

Who Sets the Price? The Vertical Origins of Uniform Pricing

Alvin Chen[†], Leandro Sanz[‡], and Michael D. Wittry[§]

June 1, 2026

Abstract

Retail prices in U.S. consumer markets are jointly produced by manufacturers and retailers, but we show that the systematic component of price-setting originates primarily upstream. Using manufacturer-product scanner data spanning approximately 5 billion UPC-store-month observations, we decompose retail price variation into manufacturer and retailer components. Manufacturer identity accounts for approximately 90 percent of explained variation in price levels and 97 percent in price changes. Price dispersion is shaped by both layers of the chain, though manufacturers remain the largest source of explained variation. We show that pricing practices change after brand acquisitions: acquired UPCs converge toward the acquirer’s incumbent pricing behavior when the acquirer already operates in the target’s category, but diverge from the acquirer’s broader pricing behavior in expansionary acquisitions. Private-label products, which compress the manufacturer-retailer information wedge, exhibit greater geographic dispersion and responsiveness to local conditions than national brands. Finally, consistent with a reduction in upstream information frictions, products sold by more AI-exposed manufacturers exhibit greater geographic dispersion, more frequent repricing, and lower prices after the introduction of scalable generative AI APIs, with stronger responsiveness to local conditions. The results indicate that retail pricing rigidities reflect upstream informational frictions that technology can relax, rather than immutable features of retail markets.

JEL Codes: L11, L81, D22, D83, O33, E31

Keywords: Uniform pricing, artificial intelligence, retail markets, price dispersion, informational frictions, vertical relationships, price rigidity, technology adoption

For their helpful comments and suggestions, we thank Aditya Chaudry, Andrea Eisfeldt, Isil Erel, Ben Matthies, Sophie Shive, René Stulz, Petra Vokata, Rafael Zambrana, and Michal Zator, and seminar participants at Ohio University, The Ohio State University, and University of Notre Dame. We thank Paul Décaire for assistance with the data. Researcher(s)’ own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[†]Stockholm School of Economics, Email: alvin.chen@hhs.se

[‡]Mendoza College of Business, University of Notre Dame, Email: lsanz@nd.edu

[§]Fisher College of Business, Ohio State University, and European Corporate Governance Institute, Email: wittry.2@osu.edu.

1 Introduction

Few corporate decisions are more important than price-setting. Yet, despite substantial variation across local markets, retail prices are remarkably uniform. DellaVigna and Gentzkow (2019) (DG) document that stores within the same retail chain charge nearly identical prices for identical products, even when serving vastly different clienteles and competitive environments.¹ Such uniformity generates significant losses—DG estimate that the median chain sacrifices \$16 million in annual profit—suggesting that the forces preventing localized pricing are substantial. Yet two important questions remain open: what are those forces and where do they arise in the vertical chain?

Because posted prices are observed and implemented at the retail level, a natural starting point for explaining uniform pricing is downstream: retailers may face menu costs, managerial bandwidth constraints, or the organizational burden of maintaining location-specific prices across large store networks (Levy, Bergen, Dutta, and Venable, 1997; Zbaracki, Ritson, Levy, Dutta, and Bergen, 2004; DellaVigna and Gentzkow, 2019). But prices are not chosen by retailers alone; they are the joint product of manufacturers upstream—who design, market, and often provide pricing recommendations for their products—and retailers downstream—who price and place them on the shelf.² Each layer possesses distinct information: retailers observe local foot traffic, demographics, and competitive conditions, while manufacturers hold product-level knowledge of demand and cross-market patterns. Translating these complementary inputs into optimal prices across thousands of SKUs is costly, and the two sides of the vertical chain rarely possess enough of each other’s information to do so independently.

Once price-setting is viewed as the outcome of this vertical chain, the explanation for uniform pricing depends on where the binding constraint lies. If that constraint lies downstream, the retail-level explanation is correct: uniformity reflects the high implementation and menu costs described

¹ We confirm that uniform pricing is pervasive in the full NielsenIQ Retail Scanner sample. Panel B of Table 1 reports that, across retailer-UPC-month observations, the median coefficient of variation in prices across stores within a retailer is 3.7%, and the 25th percentile is zero.

² Manufacturer-retailer relationships in consumer packaged goods are typically structured by a bundle of contractual instruments—trade promotions, performance rebates, pricing guidance, and category management arrangements—that tie manufacturer payoffs to retail price and volume outcomes. In our sample, 77 percent of manufacturers with available 10-K filings disclose at least one such arrangement, while only 8 percent disclose a wholesale-only relationship. Table IA.B.15 provides details on the construction of these measures; Internet Appendix IA.C discusses their implications for our interpretive framework.

above, or strategic considerations such as maintaining a consistent price image across markets (Ater and Rigbi, 2023). If, however, the relevant friction is upstream, the explanation is fundamentally different: manufacturers face high costs of observing dispersed local demand, competition, and cost signals, and of translating those signals into product-specific pricing recommendations that can be communicated through the vertical chain. Absent such information, uniform pricing is the rational default (Mankiw and Reis, 2002; Alvarez, Lippi, and Paciello, 2011). Under this view, the friction arises not from the cost of changing prices, but from the cost of knowing when and how to adjust them.³

Guided by this distinction, we show that the systematic component of retail price-setting originates primarily upstream. That is, the constraints generating price uniformity appear to arise before the point of retail execution—within manufacturers or in the communication and coordination between manufacturers and retailers—rather than in retailers’ unilateral pricing decisions.

Our empirical investigation proceeds in stages, starting with a decomposition of pricing variation across the vertical chain and then turning to settings that progressively separate upstream informational frictions from downstream implementation costs and price-image considerations. We use a bridge linking over 300 publicly traded manufacturers to product-level NielsenIQ Retail Scanner data spanning approximately 5 billion store-month observations to decompose retail pricing variation into manufacturer and retailer sources using Shapley values.⁴ We find that manufacturer identity accounts for approximately 91 percent of the explained variation in price levels and 97 percent in price changes, indicating that the systematic component of where prices sit and when they move is overwhelmingly determined upstream. For price dispersion, both layers contribute meaningfully—56 percent for manufacturers and 37 percent for retailers—with manufacturers remaining the leading source of systematic variation. Most of the cross-store variation in dispersion and repricing is, however, unsystematic—consistent with widespread uniform pricing across stores.

We next exploit ownership changes that reassign pricing authority while holding the product

³ We formalize this vertical structure in a parsimonious pricing model in Section IA.C of the Internet Appendix.

⁴ Throughout the paper, we use “manufacturer” to refer to the brand owner—the firm that designs, markets, and sets wholesale terms for the product—rather than necessarily the firm that physically produces it. In many cases these are the same entity, but where production is outsourced, the pricing decisions we study are made by the brand owner.

fixed, providing a sharper test of whether pricing behavior is governed upstream. Using brand acquisitions, we show that when ownership of a UPC changes, its pricing behavior changes with it: acquired products move away from their prior pricing patterns and converge toward the acquirer’s incumbent pricing behavior, but only when the acquirer was already operating in the target’s product category. When the acquirer had no prior presence in the category, acquired products move away from their prior patterns without converging toward the acquirer’s existing products—consistent with the absence of an incumbent category-specific pricing routine for the acquirer to extend. These patterns are difficult to reconcile with a purely retailer-driven view of pricing and instead indicate that upstream owners apply portable, category-specific pricing routines to the products they control.

Having established that pricing behavior changes with ownership, we next examine whether the relevant manufacturer-side friction is informational. We benchmark national brands, which comprise the majority of products in the scanner data and operate through the traditional manufacturer-retailer-consumer chain, against private-label brands, which are vertically integrated and controlled by the same retailer (Hoch and Banerji, 1993; Chintagunta, Bonfrer, and Song, 2002). The private-label structure substantially attenuates the vertical information wedge—pricing decisions are made by the same entity that observes local demand information (Hoch, Kim, Montgomery, and Rossi, 1995)—while much of the retailer’s pricing infrastructure and implementation costs associated with managing prices across locations (e.g., shelf tags, scanner systems, promotional calendars) remain unchanged. In many cases, private-label products are physically manufactured by the same firms that produce national brands, so production costs and supply chain logistics are also roughly similar.⁵ This comparison targets the role of information: the key difference is not how prices are implemented, but who controls the pricing decision and the flow of demand information to the shelf.

If retailer-side implementation frictions were the dominant driver of uniform pricing, private-label and national-brand products should exhibit similar uniformity. Several features of private-label markets reinforce this prediction: private labels carry the retailer’s own brand, strengthening

⁵ For example, TreeHouse Foods and Ralcorp produce private label products across many of the same categories as leading national brand manufacturers.

price-image incentives for uniformity; they are often priced as a fixed discount off the category-leading national brand, inheriting any uniformity in that benchmark; they are lower-margin products that may receive less pricing attention; and they face direct competitive pressure from other value-tier brands. If, however, manufacturer-side informational frictions are an important source of uniformity, private labels should display more localized pricing. Consistent with the latter prediction, we find that private-label products exhibit significantly greater price dispersion—even within retailer-by-product-category-by-month—and are responsive to local wage variation while national brands are not. These patterns suggest that manufacturer-side informational frictions are an important source of uniform pricing.

Finally, we examine whether reductions in upstream information-processing costs shift pricing away from uniformity. We exploit the release of OpenAI’s InstructGPT API in January 2022, which was the first generative artificial intelligence (AI) model trained via RLHF to follow natural language instructions at scale, attenuating the need for specialized prompt engineering. We argue that, following this introduction, information processing costs dropped significantly. We find that products sold by more AI-exposed manufacturers become less uniform, reprice more frequently, and respond more strongly to local economic conditions, even within county and retailer. These effects are consistent with a shift toward more localized pricing as manufacturers become better able to extract, process, and act on dispersed market signals (Agrawal, Gans, and Goldfarb, 2019; Goldfarb and Tucker, 2019; Bergemann and Bonatti, 2019).

To identify these effects, we measure firms’ AI exposure using occupation-level scores following Eisefeldt, Schubert, Taska, and Zhang (2025), aggregated to the firm level using January 2020 workforce composition, two years before the API introduction. This timing mitigates concerns that exposure reflects endogenous workforce responses to AI availability.⁶ We compare products sold by high-exposure manufacturers to a matched control group with similar pre-treatment price levels and sales volumes using a difference-in-differences event-study design. Event-study estimates show little evidence of differential pre-trends prior to the API’s introduction. After the API introduction, products sold by high-exposure manufacturers exhibit greater cross-store price dispersion, more

⁶ In Section 10, we validate this measure by showing that highly exposed firms significantly increase AI-related hiring and implementation-oriented disclosures following the API introduction.

frequent repricing, and lower average prices. Their price levels also become more sensitive to local wage conditions, with larger price decreases in lower-wage counties.

The increase in repricing is driven primarily by price cuts, and all effects are robust to retailer-by-time fixed effects, indicating that the observed changes reflect manufacturer-side rather than retailer-side forces. The effects are also concentrated where theory predicts they should be: among firms whose AI exposure loads on pricing-relevant tasks, such as demand analysis, market research, and forecasting; in occupations where AI is expected to augment rather than automate worker capabilities, consistent with AI enhancing managerial decision-making rather than mechanically automating price changes; and among firms with substantial pre-existing data assets, where the returns to improved processing capacity are highest. Taken together, these patterns are consistent with AI relaxing upstream constraints on gathering, processing, and acting on local demand signals (Zbaracki et al., 2004; Stroebel and Vavra, 2019; Brynjolfsson, Li, and Raymond, 2025; Mackowiak and Wiederholt, 2009).

The findings also survive a range of validation checks and additional robustness tests. High-exposure firms increase AI-related hiring and implementation-oriented disclosures following the API introduction, consistent with the exposure measure capturing genuine deployment rather than passive exposure. The main results are robust to controls for concurrent macroeconomic shocks, including interest-rate exposure and COVID-related disruptions, as well as to tighter fixed-effects specifications, placebo event dates, and randomization-inference exercises.

Our findings characterize the pricing of fast-moving consumer goods by large publicly traded manufacturers in brick-and-mortar retail. Whether similar patterns emerge for services, e-commerce, or smaller firms—settings where both AI adoption and pricing frictions differ markedly—remains an open question. More broadly, our results suggest that the study of price rigidity benefits from looking beyond the point of retail execution to the vertical chain that produces prices, and that technological shocks to upstream information processing can have first-order effects on how prices are set.

Related Literature. Our study connects to three strands of the literature. First, and most directly, we contribute to research on the sources of price rigidity and geographic uniformity in retail

markets. Canonical models attribute infrequent price adjustment to fixed menu costs (Goloso and Lucas, 2007; Midrigan, 2011), while a growing body of work emphasizes that the informational and organizational costs of pricing decisions can dwarf the physical costs of changing price tags (Zbaracki et al., 2004; Sims, 2003; Mackowiak and Wiederholt, 2009; Alvarez, Lippi, and Paciello, 2011). DellaVigna and Gentzkow (2019) document that retail prices are nearly uniform across locations and attribute this to the organizational costs of implementing market-specific prices. Closely related, Nakamura (2008) decomposes retail price variation in scanner data into chain-specific, store-specific, and idiosyncratic components and finds that chain effects dominate, though the absence of a manufacturer link prevents separating manufacturer- from retailer-initiated pricing behavior. We extend this line of inquiry by decomposing retail price variation into its vertical components, made possible by a bridge linking publicly traded manufacturers to product-level scanner data. The acquisition, private-label, and AI evidence then progressively distinguishes upstream informational frictions from retailer-side implementation costs and price-image considerations. Together, these results indicate that retail pricing rigidities partly reflect upstream informational frictions that technology can relax, rather than immutable features of retail markets.

This upstream interpretation also connects to research in industrial organization on vertical relationships between manufacturers and retailers, which has formally rejected linear wholesale pricing in favor of richer contract structures including nonlinear pricing, trade promotions, and resale price maintenance (Villas-Boas, 2007; Bonnet and Dubois, 2010). The pricing recommendations and contracting bundle we summarize in Section 2 and formalize in Internet Appendix IA.C are institutional channels through which this literature documents that manufacturer payoffs are tied to retail price and volume outcomes.

Second, we contribute to the rapidly growing literature on the real effects of artificial intelligence. Much of this research has focused on AI's impact on labor markets (Autor, 2015; Acemoglu and Restrepo, 2022; Webb, 2020), firm-level productivity (Brynjolfsson, Rock, and Syverson, 2021), and corporate investment and growth (Babina, Fedyk, He, and Hodson, 2024). Eisfeldt et al. (2025) develop the task-level exposure framework that we build on and show that AI exposure is priced in equity markets. Our paper extends this literature to the product market by documenting how AI

exposure translates into changes in the prices firms actually charge. The augmentation-automation decomposition aligns with [Acemoglu and Restrepo \(2022\)](#) and [Brynjolfsson, Li, and Raymond \(2025\)](#): pricing effects load primarily on augmentation exposure, consistent with AI enhancing employees’ capacity to gather and act on demand information.

Third, our findings complement a growing literature on algorithmic and automated pricing ([Calvano, Calzolari, Denicolo, and Pastorello, 2020](#); [Assad, Clark, Ershov, and Xu, 2024](#); [Dubé and Misra, 2023](#)). Our setting differs in a fundamental way: the pricing effects we document do not seem to be driven by autonomous algorithmic price-setting rules, but by AI’s effect on firms’ information environments. The increase in geographic dispersion and local demand sensitivity we find is consistent with AI enabling firms to acquire and act on demand information that was previously too costly to obtain ([Bergemann and Bonatti, 2019](#)), rather than with algorithmic optimization of existing pricing rules. This distinction matters for welfare: while algorithmic collusion can raise consumer prices, the information channel we identify increases localization and is associated with lower average prices.

Taken together, our results offer a unified account of how upstream information-processing constraints shape retail price-setting. The pricing literature asks why prices are so rigid and uniform; the AI literature studies how AI changes firm behavior. We connect the two by showing that AI relaxes upstream informational frictions that sustain uniform pricing, primarily through task-level augmentation rather than automation. By identifying upstream informational frictions as a central source of uniform pricing, our findings reframe the uniform pricing puzzle from a question about retail operations to one about the flow of information through the vertical chain.

2 Hypothesis Development

Our empirical analysis tests whether the binding friction to localized pricing originates upstream or downstream in the manufacturer-retailer ecosystem. To fix ideas, we summarize the key economic forces and resulting empirical predictions in this setting. Internet Appendix [IA.C](#) develops these arguments in a parsimonious vertical pricing framework.

Two broad classes of frictions can sustain uniform pricing across heterogeneous markets. The

first is downstream: the physical and organizational costs of maintaining localized pricing across stores—updating shelf tags, managing promotions, and coordinating price variations across a retailer’s assortment. The second is upstream: the cost of gathering local demand signals, analyzing them, and transmitting actionable guidance through the vertical chain.

These frictions generate different empirical signatures. Downstream implementation costs are features of the retail environment: they depend on the number of stores, the complexity of the product assortment, and the cost of updating price tags and promotional systems. They do not depend on how much the decision-maker knows about local demand—they may bind even when the optimal local price is known with certainty.

In contrast, upstream frictions originate within the manufacturer or in the transmission of information through the vertical chain, and operate along two distinct margins. The first is the manufacturer’s marginal cost of production, which affects the uniform default price across all markets. The second is the manufacturer’s cost of acquiring and processing demand signals and translating them into market-specific pricing recommendations.

These channels generate different predictions. A reduction in production costs lowers average prices but does not affect price dispersion or repricing behavior. A reduction in information costs increases geographic dispersion, repricing frequency, and sensitivity to local demand conditions, reflecting a shift toward more localized pricing. Implementation costs, if binding, suppress localization regardless of improvements in either upstream channel. We summarize these predictions below.

	Lower c (Production)	Lower I / Higher Φ (Information)
Average price declines	✓	×
Geographic dispersion increases	×	✓
Repricing frequency increases	×	✓
Local demand sensitivity increases	×	✓

3 Data

This section describes the sample construction, the pricing outcomes, the AI exposure measure, and other data used in the empirical analysis.

3.1 Sample construction

We assemble a panel that links NielsenIQ Retail Scanner (RMS) records to publicly traded manufacturing firms in Compustat. The RMS data is obtained from the Kilts Center at the University of Chicago and contains information on weekly prices and unit sales at the store-UPC level. Following [Décaire and Sosyura \(2024\)](#), we link RMS products to Compustat manufacturers via the GS1 barcode registry, which yields a panel that spans 2010 to 2023 and contains approximately 5 billion store-UPC-month observations across 100,250 unique UPCs, 225 retailers, 3,021 counties, 886 product modules, and approximately 350 publicly traded manufacturers. Beyond grocery, the panel covers apparel, consumer electronics, household goods, and health and beauty products. We exclude observations with product module code 9999, which corresponds to unclassified or miscellaneous products.

The empirical samples differ across stages of the analysis. The price variance decomposition in [Section 5](#) uses the full linked panel from 2010 to 2023. The brand-acquisitions analysis uses a UPC-retailer-month panel of 10 brand acquisitions by Compustat manufacturers, with 12-month event windows around each acquisition close date. The private-label test in [Section 7](#) uses the full 2010 to 2023 NielsenIQ scanner universe and does not require a manufacturer match. We construct two panels, one at the UPC-retailer-county-month level with approximately 13.2 billion observations and one at the UPC-store-month level with approximately 27.8 billion observations. The AI experiment in [Section 8](#) uses a firm-UPC-month panel for the linked Compustat manufacturers with a 6-month window around the introduction of the InstructGPT API in January 2022. This panel has 90,211 firm-UPC-month observations across 91 firms and 8,750 unique UPCs.

3.2 Pricing outcomes

We construct several pricing outcomes from the RMS data. Within each store-UPC-week, we average prices and sum unit sales. We also deflate all prices to constant dollars using the monthly Consumer Price Index from the Bureau of Labor Statistics (BLS) and winsorize the resulting real prices at the 1st and 99th percentiles.

Our primary outcomes capture price dispersion and repricing. The variance decomposition uses three variables: log prices, squared deviations from the UPC-month mean, $(\ln p_{ist} - \overline{\ln p_{it}})^2$, and month-over-month log price changes, $\Delta \ln p_{ist}$. Panel A of [Table 1](#) reports summary statistics for these variables on the linked store-UPC-month panel.

We report cross-store and cross-retailer price dispersion at three aggregation levels. Panel B of [Table 1](#) reports the coefficient of variation (CV) of prices computed across stores within a retailer, across retailer-level prices within a county, and across retailer-level prices within UPC-month. The mean CV is 0.073 within retailer, 0.138 across retailers within county, and 0.177 across retailers as a whole. Price dispersion is concentrated between retail chains rather than within them, even within the same county.

Two of our tests measure cross-store price dispersion at narrower units of observation. For the private-label test, the unit of observation is the UPC-retailer-county-month, and we compute the within-retailer cross-store standard deviation of log prices across stores in the same county. Panel C of [Table 1](#) reports the within-retailer dispersion for the private-label sample. The mean standard deviation of log prices is 0.067 for private labels and 0.057 for national brands. For the AI experiment, the unit of observation is the firm-UPC-month, and we compute the standard deviation and interquartile range of log prices across all stores at which the firm sells the UPC in that month. Panel D reports the dispersion outcomes for the AI experiment. The mean standard deviation of log prices is 0.206 and the mean interquartile range is 0.144. Both panels show substantial mass near zero, consistent with the widespread uniform pricing documented in the literature (e.g., [DellaVigna and Gentzkow, 2019](#)).

We measure repricing using an indicator for large changes in log average price at the relevant level of aggregation. For the AI experiment, the baseline outcome is $\mathbb{1}(|\Delta \ln \bar{p}_{ipt}| > 0.05)$, where \bar{p}_{ipt}

is the average price of UPC p sold by firm i in month t . In county and retailer specifications, we define the same indicator using log average prices at the firm-UPC-county-month and firm-UPC-retailer-month levels. The indicator equals one when the absolute log price change from $t - 1$ to t exceeds five percent.⁷ Panel D of [Table 1](#) reports the indicator’s distribution for the AI experiment. The mean is 0.337 for the 5 percent threshold and 0.207 for the 10 percent threshold. The median is zero in both cases, since most firm-UPC-months exhibit no large price change.

The AI experiment uses three sets of outcomes, all reported in Panel D of [Table 1](#). For cross-store dispersion, we use the standard deviation and interquartile range of log price across the stores at which a firm sells a UPC in a given month. The means of these variables are 0.206 and 0.144, respectively. For repricing, we use the 5 percent indicator defined above, which has a mean of 0.337 in this sample. For price levels, we use the natural log of the CPI-deflated monthly average price, $\ln(\textit{Average price})$, with a mean of 1.797, a standard deviation of 0.924, and a median of 1.795.⁸

3.3 AI exposure

We measure firm-level exposure to generative AI using the [Eisfeldt et al. \(2025\)](#) approach. Their approach begins with the universe of task descriptions from the Occupational Information Network (O*NET, version 27.2), which characterizes the work activities associated with each occupation in the U.S. economy. Each task is evaluated by OpenAI’s GPT-3.5 Turbo model for the extent to which generative AI is expected to reduce the time required to perform it.⁹ Tasks are classified into three exposure categories: directly exposed (scored as 1), partially exposed (scored as 0.5), and not exposed (scored as 0). The occupation-level AI exposure score X_o is then computed as the

⁷ The literature documents that price adjustments are infrequent but tend to be large when they occur, typically 10 to 15 percent in U.S. micro data (e.g., [Klenow and Kryvtsov, 2008](#); [Klenow and Malin, 2010](#)). At the same time, very small price changes may reflect measurement error ([Nakamura and Steinsson, 2008](#); [Eichenbaum, Jaimovich, and Rebelo, 2011](#)). A 5% threshold is conservative relative to the typical magnitude of price adjustments documented in the literature. It is also large enough to exclude small fluctuations that may reflect measurement error. Our results are robust to using a 10% threshold instead.

⁸ Results using $\ln(\textit{Average price})$ are robust to deflating prices using category-specific CPIs rather than the aggregate CPI. These results are reported in [Table IA.B.3](#).

⁹ We develop a similar task scoring approach where we specify the capabilities of a target model and use GPT-5-mini as the scoring model. We use this approach to obtain an AI exposure measure based on the capabilities of InstructGPT as of January 2022. We show in [Internet Appendix B](#) that all three sets of baseline results—price dispersion, repricing frequency, and price levels—are robust to using this alternative exposure measure ([Table IA.B.5](#), [Table IA.B.6](#), and [Table IA.B.7](#)).

simple average of these scores across all tasks within occupation o .

To construct firm-level AI exposure, we combine the occupation-level scores with workforce composition data from Revelio Labs, which provides individual-level employment information with Standard Occupational Classification (SOC) codes, start and end dates, and firm identifiers. Following [Eisfeldt et al. \(2025\)](#), we compute firm-level exposure as the employment-share-weighted average of occupation-level scores:

$$\text{AI Exposure}_{it} = \sum_o S_{iot} \cdot X_o, \tag{1}$$

where S_{iot} is the share of firm i 's employees in occupation o at time t and X_o is the fixed occupation-level AI score. [Figure 2](#) plots the cross-sectional distribution of firm-level AI exposure measured in January 2020. The distribution is approximately bell-shaped and concentrated in a narrow range. The interquartile range spans 0.35 to 0.42, and the standard deviation is 0.06 (Panel E of [Table 1](#)). This tight clustering reflects the large, diversified firms in our NielsenIQ sample. Since these firms employ workers across many occupations, the firm-level scores are averaged over a broad mix of tasks and converge toward the cross-firm mean. Despite this relatively tight distribution, there is meaningful variation. A firm at the 25th percentile has an AI exposure score of 0.35, compared to 0.42 at the 75th percentile, a difference of roughly one standard deviation.

We use AI exposure as a proxy for AI implementation, similar to the use of occupation-based measures as proxies for technology adoption in the broader economics literature (e.g., [Autor, Levy, and Murnane, 2003](#); [Acemoglu and Restrepo, 2022](#)). The logic is that firms whose workforces are concentrated in occupations where generative AI can meaningfully reduce task completion times are more likely to benefit from adopting these tools and, consequently, more likely to deploy them. We validate this assumption directly in [Section 10](#) by showing that high-exposure firms significantly increase their use of AI-related language in job postings and SEC 10-K filings after the InstructGPT API introduction, consistent with increased AI adoption.

We further decompose AI exposure into augmentation and automation components. We use GPT-5-mini to score each O*NET 30.1 task along two dimensions, augmentation and automation. The augmentation score captures the extent to which AI assistance increases worker productivity

while still requiring human judgment. The automation score captures the likelihood that AI completes the task independently with minimal human involvement.¹⁰ We map task scores into three categories (high, medium, and low) and aggregate them to the occupation and firm levels using the same employment-share weights as the overall exposure measure.¹¹

We also obtain a firm-level measure of data assets from [Eisfeldt et al. \(2025\)](#), which captures the extent to which a firm’s operations rely on data inputs that generative AI can process. The measure has two components. The first is the share of a firm’s workforce in occupations that work intensively with data, computed from O*NET task descriptions. The second is a GPT-scored measure of data-asset references in firms’ 10-K filings.

