Who Sets the Price? Information and the Delegation of Pricing Authority

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

August 4, 2026

Abstract

We show that price uniformity and rigidity in retail markets reflect how pricing authority is allocated between brand owners and retailers. Retailers observe local demand but do not internalize that their markdowns draw customers from others carrying the same brand; the brand, unable to verify local conditions, imposes a price floor applied broadly and revised rarely. A single instrument thus produces both uniformity across stores and rigidity over time, and both relax as the brand becomes better informed. We test the mechanism in three settings: one reassigns pricing authority, one removes the competition that motivates the floor, and one shocks the brand's information. Brand acquisitions reassign authority while holding the product and the retailer fixed: acquired products depart from their own pricing history, and converge toward the acquirer's routine only when it already prices in the same category. Private labels, which do not suffer from intrabrand competition, have prices that are 17 percent more dispersed than national brands and track local wages. After the January 2022 introduction of generative AI tools deployable at scale, products of more AI-exposed brand owners show 12 percent greater cross-store price dispersion, 19 percent more frequent repricing, lower average prices, and greater sensitivity to local wage conditions. Uniform pricing thus reflects the full vertical relationship, and not only the technology of retail execution.

JEL Codes: G34, G30, D23, L14, L11, D83, O33

Keywords: Pricing authority, decision rights, boundaries of the firm, incomplete contracts, mergers and acquisitions, vertical relationships, uniform pricing, artificial intelligence, information capital

For their helpful comments and suggestions, we thank Aditya Chaudry, Andrea Eisfeldt, Isil Erel, Naveen Khanna, 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

Local demand conditions vary widely across markets, yet retail prices do not. 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 competing in different environments.[1] These same prices are also slow to adjust over time, changing far less often than local conditions do (Bils and Klenow, 2004; Nakamura and Steinsson, 2008). Prices are thus insensitive to local demand along both dimensions at once—nearly uniform across space and rigid over time. This insensitivity is costly: DG estimate that the median chain sacrifices roughly 2 percent of annual revenue. A large literature attributes it to the real costs of pricing at the point of retail execution—the organizational burden of setting and maintaining location-specific prices (DellaVigna and Gentzkow, 2019) and the menu and managerial costs of changing them (Golosov and Lucas, 2007; Zbaracki, Ritson, Levy, Dutta, and Bergen, 2004; Midrigan, 2011). Yet losses of this magnitude are difficult to reconcile with implementation costs alone.

We take a complementary view: rather than locate the binding constraint solely within retail operations, we cast price-setting as a problem of the full vertical chain. That is, a retail price is the *joint* product of brand owners upstream—who design, market, and set the terms under which their products are sold—and retailers downstream—who price and place them on the shelf.[2] Retailers observe local foot traffic, demographics, and competitive conditions that the brands do not and that they cannot credibly convey (Dessein, 2002). Moreover, retailers compete with one another for the same brand's customers, each failing to internalize that its own markdown draws demand away from others carrying the same product. This intrabrand externality leads retailers to price below the level that maximizes profit for the channel as a whole, motivating brand owners to restrict how far retail prices may fall—the classic rationale for vertical price restraints (Telser, 1960; Rey and Tirole, 1986).[3]

The brand's problem is one of delegation. Final shelf pricing authority naturally rests with the retailer—the better-informed party—but the retailer is a conflicted agent whose incentives, through the intrabrand externality, may diverge from the brand's. Unable to observe local demand, the

[1] Hitsch, Hortaçsu, and Lin (2021) extend this finding to the full set of product categories in the Nielsen scanner data, documenting that regular prices are largely uniform within retail chains. We likewise confirm that uniform pricing is pervasive in the full NielsenIQ Retail Scanner sample: Panel B of Table 2 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] We use *brand*, *brand owner*, and *manufacturer* interchangeably; the relevant firm is the one that owns the brand and sets its commercial terms, not necessarily the one that physically produces it.

[3] In consumer packaged goods, brand owners' payoffs are tied to their products' retail sales through a bundle of contractual instruments—trade promotions, performance rebates, and consideration paid for placement and volume: in our sample, 77 percent of brand owners with available 10-K filings disclose at least one such arrangement, and under 2 percent disclose a wholesale-only relationship (Table IA.B.18). These filings reveal the financial ties between the parties, not the price-setting ones—the instruments through which brand owners shape retail prices directly, such as minimum advertised pricing and resale-price policies, are exercised through unilateral policy and private contracts that firms are not required to file, and we document them separately below.

brand owner cannot distinguish a markdown that reflects genuinely weak conditions from one that reflects competition for a rival retailer's customers, and so cannot set constraints locally. When intrabrand competition is intense, the brand's optimal response is to restrict the agent's discretion with a single instrument—applied broadly and revised rarely. Thus, price uniformity is jointly produced by the two frictions, and neither generates it alone. Absent the intrabrand externality, the brand would fully delegate, allowing retailers to price according to what they locally observe, and absent the information friction, it would restrict prices market-by-market rather than nationally. Our contribution is to show that this combination—a motive to control retail prices coupled with an inability to control them precisely—is a previously underexplored component of price uniformity and rigidity in retail markets.

To formalize this argument and guide our empirical analysis, we develop a parsimonious model of vertical pricing with the above frictions. A brand sells through many retailers, each operating in many local markets. To discipline retailer pricing, the brand sets a price floor—capturing in reduced form the bundle of contractual levers through which brands affect retail pricing in practice, including promotional funding, rebate eligibility, markdown approval, and minimum advertised pricing.[4] The brand may pay to acquire a signal of local demand. Its information-processing capacity determines how likely that signal is to detect a local demand shock. When that capacity is low, acquiring information is not worthwhile: the brand sets one floor that binds every retailer in every market and every period. Prices are uniform across stores and constant over time even in the absence of menu costs. In this framework, uniformity and rigidity arise from a single constraint rather than from separate ones, and an improvement in the brand's information relaxes both.

Once the brand's information-processing capacity clears a threshold, acquiring the signal becomes worthwhile, and uniformity breaks along each of the margins we take to the data. Prices disperse across markets, as the brand sets market-specific floors and revises them as new signals arrive, and across retailers within a market, once it lowers the floor far enough to release the least-affected retailer to price on its own information—above the floor, but below what the brand would prefer. Prices adjust more often, because the floor itself now moves with the signal. And the average price falls, while responding more strongly to local demand conditions, because the brand accepts underpricing by unconstrained retailers as the cost of better-matched prices. These margins differ in what they require. Dispersion across markets, greater sensitivity to local conditions, and more frequent repricing need only that the brand act on its own signal; dispersion across retailers and the decline in average prices require it to cede authority, and so appear together and only where delegation is optimal.

Guided by this distinction, our evidence indicates that how pricing authority is allocated between the two parties shapes retail price-setting. The constraints that generate uniformity appear to be a property of the vertical relationship—of who holds authority and what that party can

[4] Table 1 maps each primitive of the model to the instrument that implements it, with representative language from the hand-collected contract archive described in Section 2.

verify—rather than only of retailers' unilateral pricing decisions or the mechanics of retail execution.

Our empirical investigation proceeds in three stages, each varying one element of the vertical relationship: who holds pricing authority, why that party constrains it, and how well informed it is. Brands hold instruments capable of constraining retail prices: minimum-advertised-price and resale-price policies, promotional funding conditioned on price compliance, and compliance chargebacks appear in manufacturer-retailer contracts we collect and describe in Section 2. What we test is when and how those instruments bind. Throughout, we use a bridge linking over 300 publicly traded brand owners to product-level NielsenIQ Retail Scanner (RMS) data spanning roughly 5 billion store-UPC-month observations, which lets us tie the products on the shelf to the firms that own the brands.

We begin with brand acquisitions, which test the framework's premise: that pricing authority over a product rests with a single party who applies a common policy across the markets in which it is sold. An acquisition transfers that authority while leaving the product and the retailers that carry it in place, so if the premise holds, the pricing policy should travel with ownership. It does. Acquired products move away from their prior patterns and converge toward the acquirer's incumbent behavior—but only when the acquirer already operated in the target's product category. Where the acquirer had no prior presence, acquired products depart from their own history without converging on the acquirer's existing products, as there is no incumbent category-specific routine to extend. Because the product and its retailers are unchanged, the movement is attributable to the new owner, which applies a category-specific pricing routine to the products it controls.

Having established that pricing policy belongs to the brand owner, we next ask why that owner constrains it. We compare national brands, which reach consumers through the brand-retailer chain, with private labels, which the retailer owns and prices itself (Hoch and Banerji, 1993; Chintagunta, Bonfrer, and Song, 2002). No rival retailer carries a private label, so the intrabrand externality is absent, and with it the reason to constrain the price at which the label is sold—whoever prices it internalizes the effect on the brand's other sellers. In the model's terms, private labels approximate the case where the externality vanishes: no floor is worth imposing, and the brand's information becomes irrelevant because there is no constraint for it to condition. What the comparison holds roughly fixed is the machinery of retail pricing—shelf tags, scanner systems, promotional calendars—and often production itself, since private labels are frequently made by the same firms that manufacture national brands.[5] Integration also removes the information asymmetry between the two parties, so this comparison speaks to why the floor exists rather than isolating information; the later tests do that.

If uniformity were principally a matter of retail execution, the two should look alike. Several features of private-label markets push the same way: private labels carry the retailer's own name,

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

strengthening 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. We find the opposite. Private-label products are roughly 17 percent more dispersed than national brands, even within retailer-by-product-category-by-month cells, and their prices track local wages where national-brand prices do not. Prices localize once the constraint that suppressed local variation is gone.

The third setting moves information while holding the relationship and its conflict fixed. We exploit the release of OpenAI's InstructGPT API in January 2022, the first generative artificial intelligence (AI) model trained via RLHF to follow natural language instructions at scale. By removing the need for specialized prompt engineering, it made a range of information-processing tasks—summarizing unstructured local demand data, drafting category and market analyses, querying data without specialist tooling—available to non-technical staff at low marginal cost. Products sold by more AI-exposed brand owners become less uniform, reprice more frequently, and respond more strongly to local economic conditions, even within county and retailer—what the framework predicts when the brand's capacity to act on local signals rises (Agrawal, Gans, and Goldfarb, 2019; Goldfarb and Tucker, 2019; Bergemann and Bonatti, 2019).

We measure firms' AI exposure using occupation-level scores following Eisfeldt, 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. Because we observe exposure rather than deployment, the estimates that follow are intention-to-treat: they measure how pricing changed at firms whose workforces were positioned to use the technology, not the effect of putting it to work on prices. We cannot observe which firms did so, and we therefore do not scale the estimates to a treatment-on-treated effect.[6] We compare products sold by high-exposure brand owners 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.

After the API introduction, products sold by high-exposure brand owners exhibit greater dispersion in the prices charged for the same product across the stores that sell it (a 12 percent increase relative to the sample mean), more frequent repricing (a 19 percent increase), and lower average prices (a 4.3 percent decline). The rise in dispersion appears both within a chain's own stores and between competing chains serving the same county. 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 brand-side rather than retailer-side forces. The effects are also concentrated where theory predicts they should be: among firms

[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.

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 the brand's constraint 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 exposed firms acting on the technology along the margins we can observe. 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.

The same margin is visible without a shock. As a brand accumulates experience with a product, a market, or a retailer relationship, what it knows about local demand improves, and prices should localize accordingly. Measuring a product's information stock by the cumulative store-weeks over which it has been observed, we find that dispersion rises with it, and that the pattern survives holding each cell's store set fixed at the cohort it entered with—including in the most demanding comparison, an established product-retailer relationship expanding into new counties, where the estimate is larger under the fixed store set. It does not survive for entry into new counties measured on its own, which is informative in its own right: information appears to attach to the relationships through which a brand sells rather than to geography. These estimates identify no causal effect, and we read them as corroboration that the mechanism operates outside the technology event rather than as identification.

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 shocks to the information-processing capacity of the party holding pricing authority can have first-order effects on how prices are set.

Related Literature. Our study connects to five 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 (Golosov 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. Wu (2024) shows that this attention constraint interacts with where the pricing decision sits: in a model of rationally inattentive multi-region firms, headquarters economizes by ignoring regional demand differences as dispersion grows, while decentralized divisions attend to their own markets and prices separate. 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 locating the constraint in how pricing authority is allocated between brand owners and retailers, using a bridge that links publicly traded manufacturers to product-level scanner data. The acquisition, private-label, life cycle, and AI evidence then separates the brand owner's constraint from retail implementation costs and price-image considerations, indicating that these rigidities are not immutable features of retail markets.

Second, we contribute to work on the delegation of authority within firms and vertical chains. When a principal cannot verify what a better-informed agent knows, committing to a set of permitted actions rather than instructing the agent is often optimal, and under many conditions that set is an interval (Holmström, 1984; Alonso and Matouschek, 2008); ceding the decision outright can dominate communicating about it (Dessein, 2002). In finance, the same question is the allocation of residual control rights when the relevant information cannot be contracted upon (Grossman and Hart, 1986; Hart and Moore, 1990), and it recurs wherever a decision is better made by whoever holds the soft information than by whoever holds formal authority (Stein, 2002; Graham, Harvey, and Puri, 2015). The boundary we study runs between firms rather than inside one, but the object is the same: a right retained, and a range of action conceded. The price floor we study is such a restriction in contractual form, and it separates the two kinds of authority in Aghion and Tirole (1997): the brand retains the formal right to set the terms on which its product may be sold, while the retailer, better informed about its own market, exercises real authority over the price within them. The rationale for such restrictions in vertical settings is long established. Telser (1960) and Rey and Tirole (1986) attribute resale price maintenance to retailer free-riding on services and to the manufacturer's response to retail competition, and Deneckere, Marvel, and Peck (1996, 1997) provide informational rationales under demand uncertainty. Empirical work faces a common obstacle: the terms of interest are rarely observed and are typically inferred from a structural model of a single product market (Villas-Boas, 2007; Bonnet and Dubois, 2010) or documented within one industry. The contracts we collect in Section 2 let us observe the instruments directly, and show that the object we formalize—a floor, with pricing delegated within it—is one that manufacturer-retailer agreements implement.

