. Units are percentage points of annual DHS growth for the fixed exposure contrast of 0.129, relative to 2022. Horizontal lines are pointwise 95% confidence intervals with industry-clustered standard errors and $G - 1$ degrees of freedom, where G is the number of remaining industries. The dashed line marks the full-sample estimate.

Figure G.2 reports the full event-study trajectories under these exclusions. The 2023 estimate remains statistically insignificant at the 10 percent level in every reestimation, while the 2024 estimate remains positive and significant at the 5 percent level. The shaded ranges summarize variation in point estimates across exclusions; they do not measure sampling uncertainty.

Some individual pre-2023 coefficients become significant at the 5 percent level under exclusions: the 2012 coefficient when industry 518 is omitted, and the 2017 coefficient when industry 521 or family 52 is omitted. These are individual tests across many comparisons, without a multiple-testing adjustment; they are not a joint test of pre-period coefficients.




Figure G.2: Firm-Entry Event-Study Trajectories Under Sectoral Exclusions

Notes: The solid line reports $100 \times 0.129 \times \hat{\beta}_t$ from the full-sample event study, normalized to 2022. The shaded area spans the minimum and maximum point estimates after omitting each of 63 industries in panel (a), and each of 18 families in panel (b). The ranges are not confidence intervals. All exclusions remove entire histories and reestimate industry and year fixed effects and all exposure coefficients. The vertical dotted line separates 2023–2024 from earlier years.

G.3 Alternative reference year

Renormalizing to 2019, the last full pre-pandemic year, leaves the sample, fitted values, and residuals unchanged. Coefficients become $\hat{\beta}_t - \hat{\beta}_{2019}$, using the baseline estimates. Standard errors use the full covariance matrix.

Table G.3 compares both normalizations; Figure G.3 plots the trajectory relative to 2019. The 2023 estimate remains imprecise; the 2024 differential is 7.09 percentage points. Under this normalization, all industry and family exclusions yield positive 2024 estimates significant at 5 percent. Point-estimate ranges are [5.80, 7.63] and [4.97, 7.78] percentage points, respectively.

Table G.3: Firm-Entry Growth Under Alternative Reference Years

Reference year	2023	2024
2022	−1.48 [−7.43, 4.47]	6.16 [2.31, 10.01]
2019	−0.55 [−5.55, 4.44]	7.09 [2.67, 11.51]

Notes: Entries are percentage-point differences in annual DHS growth in business applications for the fixed AI-exposure contrast of 0.129, relative to the indicated reference year. Brackets report pointwise 95% confidence intervals using industry-clustered standard errors and 62 degrees of freedom. Both specifications use the same 63 industries and 882 observations over 2011–2024, with industry and year fixed effects.




Figure G.3: Firm-Entry Event Study with 2019 as the Reference Year

Notes: The baseline specification uses 63 industries and 882 observations. Points report percentage-point differences in annual DHS growth in business applications for the fixed exposure contrast of 0.129, relative to 2019 (normalized to zero). Inner thick and outer thin bars are pointwise 90% and 95% confidence intervals with industry-clustered standard errors and 62 degrees of freedom.

References

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