3.4 Other data

We obtain quarterly financial data from Compustat and data on the total number of employees from Revelio Labs. All firm-level controls are measured in the quarter before the introduction of the InstructGPT API in 2022 Q1. Panel F of [Table 1](#) reports summary statistics. Relative to the broader Compustat sample, sample firms tend to be larger, older, and more capital-intensive. For example, the median sample firm has total assets of approximately \$8 billion compared to \$1.1 billion in the unrestricted Compustat sample, and a median age of 35 years compared to 16 years. Sample firms also hold less cash (median cash-to-assets of 7% compared to 12%) and have a higher share of fixed assets (17% compared to 8%). This is not surprising given that the NielsenIQ sample consists of large firms that manufacture physical consumer products. Profitability is similar across samples, with a median return on assets (ROA) near zero in both groups.

We also obtain data on county economic conditions and firm- and county-level pandemic exposure. First, we use monthly county-level unemployment rates from the Bureau of Labor Statistics

¹⁰ Internet Appendix [IA.A.1](#) provides the system prompt and notes that the scoring was conducted in December 2025 with temperature zero. [Eisfeldt et al. \(2025\)](#) also distinguish augmentation from automation, but classify each occupation’s tasks as either core (a proxy for automation) or supplementary (a proxy for augmentation) based on O*NET categories. We instead score every task along both dimensions simultaneously.

¹¹ As an alternative, we use the Anthropic Economic Index (AEI, [Appel, Massenkoff, McCrory, McCain, Heller, Neylon, and Tamkin, 2026](#)), which maps human–AI conversations to O*NET work activities and classifies the mode of interaction between the user and the model. We proxy augmentation by the *task iteration* mode, in which users iteratively refine outputs with the AI model, and automation by the *directive* mode, in which users issue single instructions that the AI executes independently. Firm-level scores aggregate occupation-level AEI scores using the same Revelio employment-share weights. [Table IA.B.11](#) in the Internet Appendix reports the corresponding results.

(BLS) Local Area Unemployment Statistics. Second, we construct a county-level measure of interest rate (IR) exposure using the BLS Quarterly Census of Employment and Wages (QCEW). Specifically, for each county-month, we compute the share of total private employment in interest-rate-sensitive industries, including construction (NAICS 23), real estate (NAICS 531), and finance and insurance (NAICS 52).

For firm-level pandemic exposure, we follow [Barry, Campello, Graham, and Ma \(2022\)](#) and construct occupation-based measures that capture the extent to which a firm’s workforce relies on physical proximity. Specifically, we use the work-from-home feasibility index of [Dingel and Neiman \(2020\)](#) and an analogous measure constructed from the American Time Use Survey. Both measures are aggregated to the firm level using pre-pandemic employment shares and are fixed prior to the onset of COVID-19. At the county level, we measure local pandemic intensity using cumulative COVID-19 cases and deaths from the *New York Times* county-level COVID-19 database.

Finally, to assess whether AI exposure translated into actual implementation, we collect data on AI-related hiring and AI discussion in corporate disclosures. We obtain job posting data from LinkUp and textual data from annual 10-K filings from the Securities and Exchange Commission (SEC) EDGAR database.

4 Empirical Roadmap

Our empirical analysis proceeds in stages. We begin by decomposing pricing variation across manufacturers, retailers, and geography to locate the systematic component of retail price-setting within the vertical chain. We then use brand acquisitions to test whether pricing behavior changes with upstream ownership. Next, we compare private-label and national-brand pricing to examine whether attenuating the manufacturer-retailer information wedge increases price localization. Finally, we use a difference-in-differences design around the January 2022 introduction of OpenAI’s Instruct-GPT API to test whether a decrease in upstream information-processing costs shifts pricing away from uniformity.

5 Price Variance Decomposition

We begin by decomposing the variance of retail prices across manufacturers, retailers, and geography to examine where in the vertical chain the systematic component of pricing resides. We apply Shapley value decompositions of R^2 across nested fixed-effects specifications to the full manufacturer-linked NielsenIQ scanner dataset at the store \times UPC \times month level.

We estimate models with progressively richer sets of fixed effects (UPC \times month, retailer, county, and their interactions) and attribute each factor’s contribution to explained variance using the Shapley formula, which averages marginal contributions over all orderings. We report three variants: a product model, a manufacturer model, and a bilateral model following [DellaVigna and Gentzkow \(2019\)](#) that includes a retailer \times UPC interaction to capture product-specific bilateral pricing relationships. We also report a four-factor Shapley decomposition following [Couts, Gonçalves, and Loudis \(2026\)](#), which generalizes the standard Shapley R^2 decomposition to settings with interactive fixed effects. This approach treats natural interactions as single factors and preserves the standard Shapley properties.

[Table 3](#) presents results from Shapley decompositions of log prices, squared deviations from UPC \times month means, and month-over-month log price changes using the full manufacturer-linked NielsenIQ scanner dataset.

Panel A decomposes the variance of log prices. The product Shapley decomposition in column (1) uses UPC \times month as the baseline factor and attributes the remaining explained variance to retailer and county. We find that UPC \times month accounts for 99.0% of the explained variance, retailer for 0.9%, and county for 0.1%, with an overall R^2 of 0.899. Therefore, most variation in retail prices is product and time specific. The Shapley decomposition in column (2) takes month as the baseline factor and attributes the residual variance to manufacturer, retailer, and county. We find that manufacturers account for 91.4% of the explained variance, retailers 7.6%, and counties 1.0%, with an R^2 of 0.368. Column (3) shows results from a bilateral Shapley decomposition following [DellaVigna and Gentzkow \(2019\)](#), which extends the product Shapley decomposition in column (1) by adding a retailer \times UPC interaction that captures product-specific bilateral pricing relationships. The interaction absorbs 47.8% of the explained variance, and the UPC \times month share

drops to 48.4% as the bilateral term takes over part of the product-level variation. The marginal R^2 gain from adding the bilateral term is 1.0 percentage points (0.909 versus 0.899). The 47.8% Shapley share thus primarily reflects the reallocation of product-level variation rather than a large independent contribution of bilateral pricing relationships.

When manufacturers depart from a uniform per-product, per-month price, they do so at similar rates across retailers and across counties. Column (4) reports a Shapley decomposition that attributes variance to UPC \times month, manufacturer \times retailer, and manufacturer \times county fixed effects, isolating these manufacturer-specific departures. We find that UPC \times month accounts for 71.8% of the explained variance, manufacturer \times retailer for 16.0%, and manufacturer \times county for 12.2%, with an R^2 of 0.902. The remaining 28%, after accounting for the product-and-time baseline, captures the systematic component of manufacturer-specific localization. The retailer and geographic channels contribute comparable shares (16.0% and 12.2%), consistent with manufacturers responding to local market conditions through both retailer-specific and county-specific channels when they depart from uniform pricing.

Column (5) presents results from a four-factor Shapley decomposition following [Couts, Gonçalves, and Loudis \(2026\)](#). Specifically, we estimate the explanatory power of UPC, retailer, county, and month factors using fully interactive fixed effects in each subset. In this case, the R^2 is 0.954, and the UPC, retailer, month, and county shares are 88.8%, 5.9%, 3.4%, and 1.8%, respectively. Again, most variation in retail prices is product-specific.

Panel B decomposes squared deviations from UPC \times month means and thus focuses on explaining cross-store dispersion after holding the product-month price fixed. The first result is that little of this dispersion is systematic. Column (1) has an R^2 of 0.075, compared with 0.899 for price levels in Panel A, so that most cross-store deviations around a product-month price are residual to product-time, retailer, and county factors. Within the explained component, the manufacturer share remains the largest source of variation. In column (2), manufacturers account for 56.0% of the explained variance, retailers for 36.5%, and counties for 7.5%. The larger retailer share relative to the levels specification in Panel A suggests that retailer-mediated frictions are more relevant for dispersion than for the baseline price and help explain how store prices deviate from

the product-month baseline set upstream.

Columns (3) and (4) show that the retailer contribution operates through bilateral relationships rather than broad county-level price schedules. In column (3), adding retailer \times UPC increases the R^2 from 0.075 to 0.117, and the interaction accounts for 48.0% of explained variance. Column (4) separates manufacturer-specific localization across retailers and counties. The manufacturer \times retailer and manufacturer \times county factors account for 21.7% and 4.0% of explained variance, respectively, while the UPC \times month factor still accounts for 74.3%. Panel B, therefore, suggests that when prices differ across stores, the systematic part of that dispersion is mainly mediated by manufacturer-retailer relationships rather than by county-specific differentiation.

Panel C decomposes month-over-month log price changes. The systematic component of repricing is small but concentrated upstream. Column (1) shows that UPC \times month accounts for 99.9% of explained variation, with an R^2 of 0.081, which indicates that systematic price changes are product-time specific. Column (2) replaces the product factor with a manufacturer factor. Manufacturers account for 97.1% of explained variation, while retailers and counties account for 2.2% and 0.6%, respectively. Column (3) adds a retailer \times UPC factor, which accounts for only 7.4% of explained variation. This share is much smaller than in Panel B, so that retailer-product relationships explain dispersion more than they do repricing. Column (4) shows that manufacturer \times month accounts for 72.8% of explained variation, relative to 21.7% for retailer \times month and 5.5% for county \times month. Overall, the results suggest that when repricing has a systematic component, it is concentrated in manufacturer-time variation, with a smaller retailer role and little direct geographic component.

Taken together, the decompositions reveal a clear vertical structure. The baseline level of prices and the timing of price changes are primarily upstream, while the limited systematic dispersion across stores runs through manufacturer-retailer relationships. This pattern is consistent with manufacturers facing local information frictions when pricing through retailers.¹²

¹² Table IA.B.1 in the Internet Appendix reports an additive AKM decomposition that complements the Shapley results, and Table IA.B.2 replicates the decomposition separately for large and small manufacturers.

6 Brand Acquisitions

The variance decomposition shows that manufacturer identity is the largest source of systematic pricing variation. Still, it cannot rule out the possibility that this association reflects fixed product characteristics rather than active manufacturer influence on pricing. We therefore use brand acquisitions to test whether pricing behavior changes with product ownership. Acquisitions reassign upstream ownership from one manufacturer to another without altering the underlying physical product or the retail environment that sells it. If the systematic component of pricing originates upstream, acquired UPCs should depart from their pre-acquisition pricing patterns when ownership changes.

For each of 10 brand acquisitions by publicly traded manufacturers between 2012 and 2021, we form a treatment cohort consisting of the acquired brand’s UPCs and stack the cohorts on event time. The unit of observation is $\text{UPC} \times \text{retailer} \times \text{stack} \times \text{month}$, with a twelve-month event window on either side of each acquisition close date. The specification is

$$Y_{urst} = \alpha_{urs} + \alpha_t + \beta \cdot \text{Post}_{st} + \varepsilon_{urst}, \quad (2)$$

where u indexes the UPC, r the retailer, s the cohort, and t the calendar month. Post_{st} equals one in months on or after cohort s ’s acquisition close. We include $\text{UPC} \times \text{retailer} \times \text{stack}$ and calendar-month fixed effects, so identification comes from changes in the pricing behavior of the same UPC-retailer pair around the acquisition date relative to common calendar-month shocks. Standard errors are clustered at the $\text{UPC} \times \text{stack}$ level, and we also report wild cluster bootstrap p -values clustered at the deal level.

We use three pricing-distance measures as dependent variables. The first is distance from the acquired brand’s own pre-period mean, $|\text{SD}_{urt} - \overline{\text{SD}}_{ur,\text{pre}}|$ and $|\text{Repr}_{urt} - \overline{\text{Repr}}_{ur,\text{pre}}|$, where SD is the within-retailer cross-store standard deviation of log prices and Repr is the indicator for an absolute monthly log-price change above five percent. The second and third are distances from a moving acquirer benchmark, $|\text{SD}_{urt} - \overline{\text{SD}}_{rmt}^{\text{acq}}|$ and $|\text{Repr}_{urt} - \overline{\text{Repr}}_{rmt}^{\text{acq}}|$, constructed from the acquirer’s pre-existing UPCs in the same retailer r , product module m , and month t . We split deals into

within-category, or synergistic, acquisitions, in which the acquirer had a legacy presence in the acquired brand’s modules, and expansionary acquisitions, in which the acquirer’s legacy products were only in adjacent modules.

A positive β on the own-pre-period distance indicates departure from the acquired brand’s pre-acquisition pricing pattern. A negative β in the synergistic split indicates convergence to the acquirer’s incumbent pricing routine. A positive β in the expansionary split indicates divergence from the acquirer’s broader legacy pricing.

Columns (1) and (2) of Table 4 measure the absolute distance between each acquired UPC’s outcome and its own pre-acquisition mean, for the within-retailer cross-store dispersion of log prices and the indicator for an absolute monthly log-price change above five percent. A positive coefficient indicates departure from the acquired brand’s pre-acquisition pricing. The post-acquisition coefficients are 0.016 and 0.064, with deal-clustered wild-bootstrap p -values below 0.001. Acquired UPCs therefore price differently after ownership changes than they did before, on both the dispersion and repricing margins.

Columns (3) and (4) focus on within-category, or synergistic, acquisitions, in which the acquirer was already selling in the same Nielsen product modules as the acquired brand.¹³ The dependent variable is the absolute distance between the acquired UPC’s outcome and the acquirer’s legacy mean in the same retailer and product module. A negative coefficient indicates that acquired UPCs become more similar to the acquirer’s incumbent pricing in the same product module. The post-acquisition coefficients are -0.007 for cross-store dispersion and -0.021 for repricing distance, with wild-bootstrap p -values of 0.075 and 0.092, respectively. Both results are consistent with convergence to the acquirer’s pricing behavior in the same product module.

Columns (5) and (6) focus on expansionary acquisitions, in which the acquirer’s legacy products were in the same product group but distinct product modules.¹⁴ The dependent variable is the absolute distance between the acquired UPC’s outcome and the acquirer’s legacy mean in the same retailer and product group. A positive coefficient would indicate that acquired UPCs price further

¹³ A product module is the most granular category in the NielsenIQ product taxonomy, and groups products with similar use together.

¹⁴ Product groups aggregate related product modules in the NielsenIQ taxonomy. Two products in the same product group but in different modules are in adjacent subcategories within the same broad category.

from the acquirer’s adjacent-module pricing than before. The post-acquisition coefficients are 0.006 for cross-store dispersion (wild $p = 0.063$) and 0.005 for repricing distance (wild $p = 0.109$). The dispersion result is significant at the 10 percent level under wild-bootstrap inference, while the repricing estimate is positive but less precise.¹⁵

These results are consistent with manufacturers operating portable, category-specific pricing routines. Acquired UPCs move toward the acquirer’s pricing routine when the acquirer has incumbent presence in the same product module. In expansionary acquisitions, the evidence of divergence is weaker and concentrated in cross-store dispersion. This asymmetry is informative because pricing behavior moves with ownership when relevant category expertise transfers, but not when the acquirer lacks such experience. The pattern is difficult to reconcile with retailer-driven pricing, which one would not expect to respond to ownership changes that hold both the product and retailer fixed. The acquisition results thus reinforce the decomposition by showing that upstream owners bring category-specific pricing routines to the products they control.

7 Private Labels vs. National Brands

The acquisition results show that pricing behavior moves with upstream ownership. We next examine whether the relevant manufacturer-side friction is informational by exploiting the vertical structure of private label pricing. Private labels are controlled and priced by retailers themselves, which substantially attenuates the manufacturer-retailer information asymmetry that characterizes national brand pricing while holding many implementation costs, such as shelf tags, scanner systems, and promotional calendars, roughly fixed. If upstream informational frictions are a first-order driver of uniform pricing, private label products should exhibit greater geographic price dispersion than national brands sold by the same retailer in the same product module and month. Several features of private label pricing work against this prediction. Private labels are often priced as a fixed discount off the category-leading national brand, inheriting any uniformity in that benchmark. They are also lower-margin products on which retailers may invest less in price optimization,

¹⁵ Columns (3) to (6) cover 9 of the 10 deals in columns (1) and (2). For the remaining deal, the acquirer had no legacy products to serve as a benchmark.

and they face direct competitive pressure from other value-tier brands, which may compress price variation.

We test these predictions in two panels with different units of observation. Panel A of [Table 5](#) regresses the within-UPC standard deviation of log prices across stores on a private label indicator at the $\text{UPC} \times \text{retailer} \times \text{county} \times \text{month}$ level. Panel B regresses log prices on the interaction of the private label indicator with county-level log average monthly wages at the $\text{store} \times \text{UPC} \times \text{month}$ level. Both panels include progressively tighter fixed-effects structures, and standard errors are clustered at the county level.

Despite several forces that work against localized private-label pricing, we find that private labels are substantially more dispersed. As reported in Panel C of [Table 1](#), the mean within-retailer standard deviation of log prices is 0.067 for private labels versus 0.057 for national brands, a difference of approximately 17 percent relative to the national brand mean. The same pattern is present throughout the distribution. At the 75th percentile, private label dispersion is 0.059 versus 0.050 for national brands, and at the median 0.014 versus 0.009. At the 25th percentile both groups exhibit no within-retailer cross-store dispersion, consistent with widespread uniform pricing. Panel A of [Table 5](#) shows that this higher private-label dispersion survives progressively tighter fixed-effects specifications. Column (1) includes product module \times month and retailer fixed effects; column (3) includes retailer \times product module \times month fixed effects, which absorb all retailer-product-module-time variation and isolate the private label effect within the same retailer, product module, and month. The coefficient on the private label indicator ranges from 0.005 to 0.009 and is significant at the one-percent level across all three specifications. The stability of the estimate as we tighten the comparison set suggests the result reflects a genuine difference in pricing behavior rather than compositional differences across retailers or product modules.

Private-label prices also respond more strongly to local economic conditions than national-brand prices. Panel B of [Table 5](#) regresses the log of price on the interaction of the private label indicator with the log of county-level monthly wages, at the store-UPC-month level, restricting to county-months with at least 5,000 private-sector employees in the BLS Quarterly Census of Employment and Wages. The interaction coefficient is 0.003 and not significant in column (1), with

retailer \times product module \times month fixed effects. It increases to 0.019 in column (2), significant at the one-percent level, when the fixed effects absorb retailer-product-county and calendar-month heterogeneity, and remains at 0.018 in column (3), when we add retailer \times product module \times month fixed effects. With the tighter fixed effects, the wage elasticity of national-brand prices is essentially zero, while private-label prices respond positively to local wage variation. The pattern is consistent with private-label prices being shaped by local conditions in a way that national-brand prices are not.

8 AI Experiment

The decomposition, acquisition, and private-label results point to upstream informational frictions as an important source of uniform pricing. We next examine whether reducing upstream information-processing costs shifts pricing away from uniformity. We exploit the January 2022 introduction of OpenAI’s InstructGPT API, the first generative AI model available for firm-level deployment at scale, and compare pricing outcomes for manufacturers with high versus low pre-existing AI exposure. Importantly, the InstructGPT API was a firm-facing technology event. As shown in Panel A of [Figure 3](#), Google Trends search interest for “artificial intelligence” remained flat throughout the estimation window, spiking only after the consumer-facing ChatGPT release in November 2022. The pricing effects we study thus predate the period in which AI became salient to the general public. If upstream informational frictions are a binding constraint on price localization, then a reduction in the cost of acquiring and processing demand information should shift pricing away from uniformity.

Our baseline difference-in-differences specification is as follows:

$$\begin{aligned}
 Y_{ipt} = & \alpha_{ip} + \alpha_t + \beta(Treated_i \times Post_t) + \delta(\ln(\bar{Q}_{ip,pre}) \times Post_t) \\
 & + \gamma'(\mathbf{X}_{i0} \times Post_t) + \varepsilon_{ipt},
 \end{aligned}
 \tag{3}$$

where i indexes firms, p products (UPCs), and t calendar months. The dependent variable Y_{ipt} is one of the following pricing outcomes: the standard deviation or interquartile range of log prices

within a firm-UPC-month, an indicator for whether the month-over-month change in log average price exceeds a given threshold at the level of aggregation used in the specification, or the log real average price. The first two outcomes capture localization and repricing, and the price-level outcome is used to study local responsiveness and product-level heterogeneity. $Treated_i$ equals one for firms above the median of AI exposure in January 2020. $Post_t$ equals one from January 2022 onward. The term $\bar{Q}_{ip,pre}$ denotes average pre-period units sold, and its interaction with $Post_t$ allows products with different baseline sales volumes to follow different post-period changes. The vector \mathbf{X}_{i0} includes pre-event firm characteristics: $\ln(Employees)$, $\ln(Assets)$, Tobin's q , ROA, cash-to-total assets, fixed assets-to-total assets, R&D-to-total assets, market-to-book, and $\ln(Age)$. Because these controls are fixed in the pre-period and interacted with $Post_t$, they allow for differential changes in pricing outcomes based on pre-treatment characteristics, so that β is not confounded by the possibility that, for example, larger or more profitable firms would have changed prices differently in the post-period regardless of AI exposure.

The specification includes firm-UPC fixed effects α_{ip} and calendar-month fixed effects α_t . Standard errors are clustered at the firm level. The coefficient of interest, β , captures the differential change in pricing outcomes for high-exposure firms after the InstructGPT introduction. We also estimate variations of this specification at the firm-UPC-county-month and firm-UPC-retailer-month levels, which allow us to exploit within-product geographic and retailer variation in pricing outcomes and to include finer fixed effects. The identifying assumption is that, conditional on the controls, treated and control firms would have followed parallel trends in the outcomes of interest absent the introduction.

We also balance the products of treated and control firms using propensity score matching on pre-event average log prices and log units sold. This ensures that the products entering the comparison have similar pre-treatment pricing levels and sales volumes, so that pre-existing differences in these characteristics do not drive the results. Panel G of [Table 1](#) reports the resulting covariate balance. Treated firms are somewhat larger and less capital-intensive than controls ($\ln(Assets)$ 9.4 versus 8.2 and fixed-assets ratio 0.13 versus 0.26), since we match on pre-event prices and units sold rather than on firm characteristics directly. The interactions of these characteristics with $Post_t$ in

Equation 3 absorb the residual differences.

To test for differential pre-trends and plot dynamic treatment effects, we estimate

$$\begin{aligned}
 Y_{ipt} = & \alpha_{ip} + \alpha_t + \sum_{k \in \{-6, \dots, +6\}, k \neq -3} \beta_k (\mathbb{1}[Rel_t = k] \times Treated_i) \\
 & + \delta (\ln(\bar{Q}_{ip,pre}) \times Post_t) + \gamma' (\mathbf{X}_{i0} \times Post_t) + \varepsilon_{ipt},
 \end{aligned} \tag{4}$$

where Rel_t is event time in months relative to January 2022. We omit $k = -3$ (October 2021) as the reference period so that the baseline precedes the November 18, 2021 expansion of GPT-3 API access, which could have allowed early experimentation before the January 27, 2022 InstructGPT introduction (see Table 2). This event-study design is estimated on the same window used in the baseline analysis. The pre-event coefficients provide a direct test of the parallel-trends assumption, while post-event coefficients characterize the timing and persistence of treatment effects.

8.1 Price dispersion

We first test whether AI relaxes the uniform-pricing constraint by examining cross-store price dispersion within a firm-UPC-month. If AI reduces the cost of gathering and processing local demand information, firms should tailor prices more precisely to local conditions, increasing the dispersion of prices across stores. Table 6 presents the results. The dependent variable in columns (1) to (3) is the standard deviation of log prices within a firm-product-month, and in columns (4) to (6) is the interquartile range. For each measure, we compute dispersion across all stores (columns (1) and (4)), within counties (columns (2) and (5)), and within retailers (columns (3) and (6)).

The treatment effect is positive and statistically significant across all six specifications. In column (1), the standard deviation of log prices increases by 0.025 after the event for high-exposure firms relative to low-exposure firms, corresponding to approximately a 12% increase relative to the mean of 0.206. We find a larger treatment effect when dispersion is measured across all stores (columns (1) and (4)) than within counties or within retailers (columns (2) to (3) and (5) to (6)), suggesting that much of the increase in price dispersion occurs across geographic areas and retail chains. That said, the within-county and within-retailer effects are also statistically significant,

suggesting that AI implementation also enables price changes within local markets. [Figure 4](#) plots the dynamic treatment effects over the estimation window. The pre-treatment coefficients are close to zero for both the standard deviation of log prices (Panel A) and the interquartile range of log prices (Panel B), consistent with no differential pre-trends in price dispersion.

Of the two channels in our pricing framework, the information channel predicts an increase in geographic price dispersion, while the production channel does not. The information channel operates through two margins: lower costs of generating pricing recommendations expand the set of markets for which the manufacturer transmits market-specific guidance (extensive margin), and higher signal precision increases the weight placed on local information relative to priors (intensive margin). Both margins increase geographic price dispersion. In the language of [DellaVigna and Gentzkow \(2019\)](#), our results are consistent with AI eroding the informational constraints that give rise to uniform pricing across heterogeneous markets.

8.2 Price changes

We next examine whether AI implementation affects the frequency with which firms adjust prices. [Table 7](#) presents the results. In columns (1) to (3), the dependent variable is an indicator for whether the absolute month-over-month change in log average price exceeds 5% at the level of aggregation used in the column. Columns (4) to (6) and (7) to (9) isolate price increases and decreases, respectively. For each dependent variable, we report results at the firm-product, firm-product-county, and firm-product-retailer levels.

Column (1) shows that the probability of an absolute change in log average price exceeding 5% increases by 6.4 percentage points for high-AI exposure firms relative to low-AI exposure firms after the introduction of the InstructGPT API. Relative to the sample mean of 0.337 reported in Panel D of [Table 1](#), the estimated treatment effect represents approximately a 19% increase in repricing frequency. Columns (2) and (3) show consistent results when we examine changes in repricing frequency within counties and retailers. In columns (4) to (6), we decompose the effect by direction, and show that the frequency of price increases rises significantly at the county level, but the effect is small and insignificant at the firm-product and firm-product-retailer levels. By contrast,

the frequency of price decreases increases by 4.6 percentage points at the firm-product level (column (7)) and by 3.5 percentage points at the firm-product-county level (column (8)), with both effects statistically significant. The overall increase in repricing frequency is thus predominantly, though not exclusively, driven by price decreases.

Figure 5 presents the estimated dynamic treatment effects. Panel A shows that the pre-treatment coefficients for the probability of an absolute change in log average price exceeding 5% are close to zero, with a clear increase after the event. The directional decomposition reinforces the regression results. Panel B shows an increase in the frequency of price increases at the county level, though the effect is more muted at the firm-product and retailer levels. Panel C shows a broader post-event increase in the frequency of price decreases across all three levels of aggregation.