Our contribution to this literature is empirical rather than theoretical. Theory delivers the boundary of delegated authority as a function of the principal's information, but observing that boundary move requires variation in what the principal knows while the relationship is held fixed—variation that has been difficult to obtain and, to our knowledge, has not been used to estimate the comparative static at scale. Our setting supplies it, across roughly 5 billion store-UPC-month observations, and the resulting pattern is not monotone: better information leads the brand to tighten the constraint in most markets and to release it in those where it detects a shock and finds its retailers unevenly exposed. Delegation is thus not what a better-informed principal does less of; it is what it does selectively.

Third, we contribute to research on mergers and the boundaries of the firm. A large literature asks what changes when productive assets move between owners, and finds that the gains depend on the acquirer's capacity to redeploy what it buys (Maksimovic and Phillips, 2001; Maksimovic, Phillips, and Prabhala, 2011) and on how closely the two firms' product markets are related (Hoberg and Phillips, 2010; Frésard, Hoberg, and Phillips, 2020). Our acquisition evidence speaks to this literature at the level of the individual product. When a brand changes hands its pricing departs from its own history and converges toward the acquirer's routine, but only where the acquirer already prices in the same category: a pricing policy is an organizational asset that transfers with ownership and whose scope is bounded by what the acquirer already does. The closest precedent is Sheen (2014), who finds that the products of merging firms converge in quality and that their prices fall relative to competitors, but not when the acquirer is diversifying into an unfamiliar product market. We recover the same scope condition and characterize what moves with ownership: not the price level alone but the pricing policy itself—how far a product's prices are allowed to vary across stores, and how often they are revised. The private-label comparison in Section 7 is the vertical-integration margin of the same question, asking what prices look like when the retailer owns the brand and no authority need be delegated at all.

Fourth, 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. Adams, Fang, Liu, and Wang (2026) also study pricing directly, measuring adoption of AI *for pricing* from job postings and finding that adopters grow faster and earn higher markups, which they attribute to more effective learning of the demand function. The designs identify different objects: they compare firms that choose to automate pricing, whereas we hold the relationship fixed and use a common shock to the tasks a firm's existing workforce already performs. The two sets of results are not in tension. Their

adopters are firms investing in pricing capability for products they price themselves; our exposed firms are brand owners deciding how much latitude to grant the retailers that price for them. 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. Experimental evidence points the same way, with Noy and Zhang (2023) finding that generative AI raises output quality and compresses the productivity distribution on professional writing tasks. Estimated aggregate labor-market effects have been considerably more modest—Humlum and Vestergaard (2025) report precise nulls for earnings and hours in Danish administrative data—which is consistent with our reading of the mechanism, in which AI changes the quality of a decision that a small number of people were already making rather than the quantity of labor required to make it.

Fifth, 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), including evidence that pricing agents built directly on large language models can reach supracompetitive prices without being instructed to collude (Fish, Gonczarowski, and Shorrer, 2024). 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; Farboodi and Veldkamp, 2026), 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, the results connect two literatures that rarely meet. The pricing literature asks why prices are rigid and uniform; the AI literature asks how AI changes firm behavior. We join them by showing that better information relaxes a constraint the brand owner imposes for reasons of its own, primarily through task-level augmentation rather than automation. The uniform pricing puzzle then becomes a question about who holds pricing authority and what that party can verify, rather than one only about retail operations.

2 The Instruments of Brand Price Control

The owner of a national brand sells through independent retailers that ring up the final price, but it is no passive supplier at a wholesale price. Brand owners shape the prices their products carry through a bundle of contractual and policy instruments: minimum-advertised-price and unilateral resale-price policies that place a floor under advertised or transaction prices;[7] promotional funding

[7] The RMS scanner data corroborates the contractual evidence. In the sample used in Section 8, just over a fifth of chain-product-county prices sit within one percent of the lowest price recorded for that product in that county before

and performance rebates whose eligibility is conditioned on price and display compliance; data-sharing requirements that oblige retailers to transmit store-level sales data to the brand owner, together with shelf and assortment authority keyed to the brand owner's own measure of local market share; and chargebacks and termination rights that discipline retailers who deviate. The retailer chooses the price at the register, but it does so within bounds the brand shapes.

We document this arrangement with two complementary sources. First, we audit the 10-K filings of the brand owners in our sample for disclosures of their commercial relationships with retailers. Seventy-seven percent disclose at least one arrangement—trade promotions, rebates, or consideration paid for placement and volume—that ties their payoffs to their products' retail sales, and fewer than 2 percent describe a wholesale-only relationship (Table IA.B.18). These disclosures establish that brand owner payoffs move with retail price and volume, but they are nearly silent on the price-setting instruments themselves: almost none disclose minimum-advertised-price policies, resale-price terms, or category-management authority. That silence is not absence—those instruments are exercised through unilateral policy and private contracts that firms are not required to file.

Second, to examine those instruments, we hand-collect an archive of actual manufacturer-retailer contracts and pricing documents from public filings, litigation and regulatory records, and franchise registries. Much of what we find takes the form of unilateral policy rather than negotiated terms, which reflects the legal environment: under the *Colgate* doctrine a manufacturer may announce the prices at which it will continue to deal and decline to supply retailers that undercut them, whereas agreements fixing resale prices remain exposed to antitrust challenge even after Leegin Creative Leather Products, Inc. v. PSKS, Inc. (2007). The archive comprises 95 quote-verified contract and document instances spanning roughly four decades and more than a dozen consumer-goods industries; Internet Appendix IA.A.1 describes its construction, scope, and limitations. Across these documents we observe each of the instruments above in the parties' own words. The disclosures show that money flows to the brand owner with retail sales; the archive shows that the brand owner holds the instruments that constrain the retail price.

These instruments are not incidental to the model we develop in Section 3.1; they are its primitives. The price floor at the center of the model's delegation problem is a contractual object rather than an abstraction, and the monitoring and enforcement it requires is visible as compelled data feeds, compliance chargebacks, and performance rebates. Table 1 maps each primitive to the instrument that implements it, so the mechanism we formalize—a brand owner setting a floor and delegating pricing within it—is the one the contracts implement. Cross-industry records of *realized* price restraints exist: Overstreet (1983) and Ippolito (1991) assemble evidence on resale price maintenance from antitrust cases and litigation records across many industries. What those compilations document is the outcome of enforcement actions rather than the operative contract

the event, a price we treat as an estimate of the floor or reference price. Moreover, the share is almost identical for the firms later classified as treated and control (Table IA.B.12).

language, and neither is organized around the primitives of a pricing model. The archive assembled here is narrower and complementary: contract-level, quoted verbatim, and mapped to the instruments the model formalizes. The empirical literature on vertical price restraints has otherwise either inferred contract terms from a structural model of a single product market, or observed actual instruments within a single industry.[8] And because the archive records only what parties commit to writing and file, it understates rather than overstates the reach of brand owner price authority: the policies that govern the largest national brands operate, by design, outside the filing system.

3 A Model of Delegated Pricing

This section develops a parsimonious model of vertical pricing in the spirit of Deneckere, Marvel, and Peck (1996) and Deneckere, Marvel, and Peck (1997) with two frictions: an information asymmetry between the brand owner (hereafter, the brand) and its retailers, and an intrabrand externality. Retailers observe local demand conditions that the brand does not, and cannot credibly convey them. At the same time, retailers compete for the same customers and so price below the level that maximizes total channel profit. The model characterizes how the brand responds to these frictions, and how improvements in its capacity to acquire and process local information reshape retail prices. Section 3.5 summarizes the model's main empirical predictions.

3.1 Model Setup

An upstream firm (the brand) sells its product in a continuum of geographically distinct markets indexed by $m \in [0,1]$ through a continuum of retailers indexed by $j \in [0,1]$. All retailers operate in every market, so each retailer's prices are observed across markets, and each market's prices are observed across retailers. Within a market m, demand for the brand's product at retailer j is

$$q_{mj} = \alpha - z_m d_{mj} - \beta p_{mj} + \gamma \bar{p}_m, \tag{1}$$

where p_{mj} is the price at retailer j, $\bar{p}_m$ is the average price in market m, and $z_m d_{mj}$ is an adverse local demand shock specified below. The parameter $\beta > 0$ is the own-price sensitivity, and $\gamma > 0$ governs substitution across retailers carrying the brand: demand shifts toward retailer j when other retailers raise their prices. In other words, γ measures the intensity of intrabrand competition. We assume $\beta > \gamma$, so own-price effects dominate cross-price effects, and the total channel profit is strictly concave in price. To isolate the mechanism of interest, we assume that the demand

[8] Studies that infer unobserved vertical terms include Villas-Boas (2007) and Bonnet and Dubois (2010); those that observe actual instruments do so within one industry or for a single firm, including Mortimer (2008) on video rental, Asker (2016) on beer distribution, and Israeli, Anderson, and Coughlan (2016), whose data cover a single manufacturer's minimum-advertised-price program. Lafontaine and Slade (2008) survey vertical-restraint findings across industries but do not assemble a contract archive.

parameters (α, β, γ) are common across markets and retailers, and are known to all. The variation in demand across markets and retailers operates through $z_m d_{mj}$.

Local demand. In each market m, an adverse demand shock occurs with probability $\theta \in (0,1)$, independently across markets, and is indicated by $z_m \in \{0,1\}$. Conditional on $z_m = 1$, retailer j's exposure to the shock is $d_{mj} = \mu \pm \sigma$ with equal probability, drawn independently across retailers. The parameter $\mu > 0$ captures the average severity of the demand shock and $\sigma < \mu$ represents the cross-retailer heterogeneity in realized exposure. The restriction $\sigma < \mu$ ensures $d_{mj} > 0$, so the demand shock is negative for every retailer. Therefore, the realized demand intercept for retailer i in market m is $\alpha_{mj} = \alpha - z_m d_{mj}$. Because the model features a continuum of firms, we adopt the standard convention that the exact law of large numbers holds, implying that in every market m with an adverse demand shock ($z_m = 1$), half of retailers experience high exposure ($d_{mj} = \mu + \sigma$) and half experience low exposure ($d_{mj} = \mu - \sigma$). We assume α is large enough that equilibrium demand and profit are positive in every demand state.

We further assume that adverse demand events are sufficiently unlikely:

$$\theta < \bar{\theta}, \tag{2}$$

where $\bar{\theta} \in (0,1)$ depends on the remaining primitives $(\alpha, \mu, \sigma, \beta, \gamma)$ and is derived in the proof of Lemma 1. Condition (2) ensures that, absent information about the adverse demand shock, the intrabrand externality dominates the value of retailer adaptation.

Information. Retailers costlessly observe their own realized demand intercept (α_{mj}) but cannot credibly convey it to the brand. The brand observes neither the demand shock (z_m) nor any retailer's exposure (d_{mj}); it can, however, acquire costly information about the demand shock. For each market m, the brand can pay I to obtain a signal $s_m \in \{0,1\}$ with

$$\Pr(s_m = 1 \mid z_m = 0) = 0 \tag{3}$$

and

$$\Pr(s_m = 1 \mid z_m = 1) = a_m \phi, \tag{4}$$

where $a_m \in \{0,1\}$ indicates information acquisition and $\phi \in (0,1)$ captures the brand's capacity to acquire and process information about demand shocks. Simply put, acquiring information allows the brand to detect an adverse demand shock with probability ϕ; without acquisition ($a_m = 0$) the signal is uninformative, and the brand's posterior equals its prior θ. In practice, the raw material for this signal is itself contractual: the agreements in our archive obligate retailers to transmit store-level shipment and scan data to the brand, so ϕ is best read as the brand's capacity to convert those standing data feeds into a timely read on local demand—and the acquisition cost I as the

cost of processing the data, not of obtaining them. Moreover, detecting an adverse demand shock resolves only one dimension of the brand's informational disadvantage: even when the brand knows $z_m = 1$, individual retailers' realized exposures d_{mj} remain their private information.

Wholesale Pricing. The brand could in general discipline retailers with both signal-contingent wholesale prices and signal-contingent price floors. To isolate the role of information, we fix the wholesale price at the marginal cost of production, c, normalized to zero.[9] This removes the double-marginalization distortion, leaving the informational asymmetry as the only source of contractual friction.[10]

Price Floors. Because the brand cannot observe local demand, it cannot verify whether a retailer's markdown reflects genuinely weak demand or an attempt to attract customers from rival retailers carrying the same brand. We model the brand's control over markdowns as a lower bound on the retail price. Because retailers are ex-ante symmetric and the brand observes nothing retailer-specific, we restrict attention to anonymous floors conditioned on the brand's own signal,

$$p_{mj} \geq \underline{p}_m(s_m) \tag{5}$$

for all $j \in [0, 1]$. We suppress the argument s_m when doing so creates no ambiguity. We view the price floor as a reduced form for the family of contractual instruments through which brands place a lower bound on retail prices in practice—minimum-advertised-price (MAP) and unilateral resale-price policies, resale-price maintenance, and everyday-low-price commitments—backed by compliance chargebacks, suspension of promotional funding, and termination rights. Although the best-known of these, MAP, restricts only the advertised price, the stronger policies in our archive bound the transaction price itself, net of any discounts—the object our floor constrains. Table 1 maps this and the model's other primitives to the specific instruments that implement them, drawn from the contract archive of Section 2.

This formulation connects our setting to the theory of delegation in organizations. A price floor is an interval delegation scheme: the brand grants each retailer discretion over any price above

[9] Wholesale prices also divide surplus between the brand and its retailers. We abstract from this bargaining problem by assuming that lump-sum transfers implement any desired division of rents subject to retailer participation. In practice, these transfers take the form of trade promotions, promotional allowances, and performance rebates paid by the brand to the retailer; the level of such payments divides channel rents, while their conditioning on pricing compliance provides the enforcement behind the price floors introduced below.

[10] We focus on price floors rather than wholesale prices because the wholesale price does not discipline the retail pricing decision at the margin. Retailers commit to order quantities before local demand is realized, so by the time prices are set the wholesale price is sunk and the marginal cost of selling a unit already in inventory is close to zero (Deneckere, Marvel, and Peck, 1996, 1997). A wholesale markup therefore shapes how much a retailer orders rather than what it charges, while a floor constrains the retail price directly. A markup also raises prices in every state, including those in which the brand would prefer a low price, whereas a floor binds only from below and leaves intact the adaptation the brand values. Our model has no ordering stage, so we take this as a reason the restriction is reasonable rather than as a mechanism the model derives.

Table 1: Mapping model primitives to contract instruments

This table maps each primitive of the model to the contractual or policy instrument that implements it in practice, with representative evidence from the hand-collected manufacturer-retailer contract archive described in Section 2 and Internet Appendix IA.A.1. The price floor is the interval-delegation instrument at the center of the model: the brand fixes a lower bound on the retail price and delegates pricing above it to better-informed retailers. Some instruments bind only the advertised price (MAP), while others bind the net or transaction price after discounts; the model's floor, which constrains the transaction price, corresponds most literally to the latter. The enforcement, information, and transfer rows correspond, respectively, to the verification the floor requires, the brand's costly capacity ϕ to acquire local demand information, and the lump-sum transfers that divide channel rents without distorting the per-unit price. The intrabrand externality γ is the motive for the floor and appears as brand-protection and orderly-market language in the underlying policies rather than as a standalone instrument (e.g., Elevai's unilateral price policy is adopted "[i]n order to ... maintain its premium brand and product image"). Named firms are illustrative examples from the archive; Internet Appendix IA.A.1 details the full set of documents.

Model primitive	*Contractual instrument*	*Representative archive evidence*
Price floor $\underline{p}_m(s_m)$ (5)	MAP and unilateral resale-price policies; resale-price maintenance; everyday-low-price commitments	J&J ACUVUE policy: net price "at or above the UPP Minimum Retail Price, after discounts are applied"; MACH 3 uniform retail-price structure barring unauthorized discounts
Interval delegation: discretion above the floor	An enforceable minimum that leaves upward pricing free	MACH 3 distributor "allowed to market and sell product at prices that are more expensive" than the uniform structure; Remote Knowledge–Sea Smart pricing "discretion" subject to a minimum
Enforcement of the floor	Compliance chargebacks; suspension of promotional funding; termination for violation	Sears/KCD chargeback for "failure to comply with Vendor policies"; ProdecoTech MAP breach triggers "account terminated and be in default"
Information capacity ϕ	Compelled retail data feeds—store-level shipment and scan data—and manufacturer-computed local market shares	R.J. Reynolds (RMSC) "Shipment To Retail data ... through the ... ('MSAi') clearinghouse" and "Local SOM ... as determined by RMSC"; PM USA transaction-level scan data
Wholesale price $= c$ with lump-sum transfer	Trade promotions, promotional allowances, and performance rebates that move rents without distorting the per-unit price	PM USA "Retail Price Promotional Allowance paid directly by PM USA to Retailer"; RMSC everyday-low-price promotional payments conditioned on compliance

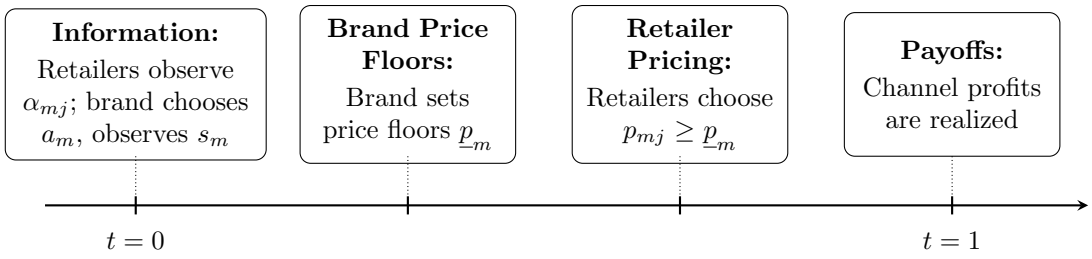


Figure 1: Model timeline. In each market m, the local demand states (α_{mj}) are drawn. Retailers observe the realized state, while the brand sees only an imperfect signal (s_m) that depends on its information acquisition decision (a_m). The brand then sets pricing floors conditional on its signals, after which the retailers set retail prices subject to the floor, as mutual best responses. Demand and channel profits are realized.

the floor, the form of authority a principal optimally grants a better-informed but biased agent when state-contingent transfers are unavailable (e.g., Holmström (1984), Alonso and Matouschek (2008)). Restricting discretion, rather than soliciting reports, is natural here because retailers cannot credibly communicate local demand to the brand (e.g., Dessein (2002)). The archive's policies take exactly this form: they fix an enforceable minimum and expressly leave the retailer free to price above it.

Retailer Problem. In market m, retailer j observes its realized demand intercept (α_{mj}) and the price floor ($\underline{p}_m$), takes the average price in the market ($\bar{p}_m$) as given, and solves

$$\max_{p_{mj} \geq \underline{p}_m} \quad p_{mj}\big(\alpha_{mj} - \beta p_{mj} + \gamma \bar{p}_m\big). \tag{6}$$

An equilibrium in market m is a price schedule $\{p_{mj}\}_{j\in[0,1]}$ and an average price $\bar{p}_m$ such that each p_{mj} solves (6) given $\bar{p}_m$, and $\bar{p}_m = \int_0^1 p_{mj} dj$.[11]

Brand's problem. In each market m, the brand maximizes total channel profits net of any information acquisition costs. First, it decides whether to acquire costly information. Second, after the signal is realized, it sets a price floor $\underline{p}_m$, anticipating its retailers' equilibrium responses. Figure 1 summarizes the model's timing.

3.2 Brand Strategy

We solve the brand's problem by backward induction: characterizing the optimal price floors conditional on the signal (Section 3.2.1), then analyzing the information-acquisition decision (Section 3.2.2).

[11] Note that retailers face no residual uncertainty: since $\mu > \sigma$, $\alpha - z_m d_{ij}$ lies strictly below α when $z_m = 1$. Hence, observing α_{mj} reveals z_m, and the exact law of large numbers renders $\bar{p}_m$ deterministic conditional on z_m and the price floor.

3.2.1 Optimal Price Floors

We begin with two benchmarks that clarify the role of price floors. The first is the unconstrained retailer-pricing equilibrium. The second is the common price the brand would choose if it could set retail prices directly but could not condition on retailer-specific exposures.

Consider retailer j's unconstrained problem. Given its own demand intercept ($\alpha_{mj} = \alpha - z_m d_{mj}$) and the average price in the market ($\bar{p}_m$), it sets a price p_{mj} to solve

$$\max_{p_{mj}} \quad p_{mj} \left(\alpha_{mj} - \beta p_{mj} + \gamma \bar{p}_m \right). \tag{7}$$

The corresponding first-order condition is

$$\alpha_{mj} + \gamma \bar{p}_m - 2\beta p_{mj} = 0. \tag{8}$$

Because there is a continuum of retailers in the market, each retailer takes the average price ($\bar{p}_m$) as given:

$$p_{mj}(\bar{p}_m) = \frac{\alpha_{mj} + \gamma \bar{p}_m}{2\beta}. \tag{9}$$

Averaging across retailers yields

$$\bar{p}_m = \mathbb{E}[p_{mj}(\bar{p}_m)] = \frac{\bar{\alpha}_m + \gamma \bar{p}_m}{2\beta}, \tag{10}$$

where $\bar{\alpha}_m = \mathbb{E}[\alpha_{mj} | z_m]$.

Absent an adverse demand shock ($z_m = 0$), all retailers share the intercept α, so $\bar{\alpha}_m = \alpha$ and every retailer charges

$$p^R_{mj} = \frac{\alpha}{2\beta - \gamma} = \bar{p}^R_m. \tag{11}$$

When an adverse event occurs ($z_m = 1$), half of retailers draw the high exposure ($d_{mj} = \mu + \sigma$) and half the low exposure ($d_{mj} = \mu - \sigma$), so $\bar{\alpha}_m = \alpha - \mu$ and

$$p^R_{mj}(d_{mj}) = \begin{cases} \frac{\alpha - \mu}{2\beta - \gamma} + \frac{\sigma}{2\beta}, & d_{mj} = \mu - \sigma, \\ \frac{\alpha - \mu}{2\beta - \gamma} - \frac{\sigma}{2\beta}, & d_{mj} = \mu + \sigma, \end{cases} \tag{12}$$

Because $\sigma < \mu$, even the least exposed retailer prices strictly below $\bar{p}^R_m$.

We now turn to the brand's preferred common price. Because floors are anonymous, the brand cannot implement retailer-specific prices. If all retailers charge p, total channel profit in market m is

$$\Pi^B(p; \bar{\alpha}_m) = p(\bar{\alpha}_m - \beta p + \gamma p), \tag{13}$$

which is strictly concave in p because $\beta > \gamma$. The associated first-order condition is

$$\bar{\alpha}_m - 2(\beta - \gamma)p = 0, \tag{14}$$

implying that the brand-preferred common price is

$$p_m^B(\bar{\alpha}_m) = \frac{\bar{\alpha}_m}{2(\beta - \gamma)} > \frac{\bar{\alpha}_m}{2\beta - \gamma} = \bar{p}_m^R. \tag{15}$$

The inequality stems from the intrabrand externality ($\gamma > 0$). Each retailer ignores that raising its price shifts demand toward other retailers carrying the same brand. The brand, therefore, prefers higher prices than retailers would choose on their own. A price floor can correct this underpricing distortion. However, a binding floor also limits retailers' ability to adapt prices to their private information about local demand. The remainder of this section characterizes how the brand's information affects this trade-off and the resulting equilibrium prices.

Uninformed Price Floor.

Lemma 1. *Assume that condition (2) holds. In the absence of information about an adverse demand shock in market m, the brand's optimal price floor is*

$$\underline{p}_m^0 = \frac{\alpha - \theta\mu}{2(\beta - \gamma)} \tag{16}$$

and it binds for all retailers in all demand states.

Lemma 1 characterizes the optimal price floor for an uninformed ($a_m = 0$) brand. In setting the floor, the brand trades off the benefit of disciplining retailers to limit intrabrand competition against the cost of preventing them from adapting prices to their privately observed demand. Under condition (2), adverse shocks are sufficiently unlikely that the benefit of discipline dominates: the brand sets the floor at its preferred price for the average market, and that floor binds every retailer in every state. Prices are consequently uniform—both across markets and across retailers within a market—which is the benchmark against which the effect of information is measured.

Informed Price Floors. Now suppose the brand acquires information ($a_m = 1$) and observes the signal s_m. The signal shifts the brand's belief in one of two ways. An affirmative signal ($s_m = 1$) allows the brand to verify an adverse demand shock: $\theta_1 = \Pr(z_m = 1 \mid s_m = 1) = 1$. A null signal ($s_m = 0$) is only partially informative—a demand shock goes undetected with probability $1 - \phi$—and lowers the posterior to $\theta_0 = \Pr(z_m = 1 \mid s_m = 0) = \frac{\theta(1-\phi)}{1-\theta\phi} < \theta$.

Lemma 2. *Assume that condition (2) holds. Suppose the brand learns $s_m = 0$. Its optimal price floor is*

$$\underline{p}_m(0)^* = \frac{\alpha - \theta_{ms}\mu}{2(\beta - \gamma)} \tag{17}$$

and it binds for all retailers and all demand states.

Lemma 2 shows that a null signal changes the level of the price floor but not its defining feature—it binds all retailers in all demand states. Because a null signal leaves the brand strictly more optimistic than its prior, $\theta_0 < \theta < \bar{\theta}$, the benefit of disciplining intrabrand competition continues to dominate that of retailer adaptation. The signal is thus informative even though it triggers no delegation—it tells the brand the market is more likely free of a demand shock, resulting in a higher price floor.

Lemma 3. *Assume that condition (2) holds. Suppose the brand learns $s_m = 1$, implying that $\theta_{ms} = 1$. Its optimal price floor is*

$$\underline{p}_m^*(\sigma) = \begin{cases} \frac{\alpha-\mu}{2(\beta-\gamma)}, & \sigma \leq \bar{\sigma}, \\ \frac{4\beta(\alpha-\mu)}{8\beta^2-8\beta\gamma+\gamma^2} - \frac{\sigma}{2\beta}, & \sigma > \bar{\sigma}, \end{cases} \tag{18}$$

where

$$\bar{\sigma} = (\alpha - \mu)\gamma\sqrt{\frac{\beta}{(\beta - \gamma)(8\beta^2 - 8\beta\gamma + \gamma^2)}}. \tag{19}$$

If $\sigma \leq \bar{\sigma}$, the optimal price floor binds all retailers in all demand states. If $\sigma > \bar{\sigma}$, the optimal price floor binds retailers with demand loss $d_{mj} = \mu + \sigma$ but not those with $d_{mj} = \mu - \sigma$:

$$p_{mj}^*(d_{mj}) = \begin{cases} \frac{4\beta(\alpha-\mu)}{8\beta^2-8\beta\gamma+\gamma^2} - \frac{\sigma}{2\beta}, & d_{mj} = \mu + \sigma, \\ \frac{2(\alpha-\mu)(2\beta-\gamma)}{8\beta^2-8\beta\gamma+\gamma^2} + \frac{\sigma}{2\beta}, & d_{mj} = \mu - \sigma. \end{cases} \tag{20}$$

Lemma 3 describes what the brand does once it has verified an adverse demand shock. The key is that detecting the shock does not merely resolve one informational problem; it also activates another. Ex ante, the retailers' private knowledge of their own exposure (d_{mj}) creates little friction, because exposure varies across retailers only when an adverse demand shock occurs ($z_m = 1$). When shocks are rare—as condition (2) requires—this heterogeneity is almost never realized, so the brand gains little by letting retailers adapt prices to their private information. Verifying an adverse demand shock changes the trade-off: the brand knows it faces the state in which exposures are dispersed, so the very heterogeneity that was negligible ex-ante becomes first-order. Retailer adaptation is now worth considering, and the brand weighs it against the loss of discipline exactly as before—but with the adaptation benefit no longer diluted by the rarity of the shock.