These results help distinguish between the two channels. A pure production-cost channel predicts changes in price levels but no change in repricing frequency, and is therefore silent on the patterns in Table 7. The information channel predicts more frequent repricing, consistent with columns (1) through (3). Moreover, the increase in repricing is concentrated in price decreases, which are statistically significant at the firm-product and firm-product-county levels. This asymmetry is consistent with better demand signals helping identify markets where firms were previously overcharging relative to local willingness to pay. That is, these results are consistent with improved information revealing downward pricing opportunities that were previously too costly for the manufacturer to identify.

8.3 Price level and local economic conditions

Having documented an increase in cross-store dispersion and repricing frequency, we now turn to the average price level and whether prices become more sensitive to local economic conditions. Table 8 presents the results. Column (1) reports results from the specification in Equation 3, with $\ln(\textit{Average price})$ as the dependent variable. The estimate is -0.043 , suggesting that high-AI exposure firms reduced prices by approximately 4.3% relative to low-AI exposure firms after the InstructGPT API introduction. Columns (2) and (3) then estimate firm-product-county-month specifications that interact the treatment effect with county-level log average monthly wages. The

coefficient on $Treated \times Post \times \ln(Wage)$ is positive and statistically significant in both specifications, including column (3), where county-by-calendar-month fixed effects absorb local time-varying shocks. Since wages enter in logs, the magnitude is easiest to interpret by comparing counties with different wage levels. In column (3), a county with wages 10 percent higher has a treatment effect that is about 0.2 percentage points less negative, while a county with wages 50 percent higher has a treatment effect that is about 0.9 percentage points less negative. These patterns suggest that treated firms’ prices load more strongly on local wage conditions after the event, with more pronounced price declines in lower-wage counties.

Figure 6 presents the corresponding event-study estimates from Equation 4. The pre-treatment coefficients are close to zero and statistically insignificant, which supports the parallel-trends assumption. The treatment effect begins to emerge in the month before the introduction, consistent with some early-adopting firms integrating AI tools following the expansion of access to the GPT-3 API in November 2021. The treatment effect persists through the end of the estimation window.

The decline in average prices is consistent with both lower marginal production costs and improved local demand signals that reveal markets where firms were previously overpricing. The wage-interaction result in columns (2) and (3) distinguishes the two. A pure production-cost channel would shift the optimal price uniformly across markets, with no interaction with local conditions, while the information channel implies that better local signals let firms tilt prices toward where consumers can afford to pay more. The positive wage loading is therefore more consistent with AI enabling localized pricing than with a uniform production-cost decline.

In Section 9, we probe this interpretation further. The pricing effects should be largest where the returns to better demand information are highest, for example, for products with more volatile or more elastic demand, where the cost of mispricing is potentially higher.

8.4 Robustness tests

We subject these results to a battery of robustness tests addressing two classes of concern. The first is whether the effects reflect manufacturer-level responses to AI or retailer-level pricing strategies. The second is that the InstructGPT API introduction in January 2022 coincided with other

macroeconomic developments that could affect retail pricing. In particular, the Federal Reserve began raising interest rates in March 2022, the economy was still adjusting to the effects of the COVID-19 pandemic, and inflation differed across product categories. If AI exposure is correlated with firms' sensitivity to these shocks, the estimated treatment effects could reflect confounding macroeconomic forces rather than AI implementation. We address the retailer concern first and then each macroeconomic concern in turn. We also conduct robustness tests using alternative price deflators, alternative AI exposure measures, a continuous treatment measure, and several placebo tests.

Retailer-specific effects. A natural concern is that our results could reflect retailer-level pricing strategies rather than manufacturer-level responses to AI. [Table 9](#) addresses this by progressively adding retailer \times UPC and retailer \times calendar month fixed effects, which absorb time-invariant retailer-product pricing differences and retailer-specific time trends such as chain-wide promotional cycles. The estimated treatment effects retain the same sign and remain statistically significant across all three outcomes. Relative to the baseline estimates, the magnitudes are smaller in the retailer-level panel, especially for price dispersion and average prices. [Figure 7](#) decomposes the price-level treatment effect by retailer. The estimate is negative for virtually every retailer in the data, consistent with a common upstream component, albeit with heterogeneous pass-through across retail chains.

COVID-19 exposure. As shown in [Figure 3](#), the InstructGPT API introduction coincides with the tail end of the Omicron wave, and COVID-19 cases decline sharply through the estimation window. If firms with high AI exposure also had workforces that were more (or less) exposed to pandemic-related workplace disruptions, COVID-19 effects could confound our estimates. We address this concern by including controls for COVID-19 exposure at the firm- and county-levels and present the results in [Table IA.B.8](#). Panel A includes occupation-level measures of workplace physical proximity aggregated to the firm-level and measured in the month before treatment, including the [Dingel and Neiman \(2020\)](#) work-from-home index and an analogous measure from the American Time Use Survey, interacted with post. Panel B estimates the same specification using the

firm-product-county-month panel and controls for log cumulative county-level COVID-19 cases and deaths. Both panels report results for price dispersion, repricing frequency, and log average prices. The estimated treatment effects remain statistically significant and consistent with the baseline results across all specifications. The COVID-19 controls are generally small and not systematically significant.

Interest rate exposure. The Federal Reserve began raising interest rates in March 2022, which increased borrowing costs and likely dampened demand in interest rate-sensitive sectors such as construction, real estate, and finance during our sample period. If products sold by high-AI exposure firms are disproportionately sold in counties with high employment shares in these sectors, the price decline we document could be driven by weaker local demand. To address this concern, we estimate [Equation 3](#) with controls for county-level interest rate (IR) exposure. We measure IR exposure as the share of private employment in the construction, real estate, and finance industries (NAICS 23, 531, and 52), which we classify as interest-rate-sensitive, using BLS data. We control for both a time-varying and an ex-ante measure of IR exposure to account for potential endogenous changes in local industry composition. [Table IA.B.9](#) presents the results. Across all specifications, the estimated treatment effects remain statistically significant and consistent with the baseline estimates. The IR exposure controls are economically small, though statistically significant in some of the columns. The stability of the treatment effects provides reassurance that the results are not driven by differential exposure to the interest rate tightening cycle.

Category-specific inflation. A related concern is that the price-level result could reflect differential inflation across product categories rather than a within-product relative price response. [Table IA.B.3](#) addresses this concern by deflating prices with category-specific CPIs from the BLS, mapping NielsenIQ product groups to 26 detailed CPI categories. The price-level effect remains negative and statistically significant, and the wage-interaction coefficients remain positive and statistically significant. These results indicate that the price-level findings are not driven by treated firms selling in categories with different inflation dynamics around the event.

Alternative AI exposure measure. We also test the robustness of the results to an alternative measure of AI exposure. As discussed in Section 3, we develop a target-judge model framework where GPT-5-mini evaluates whether the text-davinci-001 model—whose capabilities more closely match those available at the time of the InstructGPT API introduction—can perform each O*NET task. [Table IA.B.5](#), [Table IA.B.6](#), and [Table IA.B.7](#) replicate the baseline results for price dispersion, repricing frequency, and price levels using this alternative exposure measure. Across all specifications, the results are consistent with the baseline, and the corresponding event-study estimates in [Figure IA.B.1](#), [Figure IA.B.2](#), and [Figure IA.B.3](#) display similar pre-trends and post-event dynamics.

We also verify that the results are not driven by the propensity score matching procedure or the binary treatment classification. [Table IA.B.4](#) reports results using the full unmatched sample with the binary treatment indicator and the matched sample with continuous AI exposure measured in January 2020, interacted with post. The binary treatment effects in the unmatched sample are similar to those in the matched sample. Similarly, we find that all three effects are large and statistically significant using a continuous measure of AI exposure. Overall, the results confirm that the pricing effects scale with the intensity of AI exposure and are not artifacts of the propensity score matched sample or the binary treatment classification.

Placebo event dates. If our results reflect a genuine treatment effect of AI implementation, they should not appear when the event date is moved to a period with no plausible AI shock. [Table IA.B.10](#) presents placebo tests using January 2019 (36 months before the actual event) and January 2021 (12 months before) as alternative event dates. For each placebo date, we re-estimate the baseline specification for all three main outcomes: price dispersion, the probability of an absolute change in log average price exceeding 5%, and log average prices. The estimated treatment effects using these placebo dates are small, have inconsistent signs, and are statistically insignificant across all specifications. The absence of any effect at these placebo dates further supports the interpretation that our main results are not driven by pre-existing differential trends between high- and low-AI exposure firms.

We further assess the statistical significance of our results using a randomization inference

procedure. In each of 5,000 iterations, we exclude three extreme-portfolio outlier firms, randomly reassign treatment across the remaining firms, re-estimate the propensity score matching, and re-estimate the baseline specification in [Equation 3](#). [Figure IA.B.4](#) plots the distribution of placebo t -statistics for all three main outcomes. The true t -statistics, indicated by the dashed vertical lines, fall well into the tails of the permutation distributions. The two-sided permutation p -values are 0.015 for price dispersion (Panel A), 0.009 for repricing frequency (Panel B), and 0.011 for log average prices (Panel C). These results suggest that the estimated treatment effects are unlikely to be due to chance treatment assignment.

Sample composition and inference. We assess whether the results are driven by particular manufacturers or product categories. [Figure IA.B.7](#) reports leave-one-out estimates that drop each of the 91 manufacturers in turn and re-estimate the baseline specification, and [Figure IA.B.8](#) reports the analogous exercise dropping one of 100 NielsenIQ product groups at a time. The point estimates remain similar in magnitude across all leave-one-out samples for the three main outcomes, showing that the results are not driven by any single manufacturer or category. [Table IA.B.13](#) reports wild cluster bootstrap p -values clustered at the firm level, which remain consistent with the asymptotic inference.

9 Mechanism: Suggestive Evidence and Heterogeneity

The results in [Section 8](#) are consistent with a decrease in upstream informational frictions, but do not, on their own, isolate the underlying mechanism. Several of the observed patterns, including changes in average price levels, can arise through multiple channels depending on the environment. We therefore provide additional evidence by examining how the effects vary across product characteristics and across different dimensions of AI exposure.

These tests focus on cross-sectional predictions of the information channel. If AI primarily affects firms' ability to gather and process demand information, the effects should be concentrated in settings where the value of information is highest. We should also expect the effects to operate through the augmentation of worker decision-making rather than through the automation of routine

tasks. We evaluate these predictions in turn.

9.1 Product-level heterogeneity

We next investigate whether the price-level treatment effect varies across product categories. Under the information channel, the value of improved demand signals should be highest when demand conditions are more uncertain or when consumers are more price-sensitive. In such cases, the cost of mispricing is higher, and firms have stronger incentives to adjust prices in response to better information. We therefore expect the price-level effect to be concentrated in product categories where the returns to improved demand information are highest.

Specifically, we interact $Treated \times Post$ with indicator variables for discretionary products, products with high demand volatility, and products with high price elasticity. Using NielsenIQ product group codes, we classify products as discretionary when they are plausibly more substitutable or non-essential, such as alcohol and consumer electronics. The intuition is that demand for these products is more sensitive to changes in relative prices and local conditions, increasing the returns to improved demand information. We measure demand volatility using the coefficient of variation of pre-treatment units sold within each product module. We estimate price elasticities from pre-treatment regressions of log quantity changes on log price changes within product modules, and classify products as high-elasticity if the absolute value of their estimated elasticity is above the median.

Table 10 reports the results. The estimated coefficients on the triple interaction terms are negative and statistically significant in all three specifications. The estimates imply additional price declines of 3.8 percent for discretionary products, 2.5 percent for products with above-median demand volatility, and 2.3 percent for products with above-median price elasticity. Overall, these cross-sectional patterns are consistent with the information channel. Under that channel, we should expect the treatment effect to be concentrated in products with more uncertain or more elastic demand, where the cost of mispricing is higher. These are the settings where improved demand signals from AI adoption should be most valuable.

9.2 The information channel

The previous results are consistent with AI adoption improving firms’ ability to process and use demand information (Autor, 2015; Acemoglu and Restrepo, 2022; Autor and Thompson, 2025). To distinguish this information channel from alternatives in which AI directly automates pricing or simply lowers production costs, we test three additional predictions. First, the effects should concentrate in workers whose tasks involve pricing decisions rather than in workers exposed to AI in unrelated functions. Second, the effects should load on AI exposure that augments worker decision-making rather than on exposure that automates routine tasks. Third, the effects should be stronger in firms with the data assets that AI tools can leverage. Table 11 reports the three tests in three separate panels.

Panel A reports horse-race regressions that include pricing-task AI exposure and non-pricing-task AI exposure simultaneously. Pricing-task exposure is the employment-share-weighted average of each occupation’s AI score multiplied by the share of its O*NET tasks classified as pricing-related, identified by keywords such as price, cost, demand, forecast, promotion, negotiate, revenue, and margin. Non-pricing-task exposure is defined analogously using one minus the pricing-task share. Pricing-task exposure predicts more frequent repricing and lower average prices, with coefficients of 2.619 and -1.414 , respectively. The dispersion coefficient is positive but imprecise. Non-pricing-task exposure is statistically insignificant in all three specifications, and the pricing and non-pricing coefficients differ at the 10 percent level for repricing frequency and price levels. These results suggest that the repricing and price-level effects of AI exposure are concentrated in workers whose tasks involve pricing decisions.¹⁶

Panel B examines whether the effects load on augmentation exposure or automation exposure. We score each O*NET 30.1 task on both dimensions using GPT-5-mini, classify firms as having high augmentation or high automation exposure based on the median, and run separate regressions for each treatment. Firms with high augmentation exposure have 0.019 higher dispersion, a 4.7 percentage point increase in repricing frequency, and prices that are 3.9 percent lower after the

¹⁶ Table IA.B.14 in the Internet Appendix reports a placebo using general technology task intensity (software, database, network, programming, and hardware keywords, excluding pricing-related terms). Technology-task intensity is not associated with price dispersion or average prices, although it predicts repricing frequency.

event. Automation exposure has smaller and statistically insignificant coefficients in all three specifications. The augmentation result is consistent with AI strengthening employees' ability to interpret and act on demand information rather than automating routine pricing tasks.¹⁷

Panel C separates firms by their data assets exposure, a firm-level measure from [Eisfeldt et al. \(2025\)](#) that captures whether a firm's operations rely on data inputs that generative AI can process. We classify firms as high or low data-assets exposure based on the median. The dispersion and price-level effects are concentrated in high-data-assets firms. Dispersion increases by 0.025 and prices decline by 4.6 percent, while the low-data-assets coefficients are small and statistically insignificant in these specifications. Repricing frequency rises in both groups, by 6.6 percentage points for high-data-assets firms and 9.1 percentage points for low-data-assets firms. The price and dispersion patterns are consistent with AI's pricing effects being concentrated in firms whose operations generate the data inputs AI tools require.

10 Evidence of AI adoption

Our identification strategy exploits cross-sectional variation in pre-existing workforce exposure to generative AI as a proxy for the likelihood of adopting AI following the InstructGPT API introduction on January 27, 2022. It is thus important to examine whether high-AI exposure firms, as identified by our measure, actually exhibit higher AI adoption. Although AI adoption is difficult to measure, we provide evidence from several tests that is consistent with higher AI adoption among treated firms after the event.

We proxy for AI adoption in two complementary ways. First, we match firms in our sample to job posting data from LinkUp, which provides broad coverage of online vacancies. We then classify each job posting into five AI-related categories based on keywords in job titles and descriptions. Specifically, we create indicator variables that equal one for jobs we classify as broadly related to AI, applied AI, which includes jobs in predictive analytics, data science, and postings with O*NET occupational classifications identified as AI-related. We also create an indicator variable that is the

¹⁷ [Table IA.B.11](#) in the Internet Appendix reports an alternative augmentation vs. automation decomposition based on the Anthropic Economic Index, in which augmentation is proxied by the task iteration interaction mode and automation by the directive mode.

union of these four variables.¹⁸

Second, using textual data from firms' 10-K annual filings obtained from the SEC's EDGAR database, we measure the frequency of AI-related terms in corporate disclosures. Public filings provide a standardized way for firms to discuss material investments, operational changes, and emerging risks, making them a natural source of data to capture changes in AI strategy and implementation. Specifically, we count the frequency of AI-related terms such as artificial intelligence, machine learning, and automation in the full filing and separately within Item 1 (Business Description), Item 1A (Risk Factors), and Item 7 (Management's Discussion and Analysis). We also distinguish between AI adoption and implementation from discussions of AI risks by separating keywords that may reflect AI adoption and hiring from those that may reflect discussion of AI-related risks and regulation.

Table 12 reports difference-in-differences estimates of the change in AI adoption proxies for high-exposure relative to low-exposure firms after the InstructGPT API introduction. Panel A examines AI-related job postings, Panel B examines AI keyword counts in 10-K filings, and Panel C decomposes the Panel B counts by keyword type. In Panel A, we find a statistically significant increase in the monthly counts of AI-active job postings across all the different measures. High-exposure firms post approximately 8 more AI positions (column (1)), 4.4 more Applied AI and Data science positions each (columns (2) and (3)), and 22 more positions in the All AI category (column (5)) per month relative to low-exposure firms after the introduction of the InstructGPT API. The coefficients are also economically large. Overall, these findings suggest a reallocation of hiring toward AI-related skills in the first half of 2022, consistent with an increase in AI adoption among treated firms.

In column (1) of Panel B, we find that high-AI exposure firms mention AI-related keywords 1.5 more times per 10-K filing after the event. Columns (2) and (3) show that this increase is concentrated in Items 1 and 1A. The concentration of mentions in these sections suggests that AI is discussed in relation to how the firm operates and the risks it faces. However, we find no significant change in column (4) for Item 7, Management's Discussion and Analysis.

¹⁸ Appendix A describes the keywords and classification criteria for each indicator variable.

Panel C decomposes the Panel B keyword counts into implementation and hiring keywords versus risk and regulatory keywords. The implementation and hiring category captures keywords related to AI adoption, deployment, and workforce investment, whereas the risk and regulatory category captures keywords related to AI-related threats, compliance, and uncertainty. The increase in AI-related discussion in public disclosures is stronger for implementation and hiring terms compared to risk and regulatory terms. For example, the estimated treatment effect for total implementation keywords is 1.03 in column (1), compared with 0.08 for risk and regulatory keywords in column (4). The increase in AI-related discussion is therefore driven primarily by adoption and deployment language, with a much smaller increase in risk and regulatory language. Overall, we find that treated firms discuss AI more often in their public disclosures after the event, particularly in the context of adoption.¹⁹

Figure 8 presents the dynamic treatment effects corresponding to column (1) of Table 12, Panels A and B. Panel A plots the dynamic treatment effects on AI-related job postings from LinkUp at the monthly frequency, and Panel B plots the dynamic treatment effects on 10-K AI keyword counts at the annual frequency. In both panels, the pre-event coefficients are close to zero and statistically insignificant. The post-event coefficients indicate a persistent increase in AI-related job postings and AI-related disclosure following the introduction of the InstructGPT API.

To verify that these patterns reflect AI-specific adoption rather than a broader increase in technology-related discussion, we conduct a placebo test using non-AI technology keywords (e.g., cloud computing, ERP, CRM, cybersecurity, e-commerce infrastructure). Figure IA.B.6 in the Internet Appendix shows no differential change in the use of these terms in earnings calls or 10-K filings following the InstructGPT API rollout, suggesting that the increase in AI-related language is specific to artificial intelligence rather than part of a more general rise in technology-related discussion. Taken together, the results in this section are consistent with our AI exposure measure capturing meaningful differences in AI adoption.

¹⁹ We also examine the use of AI-related keywords in quarterly earnings conference call transcripts. Although conference call discussions are voluntary, they complement the 10-K analysis. Figure IA.B.5 and Table IA.B.12 in the Internet Appendix report these results. We find that treated firms significantly increase the use of AI keywords in earnings calls after the event.

11 Conclusion

The uniform pricing puzzle has traditionally been framed as a question about retail operations: why don't chains differentiate prices across locations when local conditions differ so markedly? This paper provides evidence that the answer lies upstream. Variance decompositions across approximately 5 billion store-product-month observations show that manufacturer identity accounts for approximately 90% of systematic price level variation and 97% of price change variation; for price dispersion, the manufacturer share is approximately 56%, with retailers contributing approximately 37%. Brand acquisitions—in which ownership of a UPC changes while the product and retailer remain fixed—show that pricing behavior changes with ownership, with acquired UPCs adopting the acquirer's incumbent pricing routine when the acquirer was already operating in the same product category. Private label products—which bypass the manufacturer-retailer information channel—exhibit significantly greater geographic price dispersion than national brands within the same retailer, category, and month. And when upstream informational frictions are reduced by the introduction of generative AI tools, geographic dispersion increases, repricing becomes more frequent, and prices fall, with stronger responsiveness to local economic conditions. These effects load on augmentation rather than automation exposure and survive retailer-by-time fixed effects, consistent with AI improving the flow of demand information through the vertical chain rather than automating retail execution.

The central implication is that the pricing rigidities documented in the literature are not immutable features of retail markets. They reflect, in substantial part, the cost of generating and communicating market-specific pricing information upstream—a cost that technology can reduce. The ongoing adoption of digital shelf labels by major retailers provides a natural test of this interpretation: if downstream implementation costs are the primary constraint on price localization, these technologies should substantially increase geographic price dispersion; if the bottleneck is upstream, as our evidence suggests, their effect should be more limited. Future work with richer cost data, retailer-level exposure measures, and structural estimation of the vertical pricing relationship could quantify the welfare consequences of uniform pricing and the gains from relaxing these frictions at scale.

References

- Abowd, J. M., F. Kramarz, and D. N. Margolis. 1999. High wage workers and high wage firms. *Econometrica* 67:251–333.
- Acemoglu, D., and P. Restrepo. 2022. Tasks, automation, and the rise in U.S. wage inequality. *Econometrica* 90:1973–2016.
- Agrawal, A., J. S. Gans, and A. Goldfarb. 2019. Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives* 33:31–50.
- Alvarez, F., F. Lippi, and L. Paciello. 2011. Optimal price setting with observation and menu costs. *Quarterly Journal of Economics* 126:1909–60.
- Appel, R., M. Massenkoff, P. McCrory, M. McCain, R. Heller, T. Neylon, and A. Tamkin. 2026. Anthropic economic index report: Economic primitives. Working Paper.
- Assad, S., R. Clark, D. Ershov, and L. Xu. 2024. Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *Journal of Political Economy* 132:723–71.
- Ater, I., and O. Rigbi. 2023. Price transparency, media, and informative advertising. *American Economic Journal: Microeconomics* 15:1–29.
- Autor, D., and N. Thompson. 2025. Expertise. *Journal of the European Economic Association* 23:1203–71.
- Autor, D. H. 2015. Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives* 29:3–30.
- Autor, D. H., F. Levy, and R. J. Murnane. 2003. The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics* 118:1279–333.
- Babina, T., A. Fedyk, A. He, and J. Hodson. 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151:103745–.
- Barry, J. W., M. Campello, J. R. Graham, and Y. Ma. 2022. Corporate flexibility in a time of crisis. *Journal of Financial Economics* 144:780–806.
- Bergemann, D., and A. Bonatti. 2019. Markets for information: An introduction. *Annual Review of Economics* 11:85–107.
- Bonnet, C., and P. Dubois. 2010. Inference on vertical contracts between manufacturers and retailers allowing for nonlinear pricing and resale price maintenance. *RAND Journal of Economics* 41:139–64.
- Brynjolfsson, E., D. Li, and L. Raymond. 2025. Generative AI at work. *Quarterly Journal of Economics* 140:889–942.
- Brynjolfsson, E., D. Rock, and C. Syverson. 2021. The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics* 13:333–72.

- Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello. 2020. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review* 110:3267–97.
- Card, D., J. Heining, and P. Kline. 2013. Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics* 128:967–1015.
- Chintagunta, P. K., A. Bonfrer, and I. Song. 2002. Investigating the effects of store-brand introduction on retailer demand and pricing behavior. *Management Science* 48:1242–67.
- Couts, S. J., A. S. Gonçalves, and J. A. Loudis. 2026. The subjective risk and return expectations of institutional investors. Working paper.
- Décaire, P. H., and D. Sosyura. 2024. On a spending spree: The real effects of heuristics in managerial budgets. Working paper.
- DellaVigna, S., and M. Gentzkow. 2019. Uniform pricing in U.S. retail chains. *Quarterly Journal of Economics* 134:2011–84.
- Dingel, J. I., and B. Neiman. 2020. How many jobs can be done at home? *Journal of Public Economics* 189:104235–.
- Dubé, J.-P., and S. Misra. 2023. Personalized pricing and consumer welfare. *Journal of Political Economy* 131:131–89.
- Eichenbaum, M., N. Jaimovich, and S. Rebelo. 2011. Reference prices, costs, and nominal rigidities. *American Economic Review* 101:234–62.
- Eisfeldt, A., G. Schubert, B. Taska, and M. B. Zhang. 2025. Generative AI and firm values. *Journal of Finance*, forthcoming .
- Goldfarb, A., and C. Tucker. 2019. Digital economics. *Journal of Economic Literature* 57:3–43.
- Golosov, M., and R. E. Lucas. 2007. Menu costs and Phillips curves. *Journal of Political Economy* 115:171–99.
- Hoch, S. J., and S. Banerji. 1993. When do private labels succeed? *Sloan Management Review* 34:57–67.
- Hoch, S. J., B.-D. Kim, A. L. Montgomery, and P. E. Rossi. 1995. Determinants of store-level price elasticity. *Journal of Marketing Research* 32:17–29.
- Klenow, P. J., and O. Kryvtsov. 2008. State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? *Quarterly Journal of Economics* 123:863–904.
- Klenow, P. J., and B. A. Malin. 2010. Microeconomic evidence on price-setting. In B. M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, vol. 3, 231–84. Elsevier.
- Levy, D., M. Bergen, S. Dutta, and R. Venable. 1997. The magnitude of menu costs: Direct evidence from large U.S. supermarket chains. *Quarterly Journal of Economics* 112:791–824.
- Mackowiak, B., and M. Wiederholt. 2009. Optimal sticky prices under rational inattention. *American Economic Review* 99:769–803.