The threshold $\bar{\sigma}$ is where these forces balance *conditional on a shock*. When exposures are tightly clustered ($\sigma \leq \bar{\sigma}$), even the realized dispersion is too small to justify ceding pricing authority, and

the brand binds all retailers at the common price floor $\underline{p}^N$. When they are dispersed ($\sigma > \bar{\sigma}$), the value of adaptation dominates and the brand partially delegates. It is worth noting that the two lemmas turn on different moments of demand: the uninformed brand's decision is governed by the *probability* of a shock (θ), while the informed brand's decision is determined by its *dispersion* (σ). This is the signature of the mechanism—information is valuable not because it lowers demand in expectation, but because it tells the brand when retailer-level heterogeneity is worth adapting to.

3.2.2 Optimal Information Acquisition

Lemma 4. *Suppose condition (2) holds. Moreover, suppose the cost of acquiring information is not too high ($I < \bar{I}$), then there exists a threshold $\bar{\phi} \in (0,1)$ such that the brand optimally acquires information if and only if $\phi \geq \bar{\phi}$.*[12]

The brand acquires information only when its processing capacity ϕ clears a threshold $\bar{\phi}$. A more precise signal is worth more—formally, a higher ϕ is a Blackwell improvement—but the gain is bounded by the value associated with perfectly detecting the shock, $\bar{I}$. When ϕ is low, the signal rarely detects a shock, the brand's pricing is nearly uninformed regardless, and paying I is not worthwhile; once ϕ is high enough, the improved pricing justifies the cost.

This threshold structure is what connects the model to our empirical setting. A technology that raises the brand's capacity to process demand information—the introduction of generative AI in our application—moves some brand owners across $\bar{\phi}$, shifting them discretely from uniform toward more localized pricing.

3.3 Information and Equilibrium Pricing

Price Dispersion. We study two dimensions of price dispersion. The first is within-retailer cross-market dispersion:

$$\Delta_M = \mathbb{E}_j[Var_m(p^*_{mj})], \tag{21}$$

which measures whether the same retailer charges different prices in different markets. The second is within-market cross-retailer dispersion:

$$\Delta_R = \mathbb{E}_m[Var_j(p^*_{mj})] \tag{22}$$

which measures whether the retailers charge different prices in the same local market.

Proposition 1. *Suppose condition (2) holds and $I < \bar{I}$. Then, an increase in ϕ from $\phi_0 < \bar{\phi}$ to $\phi_0 + \Delta_\phi \geq \bar{\phi}$ increases within-retailer cross-market price dispersion (Δ_M). If in addition $\sigma > \bar{\sigma}$, then it also increases within-market cross-retailer price dispersion (Δ_R).*

[12] The formulation of the thresholds ($\bar{I}$, $\bar{\phi}$) are given in the proof of the lemma in the Online Appendix.

Proposition 1 shows that improved information breaks price uniformity along two distinct margins, and that the two margins respond to different primitives. When the brand is uninformed, a single floor binds every retailer in every market, and prices are uniform in both dimensions—this is the benchmark that matches the near-uniform pricing that DellaVigna and Gentzkow (2019) document across a chain's stores. Acquiring information relaxes uniformity.

Cross-market dispersion (Δ_M) emerges as soon as the brand becomes informed. A signal lets the brand set different floors in different markets—a higher floor where a shock is unlikely, a lower one where it detects a shock—so a given retailer's price now moves with local conditions across the markets it serves. This requires only that the brand act on its signal; it does not require any retailer-level heterogeneity. Cross-retailer dispersion (Δ_R), by contrast, emerges only when the brand partially delegates the pricing decision, which it does only when exposures are dispersed enough ($\sigma > \bar{\sigma}$) to make retailer adaptation worth the loss of intrabrand discipline. Differential prices arise within a market not due to the brand's information per se, but by its decision to allow differently exposed retailers to price differently.

Price Sensitivity. In addition to price dispersion, we characterize how prices respond to the two components of local demand. We study two measures of price sensitivity. The first is price sensitivity to the market-level shock (z_m):

$$\Xi_z = \mathbb{E}\left[p^*_{mj} | z_m = 0\right] - \mathbb{E}\left[p^*_{mj} | z_m = 1\right]. \tag{23}$$

The second is price sensitivity to the retailer-specific shocks (d_{mj}):

$$\Xi_d = \mathbb{E}\left[p^*_{mj} | z_m = 1, d_{mj} = \mu - \sigma\right] - \mathbb{E}\left[p^*_{mj} | z_m = 1, d_{mj} = \mu + \sigma\right]. \tag{24}$$

Proposition 2. *Suppose condition (2) holds and $I < \bar{I}$. Then, an increase in ϕ from $\phi_0 < \bar{\phi}$ to $\phi_0 + \Delta_\phi \geq \bar{\phi}$ increases the price sensitivity to the market-level demand shock (Ξ_z). If in addition $\sigma > \bar{\sigma}$, it also increases the price sensitivity to the retailer-specific shocks (Ξ_d).*

Proposition 2 shows how the brand's information affects the way equilibrium prices track underlying demand shocks. An uninformed brand sets one floor that binds everywhere, so prices move with neither the market-level shock nor retailer-level exposure: $\Xi_z = \Xi_d = 0$. Information makes prices responsive, but to the two components in different ways and under different conditions.

Prices become sensitive to the market-level shock (Ξ_z) as soon as the brand becomes informed. Detecting a shock lets the brand lower the floor where demand is weak, so prices fall in shock markets relative to no-shock markets. Sensitivity to retailer-level exposure (Ξ_d), by contrast, requires more than information about the market-level demand shock (z_m)—it requires delegation. Even a brand that has verified a shock cannot observe which retailers are hit hardest; prices can reflect exposure only if the brand cedes authority and lets retailers price on what they privately know. That happens

only when exposures are dispersed enough ($\sigma > \bar{\sigma}$).

The distinction matters for interpretation. Sensitivity to the market-level shock reflects the brand *acting on its own information*; sensitivity to retailers' shock exposure reflects the brand *delegating to information it cannot access*. The first is a direct consequence of the brand having better information; the second is an indirect one, operating only because improved information makes surrendering partial pricing authority worthwhile to the brand.

Price Level. The preceding results concern how prices *vary*—across markets, across retailers, and with local demand. We now ask how information affects the *level* of prices. Define the average price as

$$\bar{p} = \mathbb{E}[p^*_{mj}], \tag{25}$$

where the expectation is taken over markets, retailers, and signal realizations. Whereas improved information reshapes the distribution of prices whenever the brand acts on its signal, it moves the average price only when the brand partially delegates the pricing decision.

Proposition 3. *Suppose condition (2) holds, $I \leq \bar{I}$, and $\sigma > \bar{\sigma}$. Then, an increase in ϕ from $\phi_0 < \bar{\phi}$ to $\phi_0 + \Delta_\phi \geq \bar{\phi}$ decreases the average price level ($\bar{p}$).*

Proposition 3 isolates a channel distinct from the two preceding results. In this framework with linear demand, information alone is mean-preserving. The optimal no-delegation informed floor is linear in the brand's posterior mean, implying that the informed floors average back to the uninformed floor across signal realizations. When $\sigma \leq \bar{\sigma}$, the brand never delegates, and information merely redistributes prices across markets—raising Δ_M and Ξ_z—without changing their mean.

The average price falls due to partial delegation and the intrabrand externality. When the brand partially delegates pricing authority, it allows some retailers to price without constraints. These retailers price below the level the brand would impose—each ignores that a lower price draws demand from others carrying the brand's product. The brand accepts this markdown because, conditional on a shock, matching prices to retailer exposure is worth more than the discipline it sacrifices; the lower average price is the cost of that adaptation.

The dispersion and sensitivity results describe prices moving into closer alignment with local conditions; the price-level result shows that, where this happens through delegation, consumers pay less on average. This result distinguishes the information channel we study from algorithmic or automated pricing, where improved firm technology is typically associated with *higher* prices: here, better information lowers prices precisely because it leads the brand to surrender, rather than tighten, control over retail pricing.

3.4 Repricing

Our framework is static, but we interpret it as the per-period problem of a brand and its retailers. In each period t, the demand shocks (z_{mt}) and retailer exposures (d_{mjt}) are drawn independently across periods and markets, and the brand decides whether to acquire information (a_{mt}), observes its signal (s_{mt}), and sets price floors ($\underline{p}_{mt}(s_{mt})$); retailers then price subject to their private information and the floors. When the model's primitives are constant over time, the acquisition decision of Lemma 4 is the same in every period. Without adjustment costs and other intertemporal linkages, each period is the static problem of Section 3.2.1 and Section 3.2.2, and price movements reflect only changes in the brand's signal or in retailers' realized exposures. Let the repricing frequency be defined as

$$R = \Pr(p^*_{mj,t} \neq p^*_{mj,t-1}). \tag{26}$$

Proposition 4. *Suppose condition (2) holds and $I < \bar{I}$. If the brand does not acquire information, $R = 0$. If it does,*

$$R = \begin{cases} 2\theta\phi(1-\theta\phi), & \sigma \leq \bar{\sigma}, \\ 2\theta\phi(1-\theta\phi) + \frac{1}{2}(\theta\phi)^2, & \sigma > \bar{\sigma}, \end{cases} \tag{27}$$

which is strictly positive. Hence, an increase in ϕ from $\phi_0 < \bar{\phi}$ to $\phi_0 + \Delta_\phi \geq \bar{\phi}$ raises the frequency of price adjustment.

An uninformed brand's price floor, $\underline{p}^0 = \frac{\alpha - \theta\mu}{2(\beta-\gamma)}$ depends on model parameters $(\alpha, \beta, \gamma, \theta, \mu)$. When the parameters do not change from period to period, the uninformed brand's price floor does not move, binding all retailers in all periods—despite the absence of costs to change them. Rigidity here is not the result of a menu cost but of a binding constraint set by an uninformed brand. An informed brand's floor instead tracks the realized signal, and retailers' prices move with it.

Uniformity across stores and rigidity over time are thus the same phenomenon in this framework—a floor that binds at an unchanging level—and both are relaxed by the same improvement in the brand's information-processing capacity. In this sense, the model speaks jointly to the two literatures on which we build: the geographic uniformity documented by DellaVigna and Gentzkow (2019) and the temporal rigidity that menu-cost models attribute to the physical cost of price adjustment (Golosov and Lucas, 2007; Midrigan, 2011).

3.5 Empirical Predictions

The model's main predictions revolve around two primitives. The intensity of intrabrand competition (γ) determines whether the brand imposes a binding price floor. Its information-processing capacity (ϕ) determines how well that floor can be tailored to local conditions. Below a threshold ($\bar{\phi}$), the brand does not find it worthwhile to acquire information and sets a uniform price floor that binds all retailers. Above it, the brand collects information and conditions the price floors on

its signal and—where retailer exposures to demand shocks differ greatly—partially cedes pricing authority to retailers.

Our empirical setting allows us to study variations in both primitives. Vertical integration lowers γ: a private-label product is carried by no rival retailer. A general-purpose information-processing technology increases ϕ: its availability raises the capacity of exposed brands. Because information acquisition is a threshold decision (Lemma 4), the model predicts a discrete shift in pricing behavior for brands crossing $\bar{\phi}$, which motivates our comparison of high- against low-exposure brands. Brand acquisitions vary neither primitive. They test the framework's premise: that pricing authority over a product rests with a single party who applies a common policy across the markets in which it is sold. If that party is the brand owner, transferring ownership should transfer the pricing policy, holding the product and the retailer fixed.

Prediction 1 (Intrabrand Competition). The price floor exists because retailers compete excessively for the same customers. The brand cedes pricing authority more readily as that competition weakens. In the limit $\gamma \to 0$, the prices that unconstrained retailers choose coincide with those that maximize total channel profits. Private-label products approximate this case because the intrabrand externality that motivates the floor is absent. They should therefore price more locally than national brands—greater within-retailer price dispersion, and prices that track local wages where national-brand prices do not—and their pricing should be insensitive to improvements in information-processing capacity.

Prediction 2 (Cross-Market Price Dispersion). AI-exposed brands should exhibit greater within-retailer cross-market price dispersion after the technology's introduction.

Prediction 3 (Sensitivity to Local Demand). AI-exposed brands' prices should respond more strongly to local economic conditions such as county wages.

Prediction 4 (Repricing Frequency). AI-exposed brands should reprice more frequently after the technology's introduction.

Prediction 5 (Cross-Retailer Price Dispersion.) Where retailer exposures to demand shocks vary significantly, AI-exposed brands should exhibit greater within-market cross-retailer price dispersion after the technology's introduction.

Prediction 6 (Price Level.) Where retailer exposures to demand shocks vary significantly, AI-exposed brands should exhibit lower average price levels after the technology's introduction.

Predictions 5 and 6 hold where the dispersion of retailer exposures is large enough that the brand delegates rather than simply lowering a common floor. We do not observe that dispersion,

so the tests that follow estimate unconditional average effects and should be read as consistent with these predictions rather than as conditioning on the regime in which they apply.

4 Data

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

4.1 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 over 300 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 uses the full linked panel from 2010 to 2023. The brand-acquisitions analysis uses a UPC-retailer-month panel of 68 brand acquisitions closed between 2012 and 2023, with twelve month event windows around each acquisition close date. The private-label tests in Section 7 use the full 2010 to 2023 NielsenIQ scanner universe. 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.[13] The AI experiment in Section 8 uses a firm-UPC-month panel for the sample of publicly listed manufacturing firms matched to Compustat with a six 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.

4.2 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.

[13] These samples do not require us to match UPCs to manufacturing firms.

Our primary outcomes capture price dispersion and repricing. Repricing is measured differently in the two designs that use it: the acquisition analysis uses the share of stores at which a UPC's reference price changes in a month, which filters temporary promotions, while the AI experiment uses an indicator for a month-over-month change in the average price exceeding a threshold. The first isolates changes in regular prices; the second captures any repricing, promotional or otherwise. 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 2 reports summary statistics for these variables on the linked store-UPC-month panel.

We report price dispersion at three levels of aggregation. Panel B of Table 2 reports the coefficient of variation (CV) of prices across stores within a retailer-UPC-month, across retailers within a UPC-county-month, and across retailers within a UPC-month. The mean CV is 0.073, 0.138, and 0.177, respectively. Thus, price dispersion is greater between retail chains than within them, even among chains operating in the same county.

Two of our tests measure cross-store price dispersion at narrower units of observation. For the private-label tests, 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 2 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.[14] Panel D of Table 2 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 2. For cross-

[14] 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, Eichenbaum, Jaimovich, Rebelo, and Smith (2014) find that the vast majority of very small price changes in supermarket scanner and CPI microdata reflect measurement error. 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.

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.[15]

4.3 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.[16] 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 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:

$$\text{AI Exposure}_{it} = \sum_{o} S_{iot} \cdot X_o, \tag{28}$$

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 4 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 0.36 to 0.42, and the standard deviation is 0.06 (Panel E of Table 2). 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 0.36, compared to 0.42 at the 75th percentile, a difference of roughly one standard deviation.

[15] 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.

[16] 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).

We use pre-event AI exposure to proxy for firms' likelihood of adopting and benefiting from generative AI following the InstructGPT API introduction, consistent with task-based approaches to technological change in the broader economics literature (e.g., Autor, Levy, and Murnane, 2003; Acemoglu and Restrepo, 2022).[17] Firms whose workforces are concentrated in tasks that generative AI can augment should benefit more from these tools and therefore adopt them more intensively. In our setting, greater adoption should increase firms' capacity to convert existing demand data into market-specific pricing recommendations, reducing the marginal cost of localized pricing. We provide supporting evidence for the adoption link in Section 10. After the InstructGPT API introduction, high-exposure firms increase AI-related job postings and use more implementation-oriented language in SEC 10-K filings.

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.[18] We map task scores into three categories (high, medium, and low) and aggregate them to the occupation and firm levels using the same employment shares as in the overall exposure measure.[19]

We also obtain the 10-K Data Assets measure from Eisfeldt et al. (2025), a GPT 3.5 Turbo score based on the business description in firms' 10-Ks filed in 2022. The score ranges from zero to three and captures whether the filing indicates access to data that would be valuable for large language model analysis.

4.4 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

[17] The design is therefore intent-to-treat: we observe exposure, not adoption. Systematic measurement of generative AI use postdates our event—the first nationally representative U.S. surveys were fielded in 2024—so we do not use it to date firm-level adoption. What that evidence establishes is that diffusion was rapid by historical standards, with work adoption matching the personal computer and overall adoption exceeding both the PC and the internet at a comparable horizon from first mass-market release (Bick, Blandin, and Deming, 2024), which makes a response within our post-period plausible rather than anomalous. Figure IA.B.6 supplies the firm-level counterpart: treated firms discuss AI more in earnings calls after the event, with no differential pre-trend.

[18] 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.

[19] 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. We aggregate occupation-level AEI scores to the firm level using the same Revelio employment shares. Table IA.B.14 in the Internet Appendix reports the corresponding results.

the InstructGPT API in 2022 Q1. Panel F of Table 2 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 (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 measures of limited capacity to work from home at the occupation level. The first is the complement of the Dingel and Neiman (2020) indicator for whether an occupation can be performed entirely from home. The second is the complement of the occupation-level share of American Time Use Survey respondents who report that they can work from home. We aggregate both measures to the firm level using employment shares measured before the pandemic. 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.

5 Empirical Roadmap

Our analysis asks who holds the authority to set retail prices, and why that authority produces uniformity. A variance decomposition, reported alongside the acquisitions analysis, locates the systematic component of price setting with the brand owner. To move from that correlation to the mechanism, we bring three research settings to bear, each examining one element of the brand-retailer relationship and each assigned a distinct question. The brand acquisitions analysis examines *who* holds pricing authority by reassigning ownership while holding the product and its stores fixed. In the private-label analysis, we examine whether uniformity is a property of the *relationship* or of retail execution. The premise is that vertical integration dissolves the brand-retailer relationship

entirely while leaving the shelves, scanner systems, and promotional calendars in place. Finally, we analyze product lifecycles and the generative AI shock to understand *which feature of that relationship* is the operative lever while holding the relationship fixed. The lifecycle analysis uses the gradual accumulation of information as a brand learns about its product. The AI design combines the InstructGPT API introduction with pre-existing workforce exposure to identify firms most likely to adopt and benefit from a technology that lowers the cost of converting existing demand data into market-specific pricing recommendations.

Each design isolates a different element. Acquisitions establish that pricing policy belongs to the brand owner. Private label removes the intrabrand competition that motivates the floor, and asks whether prices localize once it is gone. The third setting holds the relationship fixed and varies only the brand owner's ability to acquire, process, and act on local demand information. Identification of the information margin rests on the AI experiment, which compares changes in pricing outcomes for brand owners with high and low pre-existing workforce exposure around the InstructGPT API introduction. The product lifecycle, in which information accumulates without any shock, corroborates it. Together they point to a common mechanism—retail prices become more localized as the party holding pricing authority becomes better able to act on local demand information.

6 Brand Acquisitions

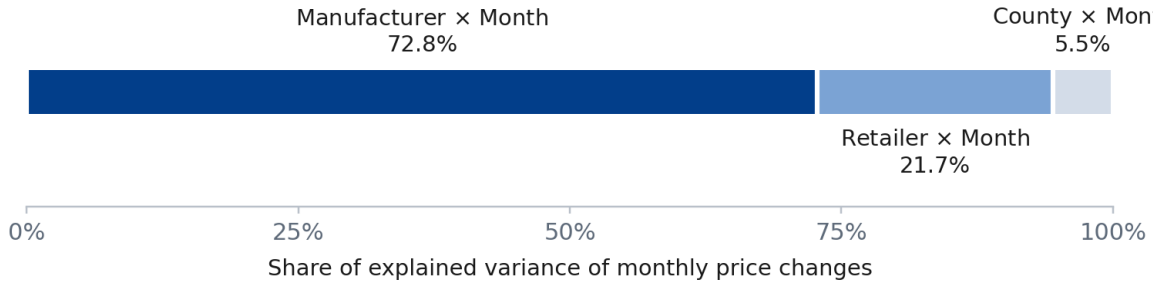
We use brand acquisitions to test whether pricing authority over a product travels with its owner. An acquisition reassigns ownership without altering the physical product or the retail environment that sells it, so if pricing policy is an attribute of the brand owner, acquired UPCs should depart from their pre-acquisition patterns when ownership changes. A variance decomposition points the same way—manufacturer identity is the dominant systematic source of price variation, and the timing of price changes is largely a manufacturer-level event (Figure 2)—but it is only a correlation, which the acquisition design sharpens into a causal test.[20]

We study 68 brand acquisitions completed between 2012 and 2023. We focus on transactions in which the acquirer already owns relevant products, allowing us to compare the acquired brand's pricing behavior with that of products already under the acquirer's control.[21] For each acquisition, we track the acquired brand's UPCs for twelve months before and after the closing date and stack the resulting panels in event time. The unit of observation is UPC $\times$ retailer $\times$ acquisition cohort $\times$ month.

[20] Full Shapley decompositions of price levels, cross-store dispersion, and price changes—including product, manufacturer, and bilateral (DellaVigna and Gentzkow, 2019) variants and a four-factor extension of the interactive fixed-effects Shapley decomposition in Couts, Gonçalves, and Loudis (2026)—appear in Table IA.B.1 in the Internet Appendix.

[21] We exclude six transactions in which the deal structure or the absence of relevant products in the acquirer's pre-acquisition portfolio makes this comparison uninformative.

Figure 2: Who synchronizes price changes?
This figure plots the Shapley decomposition of the explained variance of month-over-month log price changes into manufacturer×month, retailer×month, and county×month components, corresponding to Panel C, column (4) of Table IA.B.1. The sample is the full manufacturer-linked NielsenIQ scanner dataset from 2010 to 2023 at the store×UPC×month level (3.70 billion price-change observations).



The estimation follows Bhattacharya, Illanes, and Stillerman (2026). Because an acquisition can also affect other products in the same category, those products do not provide a clean control group. We therefore compare each acquired UPC with its own extrapolated pre-acquisition path. The first stage fits:

$$Y_{urst} = \alpha_{urs} + \gamma_s \cdot k_{st} + \delta_{m(t)} + \varepsilon_{urst} \tag{29}$$

on pre-acquisition months only. Here, u, r, s, and t index UPCs, retailers, acquisition cohorts, and calendar months, respectively. The UPC × retailer × cohort fixed effect α_{urs} absorbs persistent differences across UPC × retailer cells within each cohort. The term $\gamma_s \cdot k_{st}$ allows each acquisition cohort to follow its own linear trend in event time, while $\delta_{m(t)}$ captures recurring seasonality through indicators for January through December. We use the fitted model to extrapolate each outcome into the post-acquisition period. The second stage estimates the average difference between the realized outcome and this extrapolated path, weighting each acquisition equally. We cluster standard errors at the deal level.

We examine two dimensions of pricing behavior. The first is the within-retailer dispersion of a UPC's log price across stores. The second is regular repricing frequency, measured as the share of stores at which the UPC's reference price changes in a month. We construct the reference price as a rolling median of weekly prices, separating regular price changes from temporary discounts (Nakamura and Steinsson, 2008). For each outcome, we make two comparisons. The acquired UPC's own pre-acquisition mean allows us to test whether its pricing changes when ownership changes. The acquirer's legacy products allow us to test whether the acquired UPC adopts the new owner's pricing behavior. These acquirer benchmarks vary by retailer and month and include only products the acquirer owned before the transaction, so the acquired UPCs cannot enter their own benchmark.

For the second comparison, we divide the sample according to the acquirer's prior experience in the acquired brand's category. In 52 synergistic acquisitions, the acquirer already sold products in the same Nielsen product module as the acquired brand.[22] We therefore compare the acquired UPCs with the acquirer's legacy products in the same retailer and product module. In the remaining 16 non-synergistic acquisitions, the acquirer's legacy products belonged to adjacent modules within the same broader product group.[23] We compare the acquired UPCs in these deals with the acquirer's legacy products in the same retailer and product group. A positive coefficient relative to the acquired UPC's own history indicates departure from its prior pricing behavior. A negative coefficient relative to the acquirer's legacy products indicates convergence toward the new owner's pricing behavior.

Columns (1) and (2) of Table 4 ask whether acquired products change their pricing after ownership changes. The absolute distance from their own pre-acquisition behavior increases by 0.012 for within-retailer price dispersion and by 0.034 for regular repricing frequency. Both estimates are statistically significant at the 1 percent level.[24] Acquired UPCs therefore depart from their prior pricing patterns on both margins.

Columns (3) and (4) ask whether the changes in synergistic acquisitions move acquired products toward the new owner's established pricing behavior. The distance from the acquirer's same-module benchmark falls by 0.013 for both price dispersion and repricing frequency. Thus, when the acquirer already prices products in the same module, acquired UPCs become more similar to its legacy products on both margins. The event-time profiles in Figure 3 show the same transition. After closing, acquired products depart from their historical pricing patterns. Among synergistic acquisitions, the post-acquisition patterns of price dispersion and repricing more closely resemble those of the acquirer's legacy products.

On average, the convergence observed in synergistic acquisitions does not extend to the 16 non-synergistic acquisitions. In columns (5) and (6), the distance from the acquirer's broader product-group benchmark changes by 0.017 for price dispersion and 0.002 for repricing frequency. The same scope condition appears in price levels. Applying the design to the log of the CPI-deflated average price, prices rise by 1.5 percent after acquisition overall and by 2.2 percent among synergistic deals, while non-synergistic deals show no change; the difference between the two carries a deal-clustered wild bootstrap p-value of 0.046 (Table IA.B.2). The magnitude is close to the 1.5 percent average increase that Bhattacharya, Illanes, and Stillerman (2026) report for consummated mergers in US consumer packaged goods, and whose estimator we adopt. It runs opposite to Sheen (2014), who finds that the products of merging competitors converge in quality and that their prices fall relative to rivals. We read the difference in sign as a difference in setting rather than

[22] A product module is the most granular category in the NielsenIQ product taxonomy and groups products with similar uses.
[23] Product groups aggregate related product modules in the NielsenIQ taxonomy.
[24] The corresponding deal-clustered wild-bootstrap p-values, based on 1,000 replications, are below 0.001, below 0.001, below 0.001, 0.010, 0.093, and 0.907 for columns (1) through (6), respectively.

a contradiction: Sheen studies mergers between direct competitors, where consolidation can lower costs, whereas three-quarters of the acquisitions here are purchases of a brand by an owner already operating in its category. The boundary itself is common to both, since his price effects, like ours, are absent when the acquirer is diversifying into a product market it does not already serve.