- Mankiw, N. G., and R. Reis. 2002. Sticky information versus sticky prices: A proposal to replace the new keynesian Phillips curve. *Quarterly Journal of Economics* 117:1295–328.
- McKinsey & Company. 2019. How analytics can drive growth in consumer-packaged-goods trade promotions. <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/how-analytics-can-drive-growth-in-consumer-packaged-goods-trade-promotions>. Accessed April 2026.
- Midrigan, V. 2011. Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica* 79:1139–80.
- Nakamura, E. 2008. Pass-through in retail and wholesale. *American Economic Review* 98:430–7.
- Nakamura, E., and J. Steinsson. 2008. Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics* 123:1415–64.
- Sims, C. A. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50:665–90.
- Stroebel, J., and J. Vavra. 2019. House prices, local demand, and retail prices. *Journal of Political Economy* 127:1391–436.
- Villas-Boas, S. B. 2007. Vertical relationships between manufacturers and retailers: Inference with limited data. *Review of Economic Studies* 74:625–52.
- Webb, M. 2020. The impact of artificial intelligence on the labor market. Working paper.
- Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen. 2004. Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *Review of Economics and Statistics* 86:514–33.

Figure 1: Pricing dynamics around brand acquisitions

This figure plots monthly event-time profiles of treated-UPC outcomes around brand acquisitions. Each point is the cross-sectional mean across treated UPCs at event time k , defined as the number of months relative to the acquisition close date, with 90% confidence intervals computed from the cross-sectional standard error of the monthly mean. The vertical dashed line marks the acquisition month ($k = 0$). Panels A and B use the common set of acquisitions with usable cross-store SD data and plot deviations from the acquired brand's own pre-acquisition pricing behavior at the UPC×retailer level. In Panel A the dependent variable is $|\text{SD}_{urt} - \overline{\text{SD}}_{ur,\text{pre}}|$, where SD is the within-retailer cross-store standard deviation of log prices for UPC u sold by retailer r in month t . In Panel B the dependent variable is $|\text{Repr}_{urt} - \overline{\text{Repr}}_{ur,\text{pre}}|$, where Repr is an indicator for an absolute month-over-month log price change exceeding 5%. Panels C and D restrict to *synergistic* acquisitions, in which the acquirer was already selling products in the same Nielsen product modules as the acquired brand prior to the acquisition, and plot distance from the moving legacy-acquirer firm×retailer×module benchmark for the SD outcome (Panel C) and the repricing indicator (Panel D). [Appendix A](#) presents variable definitions.

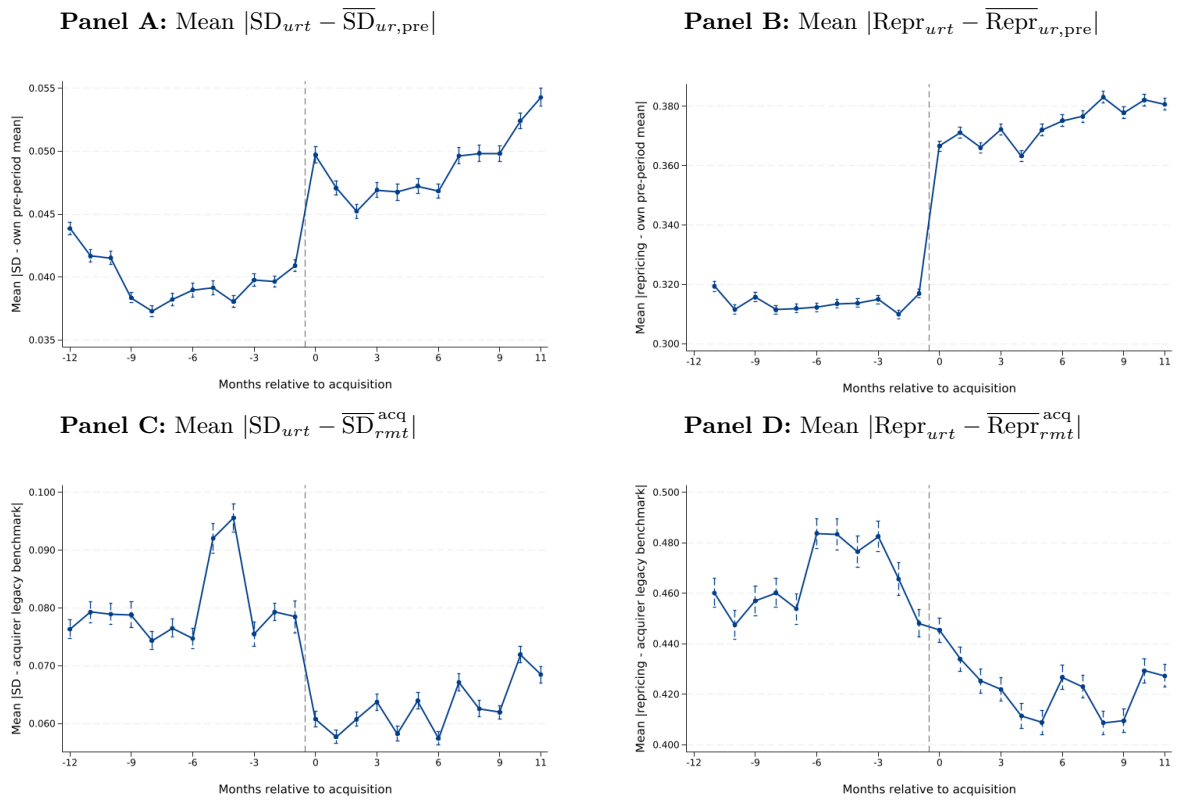


Figure 2: Distribution of AI exposure

This figure plots the cross-sectional distribution of firm-level generative AI exposure measured in January 2020. AI exposure is computed as the employment-share-weighted average of occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#), where scores capture the extent to which generative AI is expected to reduce the time required to complete the tasks typically performed by each occupation. Occupation-level scores are aggregated to the firm level using each firm's employment shares across occupations. Higher values indicate greater expected exposure to generative AI.

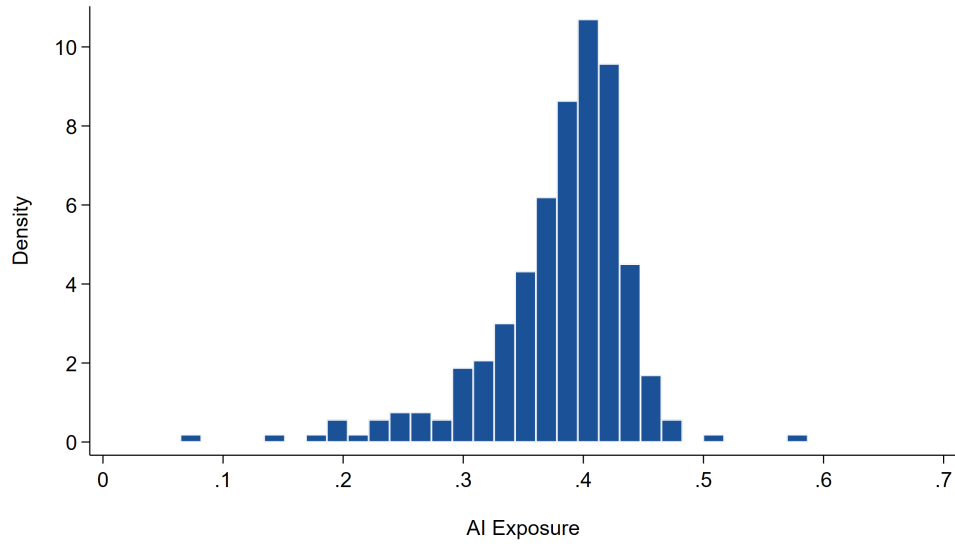


Figure 3: Macro environment around AI implementation events

This figure plots three time series around the InstructGPT API introduction in January 2022. Panel A shows Google Trends search interest for “artificial intelligence” (index, with 100 equal to the peak value over the period shown). Panel B shows the effective federal funds rate, in percent. Panel C shows new monthly COVID-19 cases in the United States, in millions. In each panel, the solid black vertical line marks the InstructGPT API introduction on January 27, 2022; the dashed gray line marks the ChatGPT public release on November 30, 2022; and the dotted gray line marks the GPT-3.5 and GPT-4 API introductions in March 2023. [Appendix A](#) presents variable definitions.

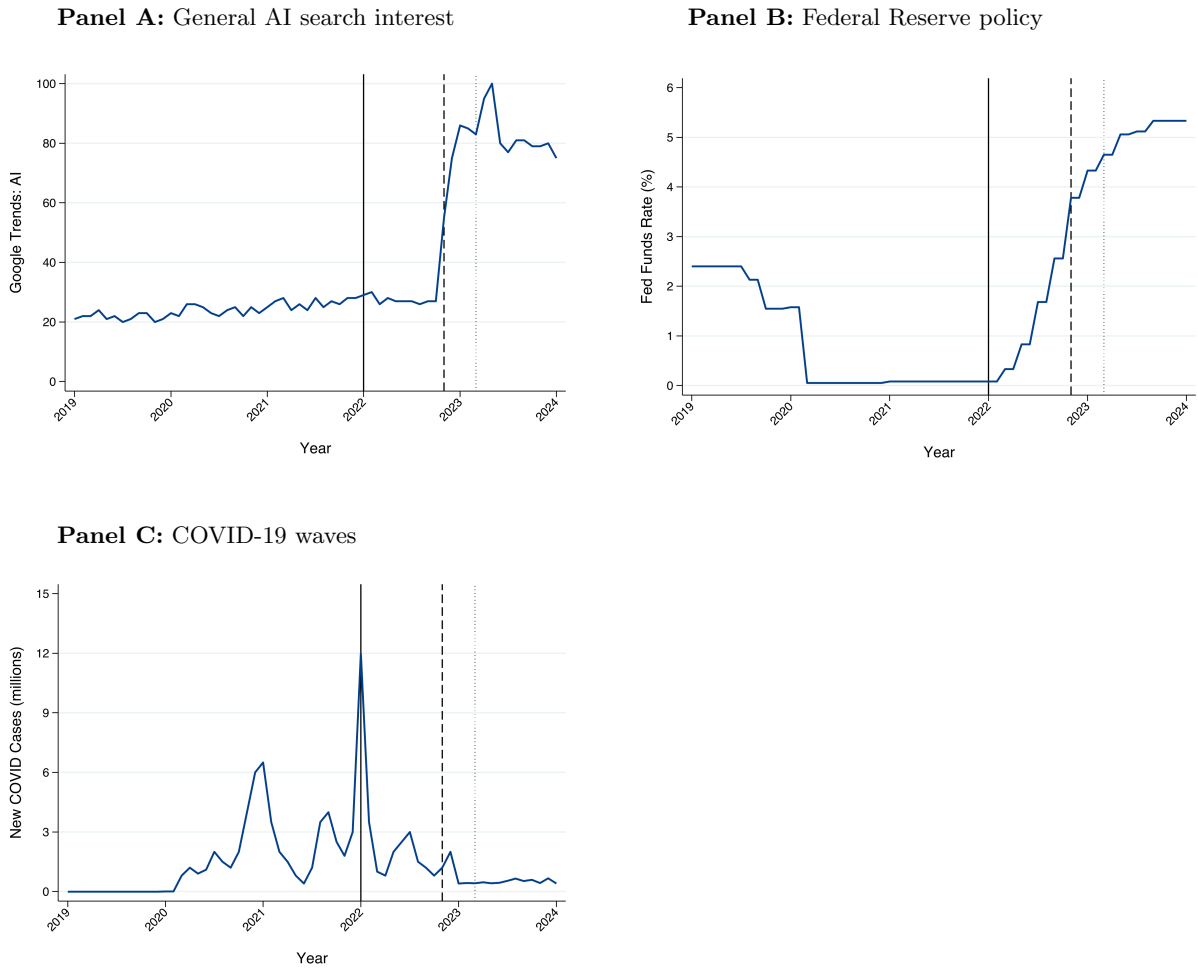


Figure 4: Price dispersion event study

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $\text{Dispersion}_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \neq -3} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), and t calendar months. In Panel A, the dependent variable is the standard deviation of log prices across stores within a firm-product-month. In Panel B, the dependent variable is the interquartile range of log prices across stores within a firm-product-month. AI exposure is computed as the employment-share-weighted average of occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#), where scores capture the extent to which generative AI is expected to reduce the time required to complete the tasks typically performed by each occupation. *Treated* equals one for firms with above-median AI exposure measured in January 2020. Event time runs from -6 to $+6$ months relative to the InstructGPT API introduction in January 2022, with $k = -3$ as the omitted base period. α_{ip} and α_t denote firm-product and calendar month fixed effects, respectively. Treated and control firm products are balanced on the pre-treatment log average price and log units sold. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.

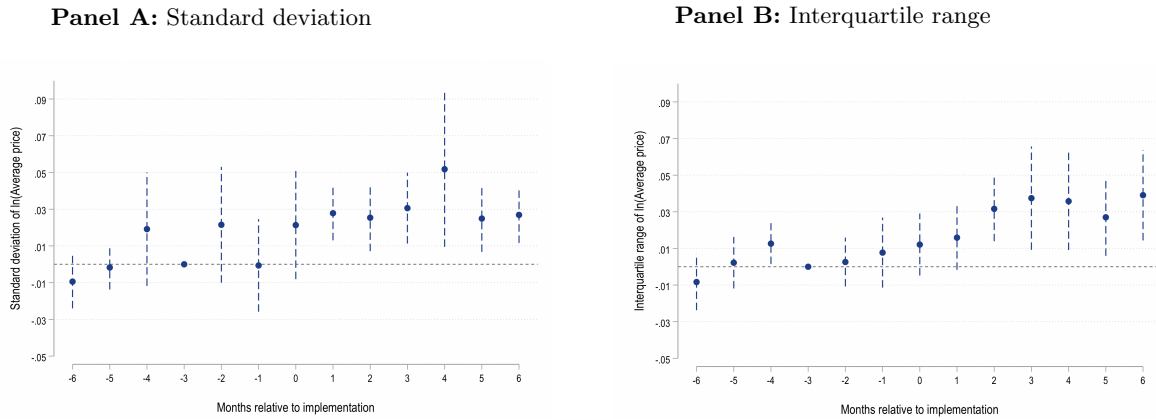


Figure 5: Price changes event study

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $Y_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \neq -3} \beta_k 1[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), and t calendar months. In Panel A, $Y_{ipt} = \mathbb{P}(|\Delta \ln \bar{P}| > 5\%)$. In Panel B, $Y_{ipt} = \mathbb{P}(\Delta \ln \bar{P} > 5\%)$. In Panel C, $Y_{ipt} = \mathbb{P}(\Delta \ln \bar{P} < -5\%)$. Within each panel, the three series correspond to specifications estimated at the firm-product level (All), the firm-product-county level (County), and the firm-product-retailer level (Retail). AI exposure is computed as the employment-share-weighted average of occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Treated* equals one for firms with above-median AI exposure measured in January 2020. Event time runs from -6 to $+6$ months relative to the InstructGPT API introduction in January 2022, with $k = -3$ as the omitted base period. α_{ip} and α_t denote firm-product and calendar month fixed effects, respectively. Treated and control firm products are balanced on the pre-treatment log average price and log units sold. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.

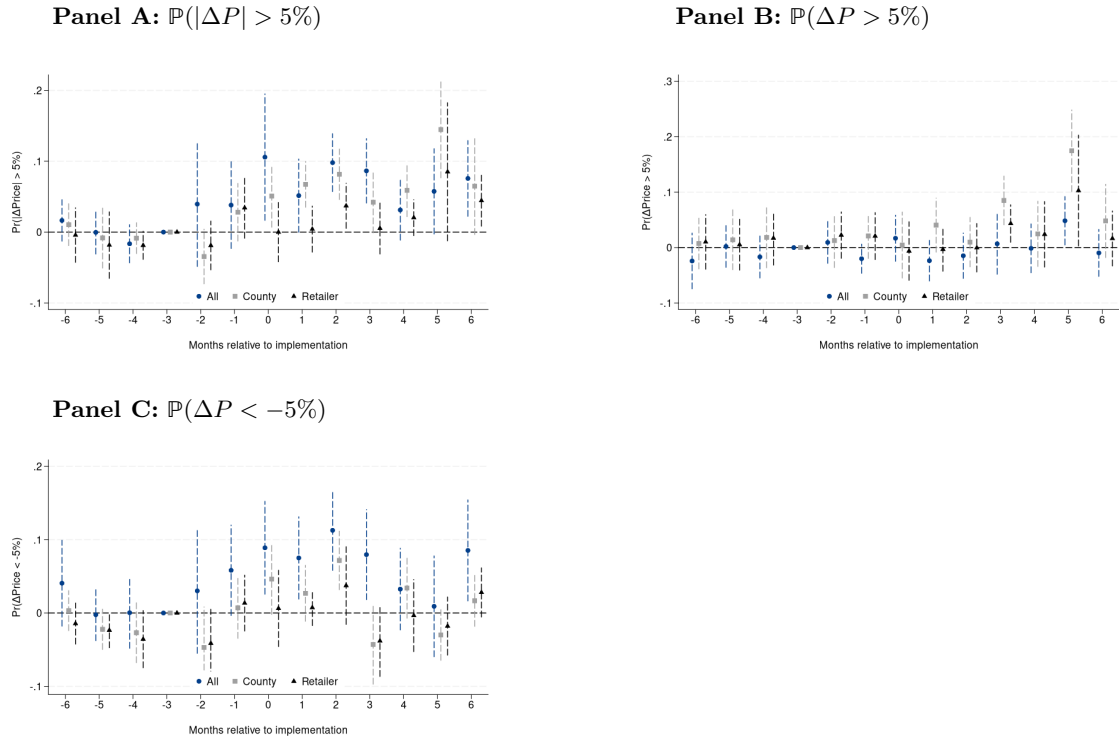


Figure 6: Price level event study

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $\ln(\text{Average price})_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \neq -3} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), and t calendar months. The dependent variable is the log of CPI-deflated average price. AI exposure is computed as the employment-share-weighted average of occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#), where scores capture the extent to which generative AI is expected to reduce the time required to complete the tasks typically performed by each occupation. *Treated* equals one for firms with above-median overall AI exposure measured in January 2020. Event time runs from -6 to $+6$ months relative to the InstructGPT API introduction in January 2022, with $k = -3$ as the omitted base period. α_{ip} and α_t denote firm-product and calendar month fixed effects, respectively. Treated and control firm products are balanced on the pre-treatment log average price and log units sold. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.




Figure 7: Treatment effect heterogeneity across retailers

This figure plots the estimated β_k coefficients and 90% confidence intervals from a regression that includes retailer-specific interactions in the baseline price-level specification: $\ln(\text{Average price})_{ipt} = \alpha_{ip} + \alpha_t + \sum_k \beta_k (\text{Treated}_i \times \text{Post}_t \times \mathbb{1}[\text{Retailer} = k]) + \mathbf{X}'_{i0} \gamma \times \text{Post}_t + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), t calendar months, and k indexes retailers. Retailers are ranked along the horizontal axis from the largest (most negative) to the smallest estimated treatment effect. The gray dashed horizontal line indicates the pooled $\text{Treated} \times \text{Post}$ estimate from the baseline specification. AI exposure is computed as the employment-share-weighted average of occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). Treated equals one for firms with above-median AI exposure measured in January 2020. Post equals one during the six months following the InstructGPT API introduction in January 2022. All specifications include firm-product and calendar month fixed effects and controls for log average pre-period units sold interacted with post, log employees interacted with post, and firm financial characteristics interacted with post. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.





Figure 8: Evidence of AI adoption

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $Y_{it} = \alpha_i + \alpha_t + \sum_{k \neq -1} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{it}$, where i denotes firms and t denotes calendar periods. In Panel A, the dependent variable is the count of active AI-related job postings from LinkUp, and event time is measured in months relative to the InstructGPT API introduction in January 2022. In Panel B, the dependent variable is the count of AI-related keywords in annual 10-K filings, and event time is measured in years relative to 2022. AI exposure is computed as the employment-share-weighted average of occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Treated* equals one for firms with above-median AI exposure measured in January 2020. α_i and α_t denote firm and calendar period fixed effects, respectively, with $k = -1$ as the omitted base period. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.

Panel A: AI-related job postings (LinkUp)



Panel B: AI keywords in 10-K filings

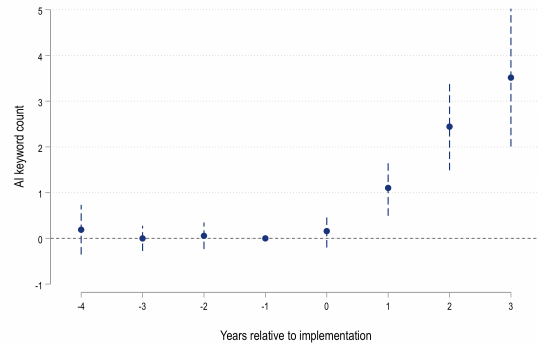


Table 1: Summary Statistics

This table reports summary statistics. Panel A reports the outcome variables used in the price variance decomposition analyses, measured at the store×UPC×month level across the full NielsenIQ Retail Scanner sample. Panel B reports the coefficient of variation of raw prices computed at three levels of aggregation in the same sample: across stores within retailer×UPC×month groups; across retailer-level value-weighted prices within UPC×county×month groups; and across retailer-level value-weighted prices within UPC×month groups. Panel C reports the outcome variable used in the private label tests. Panel D reports pricing outcomes for the AI experiment sample. Panel E reports the firm-level AI exposure measure. Panel F reports firm characteristics measured in the quarter before the InstructGPT API introduction in January 2022. Panel G reports covariate balance between treated and control firms. [Appendix A](#) presents variable definitions.

	Obs (1)	Mean (2)	Std. Dev. (3)	25th Pct. (4)	Median (5)	75th Pct. (6)
Panel A: Price variance decomposition						
ln(Average price)	4.94B	1.787	0.791	1.306	1.790	2.250
$(\ln p_{ist} - \overline{\ln p_{it}})^2$	4.94B	0.038	0.206	0.001	0.008	0.028
$\Delta \ln(\text{Price})$	3.75B	-0.004	0.257	-0.015	0.000	0.018
Panel B: Coefficient of variation						
Within retailer, across stores	35.4M	0.073	0.115	0.000	0.037	0.092
Across retailers, within county	340M	0.138	0.123	0.056	0.115	0.186
Across retailers	2.55M	0.177	0.170	0.071	0.139	0.227
Panel C: Private label sample						
SD of ln(Average price), private label	1.09B	0.067	0.241	0.000	0.014	0.059
SD of ln(Average price), national brand	12.1B	0.057	0.211	0.000	0.009	0.050
Panel D: AI tests pricing outcomes						
Ln(Average price)	90,211	1.797	0.924	1.194	1.795	2.407
SD of Ln(Average price)	78,341	0.206	0.189	0.084	0.161	0.264
IQR of Ln(Average price)	90,211	0.144	0.174	0.000	0.088	0.220
$P(\% \Delta P > 5\%)$	90,211	0.337	0.473	0.000	0.000	1.000
$P(\% \Delta P > 10\%)$	90,211	0.207	0.405	0.000	0.000	0.000
Ln(Units sold)	90,211	4.374	3.191	1.526	3.970	6.989
Panel E: AI exposure						
AI Exposure	306	0.381	0.059	0.359	0.394	0.418

	Obs (1)	Mean (2)	Std. Dev. (3)	25th Pct. (4)	Median (5)	75th Pct. (6)
Panel F: Firm characteristics						
Ln(Employees)	91	7.724	2.201	6.340	7.620	9.627
Ln(Assets)	91	8.671	2.248	7.235	8.987	10.246
Tobin's q	91	2.802	1.861	1.397	2.005	3.582
ROA	91	0.005	0.072	0.005	0.014	0.027
Cash/Assets	91	0.108	0.124	0.031	0.072	0.136
Fixed assets/Assets	91	0.212	0.151	0.102	0.174	0.283
R&D/Assets	91	0.008	0.019	0.000	0.002	0.010
Market-to-book	91	6.287	15.954	2.078	3.368	6.935
Ln(Age)	91	3.439	0.820	3.135	3.555	4.111
Panel G: Covariate balance (baseline)						
	Obs	Treated	Control	Difference	t -stat	Rank-sum
Ln(Employees)	91	8.304	7.361	0.943	2.020	[0.070]
Ln(Assets)	91	9.382	8.226	1.156	2.452	[0.029]
Tobin's q	91	2.659	2.891	-0.232	-0.576	[0.504]
ROA	91	0.018	-0.003	0.021	1.375	[0.278]
Cash/Assets	91	0.091	0.118	-0.027	-1.019	[0.602]
Fixed assets/Assets	91	0.129	0.264	-0.135	-4.633	[0.000]
R&D/Assets	91	0.011	0.007	0.004	0.866	[0.003]
Market-to-book	91	8.381	4.977	3.404	0.990	[0.138]
Ln(Age)	91	3.555	3.367	0.188	1.068	[0.050]

Table 2: OpenAI API release schedule

This table lists the OpenAI large language model families and their API availability dates. Each row corresponds to a model version made available through the OpenAI API. The *Event Date* column indicates the date on which the model became available for commercial use via the API.

Family	Model API Name	Event Date
GPT-3	davinci (beta)	2020-06-11
GPT-3	davinci (general release)	2021-11-18
InstructGPT	text-davinci-001	2022-01-27
InstructGPT	text-davinci-002	2022-03-15
InstructGPT	text-davinci-003	2022-11-28
GPT-3.5	gpt-3.5-turbo-0301	2023-03-01
GPT-3.5	gpt-3.5-turbo-0613	2023-06-13
GPT-4	gpt-4	2023-03-14
GPT-4	gpt-4-32k	2023-03-14
GPT-4 Turbo	gpt-4-1106-preview	2023-11-06
GPT-4o	gpt-4o	2024-05-13

Table 3: Price variance decomposition

This table decomposes the variance of three pricing outcomes at the store×UPC×month level using the full manufacturer-linked NielsenIQ scanner dataset from 2010 to 2023. Panel A decomposes the variance of log prices. Panel B decomposes the variance of squared deviations from the UPC×month mean. Panel C decomposes the variance of month-over-month log price changes. Each panel reports Shapley value decompositions of the explanatory power of different fixed-effect groupings. Column (1) uses UPC×month as the baseline and attributes the marginal R^2 to retailer and county fixed effects. Column (2) uses manufacturer as the baseline and attributes variance to retailer and county. Column (3) extends the decomposition following DellaVigna and Gentzkow (2019) to include a bilateral retailer×UPC pricing term. In Panels A and B, column (4) reports a Shapley decomposition that attributes variance to UPC×month, manufacturer×retailer, and manufacturer×county fixed effects, isolating manufacturer-specific pricing patterns across retailers and counties on top of UPC×month variation. In Panel C, column (4) reports a Shapley decomposition that attributes variance to manufacturer×month, retailer×month, and county×month fixed effects, isolating who synchronizes price changes over time. In Panel A, column (5) reports the four-factor Shapley decomposition of Couts, Gonçalves, and Loudis (2026) that splits explanatory power across UPC, retailer, county, and month using fully interactive fixed effects in each subset. All decompositions are estimated via iterative demeaning. Appendix A presents variable definitions.

Panel A: Variance of log prices					
	Shapley			Sources	CGL
	Product	Manufacturer	Bilateral	Geographic	4-factor
	(1)	(2)	(3)	(4)	(5)
UPC					88.8%
UPC × Month	99.0%		48.4%	71.8%	
Manufacturer		91.4%			
Manufacturer × Retailer				16.0%	
Manufacturer × County				12.2%	
Retailer	0.9%	7.6%	3.3%		5.9%
Retailer × UPC			47.8%		
Month					3.4%
County	0.1%	1.0%	0.4%		1.8%
R^2	0.899	0.368	0.909	0.902	0.954
Observations (billion)	4.94	4.94	4.94	4.94	4.94

Panel B: Variance of squared deviations from UPC×month mean				
	Shapley			Sources
	Product	Manufacturer	Bilateral	Geographic
	(1)	(2)	(3)	(4)
UPC × Month	96.2%		50.5%	74.3%
Manufacturer		56.0%		
Manufacturer × Retailer				21.7%
Manufacturer × County				4.0%
Retailer	3.0%	36.5%	1.1%	
Retailer × UPC			48.0%	
County	0.8%	7.5%	0.4%	
R^2	0.075	0.007	0.117	0.092
Observations (billion)	4.94	4.94	4.94	4.94

Panel C: Variance of log price changes

	Shapley			Sources
	Product	Manufacturer	Bilateral	Repricing
	(1)	(2)	(3)	(4)
UPC \times Month	99.9%		92.5%	
Manufacturer		97.1%		
Manufacturer \times Month				72.8%
Retailer \times Month				21.7%
Retailer	0.1%	2.2%	0.1%	
Retailer \times UPC			7.4%	
County	0.0%	0.6%	0.0%	
County \times Month				5.5%
R^2	0.081	0.002	0.083	0.016
Observations (billion)	3.75	3.75	3.75	3.70

Table 4: Pricing behavior around brand acquisitions

This table examines changes in pricing outcomes around brand acquisitions. The sample consists of UPCs belonging to the acquired brands. The estimation window is ± 12 months around each acquisition close date, and the unit of observation is $\text{UPC} \times \text{retailer} \times \text{stack} \times \text{month}$. SD denotes the within-retailer cross-store standard deviation of log prices for UPC u sold by retailer r in month t , and $Repr$ is an indicator for an absolute month-over-month log price change exceeding 5%. Columns (1) and (2) use the common set of acquisitions with usable cross-store SD data and test whether acquired brands depart from their own pre-acquisition pricing behavior. The outcomes are $|SD_{urt} - \overline{SD}_{ur,pre}|$ and $|Repr_{urt} - \overline{Repr}_{ur,pre}|$, where the benchmark is the acquired UPC-retailer pair's own pre-period mean. Columns (3) and (4) restrict to *synergistic* acquisitions, those in which the acquirer was already selling products in the same Nielsen product modules as the acquired brand prior to the acquisition, and test alignment with the acquirer's legacy pricing behavior. The outcomes are absolute distances from moving legacy-acquirer firm \times retailer \times module benchmarks, constructed from the acquirer's pre-existing UPCs in the same retailer, month, and module. Columns (5) and (6) restrict to *expansionary* acquisitions, those in which the acquirer had identifiable legacy products but not in the acquired brand's modules, and test divergence from the acquirer's legacy pricing behavior. The outcomes are the corresponding absolute distances from moving legacy-acquirer firm \times retailer \times product-group benchmarks. In columns (1) and (2), a positive coefficient indicates departure from the acquired brand's own pre-acquisition pricing. In columns (3) and (4), a negative coefficient indicates the acquired brand becoming more similar to the acquirer's legacy pricing in the same module. In columns (5) and (6), a positive coefficient indicates divergence from the acquirer's legacy pricing in adjacent product groups. All columns include UPC \times retailer \times stack fixed effects and cluster standard errors at the UPC \times stack level. Asymptotic standard errors are reported in parentheses. Wild bootstrap p -values with 1,000 replications are clustered at the deal level. Statistical significance at the 1, 5, and 10 percent significance levels based on asymptotic standard errors is denoted by ***, **, and *, respectively.

	Distance from own-UPC pre-acquisition pricing		Distance from acquirer's legacy pricing			
	SD(ln price)	Repricing frequency	Synergistic acquisitions		Expansionary acquisitions	
			SD(ln price)	Repricing frequency	SD(ln price)	Repricing frequency
	(1)	(2)	(3)	(4)	(5)	(6)
Post acquisition	0.016*** (0.000)	0.064*** (0.001)	-0.007*** (0.001)	-0.021*** (0.003)	0.006*** (0.001)	0.005*** (0.001)
Wild bootstrap p -value	<0.001	<0.001	0.075	0.092	0.063	0.109
Fixed effects						
UPC \times retailer \times stack	✓	✓	✓	✓	✓	✓
Deals	10	10	4	4	5	5
Adj. R^2	0.375	0.393	0.262	0.164	0.268	0.142
Observations	1,720,245	1,572,497	330,966	350,028	826,043	814,767

Table 5: Private label products, price dispersion, and local wages

This table examines whether private-label products exhibit greater price localization than national brands and whether their prices respond more strongly to local wage variation. Panel A regresses the within-UPC standard deviation of log prices across stores on a private label indicator. The unit of observation is UPC \times retailer \times county \times month. Panel B regresses the log of price on the interaction of the private label indicator with county-level log average monthly wages from the BLS Quarterly Census of Employment and Wages, restricting to county-months with at least 5,000 private-sector employees. The unit of observation is a store \times UPC \times month. In both panels, *Private Label* equals one for store brand products. The fixed effects used in each specification are listed at the bottom of each panel. Standard errors clustered at the county level are reported in parentheses. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

Panel A: Price dispersion			
	Standard deviation ln(Price)		
	(1)	(2)	(3)
Private label	0.009*** (0.000)	0.005*** (0.000)	0.006*** (0.000)
ln(Units sold)	-0.004*** (0.000)	0.001*** (0.000)	0.001*** (0.000)
Fixed effects			
Month \times Product module	✓		
Retailer	✓		
Retailer \times Product module		✓	
Retailer \times Month		✓	
Retailer \times Product module \times Month			✓
Observations (billion)	13.2	13.2	13.2
Adjusted R^2	0.04	0.08	0.10
Panel B: Private label prices and local wages			
	ln(Price)		
	(1)	(2)	(3)
Private label \times ln(Wage)	0.003 (0.006)	0.019*** (0.003)	0.018*** (0.003)
Private label	-0.220*** (0.053)	-0.387*** (0.026)	-0.378*** (0.027)
ln(Wage)	0.362*** (0.015)	-0.005 (0.004)	-0.004 (0.004)
ln(Units sold)	-0.174*** (0.002)	-0.158*** (0.001)	-0.159*** (0.001)
Fixed effects			
Retailer \times Product module \times Month	✓		
Retailer \times Product module \times County		✓	✓
Calendar month		✓	
Retailer \times Product module \times Month			✓
Observations (billion)	25.7	25.7	25.7
Adjusted R^2	0.21	0.47	0.48

Table 6: Price dispersion

This table presents results from difference-in-differences regressions of price dispersion around the introduction of the InstructGPT API. The sample is a firm-product-month panel of publicly listed U.S. firms. The dependent variable in columns (1) to (3) is the standard deviation of log prices, and in columns (4) to (6) is the interquartile range of log prices, both computed across stores within a firm-product-month. Columns (1) and (4) compute dispersion across all stores; columns (2) and (5) compute dispersion within counties; and columns (3) and (6) compute dispersion within retailers. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. All specifications include controls for the log of average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. All specifications include firm-product and calendar month fixed effects. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	Standard deviation			Interquartile range		
	All	County	Retailer	All	County	Retailer
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	0.025*** (0.006)	0.008** (0.003)	0.009** (0.004)	0.020*** (0.006)	0.001* (0.001)	0.007** (0.003)
Ln(Avg. pre-period units) \times Post	0.004*** (0.001)	0.002*** (0.001)	0.002*** (0.001)	0.005*** (0.001)	0.000* (0.000)	0.001*** (0.000)
Ln(Employees) \times Post	-0.001 (0.006)	0.002 (0.003)	0.005 (0.005)	0.008 (0.007)	0.001 (0.001)	0.006* (0.004)
Ln(Assets) \times Post	-0.000 (0.006)	-0.003 (0.003)	-0.007 (0.005)	-0.012 (0.007)	-0.001** (0.001)	-0.008** (0.004)
Tobin's q \times Post	0.000 (0.002)	-0.002* (0.001)	0.000 (0.001)	0.002 (0.002)	0.000 (0.000)	0.002 (0.001)
ROA \times Post	-0.030 (0.160)	-0.054 (0.064)	0.011 (0.105)	0.208 (0.157)	0.036*** (0.012)	-0.006 (0.098)
Cash/assets \times Post	-0.004 (0.040)	-0.006 (0.017)	-0.042 (0.028)	-0.069* (0.040)	-0.003 (0.004)	-0.040* (0.021)
Fixed assets/assets \times Post	0.002 (0.019)	0.013 (0.010)	-0.033** (0.015)	-0.039 (0.024)	-0.005* (0.003)	-0.067*** (0.015)
R&D/assets \times Post	0.157 (0.237)	-0.037 (0.127)	0.256 (0.167)	-0.069 (0.258)	-0.080** (0.031)	-0.021 (0.111)
Market-to-book \times Post	-0.000 (0.000)	-0.000* (0.000)	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Ln(Age) \times Post	0.002 (0.005)	0.003 (0.003)	0.003 (0.004)	0.000 (0.005)	0.000 (0.001)	0.003 (0.003)
Fixed effects						
UPC	✓	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓	✓
Observations	77,607	63,304	75,242	77,607	63,304	75,242
Adjusted R^2	0.46	0.46	0.46	0.50	0.78	0.45

Table 7: Price changes

This table presents results from difference-in-differences regressions of the frequency of price changes around the introduction of the InstructGPT API. The dependent variable is an indicator for whether the month-over-month change in log average price exceeds 5% in magnitude at the level of aggregation used in the column. Columns (1) to (3) use the absolute value of that change; columns (4) to (6) isolate price increases; and columns (7) to (9) isolate price decreases. Within each group, columns are estimated at the firm-product level (All), the firm-product-county level (County), and the firm-product-retailer level (Retail). The All specifications include firm-product and calendar month fixed effects; the County specifications include UPC×county and calendar month fixed effects; and the Retail specifications include UPC×retailer and calendar month fixed effects. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from Eisfeldt et al. (2025). *Post* equals one during the event month and the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. All specifications include controls for the log of average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin’s q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	$\mathbb{P}(\% \Delta P > 5\%)$			$\mathbb{P}(\% \Delta P > 5\%)$			$\mathbb{P}(\% \Delta P < -5\%)$		
	All	County	Retailer	All	County	Retailer	All	County	Retailer
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated × Post	0.064*** (0.018)	0.064*** (0.017)	0.033** (0.013)	0.018 (0.012)	0.029** (0.012)	0.016 (0.011)	0.046*** (0.012)	0.035*** (0.013)	0.017 (0.011)
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Fixed effects									
UPC	✓			✓			✓		
UPC×County		✓			✓			✓	
UPC×Retailer			✓			✓			✓
Calendar month	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	89,168	16,509,927	960,796	89,168	16,509,927	960,796	89,168	16,509,927	960,796
Adjusted R^2	0.31	0.20	0.25	0.08	0.04	0.07	0.12	0.06	0.10

Table 8: Price level and local economic conditions

This table examines whether average prices fall and become more sensitive to local economic conditions after the introduction of the InstructGPT API. The dependent variable is the log of CPI-deflated average price. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. Column (1) reports the baseline price-level treatment effect on the firm-product-month panel, with firm-product and calendar month fixed effects. Columns (2) and (3) extend the design to a firm-product-county-month panel and interact the treatment indicator with county-level log average monthly wages from the BLS Quarterly Census of Employment and Wages. Column (2) includes UPC, county, and calendar month fixed effects. Column (3) replaces the separate county and calendar month fixed effects with county \times calendar month fixed effects, which absorb local time-varying shocks. All specifications include controls for the log of average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	ln(Average price)		
	(1)	(2)	(3)
Treated \times Post	-0.043*** (0.011)	-0.219** (0.098)	-0.223** (0.100)
Treated \times Post \times ln(Wage)		0.023** (0.011)	0.023** (0.012)
Treated \times ln(Wage)		0.004 (0.010)	0.004 (0.009)
ln(Wage) \times Post		-0.020* (0.011)	
ln(Wage)		-0.012 (0.013)	
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
County		✓	
Calendar month	✓	✓	
County \times Calendar month			✓
Observations	89,168	12,682,944	12,682,944
Adjusted R^2	0.95	0.95	0.95

Table 9: Controlling for retailer-specific effects

This table presents results from difference-in-differences regressions of prices around the introduction of the InstructGPT API, controlling for retailer-specific fixed effects. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. The dependent variables are the standard deviation of log prices (columns (1)–(3)), an indicator for absolute month-over-month changes in log average price exceeding 5% (columns (4)–(6)), and the log of CPI-deflated average price (columns (7)–(9)). Columns (1), (4), and (7) include UPC and calendar month fixed effects (baseline). Columns (2), (5), and (8) replace UPC fixed effects with UPC×retailer fixed effects, which absorb any time-invariant retailer-product pricing differences. Columns (3), (6), and (9) further add retailer×calendar month fixed effects, which absorb retailer-specific time trends such as chain-wide promotional cycles and seasonal pricing strategies. All specifications include controls for the log of average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin’s q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	Standard deviation			$\mathbb{P}(\% \Delta P > 5\%)$			ln(Average price)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated × Post	0.011*** (0.003)	0.007*** (0.003)	0.005* (0.003)	0.027** (0.012)	0.033** (0.013)	0.028** (0.012)	−0.036*** (0.007)	−0.033*** (0.008)	−0.024*** (0.006)
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Fixed effects									
UPC	✓			✓			✓		
UPC × Retailer		✓	✓		✓	✓		✓	✓
Retailer × Month			✓			✓			✓
Calendar month	✓	✓		✓	✓		✓	✓	
Observations	856,345	853,644	853,644	963,500	960,781	960,781	963,500	960,781	960,781
Adjusted R^2	0.17	0.41	0.43	0.12	0.25	0.26	0.91	0.96	0.96

Table 10: Product-level heterogeneity

This table presents results from difference-in-differences regressions of prices around the introduction of the Instruct-GPT API, including triple interactions with product-level characteristics. The sample is a firm-product-month panel of publicly listed U.S. firms. The dependent variable is the log of CPI-deflated average price. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from Eisfeldt et al. (2025). *Product indicator* is a binary variable capturing a product-level characteristic interacted with *Treated* \times *Post*. In column (1), the product indicator equals one for discretionary products, classified based on NielsenIQ product group codes. In column (2), the product indicator equals one for products with above-median demand volatility, measured as the coefficient of variation of pre-treatment units sold within each product module. In column (3), the product indicator equals one for products with above-median price elasticity, estimated from pre-treatment regressions of log quantity changes on log price changes. *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. All specifications include controls for the log of average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. All specifications include UPC and product group \times calendar month fixed effects. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. Appendix A presents variable definitions.

	ln(Average price)		
	Discretionary	Demand volatility	Price elasticity
	(1)	(2)	(3)
Treated \times Post	-0.003 (0.009)	-0.011 (0.010)	-0.008 (0.011)
Treated \times Post \times Product indicator	-0.038*** (0.013)	-0.025** (0.011)	-0.023** (0.011)
Product indicator \times Post		0.015* (0.008)	-0.003 (0.008)
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
Product group \times Calendar month	✓	✓	✓
Observations	89,063	89,050	89,038
Adjusted R^2	0.95	0.95	0.95

Table 11: The information channel

This table presents results from difference-in-differences regressions of pricing outcomes on three decompositions of firm-level AI exposure around the introduction of the InstructGPT API. The sample is a firm-product-month panel of publicly listed U.S. firms. The dependent variables are the standard deviation of log prices (column 1), an indicator for absolute month-over-month changes in log average price exceeding 5% (column 2), and the log of CPI-deflated average price (column 3). Panel A reports horse-race regressions that include pricing-task AI exposure and non-pricing-task AI exposure simultaneously. Pricing-task exposure is the employment-share-weighted average of each occupation’s AI exposure score multiplied by the share of its O*NET tasks classified as pricing-related, identified by keywords such as price, cost, demand, forecast, promotion, negotiate, revenue, and margin. Non-pricing-task exposure is defined analogously using one minus the pricing-task share. Panel B reports separate regressions for above-median augmentation and above-median automation exposure. Augmentation scores capture the extent to which AI assistance increases worker productivity while human judgment remains essential. Automation scores capture the likelihood that AI could complete the task independently with minimal human involvement. Both measures are scored at the O*NET 30.1 task level using GPT-5-mini, aggregated to the occupation level, and then to the firm level using employment shares. Panel C reports separate regressions for above-median and below-median data assets exposure, a firm-level measure from [Eisfeldt et al. \(2025\)](#) that captures whether a firm’s operations rely on data inputs that generative AI can process. *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. All specifications include firm-product and calendar month fixed effects, controls for the log of average pre-period units sold interacted with *post*, log employees measured in the month before the event and interacted with *post*, and firm financial controls (log assets, Tobin’s *q*, ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with *post*. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

Panel A: Pricing-task vs. non-pricing-task AI exposure			
	SD of ln(price)	Pr(ΔP > 5%)	ln(Average price)
	(1)	(2)	(3)
Pricing-task exposure \times Post	0.493 (0.531)	2.619** (1.010)	-1.414* (0.740)
Non-pricing-task exposure \times Post	-0.057 (0.182)	0.073 (0.396)	0.370 (0.294)
$p(\text{pricing} = \text{non-pricing})$	0.410	0.060	0.062
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
Calendar month	✓	✓	✓
Observations	77,976	88,685	89,589

Panel B: Augmentation vs. automation exposure			
	SD of ln(price)	Pr(ΔP > 5%)	ln(Average price)
	(1)	(2)	(3)
Augmentation exposure \times Post	0.019** (0.007)	0.047** (0.020)	-0.039*** (0.010)
Automation exposure \times Post	0.015 (0.012)	0.022 (0.022)	-0.014 (0.016)
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
Calendar month	✓	✓	✓
Observations (Augmentation)	65,659	73,552	73,981
Observations (Automation)	55,774	64,154	64,903

Panel C: High vs. low data assets			
	SD of ln(price)	Pr(ΔP > 5%)	ln(Average price)
	(1)	(2)	(3)
High-data-assets exposure \times Post	0.025*** (0.006)	0.066*** (0.016)	-0.046*** (0.012)
Low-data-assets exposure \times Post	0.005 (0.008)	0.091** (0.044)	-0.004 (0.013)
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
Calendar month	✓	✓	✓
Observations (High data assets)	74,138	74,138	74,138
Observations (Low data assets)	10,482	10,482	10,482

Table 12: Evidence of AI adoption

This table presents results from difference-in-differences regressions validating that firms with high AI exposure increase AI-adoption activity after the introduction of the InstructGPT API. Panel A reports regressions of AI-related job posting counts. The dependent variables are monthly counts of active job postings in five AI categories: AI, Applied AI, Data science, O*NET AI, and All AI. Panel B reports regressions of AI-related keyword counts in annual 10-K filings. The dependent variable is the count of AI-related keywords measured across the full filing (Total) and separately for Item 1 (Business Description), Item 1A (Risk Factors), and Item 7 (Management’s Discussion and Analysis). Panel C decomposes 10-K AI keywords into implementation and hiring and risk and regulatory categories. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Post* equals one after the introduction of the InstructGPT API. All specifications include firm and calendar period fixed effects. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

Panel A: Job postings (LinkUp)					
	Job postings				
	AI	Applied AI	Data science	O*NET AI	All AI
	(1)	(2)	(3)	(4)	(5)
Treated × Post	7.952** (3.381)	4.448** (1.771)	4.385** (1.773)	16.014** (7.247)	22.451** (9.200)
Fixed effects					
Firm	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓
Observations	1,586	1,586	1,586	1,586	1,586
Adjusted R^2	0.96	0.93	0.93	0.94	0.95

Panel B: 10-K AI keywords

	AI keyword count			
	Total	Item 1	Item 1A	Item 7
	(1)	(2)	(3)	(4)
Treated \times Post	1.516*** (0.392)	0.498** (0.215)	0.708*** (0.218)	0.038 (0.058)
Fixed effects				
Firm	✓	✓	✓	✓
Year	✓	✓	✓	✓
Observations	2,384	2,384	2,384	2,384
Adjusted R^2	0.68	0.73	0.43	0.61

Panel C: 10-K implementation vs. risk

	Implementation/hiring			Risk/regulatory		
	Total	Item 1	Item 1A	Total	Item 1A	Item 7
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	1.026*** (0.283)	0.405** (0.174)	0.433*** (0.135)	0.079** (0.037)	0.089** (0.035)	0.002 (0.002)
Fixed effects						
Firm	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓
Observations	2,384	2,384	2,384	2,384	2,384	2,384
Adjusted R^2	0.69	0.73	0.43	0.28	0.24	-0.01

Appendix A: Variable Definitions

Variable	Definition	Source
<i>Panel A: Dependent variables</i>		
$\ln(\text{Price})$	The natural log of the store-level unit price for a UPC in a store-month.	NielsenIQ Retail Panel
$\ln(\text{Average price})$	The natural log of the CPI-deflated average unit price of a firm-product (UPC) in a given month.	NielsenIQ Retail Panel; BLS Consumer Price Index (CPI)
$(\ln p_{ist} - \overline{\ln p_{it}})^2$	The squared deviation of store-level log price from the UPC-month mean log price, where i indexes UPCs, s stores, and t months.	NielsenIQ Retail Panel
$\Delta \ln(\text{Price})$	The month-over-month change in log price within the same unit of observation. In the variance decomposition, the unit is store-UPC; in the AI tests, it is firm-product, firm-product-county, or firm-product-retailer.	NielsenIQ Retail Panel
SD of $\ln(\text{Average price})$	The standard deviation of log prices across stores within a firm-product-month. Variants are computed across all stores, within counties, and within retailers.	NielsenIQ Retail Panel
IQR of $\ln(\text{Average price})$	The interquartile range (75th minus 25th percentile) of log prices across stores within a firm-product-month. Variants are computed across all stores, within counties, and within retailers.	NielsenIQ Retail Panel
Coefficient of variation	The standard deviation of prices divided by the mean price, computed across stores within a retailer-UPC-month, across retailer-level prices within a UPC-county-month, or across retailer-level prices within a UPC-month.	NielsenIQ Retail Panel
$\mathbb{P}(\% \Delta P > 5\%)$	An indicator variable equal to one if the absolute month-over-month change in log average price exceeds 5% at the relevant level of aggregation.	NielsenIQ Retail Panel
$\mathbb{P}(\% \Delta P > 5\%)$	An indicator variable equal to one if the month-over-month change in log average price exceeds +5% at the relevant level of aggregation.	NielsenIQ Retail Panel
$\mathbb{P}(\% \Delta P < -5\%)$	An indicator variable equal to one if the month-over-month change in log average price is less than -5% at the relevant level of aggregation.	NielsenIQ Retail Panel
Own pre-acquisition pricing distance	The absolute distance between an acquired UPC-retailer outcome and its own pre-acquisition mean, computed for within-retailer price dispersion and repricing frequency.	NielsenIQ Retail Panel; hand-collected brand acquisitions

Variable	Definition	Source
Acquirer legacy pricing distance	The absolute distance between an acquired UPC-retailer outcome and the acquirer’s legacy benchmark in the same retailer-month and product module or product group.	NielsenIQ Retail Panel; hand-collected brand acquisitions

Panel B: Treatment and exposure variables

AI Exposure	The employment-share-weighted average of occupation-level generative AI scores from Eisfeldt et al. (2025) . Scores capture the extent to which generative AI is expected to reduce the time required to complete the tasks typically performed by each occupation.	ONET (v27.2); Eisfeldt et al. (2025) ; Revelio Labs
Treated	An indicator variable equal to one for firms with above-median AI exposure measured in January 2020.	ONET (v27.2); Eisfeldt et al. (2025) ; Revelio Labs
Post	An indicator variable equal to one during the six months following the AI implementation event.	OpenAI API release schedule
AI Exposure (Davinci-01)	An alternative AI exposure measure derived from a target-judge model framework, where GPT-5-mini evaluates whether the text-davinci-001 model can perform O*NET occupational tasks. Aggregated to the firm level using employment shares.	ONET (v30.1); Eisfeldt et al. (2025) ; Revelio Labs
Private label	An indicator variable equal to one for store-brand products and zero for national-brand products.	NielsenIQ Retail Panel
Post acquisition	An indicator variable equal to one in months on or after the acquisition close date for the acquired-brand cohort.	Hand-collected brand acquisitions
Synergistic acquisition	An acquisition in which the acquirer sold products in the acquired brand’s NielsenIQ product modules before the acquisition.	NielsenIQ Retail Panel; hand-collected brand acquisitions
Expansionary acquisition	An acquisition in which the acquirer had identifiable legacy products in the same NielsenIQ product group but not in the acquired brand’s product modules.	NielsenIQ Retail Panel; hand-collected brand acquisitions

Panel C: Information-channel variables

Pricing-task AI exposure	The employment-share-weighted average of each occupation’s AI exposure score multiplied by the share of the occupation’s O*NET tasks classified as pricing-related. Pricing-related tasks are identified using keywords such as price, cost, demand, forecast, promotion, negotiate, revenue, and margin.	ONET (v27.2); Eisfeldt et al. (2025) ; Revelio Labs
--------------------------	---	---

Variable	Definition	Source
Augmentation exposure	An indicator variable equal to one for firms with above-median firm-level augmentation exposure. The firm-level score is the employment-share-weighted average of GPT-5-mini task-level scores capturing the extent to which AI assistance increases worker productivity while human judgment remains essential.	ONET (v30.1); GPT-5-mini; Revelio Labs
Automation exposure	An indicator variable equal to one for firms with above-median firm-level automation exposure. The firm-level score is the employment-share-weighted average of GPT-5-mini task-level scores capturing the likelihood that AI could complete the task independently with minimal human involvement.	ONET (v30.1); GPT-5-mini; Revelio Labs
Data-assets exposure	A firm-level measure from Eisfeldt et al. (2025) capturing the extent to which a firm's operations rely on data inputs that generative AI can process. The measure combines the firm's employment share in data-intensive occupations with GPT-scored references to data assets in 10-K filings.	ONET; SEC EDGAR; Eisfeldt et al. (2025)
Technology-task intensity	The employment-share-weighted average of each occupation's share of non-AI technology tasks, identified using O*NET task statements containing software, database, network, programming, or hardware terms and excluding pricing-related terms.	ONET (v27.2); Revelio Labs
Task-iteration exposure (AEI)	An employment-share-weighted firm-level exposure measure based on Anthropic Economic Index interactions classified as task iteration, in which users iteratively refine outputs with an AI model.	Anthropic Economic Index; O × NET; Revelio Labs
Directive exposure (AEI)	An employment-share-weighted firm-level exposure measure based on Anthropic Economic Index interactions classified as directive, in which users issue a single instruction that the AI model executes independently.	Anthropic Economic Index; O × NET; Revelio Labs