For our purposes the policy margins are the more informative ones, precisely because they are not price levels. How widely a product's prices vary across stores, and how often they are revised, is not what a change in marginal cost alone would deliver. Convergence on those margins points instead to a category-specific decision rule that an acquirer can apply where it already operates and cannot apply where it does not.

Taken together, the results show that acquired products change how they price when ownership changes and, when the acquirer has relevant category experience, move toward the new owner's established pricing routine. This pattern is difficult to reconcile with pricing determined solely by retailers because the product and retailer remain fixed while the brand owner changes. Instead, it supports the interpretation that brand owners operate category-specific routines that transfer with ownership. The variance decomposition corroborates this reading: it attributes 91% of the systematic variation in price levels, 72.8% of the systematic timing of price changes, and 56% of the systematic cross-store dispersion (against 37% for retailers) to manufacturer identity—an accounting correlation for which the acquisition results supply the causal counterpart.

7 Private Labels vs. National Brands

The acquisition results show that pricing policy belongs to the brand owner. We next ask why that owner withholds discretion, by exploiting the vertical structure of private-label pricing. A private label is carried by no rival retailer, so the intrabrand externality that motivates a price floor is absent: whoever sets the price internalizes its full effect on the brand. The comparison holds the machinery of retail pricing—shelf tags, scanner systems, promotional calendars—roughly fixed. If the floor is what suppresses local price variation, private-label products should be more dispersed 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 lower-margin products on which retailers may invest less in price optimization; 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 UPC $\times$ retailer $\times$ county $\times$ month level. Panel B regresses log prices on the interaction of the private-label indicator with county-level log average monthly wages at the store $\times$ UPC $\times$ 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 2, 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×month and retailer fixed effects and column (3) includes retailer×product module×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 × product module × 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 × product module × 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 acquisition and private-label results place pricing policy with the brand owner and tie uniformity to the floor it imposes. We next ask what happens when the brand's capacity to process local demand information rises. 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 brand owners with high versus low pre-existing AI exposure. The event was firm-facing: Google Trends search interest for "artificial intelligence" remained flat throughout the estimation window, spiking only after the consumer-facing ChatGPT release in November 2022 (Panel A of Figure 5), so the effects we study predate the period in which AI became salient to the

general public.

The premise of this design is that the binding friction is the cost of *processing* demand information into market-level pricing decisions, not access to the data. Kanazawa, Kawaguchi, Shigeoka, and Watanabe (2026) provide direct evidence that this is the margin the technology moves: an AI system that predicts where demand will arise raises taxi drivers' productivity by shortening search, and the gains accrue to the least skilled drivers, who had the least capacity to process that information unaided. The pricing analogue is a brand owner that can now identify which of its markets warrant a departure from a uniform price. Verified local demand data have been available to brand owners for decades: syndicated scanner data date to the 1980s, and electronic data interchange and point-of-sale sharing diffused through the 1990s and 2000s without a corresponding move away from uniform pricing (DellaVigna and Gentzkow, 2019; Hitsch, Hortaçsu, and Lin, 2021). What a per-market pricing decision required before 2022 was analyst attention—a cost that scales with the number of markets while the value of localization per market is small—so state-contingent pricing was reserved for the largest accounts and the uniform default governed the long tail. Scalable generative AI APIs collapsed the marginal cost of producing a decision-grade recommendation for the marginal market. Because these tools augment the analysts and category managers who already set prices—raising their throughput rather than requiring new pricing systems—the response can plausibly appear within months of the API becoming available rather than over a multi-year adoption cycle. In the model's terms, the InstructGPT introduction is a shock to the brand's information capacity ϕ (equivalently, to the acquisition cost I) on the extensive margin of markets, which is why the predicted response is greater dispersion, more frequent repricing, and greater sensitivity to local conditions rather than a level shift in any single market's price rule.

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

$$\begin{aligned} Y_{ipt} = \alpha_{ip} + \alpha_t + \beta(\mathit{Treated}_i \times \mathit{Post}_t) + \delta\left(\ln(\overline{Q}_{ip,pre}) \times \mathit{Post}_t\right) \\ + \gamma'\left(\mathbf{X}_{i0} \times \mathit{Post}_t\right) + \varepsilon_{ipt}, \end{aligned} \tag{30}$$

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. $\mathit{Treated}_i$ equals one for firms above the median of AI exposure in January 2020. Post_t equals one from January 2022 onward. The term $\overline{Q}_{ip,pre}$ 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: $\ln(\mathit{Employees})$, $\ln(\mathit{Assets})$, Tobin's q, ROA, cash-to-total assets, fixed assets-to-total assets, R&D-to-total assets, market-to-book, and $\ln(\mathit{Age})$.

Interacted with $Post_t$, these pre-period controls absorb differential post-period trends driven by pre-treatment characteristics, so that β is not confounded by, for example, larger or more profitable firms repricing differently 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 2 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 30 absorb the residual differences.

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

$$\begin{aligned} Y_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \in \{-6,\ldots,+6\},\, k \neq -3} \beta_k \left(\mathbb{1}[Rel_t = k] \times Treated_i \right) \\ + \delta \left(\ln(\overline{Q}_{ip,pre}) \times Post_t \right) + \gamma' \left(\mathbf{X}_{i0} \times Post_t \right) + \varepsilon_{ipt}, \end{aligned} \tag{31}$$

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 3). 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 the dispersion of prices across stores 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 (4) is the standard deviation of log prices within a firm-product-month, and in columns (5) to

(8) is the interquartile range. For each measure, we compute dispersion across all stores (columns (1) and (5)), within counties (columns (2) and (6)), within retailers (columns (3) and (7)), and across retailers within the same county (columns (4) and (8)). The cross-retailer measure holds the county fixed and asks whether different chains charge different prices for the same product.

The treatment effect is positive in all eight specifications, and significant at the 5% level in seven of the eight. 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. The effect is larger for dispersion measured across all stores (columns (1) and (5)) than for either of the two components the model distinguishes: within-retailer cross-market dispersion, the Δ_M of 1 (columns (3) and (7)), and within-market cross-retailer dispersion, its Δ_R (columns (4) and (8)). Dispersion within counties (columns (2) and (6)) pools both and is reported as a supplementary cut. Much of the increase in the pooled measure therefore occurs across geographic areas and across retail chains rather than within a single chain's stores in a single market. That said, the within-county and within-retailer effects are also statistically significant, suggesting that AI adoption also enables price changes within local markets. Figure 6 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.

Columns (4) and (8) test the margin the model ties most directly to delegation: dispersion across retailers within the same market. The standard deviation of log prices across retailer-level prices within a county increases by 0.017 for high-exposure firms after the event, and the interquartile range increases by 0.013—both exceeding the corresponding within-retailer estimates in columns (3) and (7). In the corresponding event-study specification, the pre-treatment coefficients are jointly indistinguishable from zero ($p = 0.66$ for the standard deviation, 0.70 for the interquartile range). The cross-retailer margin also separates the information channel from a cost-based alternative, and Table IA.B.11 sharpens the separation: the entire increase in dispersion occurs below a pre-period reference price, with the above-reference margin flat, which is the one-sided signature of a floor being relaxed rather than of a common cost change passed through unevenly. A common input-cost shock moves the prices of every chain carrying a UPC and leaves relative prices across chains unchanged, so it cannot generate rising cross-retailer dispersion within a county. The pattern instead matches the model's delegation prediction, in which price dispersion across retailers within a market appears together with the decline in average prices.

The increase in price dispersion is difficult to explain through a uniform decline in production costs but follows directly from the model's information mechanism. When the brand becomes better able to process information, it can condition its price floors on market-level signals, increasing dispersion across markets. Where retailer exposures are sufficiently heterogeneous, the brand may also relax a binding floor and allow retailers to price on their private information, generating

dispersion within markets. Better information lets the brand tailor, and ultimately relax, the floor that holds prices together across the heterogeneous markets DellaVigna and Gentzkow (2019) document.

8.2 Price changes

We next examine whether AI adoption 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-exposure firms relative to low-exposure firms after the introduction of the InstructGPT API. Relative to the sample mean of 0.337 reported in Panel D of Table 2, 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 7 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 are more consistent with the information mechanism than with a uniform decline in production costs. Once information acquisition becomes worthwhile, the brand revises its market-specific price floors as its signals change, increasing repricing frequency even without delegating pricing authority. A uniform cost decline may generate a one-time price adjustment, but it does not naturally produce a sustained increase in repricing. The concentration in price *decreases* is also consistent with the model's delegation margin, under which the brand accepts lower prices from unconstrained retailers so that prices can better reflect local demand. Table IA.B.11 shows that the asymmetry is not confined to the frequency of adjustment: the dispersion increase is itself

one-sided, occurring entirely below a pre-period reference price.

8.3 Price level and local economic conditions

Having documented an increase in price 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 30, with $\ln(\textit{Average price})$ as the dependent variable. The estimate is -0.043, suggesting that high-exposure firms reduced prices by approximately 4.3% relative to low-exposure firms after the InstructGPT API introduction. Decomposing deviations from a pre-period reference price attributes roughly five-sixths of the movement in the mean to mass opening below the reference rather than to a shift in the whole distribution (Table IA.B.11). 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 $\textit{Treated} \times \textit{Post} \times \ln(\textit{Wage})$ 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 high-exposure firms' prices load more strongly on local wage conditions after the event, with more pronounced price declines in lower-wage counties.

Figure 8 presents the corresponding event-study estimates from Equation 31. 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 average-price decline and wage interaction map to distinct predictions of the model. Greater wage sensitivity indicates that treated brands condition prices more closely on local demand, as expected when better-informed brands act on market-level signals. The decline in average prices points to the model's delegation margin. When the brand relaxes a binding floor, unconstrained retailers price on their private information and charge less than the brand would otherwise impose. A common decline in production costs could lower average prices, but it would not explain why high-exposure firms' prices become more responsive to local wages. Together, the results are consistent with AI improving the brand owner's information and allowing greater adaptation to local demand.

In Section 9, we probe this interpretation by examining whether the pricing effects are concentrated in products with more volatile or elastic demand, where local demand information is especially valuable.

9 The Information Mechanism

The results in Section 8 are consistent with the brand's information constraint relaxing, but do not, on their own, isolate the mechanism: several of the patterns, including the change in average prices, can arise through other channels. We therefore examine how pricing changes as brand owners accumulate product-market information, and how the AI effects vary across product characteristics and dimensions of AI exposure.

If information is costly to process, prices should become more localized as brand owners accumulate experience with a product, market, or retailer relationship. The AI effects should concentrate where the value of information is highest, and should operate through the augmentation of worker decision-making rather than the automation of routine tasks.

9.1 Information without a shock: the product lifecycle

The AI experiment identifies a shift in the brand owner's information capacity. If that mechanism operates outside the technology event, the same margin should be visible wherever information accumulates naturally—as a brand gains experience with a product, a market, or a retailer relationship. We test that here, observationally: there is no shock and no exposure measure, and we read the estimates as corroboration rather than identification. We measure information stock as the log of one plus cumulative store-weeks observed through month $t-3$, using the relevant history for each specification.

Table 10 reports four specifications. Column (1) follows new UPCs after launch. Column (2) follows established UPCs entering new counties. Column (3) follows established UPCs entering new retailers. Column (4) uses the most restrictive comparison, following an existing UPC-retailer relationship as it expands into new counties while absorbing UPC×retailer×month fixed effects. Across all four specifications, the coefficient on lagged information stock is positive and statistically significant. The magnitudes are smaller in the most restrictive specification, as expected, but the sign is unchanged. These patterns are consistent with pricing becoming more locally differentiated as brand owners accumulate information about where and through whom a product is sold.

A concern with these estimates is that products enter a broader and more heterogeneous set of stores as they age. Because information stock also increases with product age, the baseline estimates may capture changes in store composition rather than greater price differentiation within the same stores. Panel B addresses this concern in two steps. The first row re-estimates the Panel A specification on the restricted sample for which both dispersion measures are available, allowing the store set to change over time as in Panel A. The second row uses the identical sample and specification but recomputes dispersion using only the stores present in each cell's entry cohort. Comparing the two rows therefore isolates the effect of changes in the stores over which dispersion is measured.

The most demanding comparison is column (4), which follows an established UPC-retailer relationship as it expands into new counties. There the coefficient survives the fixed store set and is larger under it. That direction is what the mechanism predicts: stores added during a rollout enter at or near the chain's prevailing price, so their inclusion compresses measured dispersion, and holding the cohort fixed removes an attenuation rather than manufacturing an effect. Launches and entry into new retailers also survive, at roughly one-half and one-quarter of their restricted-sample magnitudes.

The exception is county entry measured on its own, where the fixed-store coefficient falls to zero. We read this as informative rather than as a casualty. Were the gradient produced by something mechanical—products aging into markdown cycles, or secular trends in the panel—it would appear in the county specification as well. Instead it survives precisely where the framework locates pricing authority, in the relationships through which a brand sells, and vanishes where the framework gives it no purchase, in geography as such. Information appears to attach to the contracting parties rather than to the map.

Two limits deserve emphasis. These are gradients rather than treatment effects: information stock is an outcome of a product's own rollout, not an assigned variation, and the estimates are modest in economic terms even where precisely estimated. And cumulative observation accrues to both parties, so the design cannot by itself attribute the accumulated information to the brand rather than to the retailer; that attribution rests on the AI experiment, where exposure is measured on the brand owner's workforce.

9.2 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 $\textit{Treated} \times \textit{Post}$ 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 11 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.

9.3 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.

Panel A of Table 12 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.[25]

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 0.017 higher dispersion, a 4.7 percentage point increase in repricing frequency, and prices that are 3.9 percent lower after the

[25] Table IA.B.17 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.[26]

Panel C separates firms using the 10-K Data Assets measure from Eisfeldt et al. (2025), a GPT-based score of whether business descriptions in firms' 10-Ks filed in 2022 indicate access to data valuable for large language model analysis. We classify firms as having high or low data assets based on the median score. The dispersion and price-level effects are concentrated among high-data-assets firms. Dispersion increases by 0.025 and prices decline by 4.6 percent, while the coefficients for low-data-assets firms 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 effects are concentrated among firms with high data-assets scores, consistent with AI being more useful when firms have access to data that these tools can process.

10 Robustness and Validation

Our design treats pre-existing workforce exposure as a proxy for the likelihood of adopting generative AI after January 2022, so we first ask whether high-exposure firms visibly adopted it. Adoption is difficult to observe directly, so we use two indirect proxies: AI-related job postings from LinkUp, classified into five categories from job titles and descriptions, and the frequency of AI terms in 10-K filings, measured for the full filing and separately by item, and split between adoption and hiring language and risk and regulatory language.[27] Both proxies move (Table 13). After the introduction, high-exposure firms post roughly 8 more AI positions per month and 22 more across all AI categories than low-exposure firms (Panel A), and mention AI 1.5 more times per filing, with the increase concentrated in Items 1 and 1A rather than in the MD&A (Panel B). It is driven by adoption and deployment language rather than risk and compliance language: 1.03 for implementation and hiring keywords against 0.08 for risk and regulatory terms (Panel C). Pre-event coefficients are close to zero in both series and the post-event increase persists (Figure 10), the pattern extends to quarterly earnings calls (Figure IA.B.6 and Table IA.B.15), and a placebo using non-AI technology keywords such as cloud computing, ERP, and cybersecurity shows no differential change (Figure IA.B.7). This suggests that the shift is specific to artificial intelligence rather than part of a general rise in technology-related discussion, and that our exposure measure captures meaningful differences in AI adoption.

[26] Table IA.B.14 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. The results are robust to this alternative decomposition. Task iteration exposure predicts significantly greater price dispersion, more frequent repricing, and lower prices, whereas directive exposure is not significantly associated with dispersion or price levels and is only marginally associated with repricing.
[27] Appendix A describes the keywords and classification criteria for each indicator variable.

The pricing estimates then survive two classes of concern: that they reflect retailer rather than brand behavior, and that January 2022 coincided with other macroeconomic developments. Adding retailer×UPC and retailer×calendar month fixed effects, which absorb chain-level pricing differences and chain-wide promotional cycles, leaves all three effects signed and significant, though smaller in the retailer panel (Table 9); the price-level effect is negative for virtually every retailer in the data, consistent with a common brand-side component passed through heterogeneously (Figure 9). The event falls in the tail of the Omicron wave (Figure 5), but controlling for limited capacity to work from home—the Dingel and Neiman (2020) measure and an American Time Use Survey analogue—and for county COVID-19 cases and deaths leaves all three effects close to baseline (Table IA.B.8), as does controlling for county exposure to rate-sensitive employment around the Federal Reserve's March 2022 tightening (Table IA.B.9). The window also contains the February 2022 invasion of Ukraine and the input-cost shock that followed, but treated and control firms show no differential change in margins or costs over the period (Table IA.B.10). Deflating with category-specific CPIs, mapping NielsenIQ product groups to 26 detailed BLS categories, leaves both the price decline and the wage interactions intact, so the result is not differential inflation across categories (Table IA.B.3).

Nor do the findings depend on how we measure exposure, form the sample, or conduct inference. Exposure scores generated by large language models can depend materially on which model generates them (Yin, Vu, and Persico, 2026). An alternative exposure measure, built from the target-judge framework in which GPT-5 mini evaluates whether text-davinci-001—closer to the capabilities available at the time of the event—can perform each O*NET task, reproduces all three headline results with similar event-study dynamics, though the directional repricing estimates and the wage interactions are less precisely estimated (Table IA.B.5, Table IA.B.6, and Table IA.B.7; Figure IA.B.2, Figure IA.B.3, and Figure IA.B.4). The effects hold in the full unmatched sample under both binary and continuous treatment, and rise with exposure intensity (Table IA.B.4). Placebo events in January 2019 and January 2021 yield small, insignificant, sign-varying coefficients (Table IA.B.13). Randomization inference that excludes three firms with outsized product portfolios and reassigns treatment across the remainder 5,000 times, preserving the number treated and rematching products at each iteration, places the baseline statistics well into the tails of the placebo distributions, with two-sided p-values of 0.015 for price dispersion, 0.009 for repricing frequency, and 0.011 for log average prices (Figure IA.B.5). The exclusion is specific to this exercise rather than to the results: dropping the same three firms from the baseline specifications leaves all three effects statistically significant, with dispersion and repricing essentially unchanged and the price decline attenuating by roughly a third.[28] Leave-one-out estimates dropping each of the 91 manufacturers (Figure IA.B.8) or each of the 100 product groups (Figure IA.B.9) leave all three

[28] These three firms account for a quarter of the observations in the matched sample, which is why they are excluded here: any permutation draw that assigns one of them to treatment moves a large share of the sample's products at once, leaving the placebo distributions highly irregular. Retaining them leaves the permutation test rejecting as well.

effects similar in magnitude, and wild cluster bootstrap p-values clustered at the firm level agree with the asymptotic inference (Table IA.B.16).

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 offers a complementary explanation. Price-setting in these markets is a delegation problem. The party with the local information is not the party whose payoff the price is designed to protect, and the brand owner cannot verify whether a proposed markdown reflects genuinely weak local demand or competition for a rival retailer's customers. Withholding discretion through a broadly applied, rarely revised price floor is the brand's optimal response, and uniformity across stores and rigidity over time are its joint consequence. Three designs support this reading. Brand acquisitions—in which ownership of a UPC changes while the product and retailer remain fixed—show that pricing behavior travels with ownership: acquired products move away from their own pricing history and converge toward the acquirer's routine only when the acquirer already prices in the same category. Private-label products, for which vertical integration dissolves the brand-retailer relationship altogether, are substantially more dispersed than national brands within the same retailer, category, and month, and their prices track local wages where national-brand prices do not. And when generative AI relaxes the brand's information constraint, geographic dispersion rises—both within a chain's own stores and between competing retailers in the same county—repricing becomes more frequent, prices fall, and prices respond more strongly 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 same gradient appears without any shock, as brands accumulate selling history with a product and a retailer.

Read together, these results recast two facts that are usually explained separately. Geographic uniformity and temporal rigidity are the same object here: a constraint that binds across a chain's stores also binds across months, so a poorly informed brand owner produces both at once, without any appeal to the physical cost of changing a price. The binding friction is also not access to data. Syndicated scanner data have been available for decades, and electronic data interchange diffused through the 1990s without a move away from uniform pricing. What was scarce was the capacity to process that data into a decision for each of thousands of markets, and the ability to verify a local claim well enough to relax a constraint that exists precisely because the agent's incentives differ.

The constraint is imposed because of the intrabrand externality and is blunt because of the information friction, and its consequences reach beyond the level of any single price. The frequency and state dependence of price adjustment—primitives in models of monetary transmission—would

then depend in part on how well informed the party holding pricing authority happens to be, and would move as that information improves rather than staying fixed by the technology of changing a price tag. Cost shocks would pass through more locally as the constraint relaxes. And the distributional consequence emphasized by DellaVigna and Gentzkow (2019)—that uniform pricing holds prices in low-income markets above what local conditions alone would support—should attenuate, a pattern consistent with our finding that the AI-induced price declines are concentrated in lower-wage counties. We note these implications without developing them; quantifying them at the aggregate level requires a different exercise than the one we conduct here.

The central implication for our setting is that the pricing rigidities documented in the literature are not immutable features of retail markets. They reflect, at least in part, the cost of producing and acting on market-specific information at the level of the chain where pricing authority sits—a cost that technology can reduce. The ongoing adoption of digital shelf labels provides a natural test: if the constraint on localization lies in retail execution, these technologies should substantially increase geographic price dispersion; if it lies with the brand owner, 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

Acemoglu, D., and P. Restrepo. 2022. Tasks, automation, and the rise in U.S. wage inequality. *Econometrica* 90:1973–2016.

Adams, J. J., M. Fang, Z. Liu, and Y. Wang. 2026. The rise of AI pricing: Trends, driving forces, and implications for firm performance. *Journal of Monetary Economics* 157:103875–.

Aghion, P., and J. Tirole. 1997. Formal and real authority in organizations. *Journal of Political Economy* 105:1–29.

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.

Alonso, R., and N. Matouschek. 2008. Optimal delegation. *Review of Economic Studies* 75:259–93.

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, Anthropic.

Asker, J. 2016. Diagnosing foreclosure due to exclusive dealing. *Journal of Industrial Economics* 64:375–410.

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.

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.

Bhattacharya, V., G. Illanes, and D. Stillerman. 2026. Merger effects and antitrust enforcement: Evidence from US consumer packaged goods. *American Economic Review* Forthcoming.

Bick, A., A. Blandin, and D. J. Deming. 2024. The rapid adoption of generative AI. Working Paper 32966, National Bureau of Economic Research.

Bils, M., and P. J. Klenow. 2004. Some evidence on the importance of sticky prices. *Journal of Political Economy* 112:947–85.

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.

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.

Deneckere, R., H. P. Marvel, and J. Peck. 1996. Demand uncertainty, inventories, and resale price maintenance. *Quarterly Journal of Economics* 111:885–913.

———. 1997. Demand uncertainty and price maintenance: Markdowns as destructive competition. *American Economic Review* 87:619–41.

Dessein, W. 2002. Authority and communication in organizations. *Review of Economic Studies* 69:811–38.

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, S. Rebelo, and J. Smith. 2014. How frequent are small price changes? *American Economic Journal: Macroeconomics* 6:137–55.

Eisfeldt, A., G. Schubert, B. Taska, and M. B. Zhang. 2025. Generative AI and firm values. *Journal of Finance, forthcoming* .

Farboodi, M., and L. Veldkamp. 2026. A model of the data economy. *Review of Economic Studies* Advance access, doi:10.1093/restud/rdag074.

Fish, S., Y. A. Gonczarowski, and R. I. Shorrer. 2024. Algorithmic collusion by large language models. Working Paper 2404.00806, arXiv.

Frésard, L., G. Hoberg, and G. Phillips. 2020. Innovation activities and integration through vertical acquisitions. *Review of Financial Studies* 33:2937–76.

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.

Graham, J. R., C. R. Harvey, and M. Puri. 2015. Capital allocation and delegation of decision-making authority within firms. *Journal of Financial Economics* 115:449–70.

Grossman, S. J., and O. D. Hart. 1986. The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy* 94:691–719.

Hart, O. D., and J. Moore. 1990. Property rights and the nature of the firm. *Journal of Political Economy* 98:1119–58.

Hitsch, G. J., A. Hortaçsu, and X. Lin. 2021. Prices and promotions in U.S. retail markets: Evidence from big data. *Quantitative Marketing and Economics* 19:289–368.

Hoberg, G., and G. Phillips. 2010. Product market synergies and competition in mergers and acquisitions: A text-based analysis. *Review of Financial Studies* 23:3773–811.

Hoch, S. J., and S. Banerji. 1993. When do private labels succeed? *Sloan Management Review* 34:57–67.

Holmström, B. 1984. On the theory of delegation. In M. Boyer and R. E. Kihlstrom, eds., *Bayesian Models in Economic Theory*, 115–41. Amsterdam: North-Holland.

Humlum, A., and E. Vestergaard. 2025. Large language models, small labor market effects. Working Paper 33777, National Bureau of Economic Research.

Ippolito, P. M. 1991. Resale price maintenance: Empirical evidence from litigation. *Journal of Law and Economics* 34:263–94.

Israeli, A., E. T. Anderson, and A. T. Coughlan. 2016. Minimum advertised pricing: Patterns of violation in competitive retail markets. *Marketing Science* 35:539–64.

Kanazawa, K., D. Kawaguchi, H. Shigeoka, and Y. Watanabe. 2026. AI, skill, and productivity: The case of taxi drivers. *Management Science* 72.

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.

Lafontaine, F., and M. E. Slade. 2008. Exclusive contracts and vertical restraints: Empirical evidence and public policy. In P. Buccirossi, ed., *Handbook of Antitrust Economics*, 391–414. Cambridge, MA: MIT Press.

Leegin Creative Leather Products, Inc. v. PSKS, Inc. 2007. 551 u.s. 877. *Supreme Court of the United States* .

Mackowiak, B., and M. Wiederholt. 2009. Optimal sticky prices under rational inattention. *American Economic Review* 99:769–803.

Maksimovic, V., and G. Phillips. 2001. The market for corporate assets: Who engages in mergers and asset sales and are there efficiency gains? *Journal of Finance* 56:2019–65.

Maksimovic, V., G. Phillips, and N. R. Prabhala. 2011. Post-merger restructuring and the boundaries of the firm. *Journal of Financial Economics* 102:317–43.

Midrigan, V. 2011. Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica* 79:1139–80.

Mortimer, J. H. 2008. Vertical contracts in the video rental industry. *Review of Economic Studies* 75:165–99.

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.

Noy, S., and W. Zhang. 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381:187–92.

Overstreet, T. R. 1983. Resale price maintenance: Economic theories and empirical evidence. Working Paper, Bureau of Economics Staff Report, Federal Trade Commission.

Rey, P., and J. Tirole. 1986. The logic of vertical restraints. *American Economic Review* 921–39.

Sheen, A. 2014. The real product market impact of mergers. *Journal of Finance* 69:2651–88.

Sims, C. A. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50:665–90.

Stein, J. C. 2002. Information production and capital allocation: Decentralized versus hierarchical firms. *Journal of Finance* 57:1891–921.

Stroebel, J., and J. Vavra. 2019. House prices, local demand, and retail prices. *Journal of Political Economy* 127:1391–436.

Telser, L. G. 1960. Why should manufacturers want fair trade? *The Journal of Law and Economics* 3:86–105.

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.

Wu, E. 2024. Rationally inattentive seller: The geographic dispersion of U.S. retail chains and its aggregate implications. Working paper, University of Texas at Austin.