Panel D: Control variables

$\ln(\text{Units sold})$	The natural log of total units sold for a firm-product in a given month.	NielsenIQ Retail Panel
$\ln(\text{Avg. pre-period units}) \times \text{Post}$	The natural log of average pre-period units sold for the relevant firm-product pair, interacted with the post-event indicator.	NielsenIQ Retail Panel
$\ln(\text{Employees})$	The natural log of the number of employees, measured in the month before the event.	Revelio Labs
$\ln(\text{Assets})$	The natural log of total assets (atq), measured in the quarter before the event.	Compustat
Tobin's q	Tobin's q , measured in the quarter before the event.	Compustat

Variable	Definition	Source
ROA	Return on assets, computed as income before extraordinary items divided by total assets (ibq/atq), measured in the quarter before the event.	Compustat
Cash/Assets	Cash and short-term investments divided by total assets (chq/atq), measured in the quarter before the event.	Compustat
Fixed assets/Assets	Net property, plant, and equipment divided by total assets ($ppentq/atq$), measured in the quarter before the event.	Compustat
R&D/Assets	Research and development expenditures divided by total assets ($xrdq/atq$), measured in the quarter before the event. Missing values set to zero.	Compustat
Market-to-book	Shares outstanding times stock price divided by common equity ($cshoq \times prccq/ceqq$), measured in the quarter before the event.	Compustat
ln(Age)	The natural log of firm age.	CRSP
<i>Panel E: Local conditions and robustness variables</i>		
ln(Wage)	The natural log of county-level average monthly wages, computed as total quarterly wages divided by average monthly employment.	BLS Quarterly Census of Employment and Wages
County IR exposure	The county-month share of private employment in interest-rate-sensitive industries: construction (NAICS 23), real estate (NAICS 531), and finance and insurance (NAICS 52).	BLS Quarterly Census of Employment and Wages
County IR exposure (ex-ante)	County IR exposure fixed at its pre-treatment value before the InstructGPT API event.	BLS Quarterly Census of Employment and Wages
Category-specific CPI	Product-category CPI used to deflate prices in the category-CPI robustness test, constructed by mapping NielsenIQ product groups to detailed BLS CPI categories.	NielsenIQ Retail Panel; BLS Consumer Price Index (CPI)
Workplace Covid exposure (DN)	An occupation-level measure of physical proximity at work from Dingel and Neiman (2020) , aggregated to the firm level using employment shares and measured before the event.	Dingel and Neiman (2020) ; Revelio Labs
Workplace Covid exposure (ATUS)	An occupation-level measure of physical proximity at work from the American Time Use Survey, aggregated to the firm level using employment shares and measured before the event.	American Time Use Survey (ATUS); Revelio Labs
ln(County Covid cases)	The natural log of one plus monthly county-level cumulative COVID-19 cases.	NYT County-Level COVID-19 Database

Variable	Definition	Source
$\ln(\text{County Covid deaths})$	The natural log of one plus monthly county-level cumulative COVID-19 deaths.	NYT County-Level COVID-19 Database
<i>Panel F: Product-level heterogeneity variables</i>		
Product module	The most granular product category in the NielsenIQ product hierarchy used in the paper.	NielsenIQ Retail Panel
Product group	A broader NielsenIQ product category that aggregates related product modules.	NielsenIQ Retail Panel
Discretionary	An indicator variable equal to one for discretionary products, classified based on NielsenIQ product group codes.	NielsenIQ Retail Panel
Demand volatility	An indicator variable equal to one for products with above-median demand volatility, measured as the coefficient of variation (standard deviation divided by mean) of pre-treatment units sold within each product module.	NielsenIQ Retail Panel
Price elasticity	An indicator variable equal to one for products with above-median price elasticity, estimated from pre-treatment regressions of log quantity changes on log price changes.	NielsenIQ Retail Panel
<i>Panel G: AI adoption variables</i>		
AI job postings	Monthly count of active job postings classified from job titles and descriptions into broad AI, applied AI, data science, O*NET AI-related occupations, and the union of these categories.	LinkUp
AI keyword count (10-K)	The count of AI-related keywords (e.g., “artificial intelligence,” “machine learning,” “automation”) in firm i 's annual 10-K filing, measured across the full filing (Total) and separately for Item 1 (Business Description), Item 1A (Risk Factors), and Item 7 (Management's Discussion and Analysis).	SEC EDGAR
AI keyword count (earnings calls)	The count of AI-related keywords in firm i 's quarterly earnings conference call transcript, measured across the full transcript and separately within prepared remarks, management Q&A, and analyst Q&A.	Earnings conference call transcripts
Implementation/hiring keywords	A subset of 10-K AI keywords capturing references to AI adoption, deployment, and workforce investment.	SEC EDGAR
Risk/regulatory keywords	A subset of 10-K AI keywords capturing references to AI-related threats, compliance, and uncertainty.	SEC EDGAR

Variable	Definition	Source
Placebo keyword count	The count of non-AI technology keywords, including cloud, ERP, CRM, digital transformation, digitization, cybersecurity, data warehouse, software modernization, omnichannel, e-commerce infrastructure, and automation excluding AI contexts.	SEC EDGAR; earnings conference call transcripts

Internet Appendix for
Who Sets the Price? The Vertical Origins of Uniform Pricing

Alvin Chen, Leandro Sanz, and Michael Wittry

IA.A Data

IA.A.1 AI implementation exposure decomposition

The original measure of workforce exposure to generative AI developed by [Eisfeldt et al. \(2025\)](#) can manifest through two distinct channels: augmenting worker productivity while preserving human involvement, or automating the completion of tasks with minimal or no human involvement. We extend the [Eisfeldt et al. \(2025\)](#) measure by decomposing it into separate measures of augmentation and automation. We conducted this categorization in December 2025, using GPT model 5-mini, and a temperature parameter of zero.

For each occupation, we use the O*NET 30.1 database task descriptions and evaluate each task description along an augmentation and automation dimension. Our augmentation score ranges from 0 to 10 and captures the extent by which AI assistance increases worker productivity while human judgment and oversight remains important for task completion. High scores (8–10) indicate tasks where AI is expected to reduce the time to complete a task by 50% or more, or the use of AI significantly improves quality. Examples of tasks that have high augmentation scores include writing and data analysis, where AI drafts text or processes information but human oversight is important to interpret and make final decisions. Similarly, our automation score ranges from 0 to 10, and measures the likelihood that AI could complete the task independently with minimal human involvement while achieving acceptable quality. Tasks with high automation scores include structured data entry and the creation of standardized reports.

For each occupation, we aggregate task-level augmentation and automation scores by taking an equal-weighted average across tasks within an occupation, and then weighting each occupation score by the share of workers in that occupation within a firm. This decomposition allows us to distinguish firms whose workforces face primarily augmentation effects from those facing automation effects. The system prompt submitted to the GPT model is as follows:

Consider the most powerful large language models (LLMs) and AI systems currently available. These models can complete many tasks with text input and output where context can be captured in ~100,000 words. They can access up-to-date information through web search and can generate and analyze images and code. Assume you are a worker with average expertise in your role trying

to complete the given task. You have access to AI systems as well as any commonly available software, computer hardware, and technical tools accessible via a laptop. You do not have access to specialized physical tools, materials, or equipment beyond standard office technology.

Your job is to evaluate two separate dimensions for each task:

Augmentation score (0–10): Rate how much AI assistance increases worker productivity while the human remains essential to task completion.

High augmentation (8–10): AI significantly enhances human capability, reducing time by 50% or more, or substantially improving quality, but human judgment, creativity, or oversight remains critical:

- *Complex writing where AI drafts but human provides strategic direction and refinement*
- *Data analysis where AI processes information but human interprets implications and makes decisions*
- *Research where AI gathers and synthesizes information but human evaluates relevance and draws conclusions*
- *Creative work where AI generates options but human provides vision and final selection*
- *Technical work where AI handles routine coding and calculations but human designs architecture*
- *Advisory work where AI provides information and recommendations but human provides contextualized counsel*
- *Planning where AI analyzes scenarios but human makes strategic choices*

Medium augmentation (4–7): AI provides moderate assistance but human expertise drives the majority of value:

- *Editing and proofreading where AI catches errors but human ensures meaning*
- *Routine communication where AI assists with drafting but human handles nuanced interactions*
- *Information organization where AI structures data but human ensures accuracy*
- *Basic analysis where AI highlights patterns but human validates findings*

Low augmentation (1–3): AI provides minimal assistance; task primarily relies on human skills:

- *Tasks requiring real-time human interaction or relationship building*
- *Physical manipulation or in-person demonstration*
- *Work requiring tacit knowledge not easily articulated*
- *Tasks where AI assistance does not meaningfully reduce time or improve quality*

No augmentation (0): AI cannot assist meaningfully with this task.

Automation score (0–10): Rate the likelihood that AI could complete this task independently with minimal human involvement, achieving acceptable quality.

High automation (8–10): AI can complete the task end-to-end with minimal human oversight:

- *Highly structured data entry, transformation, or formatting*
- *Routine documentation following clear templates*
- *Straightforward text translation or summarization*
- *Basic code generation for well-defined specifications*
- *Repetitive communication following standard scripts*
- *Simple scheduling or record-keeping*
- *Standardized report generation from structured inputs*

Medium automation (4–7): AI can complete substantial portions independently but requires human review or input at key decision points:

- *Complex documentation where AI generates drafts requiring human validation*
- *Analysis with clear methodologies but outputs need expert interpretation*
- *Customer service where AI handles routine queries but escalates complex issues*
- *Research where AI gathers information but human validates sources and conclusions*

Low automation (1–3): AI can only handle small, discrete sub-tasks; human drives most of the work:

- *Work requiring situational judgment or contextual expertise*
- *Tasks involving negotiation or persuasion*
- *Creative problem-solving without clear parameters*

- *Processes requiring real-time adaptation based on feedback*

No automation (0): Human must perform this task; AI cannot automate it:

- *Tasks legally requiring human performance*
- *Work requiring physical presence or manipulation*
- *Tasks requiring precise measurements or detailed visual inspection*
- *Decisions directly affecting human welfare (hiring, medical diagnosis, legal judgments)*
- *Work requiring trust, empathy, or human connection*
- *Tasks where accountability must rest with a human*

Relationship between scores:

- *High augmentation + low automation = AI augments human expert (e.g., strategic analysis).*
- *High augmentation + high automation = Transitional state (AI increasingly capable, human adds value but may become optional).*
- *Low augmentation + high automation = Full automation potential (e.g., data entry).*
- *Low augmentation + low automation = AI-resistant work (e.g., skilled trades, in-person services).*

Evaluation guidelines:

- *Augmentation asks: “How much does AI help a human do this better or faster?”*
- *Automation asks: “Could AI do this acceptably without the human?”*
- *Consider both current AI capabilities and easily imaginable near-term applications.*
- *Consider whether quality or judgment trade-offs are acceptable for automation.*
- *Default to lower scores when uncertain.*
- *Account for regulatory, ethical, or practical constraints on AI autonomy.*

Output format: Return valid JSON with:

- *Augmentation: score 0–10*
- *Automation: score 0–10*

- *Augmentation_reason: ≤ 30 words explaining why AI helps humans perform this task*
- *Automation_reason: ≤ 30 words explaining whether and why AI could perform this independently*

IA.B Additional results

This section presents additional results and robustness tests that are described and referenced in the manuscript.

IA.B.1 AKM variance decomposition

Table IA.B.1 reports variance decompositions in the spirit of [Abowd, Kramarz, and Margolis \(1999\)](#) and [Card, Heining, and Kline \(2013\)](#), corresponding to the Shapley decompositions in [Table 3](#). In each panel, column (1) decomposes variance into additive UPC, retailer, county, and month fixed effects; column (2) replaces UPC with manufacturer fixed effects. $\text{Cov}(\text{UPC}, \text{Retailer})$ denotes $2 \cdot \text{Cov}(\hat{\alpha}, \hat{\psi})$, where a positive value indicates positive assortative matching between product (or manufacturer) and retailer contributions to the dependent variable.

Table IA.B.1: AKM variance decomposition

This table reports variance decompositions corresponding to the Shapley decompositions in Table 3. In each panel, column (1) decomposes variance into additive UPC, retailer, county, and month fixed effects; column (2) replaces UPC with manufacturer. $\text{Cov}(\text{UPC}, \text{Retailer})$ denotes $2 \cdot \text{Cov}(\hat{\alpha}, \hat{\psi})$. All regressions are estimated via iterative demeaning. Appendix A presents variable definitions.

Panel A: Variance of log prices		
	Product	Manufacturer
	(1)	(2)
UPC	98.2%	
Manufacturer		77.7%
Retailer	1.3%	8.6%
County	0.1%	0.3%
Month	0.8%	6.6%
Cov(UPC, Retailer)	-1.5%	8.4%
R^2	0.880	0.368
Observations (billion)	4.94	4.94

Panel B: Variance of squared deviations		
	Product	Manufacturer
	(1)	(2)
UPC	86.0%	
Manufacturer		40.4%
Retailer	8.7%	39.8%
County	1.5%	5.3%
Month	6.0%	15.6%
Cov(UPC, Retailer)	1.6%	-0.8%
R^2	0.028	0.007
Observations (billion)	4.94	4.94

Panel C: Variance of log price changes		
	Product	Manufacturer
	(1)	(2)
UPC	88.0%	
Manufacturer		41.7%
Retailer	0.5%	2.6%
County	0.1%	0.4%
Month	11.6%	52.2%
Cov(UPC, Retailer)	-0.1%	1.7%
R^2	0.008	0.002
Observations (billion)	3.75	3.75

IA.B.2 Price variance decomposition by firm size

Table IA.B.2 replicates the price variance decomposition of Table 3 separately for large and small manufacturers, split at the median of total revenue. Panel A decomposes the variance of log price levels and Panel B decomposes the variance of squared deviations from the UPC \times month mean. Large manufacturers exhibit price structures similar to the full sample, with UPC \times month effects explaining the vast majority of level variance and bilateral retailer \times UPC terms dominating the dispersion decomposition. For small manufacturers, the manufacturer fixed effect absorbs a larger share of price-level variance, and retailer and county components are proportionally larger in both panels, consistent with smaller firms facing greater cross-retailer pricing heterogeneity.

Panel B: Variance of squared deviations from UPC×month mean

	Large Manufacturers					Small Manufacturers				
	Shapley			AKM		Shapley			AKM	
	Product	Mfr	Bilateral	Product	Mfr	Product	Mfr	Bilateral	Product	Mfr
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
UPC × Month	97.6%		44.0%			79.0%		22.9%		
UPC				2.6%					4.0%	
Manufacturer		73.0%			0.3%		53.0%			1.5%
Retailer	2.2%	24.7%	0.8%	0.1%	0.1%	14.8%	33.1%	4.1%	1.6%	1.3%
Retailer × UPC			55.1%					70.2%		
County	0.2%	2.3%	0.1%	0.0%	0.0%	6.2%	13.9%	2.8%	0.3%	0.3%
Month				0.0%	0.0%				0.0%	0.0%
Cov(UPC, Retailer)				0.0%	−0.0%				−1.6%	−0.8%
R^2	0.049	0.005	0.079	0.028	0.005	0.048	0.023	0.095	0.043	0.023
Firms	175	175	175	175	175	175	175	175	175	175

IA.B.3 Category-specific CPI deflation

A potential concern is that the baseline price-level result could be driven by sectoral inflation differences: if treated firms disproportionately sell in categories that had different inflation around the event (e.g., food at home versus discretionary goods in 2022), the estimated gap could reflect composition and sectoral price dynamics rather than a within-category relative price response. To address this, we deflate prices using category-specific CPIs from the BLS. We use 26 detailed expenditure categories (e.g., cereals and bakery, dairy, meats, fruits and vegetables, beverages, apparel, personal care, pets) and map each NielsenIQ product group to the corresponding BLS category. [Table IA.B.3](#) replicates the price-level and local-wage specifications in [Table 8](#), with the dependent variable defined as the log of average price deflated by the product’s category CPI rather than the national CPI. The treatment effect remains negative and statistically significant in column (1), and the wage-interaction coefficients remain positive and statistically significant in columns (2) and (3).

Table IA.B.3: Price level with category-specific CPI deflation

This table presents results from difference-in-differences regressions of prices around the introduction of the Instruct-GPT API, where the dependent variable is the log of average price deflated by category-specific CPIs from the BLS rather than the national CPI. Prices are deflated using 26 detailed expenditure categories (e.g., cereals and bakery, dairy, meats, fruits and vegetables, nonalcoholic beverages, other food at home, apparel, personal care products, pets); each NielsenIQ product group is mapped to the corresponding BLS category. The sample, treatment definition, event window, matching, and specifications follow the baseline price-level and local-wage regressions in Table 8. Column (1) estimates the baseline price-level specification; columns (2) and (3) estimate local-wage interaction specifications. All specifications include controls for log average pre-period units interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. Fixed effects are reported in the table. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. Appendix A presents variable definitions.

	ln(Average price), category-CPI deflated		
	(1)	(2)	(3)
Treated \times Post	-0.030*** (0.010)	-0.153* (0.082)	-0.158* (0.085)
Treated \times Post \times ln(Wage)		0.017* (0.010)	0.017* (0.010)
Treated \times ln(Wage)		0.011 (0.010)	0.010 (0.010)
ln(Wage) \times Post		-0.014 (0.009)	
ln(Wage)		-0.017 (0.013)	
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
County		✓	
Calendar month	✓	✓	
County \times Calendar month			✓
Observations	89,168	12,682,944	12,682,944
Adjusted R^2	0.95	0.95	0.95

IA.B.4 Unmatched sample and continuous treatment

Our baseline results use a matched sample constructed via propensity-score matching on pre-treatment price levels and sales volumes. To assess whether the matching procedure drives the results, [Table IA.B.4](#) reports two sets of robustness tests. Columns (1), (3), and (5) re-estimate the binary treatment specification on the full unmatched sample. The unmatched estimates are close to the matched-sample estimates across all three outcomes. Price dispersion increases by 0.024, nearly identical to the matched-sample estimate of 0.025. Repricing frequency increases by 5.6 percentage points, compared with 6.4 percentage points in the matched sample. Prices decline by approximately 4.6%, close to the baseline decline of 4.3%. All three estimates are significant at the 1% level.

Columns (2), (4), and (6) replace the binary treatment indicator with continuous AI exposure, measured as the employment-share-weighted average of occupation-level generative AI scores. The signs are the same as in the binary specifications. Higher AI exposure predicts greater price dispersion, more frequent repricing, and lower average prices, with all three estimates significant at the 1% level. The consistency of the unmatched and continuous specifications suggests that the results are not sensitive to the propensity-score matched sample or to the binary treatment classification.

Table IA.B.4: Unmatched sample and continuous treatment

This table reports two alternative difference-in-differences specifications for the AI experiment around the introduction of the InstructGPT API. The dependent variables are the standard deviation of log prices across stores within a firm-product-month (columns (1) and (2)), an indicator for absolute month-over-month changes in log average price exceeding 5% (columns (3) and (4)), and the log of CPI-deflated average price (columns (5) and (6)). The sample is a firm-product-month panel of publicly listed U.S. firms. Odd columns re-estimate the baseline binary-treatment specification on the full unmatched sample. Even columns replace the binary treatment indicator with continuous AI exposure, measured as the employment-share-weighted average of occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#), on a Mahalanobis-matched sample. *Treated* equals one for firms with above-median AI exposure measured in January 2020. *Post* equals one during the six months following the event. All specifications include controls for log average pre-period units interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin’s q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. All specifications include UPC and calendar month fixed effects. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	Standard deviation		Pr($ \Delta P > 5\%$)		ln(Average price)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	0.024*** (0.005)		0.056*** (0.016)		-0.046*** (0.012)	
Exposure \times Post		1.324*** (0.356)		2.919*** (0.938)		-2.621*** (0.847)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC	✓	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓	✓
Observations	113,987	57,390	133,457	63,499	133,457	63,499
Adjusted R^2	0.48	0.45	0.32	0.31	0.96	0.93

IA.B.5 Specifying the target and judge model

Our baseline AI exposure measure uses the [Eisfeldt et al. \(2025\)](#) scores, which evaluate each O*NET task against the capabilities of GPT 3.5 Turbo. An alternative approach is to specify the capabilities of a particular target model and use a separate judge model to evaluate whether the target model can perform each task. We implement this approach using GPT-5-mini as the judge model and text-davinci-001 as the target model, whose capabilities more closely match those available at the time of the InstructGPT API introduction. Scores are aggregated to the firm level using employment shares, and treatment is defined as above-median exposure, identical to the baseline.

[Table IA.B.5](#), [Table IA.B.6](#), and [Table IA.B.7](#) replicate the baseline price dispersion, price changes, and price level results using the Davinci-01 exposure measure. Across all specifications, the results are consistent with the baseline results. That is, we find that price dispersion increases, repricing frequency increases, and prices decline. The corresponding event-study estimates ([Figure IA.B.1](#), [Figure IA.B.2](#), and [Figure IA.B.3](#)) display parallel pre-trends and post-event treatment effects similar to those in the main analysis.

Figure IA.B.1: Price dispersion event study, using AI exposure based on Davinci-01

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $\text{Dispersion}_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \neq -3} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), and t calendar months. In Panel A, the dependent variable is the standard deviation of log prices across stores within a firm-product-month. In Panel B, the dependent variable is the interquartile range of log prices across stores within a firm-product-month. AI exposure is derived from a target-judge model framework, where GPT-5-mini evaluates whether the text-davinci-001 model can perform O*NET occupational tasks. Scores are aggregated to the firm level using employment shares. *Treated* equals one for firms with above-median AI exposure measured in January 2020. Event time runs from -6 to $+6$ months relative to the implementation event, with $k = -3$ as the omitted base period. α_{ip} and α_t denote firm-product and calendar month fixed effects, respectively. Treated and control firms are balanced on pre-treatment log average price and log units sold. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.

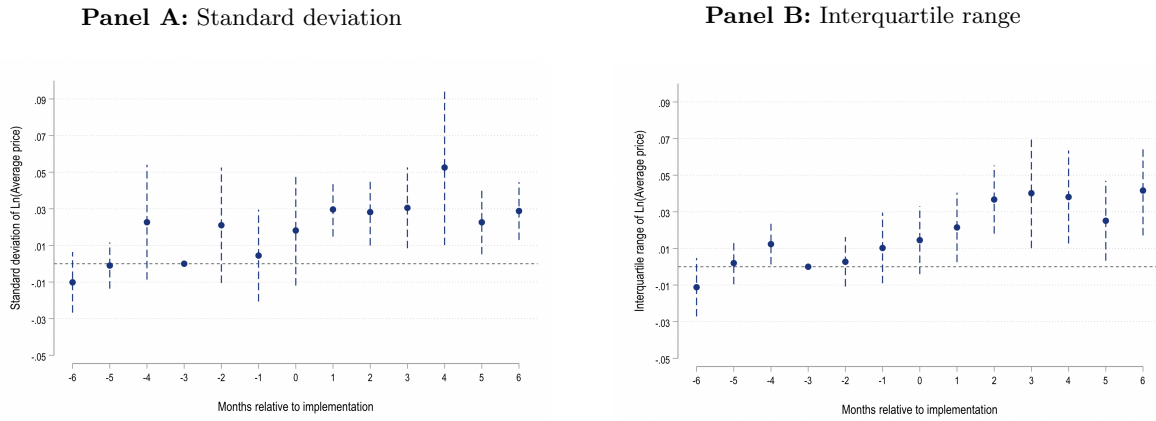


Figure IA.B.2: Price changes event study, using AI exposure based on Davinci-01

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $\mathbb{P}(|\Delta \ln \bar{P}| > 5\%)_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \neq -3} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), and t calendar months. The dependent variable is an indicator for whether the month-over-month change in log average price exceeds a given threshold. In Panel A, the dependent variable is an indicator for whether the absolute change in log average price exceeds 5%. Panels B and C focus on price increases and decreases, respectively. AI exposure is derived from a target-judge model framework, where GPT-5-mini evaluates whether the text-davinci-001 model can perform O*NET occupational tasks. Scores are aggregated to the firm level using employment shares. *Treated* equals one for firms with above-median AI exposure measured in January 2020. Event time runs from -6 to +6 months relative to the InstructGPT API introduction in January 2022, with $k = -3$ as the omitted base period. α_{ip} and α_t denote firm-product and calendar month fixed effects, respectively. Treated and control firm products are balanced on the pre-treatment log average price and log units sold. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.

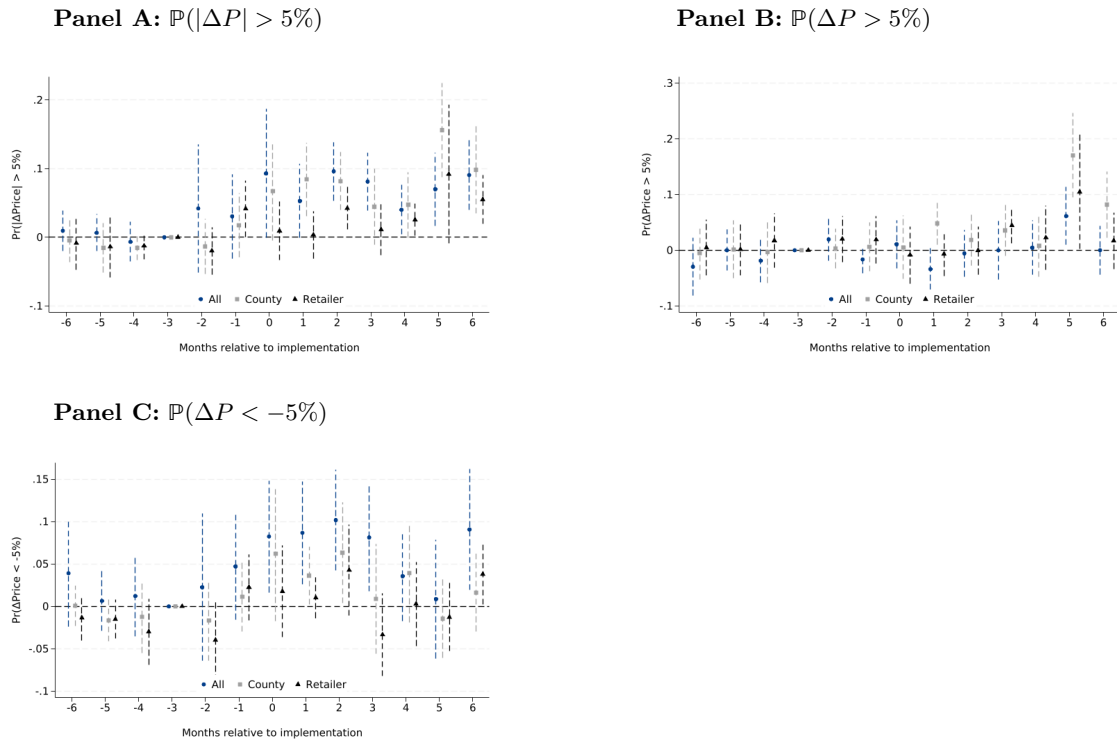


Figure IA.B.3: Price level event study, using AI exposure based on Davinci-01

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $\ln(\text{Average price})_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \neq -3} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), and t calendar months. The dependent variable is the log of CPI-deflated average price. AI exposure is derived from a target-judge model framework, where GPT-5-mini evaluates whether the text-davinci-001 model can perform O*NET occupational tasks. Scores are aggregated to the firm level using employment shares. *Treated* equals one for firms with above-median AI exposure measured in January 2020. Event time runs from -6 to $+6$ months relative to the event, with $k = -3$ as the omitted base period. α_{ip} and α_t denote firm-product and calendar month fixed effects, respectively. Treated and control firms are balanced on pre-treatment log average price and log units sold. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.

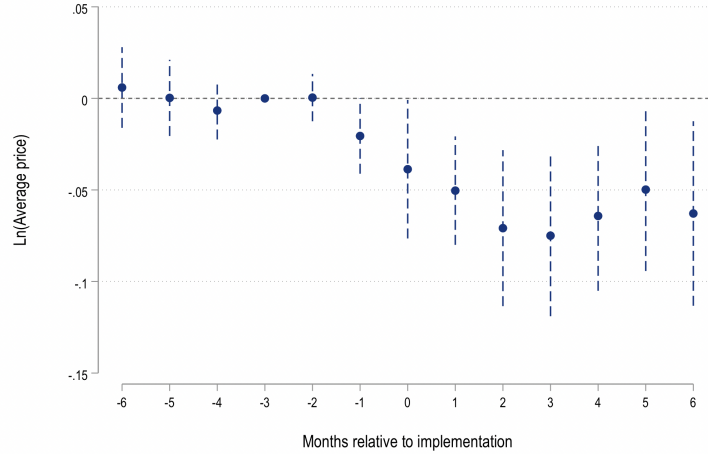


Table IA.B.5: Price dispersion, using AI exposure based on Davinci-01

This table presents results from difference-in-differences regressions of price dispersion around the introduction of the InstructGPT API. The sample is a firm-product-month panel of publicly listed U.S. firms. The dependent variable in columns (1) to (3) is the standard deviation of log prices, and in columns (4) to (6) is the interquartile range of log prices, both computed across stores within a firm-product-month. Columns (1) and (4) compute dispersion across all stores; columns (2) and (5) compute dispersion within counties; and columns (3) and (6) compute dispersion within retailers. AI exposure is derived from a target-judge model framework, where GPT-5-mini evaluates whether the text-davinci-001 model can perform O*NET occupational tasks. Scores are aggregated to the firm level using employment shares. *Treated* equals one for firms with above-median AI exposure measured in January 2020. *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. All specifications include controls for log average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. All specifications include firm-product and calendar month fixed effects. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	Standard deviation			Interquartile range		
	All	County	Retailer	All	County	Retailer
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	0.023*** (0.005)	0.010*** (0.003)	0.009** (0.004)	0.022*** (0.006)	0.002** (0.001)	0.007** (0.003)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC	✓	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓	✓
Observations	77,912	63,354	75,528	77,912	63,354	75,528
Adjusted R^2	0.47	0.46	0.46	0.51	0.79	0.45

Table IA.B.6: Price changes, using AI exposure based on Davinci-01

This table presents results from difference-in-differences regressions of the frequency of price changes around the introduction of the InstructGPT API. The dependent variable is an indicator for whether the month-over-month change in log average price exceeds 5% at the level of aggregation used in the column. Columns (1) to (3) use the absolute value of that change; columns (4) to (6) isolate price increases; and columns (7) to (9) isolate price decreases. Within each group, columns are estimated at the firm-product level (All), the firm-product-county level (County), and the firm-product-retailer level (Retail). AI exposure is derived from a target-judge model framework, where GPT-5-mini evaluates whether the text-davinci-001 model can perform O*NET occupational tasks. Scores are aggregated to the firm level using employment shares. *Treated* equals one for firms with above-median AI exposure measured in January 2020. *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. All specifications include controls for log average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. The All specifications include firm-product and calendar month fixed effects; the County specifications include UPC \times county and calendar month fixed effects; and the Retail specifications include UPC \times retailer and calendar month fixed effects. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	$\mathbb{P}(\% \Delta P > 5\%)$			$\mathbb{P}(\% \Delta P > 5\%)$			$\mathbb{P}(\% \Delta P < -5\%)$		
	All	County	Retail	All	County	Retail	All	County	Retail
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated \times Post	0.061*** (0.015)	0.087*** (0.028)	0.035*** (0.013)	0.013 (0.010)	0.051*** (0.012)	0.013 (0.011)	0.049*** (0.013)	0.036 (0.024)	0.022* (0.012)
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Fixed effects									
UPC	✓			✓			✓		
UPC \times County		✓			✓			✓	
UPC \times Retailer			✓			✓			✓
Calendar month	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	89,776	13,296,290	941,982	89,776	13,296,290	941,982	89,776	13,296,290	941,982
Adjusted R^2	0.31	0.20	0.25	0.09	0.05	0.07	0.13	0.06	0.10

Table IA.B.7: Price level and local economic conditions, using AI exposure based on Davinci-01

This table examines whether average prices fall and become more sensitive to local economic conditions after the introduction of the InstructGPT API, using the Davinci-01 AI exposure measure. The dependent variable is the log of CPI-deflated average price. AI exposure is derived from a target-judge model framework, where GPT-5-mini evaluates whether the text-davinci-001 model can perform O*NET occupational tasks. Scores are aggregated to the firm level using employment shares. *Treated* equals one for firms with above-median AI exposure measured in January 2020. *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. Column (1) reports the baseline price-level treatment effect on the firm-product-month panel, with firm-product and calendar month fixed effects. Columns (2) and (3) extend the design to a firm-product-county-month panel and interact the treatment indicator with county-level log average monthly wages from the BLS Quarterly Census of Employment and Wages. Column (2) includes UPC, county, and calendar month fixed effects. Column (3) replaces the separate county and calendar month fixed effects with county×calendar month fixed effects, which absorb local time-varying shocks. All specifications include controls for the log of average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin’s q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	ln(Average price)		
	(1)	(2)	(3)
Treated × Post	−0.056*** (0.011)	−0.189* (0.102)	−0.194* (0.104)
Treated × Post × ln(Wage)		0.019 (0.012)	0.020* (0.012)
Treated × ln(Wage)		0.006 (0.011)	0.006 (0.011)
ln(Wage) × Post		−0.018 (0.012)	
ln(Wage)		−0.017 (0.011)	
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
County		✓	
Calendar month	✓	✓	
County × Calendar month			✓
Observations	89,255	11,750,846	11,750,846
Adjusted R^2	0.95	0.95	0.95

IA.B.6 COVID-19 exposure

The InstructGPT API introduction in January 2022 occurred while the economy was still adjusting to the effects of the COVID-19 pandemic. If firms with high AI exposure also had workforces that were more exposed to pandemic-related workplace disruptions, COVID-19 effects could confound the estimated treatment effects. [Table IA.B.8](#) addresses this concern by augmenting the baseline specification with direct controls for pandemic exposure. Panel A adds firm-level measures of workplace physical proximity—the [Dingel and Neiman \(2020\)](#) work-from-home index and an analogous measure from the American Time Use Survey—interacted with post. Panel B moves to the firm-product-county-month panel and controls for cumulative county-level COVID-19 cases and deaths. Across all specifications, the treatment effects remain statistically significant and consistent with the baseline, and the COVID-19 controls are generally small.

Table IA.B.8: Controlling for COVID-19 exposure

This table presents results from difference-in-differences regressions of prices around the introduction of the InstructGPT API, controlling for COVID-19 exposure. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Post* equals one during the six months following the event. Treated and control firms are balanced on pre-treatment log average price and log units sold. The dependent variables are the standard deviation of log prices (columns (1) and (2)), an indicator for absolute month-over-month changes in log average price exceeding 5% (columns (3) and (4)), and the log of CPI-deflated average price (columns (5) and (6)). Panel A uses the firm-product-month panel with firm-product and calendar month fixed effects. *Workplace Covid exposure (DN)* and *Workplace Covid exposure (ATUS)* are occupation-level measures of physical proximity at work from [Dingel and Neiman \(2020\)](#) and the American Time Use Survey, respectively, aggregated to the firm level using employment shares and interacted with *post*. Odd columns control for the [Dingel and Neiman \(2020\)](#) measure and even columns control for the ATUS measure. Panel B uses the firm-product-county-month panel with firm-product \times county and calendar month fixed effects. Odd columns control for the log of one plus cumulative county-level COVID-19 cases and even columns control for the log of one plus cumulative county-level COVID-19 deaths. All specifications include controls for log average pre-period units sold interacted with *post*, log employees measured in the month before the event and interacted with *post*, and firm financial controls (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with *post*. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

Panel A: Firm-level COVID-19 exposure						
	Standard deviation		$\mathbb{P}(\% \Delta P > 5\%)$		ln(Average price)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	0.023** (0.009)	0.027*** (0.008)	0.045** (0.017)	0.060*** (0.022)	-0.051*** (0.012)	-0.047*** (0.012)
Workplace Covid exposure (DN) \times Post	0.000 (0.044)		-0.141 (0.085)		-0.060 (0.067)	
Workplace Covid exposure (ATUS) \times Post		0.091 (0.075)		-0.047 (0.198)		-0.091 (0.128)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC	✓	✓	✓	✓	✓	✓
UPC \times County						
Calendar month	✓	✓	✓	✓	✓	✓
Observations	77,604	77,604	89,147	89,147	89,147	89,147
Adjusted R^2	0.44	0.44	0.31	0.31	0.95	0.95

Panel B: County-level COVID-19 exposure

	Standard deviation		$\mathbb{P}(\% \Delta P > 5\%)$		$\ln(\text{Average price})$	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	0.011*** (0.003)	0.011*** (0.003)	0.064*** (0.017)	0.064*** (0.017)	-0.020*** (0.005)	-0.020*** (0.005)
$\ln(\text{County Covid cases})$	-0.001 (0.003)		-0.006 (0.008)		-0.006** (0.003)	
$\ln(\text{County Covid deaths})$		0.001 (0.003)		0.001 (0.004)		0.003 (0.002)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC						
UPC \times County	✓	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓	✓
Observations	11,118,967	11,118,967	16,456,600	16,456,600	16,456,600	16,456,600
Adjusted R^2	0.31	0.31	0.20	0.20	0.97	0.97

IA.B.7 Interest rate exposure

The Federal Reserve began raising interest rates in March 2022, shortly after the InstructGPT API introduction. If AI exposure is correlated with firms' geographic concentration in interest-rate-sensitive sectors, the price decline could reflect local demand contractions rather than AI implementation. [Table IA.B.9](#) augments the baseline specification with controls for county-level interest rate exposure, measured as the share of private employment in interest-rate-sensitive industries (NAICS 23, 531, and 52). Odd columns control for the time-varying monthly IR exposure measure, and even columns use the baseline (January 2020) value to rule out endogenous changes in local industry composition. The treatment effects are virtually identical to the corresponding county-level baseline estimates across all six specifications, indicating that differential exposure to the interest rate tightening cycle does not drive the results.

Table IA.B.9: Controlling for interest rate exposure

This table presents results from difference-in-differences regressions of prices around the introduction of the InstructGPT API, controlling for county-level interest rate exposure. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Post* equals one during the six months following the implementation event. Treated and control firms are balanced on pre-treatment log average price and log units sold. The dependent variables are the standard deviation of log prices (columns (1) and (2)), an indicator for absolute month-over-month changes in log average price exceeding 5% (columns (3) and (4)), and the log of CPI-deflated average price (columns (5) and (6)). *County IR exposure* is the share of private employment in interest-rate-sensitive industries (NAICS 23: construction, NAICS 531: real estate, and NAICS 52: finance and insurance), constructed using data from the BLS Quarterly Census of Employment and Wages. Odd columns control for the time-varying monthly county IR exposure measure interacted with post. Even columns control for county IR exposure fixed at its pre-treatment (January 2020) value interacted with post. All specifications include UPC \times county and calendar month fixed effects. Controls include log average pre-period units sold interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	Standard deviation		$\mathbb{P}(\% \Delta P > 5\%)$		ln(Average price)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	0.011*** (0.003)	0.011*** (0.003)	0.064*** (0.017)	0.064*** (0.017)	-0.020*** (0.005)	-0.020*** (0.005)
County IR exposure \times Post	-0.000 (0.000)		-0.001 (0.001)		-0.002*** (0.000)	
County IR exposure (ex-ante) \times Post		-0.000 (0.000)		-0.001 (0.001)		-0.002*** (0.000)
County IR exposure	-0.000 (0.000)		-0.000 (0.001)		0.001*** (0.000)	
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC						
UPC \times County	✓	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓	✓
Observations	11,079,256	11,092,395	16,399,494	16,420,473	16,399,494	16,420,473
Adjusted R^2	0.31	0.31	0.20	0.20	0.97	0.97

IA.B.8 Placebo tests

We conduct two types of placebo tests: date-based placebos that move the event window away from the actual API introduction, and randomization-based placebos that reassign treatment status across firms.

Placebo event dates. If the estimated treatment effects reflect a genuine response to the InstructGPT API introduction, they should not appear when the event date is moved to a period with no plausible AI shock. [Table IA.B.10](#) presents placebo tests using January 2019 (36 months before the actual event) and January 2021 (12 months before) as alternative event dates. For each placebo date, we re-estimate the baseline specification for all three main outcomes. Across all six specifications, the placebo coefficients are small, of inconsistent sign, and statistically insignificant, supporting the interpretation that the main results are driven by the actual AI implementation event.

Randomization inference. As an alternative to conventional asymptotic inference, we conduct a permutation test that directly evaluates the likelihood of observing our treatment effects under random assignment. In each of 5,000 iterations, we randomly reassign the treatment label across firms—preserving the number of treated firms—drop three outlier firms whose product portfolios are sufficiently large to exert disproportionate influence on the permutation distribution, re-match treated and control products using propensity-score matching on pre-treatment price levels and sales volumes, and re-estimate the full baseline specification. For each regression, we store the t -statistic on the treatment interaction. [Figure IA.B.4](#) plots the resulting placebo distributions alongside the t -statistics from the baseline specification estimated on the restricted sample (dashed vertical lines). Across all three outcomes, the true t -statistics fall in the extreme tails of the placebo distributions: the two-sided permutation p -values are 0.015 for price dispersion, 0.009 for repricing frequency, and 0.011 for average prices. These results provide nonparametric evidence that the estimated treatment effects are unlikely to arise from any arbitrary partition of firms into treated and control groups.

Table IA.B.10: Placebo tests using January 2019 and January 2021 as event dates

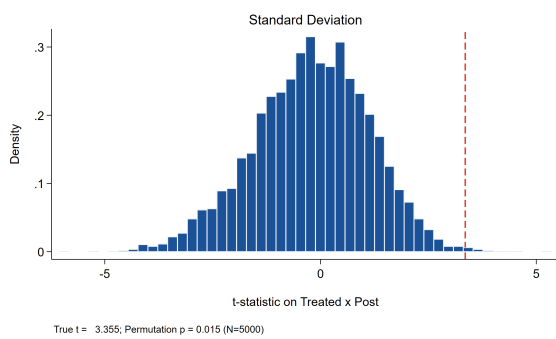
This table presents results from placebo tests that move the event date away from the actual InstructGPT API introduction in January 2022. Columns (1) to (3) use January 2019 (36 months before) and columns (4) to (6) use January 2021 (12 months before) as placebo event dates. Within each placebo date, the dependent variables are within-product price standard deviation, an indicator for whether the absolute month-over-month change in log average price exceeds 5%, and the log of CPI-deflated average price. *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Post* equals one during the six months following the placebo event date. All specifications include UPC and calendar month fixed effects. Controls include log average pre-period units sold interacted with post, log employees interacted with post, and firm financial controls measured in the quarter before the placebo event and interacted with post. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	January 2019			January 2021		
	Standard deviation	$\mathbb{P}(\% \Delta P > 5\%)$	$\ln(\text{Average price})$	Standard deviation	$\mathbb{P}(\% \Delta P > 5\%)$	$\ln(\text{Average price})$
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	0.010 (0.007)	-0.007 (0.012)	-0.008 (0.011)	0.002 (0.009)	0.015 (0.014)	-0.015 (0.019)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC	✓	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓	✓
Observations	91,843	95,254	95,254	97,240	109,531	109,531
Adjusted R^2	0.49	0.27	0.96	0.39	0.28	0.95

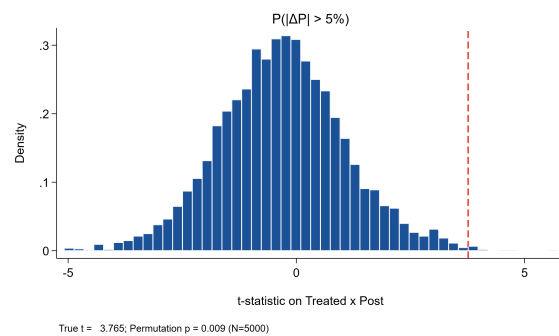
Figure IA.B.4: Placebo tests randomizing treatment status

This figure presents the distribution of placebo t -statistics from 5,000 random permutations of the treatment indicator. In each iteration, we randomly reassign the treated label across firms (preserving the number of treated firms), drop three outlier firms whose product portfolios are sufficiently large to exert disproportionate influence on the permutation distribution, re-match treated and control products using propensity-score matching on pre-treatment price levels and sales volumes, and re-estimate the baseline difference-in-differences specification. Panel A reports results for the standard deviation of log price, Panel B for the frequency of absolute changes in log average price exceeding 5%, and Panel C for log average price. The dashed vertical line indicates the true estimated t -statistic. Two-sided permutation p -values are 0.015, 0.009, and 0.011 for Panels A, B, and C, respectively. All specifications include firm-product and calendar month fixed effects, with controls for log average pre-period units sold interacted with post, log employees interacted with post, and firm financial characteristics (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) interacted with post. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.

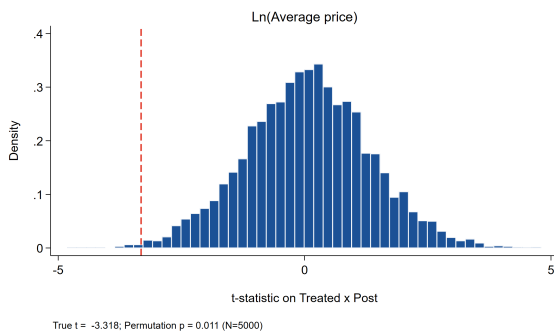
Panel A: Standard deviation



Panel B: $\mathbb{P}(|\Delta P| > 5\%)$



Panel C: $\ln(\text{Average price})$



IA.B.9 Alternative augmentation vs. automation decomposition

As a complement to the augmentation vs. automation decomposition in Table 11, we examine an alternative decomposition based on the Anthropic Economic Index (Appel et al., 2026). The AEI maps human–AI conversations to O*NET work activities and classifies the mode of interaction between the user and the model. This allows us to construct a second set of occupation-level measures without relying on our GPT-5-mini task scores. We proxy augmentation with the *task iteration* mode, in which users iteratively refine outputs with the AI model. We proxy automation with the *directive* mode, in which users issue a single instruction and the AI executes independently.

Table IA.B.11 reports the results. Columns (1) to (3) use task iteration exposure. Firms with above-median task iteration exposure exhibit higher price dispersion, more frequent repricing, and lower average prices after the event. The estimates imply a 0.020 increase in the standard deviation of log prices, a 5.5 percentage point increase in repricing frequency, and a price decline of 4.2%. Columns (4) to (6) use directive exposure. The corresponding estimates are smaller and less precisely estimated. Directive exposure is not significantly associated with price dispersion or average prices, and the repricing effect is only marginally significant. Overall, these results mirror the augmentation vs. automation pattern and provide independent corroboration using an external data source. The pricing effects of AI therefore load on occupations where workers interact iteratively with AI tools, not where AI executes tasks autonomously.

Table IA.B.11: Augmentation vs. automation (Anthropic Economic Index)

This table presents results from difference-in-differences regressions that decompose AI exposure into augmentation and automation components using the Anthropic Economic Index (AEI). The AEI maps human–AI conversations to O*NET work activities and classifies the mode of interaction between the user and the model. Augmentation is proxied by the *task iteration* mode, in which users iteratively refine outputs with the AI model. Automation is proxied by the *directive* mode, in which users issue single instructions that the AI executes independently. Firm-level scores are computed as employment-share-weighted averages of occupation-level AEI scores using employment shares from Revelio Labs. Treatment is defined as above-median exposure measured in January 2020, with propensity score matching on pre-event log average price and log units sold. *Post* equals one during the six months following the InstructGPT API introduction in January 2022. Within each block, columns report standard deviation, repricing, and log average price. All specifications include firm-product and calendar month fixed effects and controls for log average pre-period units sold interacted with post, log employees interacted with post, and firm financial characteristics interacted with post. Standard errors clustered at the firm level are reported in parentheses. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	Augmentation (task iteration)			Automation (directive)		
	Standard deviation	$\mathbb{P}(\% \Delta P > 5\%)$	$\ln(\text{Average price})$	Standard deviation	$\mathbb{P}(\% \Delta P > 5\%)$	$\ln(\text{Average price})$
	(1)	(2)	(3)	(4)	(5)	(6)
Treated \times Post	0.020*** (0.006)	0.055*** (0.019)	-0.042** (0.017)	0.014 (0.010)	0.040* (0.021)	-0.014 (0.017)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC	✓	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓	✓
Observations	61,580	70,405	71,179	70,192	80,265	81,110
Adjusted R^2	0.44	0.31	0.94	0.44	0.30	0.95

IA.B.10 AI adoption: Earnings conference calls

As a complement to the job posting and 10-K analyses in the main text, we examine AI-related keywords in quarterly earnings conference call transcripts, which provide more timely information about corporate strategy. Table IA.B.12 reports PPML regressions of AI keyword counts on the $Treated \times Post$ interaction, with firm and calendar quarter fixed effects. High-exposure firms significantly increase AI keyword usage in earnings calls after the event, both in total and across several of the individual sections (prepared remarks and management Q&A). Figure IA.B.5 presents the corresponding event-study estimates, which show parallel pre-trends and a sustained increase in AI mentions beginning at the event date, consistent with the job posting and 10-K evidence.

Figure IA.B.5: AI keyword event study using earnings conference calls

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $AI\ Keywords_{it} = \alpha_i + \alpha_t + \sum_{k \neq -1} \beta_k \mathbb{1}[\text{Event time} = k] \times Treated_i + \epsilon_{it}$, where i denotes firms and t calendar quarters. The dependent variable is the count of AI-related keywords (artificial intelligence, machine learning, ChatGPT, LLM, neural network, predictive analytics, pricing algorithm) in firm i 's earnings call transcript in quarter t . AI exposure is computed as the employment-share-weighted average of occupation-level generative AI scores from Eisfeldt et al. (2025), where scores capture the extent to which generative AI is expected to reduce the time required to complete the tasks typically performed by each occupation. $Treated$ equals one for firms with above-median AI exposure measured in January 2020. Event time runs from -4 to $+4$ quarters relative to the InstructGPT API introduction in 2022 Q1, with $k = -1$ as the omitted base period. α_i and α_t denote firm and calendar quarter fixed effects, respectively. Standard errors are clustered at the firm level. Appendix A presents variable definitions.




Table IA.B.12: AI keyword counts in earnings conference calls

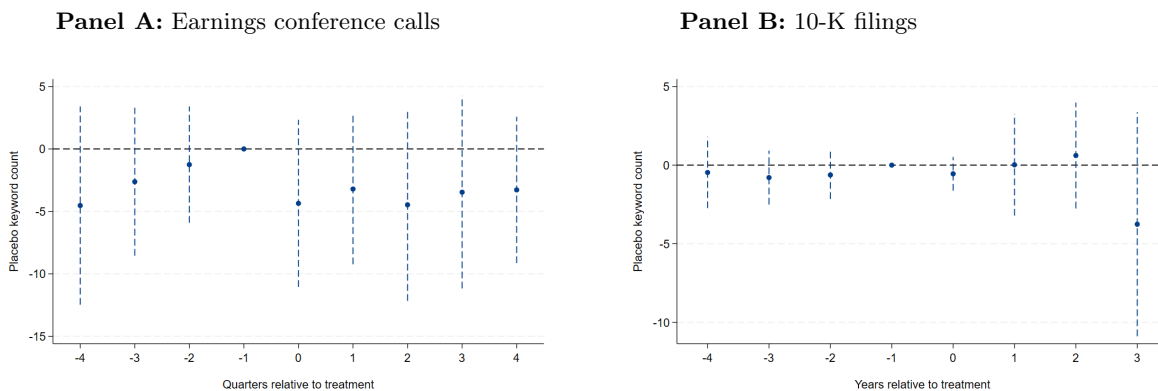
This table presents results from Poisson pseudo-maximum-likelihood (PPML) regressions of AI-related keyword counts in quarterly earnings conference call transcripts around AI implementation events. The dependent variable is the count of AI-related keywords (artificial intelligence, machine learning, ChatGPT, LLM, neural network, predictive analytics, pricing algorithm) in firm i 's earnings call transcript in quarter t . Column (1) uses the total keyword count across the full transcript. Columns (2)–(4) decompose the count by transcript section: prepared remarks (column (2)), management Q&A (column (3)), and analyst Q&A (column (4)). *Treated* equals one for firms with above-median AI exposure, measured in January 2020 using employment-share-weighted occupation-level generative AI scores from Eisfeldt et al. (2025). *Post* equals one after the InstructGPT API introduction 2022 Q1. All specifications include firm and calendar quarter fixed effects. Standard errors clustered at the firm level are reported in parentheses below coefficient estimates. Statistical significance at the 1, 5, and 10 percent significance levels is denoted by ***, **, and *, respectively. Appendix A presents variable definitions.

	AI keyword count			
	Total	Prepared remarks	Management Q&A	Analyst Q&A
	(1)	(2)	(3)	(4)
Treated \times Post	1.800*** (0.609)	1.629** (0.746)	2.265*** (0.740)	0.348 (2.092)
Fixed effects				
Firm	✓	✓	✓	✓
Quarter	✓	✓	✓	✓
Observations	297	238	155	30

The placebo keyword event studies provide a falsification test for the earnings-call evidence. We replace AI keywords with non-AI technology terms, including cloud computing, ERP, CRM, cybersecurity, e-commerce infrastructure, and enterprise systems. If high-exposure firms were simply increasing all technology-related discussion after 2022, the placebo series should exhibit a similar post-event increase. [Figure IA.B.6](#) shows no differential increase in either earnings calls (Panel A) or 10-K filings (Panel B). These patterns support the interpretation that the validation results capture AI adoption rather than a broad shift in technology language.

Figure IA.B.6: Placebo keyword event studies

This figure plots the β_k coefficients and 90% confidence intervals from event study regressions of the form: $\text{Placebo Keywords}_{it} = \alpha_i + \alpha_t + \sum_{k \neq -1} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{it}$, where i denotes firms and t denotes calendar periods. The dependent variable is the count of non-AI technology keywords (cloud, ERP, CRM, digital transformation, digitization, cybersecurity, data warehouse, software modernization, omnichannel, e-commerce infrastructure, automation excluding AI contexts, enterprise systems). In Panel A, the dependent variable is measured in quarterly earnings conference call transcripts, and event time runs from -4 to $+4$ quarters relative to the Instruct-GPT API introduction in 2022 Q1. In Panel B, the dependent variable is measured in annual 10-K filings, and event time runs from -4 to $+3$ years relative to 2022. AI exposure is computed as the employment-share-weighted average of occupation-level generative AI scores from [Eisfeldt et al. \(2025\)](#). *Treated* equals one for firms with above-median AI exposure measured in January 2020. α_i and α_t denote firm and calendar period fixed effects, respectively, with $k = -1$ as the omitted base period. Standard errors are clustered at the firm level. All coefficients are statistically indistinguishable from zero, indicating no differential trend in non-AI technology discussion between treated and control firms. [Appendix A](#) presents variable definitions.



IA.B.11 Sample composition and inference

Leave-one-out by manufacturer and category. The matched sample contains 91 manufacturers and roughly 100 product categories, raising the concern that the estimates could be driven by a small number of firms or categories. We address this by re-estimating the baseline specification 91 times, dropping each manufacturer in turn, and a further 100 times, dropping each product category in turn. [Figure IA.B.7](#) ranks the resulting coefficients for the standard deviation of log prices (Panel A), the probability of an absolute change in log average price exceeding 5% (Panel B), and log average prices (Panel C); [Figure IA.B.8](#) reports the analogous decomposition by product category. Across all three outcomes and both exercises, the estimated treatment effect retains its sign relative to the full-sample estimate (dashed horizontal line) in every iteration. No single manufacturer and no single product category reverses the headline pattern of lower prices, greater dispersion, and more frequent repricing among AI-exposed firms.

Wild cluster bootstrap inference. The baseline AI regressions cluster standard errors at the firm level, and the matched sample contains fewer than 100 firm clusters. To address finite-sample concerns, [Table IA.B.13](#) reports wild cluster bootstrap t -test p -values for the three main outcomes. The bootstrap imposes the null hypothesis and uses 9,999 replications with Webb six-point weights. The resulting p -values are 0.003 for the standard deviation of log prices, 0.001 for the probability of an absolute change in log average price exceeding 5%, and below 0.001 for log average prices. The main estimates therefore remain significant at the 1% level under finite-cluster inference, suggesting that the results are not driven by reliance on asymptotic cluster-robust standard errors.

Figure IA.B.7: Leave-one-out manufacturer analysis

This figure presents the estimated difference-in-differences coefficient on $Treated \times Post$ when iteratively dropping each manufacturer from the sample. For each of the 91 manufacturers in the matched sample, we exclude that manufacturer and re-estimate the baseline specification. Coefficients are ranked from smallest to largest, plotted with 90% confidence intervals. The dashed horizontal lines mark the full-sample estimate and zero. Panel A reports results for the standard deviation of log prices, Panel B for the probability of an absolute change in log average price exceeding 5%, and Panel C for log average price. All specifications include UPC and calendar month fixed effects, with controls for log average pre-period units sold interacted with post, log employees interacted with post, and firm financial characteristics (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) interacted with post. Standard errors are clustered at the firm level. [Appendix A](#) presents variable definitions.

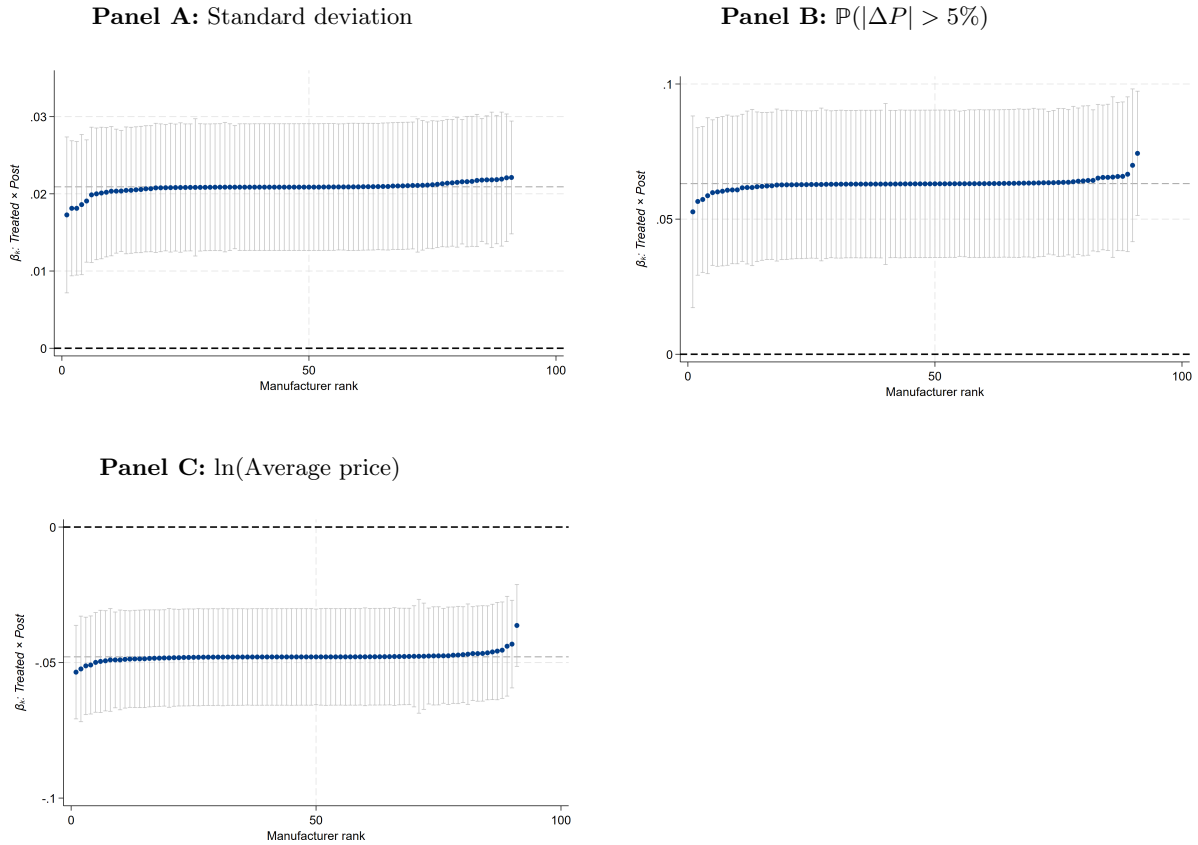


Figure IA.B.8: Leave-one-out product category analysis

This figure presents the estimated difference-in-differences coefficient on $Treated \times Post$ when iteratively dropping each product category from the sample. For each of the 100 NielsenIQ product groups in the matched sample, we exclude all products in that category and re-estimate the baseline specification. Coefficients are ranked from smallest to largest, plotted with 90% confidence intervals. The dashed horizontal lines mark the full-sample estimate and zero. Panel A reports results for the standard deviation of log prices, Panel B for the probability of an absolute change in log average price exceeding 5%, and Panel C for log average price. All specifications include UPC and calendar month fixed effects, with controls for log average pre-period units sold interacted with post, log employees interacted with post, and firm financial characteristics (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) interacted with post. Standard errors are clustered at the firm level. Appendix A presents variable definitions.

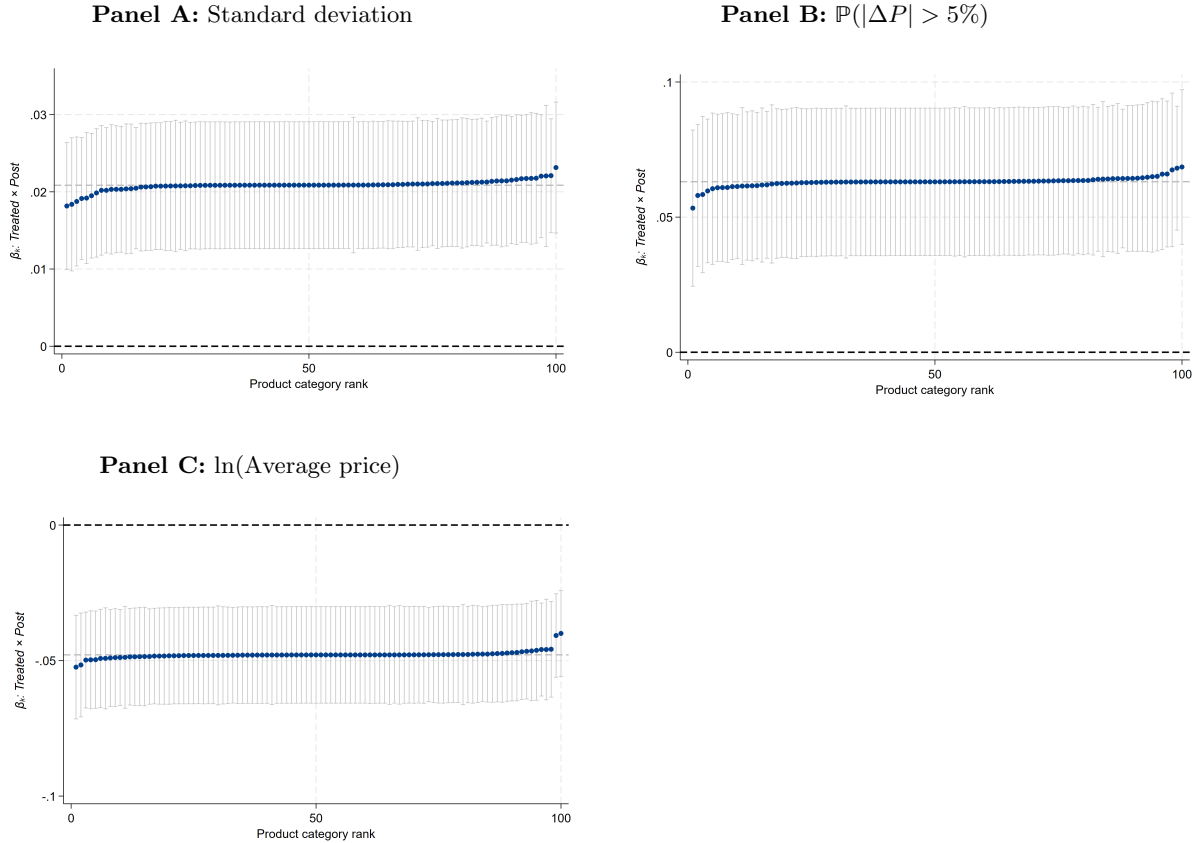


Table IA.B.13: Wild cluster bootstrap inference

This table reports finite-cluster inference for AI-experiment regressions covering the three main outcomes. Columns (1) to (3) report estimates for the standard deviation of log prices, the probability of an absolute change in log average price exceeding 5%, and $\ln(\text{Average price})$, respectively. Each column reports the coefficient on $Treated \times Post$, firm-clustered standard errors in parentheses, and the wild cluster bootstrap t -test p -value below. The bootstrap imposes the null hypothesis and uses 9,999 replications with Webb six-point weights. Significance stars are based on the wild cluster bootstrap p -values. All specifications include UPC and calendar month fixed effects, with controls for log average pre-period units sold interacted with post, log employees interacted with post, and firm financial characteristics (log assets, Tobin's q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) interacted with post. [Appendix A](#) presents variable definitions.

	Standard deviation	$\mathbb{P}(\% \Delta P > 5\%)$	$\ln(\text{Average price})$
	(1)	(2)	(3)
Treated \times Post	0.025*** (0.006)	0.071*** (0.018)	-0.043*** (0.010)
Wild bootstrap p -value	0.003	0.001	<0.001
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
Calendar month	✓	✓	✓
Observations	77,614	80,782	89,195
Adjusted R^2	0.46	0.31	0.95

IA.B.12 Technology-task intensity placebo

The main text shows that the pricing effects are concentrated in AI exposure attached to pricing-related tasks. To assess whether those results instead reflect general technological sophistication, [Table IA.B.14](#) reports a placebo test using technology-task intensity, constructed from O*NET tasks containing non-AI technology keywords (software, database, network, programming, hardware) but excluding pricing-related terms. Technology-task intensity is not associated with price dispersion or average prices, although it predicts repricing frequency.

Table IA.B.14: Technology-task intensity placebo

This table reports placebo difference-in-differences regressions using technology-task intensity, defined as the employment-share-weighted average of each occupation’s share of non-AI technology tasks (software, database, network, programming, hardware), excluding pricing-related terms. Columns (1) to (3) report estimates for the standard deviation of log prices, the probability of an absolute change in log average price exceeding 5%, and $\ln(\text{Average price})$. All specifications include controls for log average pre-period units sold and log average pre-period prices interacted with post, log employees measured in the month before the event and interacted with post, and firm financial controls (log assets, Tobin’s q , ROA, cash-to-assets, fixed assets-to-assets, R&D-to-assets, market-to-book, and log firm age) measured in the quarter before the event and interacted with post. All specifications include UPC and calendar month fixed effects. Standard errors clustered at the firm level are reported in parentheses. Statistical significance at the 1, 5, and 10 percent levels is denoted by ***, **, and *, respectively. [Appendix A](#) presents variable definitions.

	Standard deviation	$\mathbb{P}(\% \Delta P > 5\%)$	$\ln(\text{Average price})$
	(1)	(2)	(3)
Tech-task intensity \times Post	0.060 (0.439)	2.305** (1.117)	0.210 (0.676)
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
Calendar month	✓	✓	✓
Observations	77,967	81,036	89,542
Adjusted R^2	0.47	0.31	0.95

IA.B.13 Manufacturer-retailer contract disclosures

To provide institutional context on the vertical pricing relationships in our sample, we audit the 10-K filings of all 229 sample manufacturers for disclosures related to manufacturer-retailer contracting. We classify disclosures into four categories: (1) trade promotions and allowances, (2) pricing policies and guidance, (3) category management and shelf arrangements, and (4) rebates and performance-based payments. We also identify whether the firm discloses a wholesale or contract pricing relationship as a baseline control category. Panel A of [Table IA.B.15](#) reports the prevalence of each disclosure type. Approximately 77% of sample firms disclose at least one of the four contract categories, with rebates and performance payments (75%) and trade promotions (70%) being the most common. Roughly half of firms disclose pricing policies or guidance, and nearly half disclose three or more categories. Panel B provides illustrative excerpts spanning all four disclosure categories, demonstrating the range of contractual arrangements through which manufacturers influence retail pricing.

Table IA.B.15: Manufacturer-retailer contract disclosures

This table summarizes the prevalence and nature of manufacturer-retailer contract disclosures in the 10-K filings of 229 sample manufacturers. Panel A reports the share of firms disclosing each of four contract categories: (1) trade promotions and allowances, (2) pricing policies and guidance, (3) category management and shelf arrangements, and (4) rebates and performance-based payments. “Wholesale only” identifies firms that disclose only a wholesale or contract pricing relationship without any of the four contract categories. “Broad contract disclosure” identifies firms disclosing three or more of the four contract categories. Panel B provides illustrative excerpts spanning all four categories.

Panel A: Contract disclosure rates		
	<i>N</i>	Share (%)
Trade promotions & allowances	161	70.3
Pricing policies & guidance	119	52.0
Category management & shelf	11	4.8
Rebates & performance payments	172	75.1
Wholesale/contract (control)	184	80.3
Any contract disclosure (cats 1–4)	177	77.3
Consideration to customers disclosed	36	15.7
Wholesale only (no cats 1–4)	18	7.9
Broad contract disclosure (3+ categories)	113	49.3
Mean contract categories (0–4)	2.02 (SD = 1.26)	
Median contract categories	2	
Firms in sample	229	

Panel B: Illustrative excerpts from 10-K filings

Firm	Category	Excerpt
Procter & Gamble	Trade promo	“Trade promotions, consisting primarily of customer pricing allowances, in-store merchandising funds, advertising and other promotional activities, and coupons...[are] recognized as incurred at the time of the sale.”
PepsiCo	Rebates	“For our independent distributors and retailers, these incentives include volume-based rebates, product placement fees, promotions and displays.”
General Mills	Trade promo	“Trade promotions are recorded using significant judgment of estimated participation and performance levels for offered programs at the time of sale.”
Fossil Group	Pricing guidance	“The Company has established minimum advertised price policies to protect brand integrity.”
Alkaline Water	Pricing guidance	“We currently offer our product in a one-gallon bottle for a suggested resale price or a SRP of \$4.99, three-liter bottle for a SRP of \$3.99...”
Hershey	Trade promo	“We promote our products with advertising, trade promotions and consumer incentives. These programs include, but are not limited to, discounts, coupons, displays, and other promotional activities.”
Central Garden & Pet	Category mgmt	“[Our field sales force provides] acquisition and maintenance of premium shelf placement, prompt product replenishment, customization of retailer programs”...and “in-store training and merchandising for our customers.”
Abbott Labs	Rebates	“Rebates and chargebacks charged against gross sales in 2021, 2020 and 2019 amounted to approximately \$3.9 billion, \$3.3 billion and \$3.1 billion, respectively.”
Coca-Cola	Trade promo	“The aggregate amount provided by our Company to bottlers, resellers and other customers...principally for participation in promotional and marketing programs, was \$4.7 billion in 2021.”
Mattel	Rebates	“Sales adjustments for [customer incentive] programs totaled \$623.9 million or 11.4% as a percent of net sales in 2021.”

IA.C A Parsimonious Vertical Pricing Framework

The empirical results point to a common mechanism: pricing decisions are shaped by upstream constraints in generating and transmitting market-specific information. Manufacturer identity accounts for the majority of systematic pricing variation, and pricing becomes more localized when the vertical information wedge is attenuated. Reductions in information-processing costs shift pricing away from uniformity and toward localization.

To synthesize these findings, we develop a parsimonious vertical pricing framework that emphasizes two sequential decisions. First, the manufacturer generates pricing guidance based on its information about demand. Second, the retailer implements that guidance, subject to operational constraints. This interaction highlights a central asymmetry: manufacturers observe demand across many geographically dispersed markets and can, in principle, aggregate this information into pricing recommendations, while retailers observe local conditions but rely on upstream guidance to coordinate pricing across products and locations.

Translating this demand intelligence into actionable, market-specific pricing recommendations is nonetheless costly. It requires gathering and processing local demand signals, then transmitting them through organizational or contractual channels to retailers, who must then act on them. The private-label evidence is consistent with this view: retailers are fully capable of implementing localized pricing when they possess the relevant demand information directly, suggesting that the friction lies upstream in generating market-specific pricing recommendations rather than downstream in retail implementation. The model formalizes this upstream information friction and its implications for pricing behavior across markets.

Model Setup. A manufacturer sells a product through retailers operating across N segmented markets. In market $i \in \{1, \dots, N\}$, demand is $Q_i(P_i) = (\alpha_i - P_i)/\beta_i$, where α_i summarizes local demand conditions and $\beta_i > 0$ parameterizes demand-price sensitivity. For simplicity, we assume that β_i is commonly known, but there is uncertainty about $\alpha_i \sim N(\mu_\alpha, \sigma_\alpha^2)$, which is the source of the information friction. The manufacturer produces at a constant marginal cost c .

Manufacturer’s Problem. The manufacturer can generate pricing recommendations by incurring a fixed cost I , which captures broadly the cost of gathering and processing information into pricing recommendations that can be communicated through the vertical chain—via wholesale pricing schedules, trade promotions, minimum advertised price policies, or category captain arrangements.²⁰ When the manufacturer pays I , the retailer receives a signal $s_i = \alpha_i + \eta_i$, where $\eta_i \stackrel{i.i.d.}{\sim} N(0, \frac{1}{\phi}\sigma_\alpha^2)$. The retailer’s posterior mean becomes

$$\hat{\alpha}_i = \mu_\alpha + \Phi(s_i - \mu_\alpha), \quad \Phi = \frac{\phi}{1 + \phi}, \quad (\text{IA.C.1})$$

where Φ summarizes the informativeness of the manufacturer’s recommendation. When the manufacturer does not pay I , no signal is transmitted and the retailer prices on the prior alone. The manufacturer pays the cost I whenever the resulting expected gains exceed the cost.

Retailer’s Problem. The retailer sets prices in two stages to maximize the joint manufacturer-retailer surplus. First, prior to receiving information, the retailer sets a uniform price:

$$\bar{P} = \frac{\mu_\alpha + c}{2}, \quad (\text{IA.C.2})$$

which is the standard monopoly price maximizing the expected joint manufacturer-retailer surplus given the prior. This uniform price serves as the retailer’s default across all markets in the absence of communication from the manufacturer.

When the manufacturer transmits signals s_i , the retailer updates its beliefs about α_i and faces a trade-off between the gains from more efficient pricing and the cost of implementing price changes. Deviating from the uniform price \bar{P} incurs a quadratic implementation cost $\frac{k}{2}(P_i - \bar{P})^2$, capturing menu costs, organizational inertia, and the costs of managing price variation across a store’s assortment. The retailer’s optimal shelf price upon receiving signal s_i is

$$P_i^* = (1 - \lambda_i)\bar{P} + \lambda_i \left(\frac{\hat{\alpha}_i + c}{2} \right), \quad (\text{IA.C.3})$$

²⁰ This maps onto the “managerial costs” documented by Zbaracki et al. (2004), who find that the informational and deliberation costs of price adjustment are roughly six times larger than the physical costs of changing price tags.

where $\hat{\alpha}_i$ is the posterior mean and $\lambda_i = \frac{2}{2+k\beta_i}$ is the optimal weight placed on the updated information relative to the uniform benchmark. A higher implementation cost (k) lowers the optimal weight (λ). In the extreme ($k \rightarrow \infty$), implementation costs drive the optimal shelf price (P_i^*) to \bar{P} , regardless of the signal. A lower implementation cost increases the optimal weight. In the extreme ($k \rightarrow 0$), the optimal shelf price fully adjusts to reflect the updated posterior.

Contracting Environment. The assumption that the retailer prices to maximize joint manufacturer-retailer surplus deserves justification. Standard double-marginalization models treat manufacturers and retailers as independent profit-maximizers connected only by a wholesale price, under which the manufacturer would have limited incentive to invest in costly pricing recommendations beyond setting an appropriate wholesale schedule. In CPG retail, this representation is not borne out in the data (Villas-Boas, 2007; Bonnet and Dubois, 2010). Manufacturer-retailer relationships in CPG are structured by a bundle of contractual instruments—trade promotions, scan-back and bill-back allowances, off-invoice discounts, slotting fees, category captain arrangements, and minimum advertised price (MAP) policies—that tie manufacturer net revenue directly to retail price and volume behavior. The magnitudes are substantial: industry estimates place trade promotion spending at approximately 15 to 20 percent of CPG manufacturer revenue, making it the second-largest line item on the manufacturer P&L after cost of goods sold (McKinsey & Company, 2019).

For the 229 manufacturers in our sample with available 10-K filings, 77 percent of filings disclose at least one form of contractual arrangement tying payoffs to retail behavior, 75 percent disclose performance-based rebates or payments, 70 percent disclose trade promotions or allowances, and 52 percent disclose pricing policies or guidance to retailers; only 8 percent disclose a wholesale-only relationship without any of these additional arrangements. Nearly half of sample manufacturers disclose three or more categories of contractual arrangements, indicating that the contracting bundle is both pervasive and deep.

In this contracting environment, modeling the retailer as maximizing joint surplus is a useful reduced form: the contracting bundle aligns incentives sufficiently that the binding friction we model is informational rather than strategic, and the cost I in our model captures the real and substantial expense of generating market-specific recommendations through this machinery.

$$\frac{\partial P_i^*}{\partial \alpha_i} = \lambda_i \cdot \frac{\Phi}{2}. \quad (\text{IA.C.4})$$

The manufacturer’s decisions thus affect pricing sensitivity through two margins: the extensive margin, whether to transmit a recommendation at all, and the intensive margin, the informativeness of pricing recommendations. Both are upstream decisions that determine how much local demand information reaches the shelf, and both are affected by AI adoption.

Although the framework is static, we interpret it as the per-period problem of the manufacturer and retailer: each period, market i draws a new α_{it} , and the manufacturer decides whether to transmit a signal, and the retailer decides how much to update prices. Under this interpretation, the within-market price variance $Var(P_i^*) = (\frac{\lambda_i}{2})^2 \Phi \sigma_\alpha^2$ corresponds to time-series price variability, and changes that increase responsiveness to signals also increase repricing frequency.

IA.C.1 AI Adoption and Empirical Predictions

AI can affect pricing through two channels. The following table summarizes the empirical predictions across these channels.

	Lower c (Production)	Lower I / Higher Φ (Information)
Average price declines	✓	×
Geographic dispersion increases	×	✓
Repricing frequency increases	×	✓
Local demand sensitivity increases	×	✓

No single channel accounts for the full set of empirical results. A reduction in marginal production costs ($c \downarrow$) shifts the uniform default $\bar{P} = (\mu_\alpha + c)/2$ downward, reducing average prices uniformly across markets, but leaving dispersion, repricing frequency, and local price-demand sensitivity unchanged. Lower costs of generating pricing recommendations ($I \downarrow$) expand the set of markets for which the manufacturer transmits a recommendation, increasing dispersion and repricing frequency on the extensive margin, but leave the uniform default \bar{P} unchanged and therefore do not affect average prices. Higher signal precision ($\Phi \uparrow$) increases the weight the retailer places

on new information, further raising dispersion and repricing frequency on the intensive margin, and increasing sensitivity to local demand conditions. The full set of empirical patterns—lower average prices, greater geographic dispersion, higher repricing frequency, and increased sensitivity to local economic conditions—is therefore consistent with AI operating through both channels simultaneously.