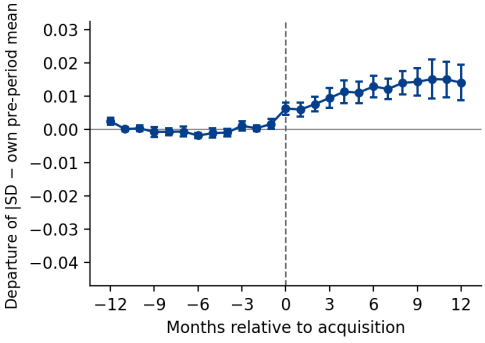
Yin, M., H. Vu, and C. Persico. 2026. How (un)stable are LLM occupational exposure scores? evidence from multi-model replication. Working Paper 35110, National Bureau of Economic Research.

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.

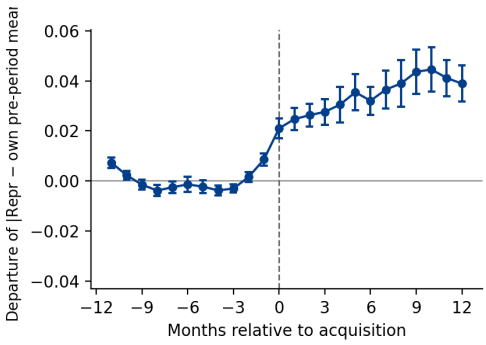
Figure 3: Pricing dynamics around brand acquisitions: expanded sample

This figure plots monthly event-time profiles of pricing behavior around the 68 brand acquisitions in Table 4. For each outcome, we use the available observations from the twelve months before closing to estimate a counterfactual path that accounts for persistent differences across UPC×retailer combinations within acquisition cohorts, acquisition-specific linear trends, and recurring seasonality through indicators for January through December. Each point is the difference between the observed outcome and this counterfactual path in a given event month, averaged first across UPC×retailer observations within each acquisition and then across acquisitions, with each acquisition receiving equal weight. Whiskers show 90% pointwise confidence intervals based on variation across acquisitions, and the vertical dashed line marks the acquisition month ($k = 0$). Panels A and B use all 68 acquisitions. They measure the absolute distance of cross-store price dispersion and regular repricing frequency, respectively, from the acquired UPC×retailer pair's own pre-acquisition mean. Panels C and D restrict the sample to synergistic acquisitions and measure the absolute distance of the same outcomes from the acquirer's legacy products in the same retailer and product module. Panels E and F report the corresponding comparisons for non-synergistic acquisitions using the acquirer's legacy products in the same retailer and product group. The acquirer benchmarks vary by month and include only products owned before closing. *SD* denotes the within-retailer cross-store standard deviation of log prices, and *Repr* denotes the share of stores where the UPC's regular price changes in a month. The repricing profiles begin at $k = -11$ because calculating a price change requires a lagged reference price. In Panels E and F, arrowheads indicate that some confidence intervals extend beyond the common vertical scale. Table 4 reports the corresponding formal tests. Appendix A presents variable definitions.

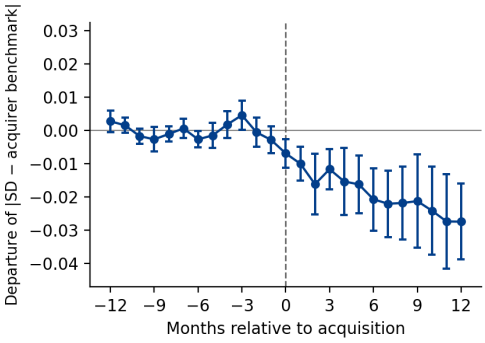
Panel A: Departure of $|\mathrm{SD}_{urt} - \overline{\mathrm{SD}}_{ur,\mathrm{pre}}|$



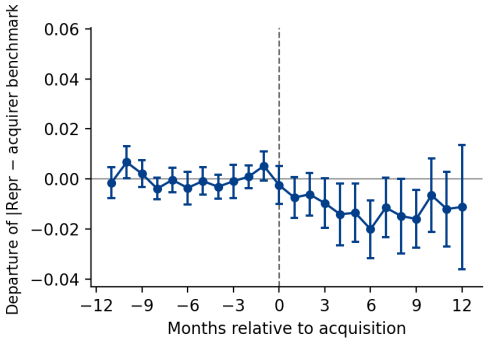
Panel B: Departure of $|\mathrm{Repr}_{urt} - \overline{\mathrm{Repr}}_{ur,\mathrm{pre}}|$



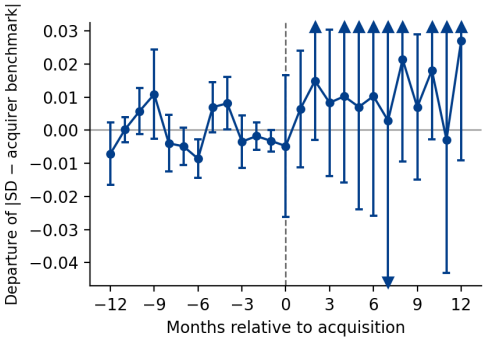
Panel C: Departure of $|\mathrm{SD}_{urt} - \overline{\mathrm{SD}}^{\mathrm{acq}}_{rmt}|$



Panel D: Departure of $|\mathrm{Repr}_{urt} - \overline{\mathrm{Repr}}^{\mathrm{acq}}_{rmt}|$



Panel E: Departure of $|\mathrm{SD}_{urt} - \overline{\mathrm{SD}}^{\mathrm{acq}}_{rgt}|$



Panel F: Departure of $|\mathrm{Repr}_{urt} - \overline{\mathrm{Repr}}^{\mathrm{acq}}_{rgt}|$

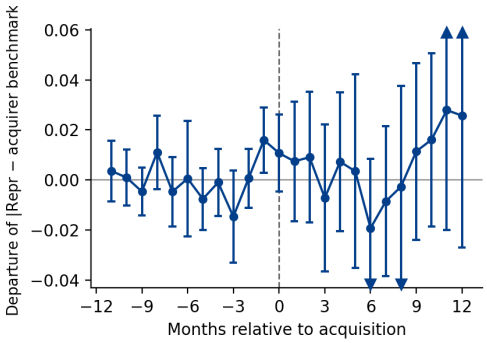


Figure 4: Distribution of AI exposure

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. Appendix A presents variable definitions.

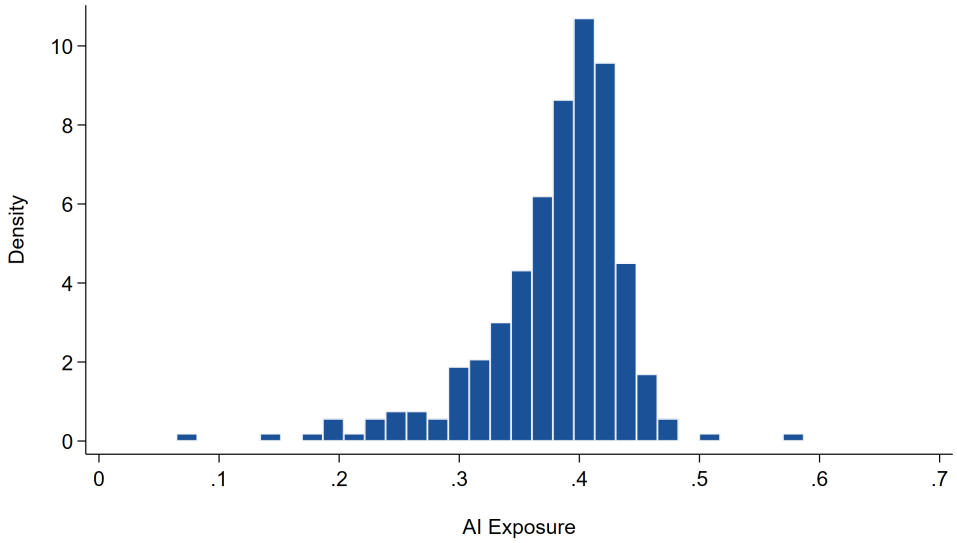
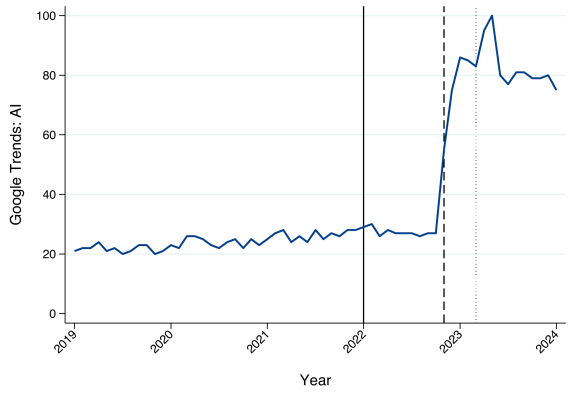


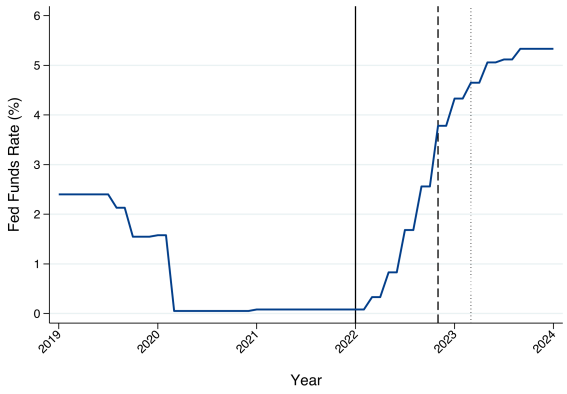
Figure 5: Macro environment around the AI implementation event

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.

Panel A: General AI search interest



Panel B: Federal Reserve policy



Panel C: COVID-19 waves

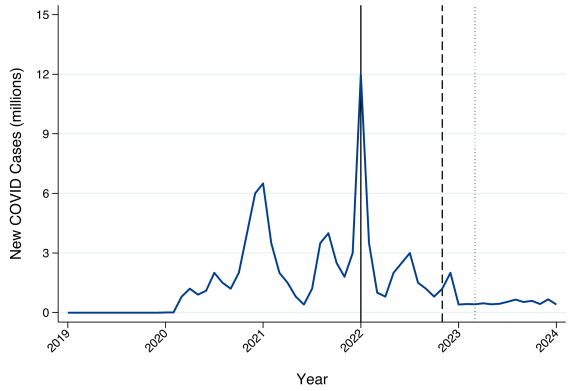
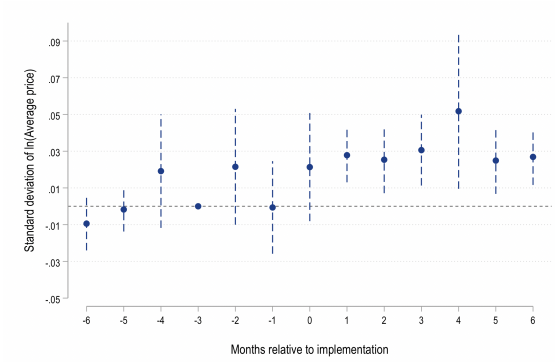


Figure 6: Price dispersion event study

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.

Panel A: Standard deviation



Panel B: Interquartile range

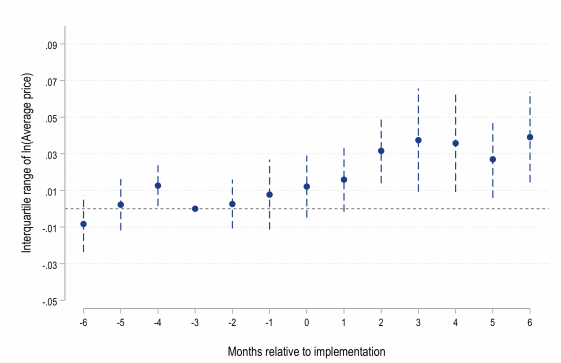
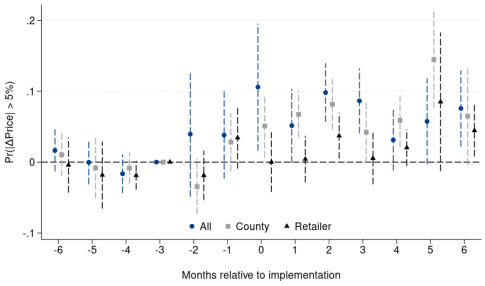


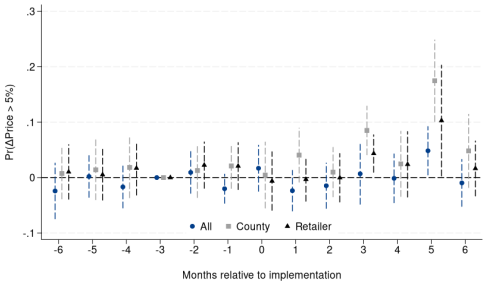
Figure 7: Price changes event study

This figure plots the β_k coefficients and 90% confidence intervals from event-study regressions of the form: $Y_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \neq -3} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{ipt}$, where i denotes firms, p products (UPCs), and t calendar months. In Panel A, $Y_{ipt} = \mathbb{P}(|\Delta \ln \bar{P}| > 5\%)$. 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.

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



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



Panel C: $\mathbb{P}(\Delta P < -5\%)$

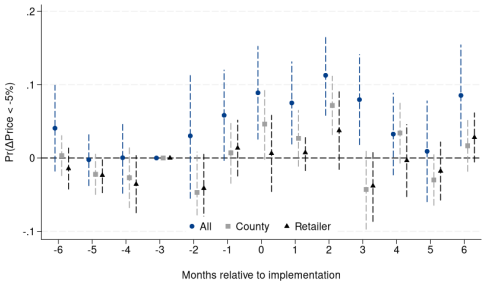


Figure 8: Price level event study

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.

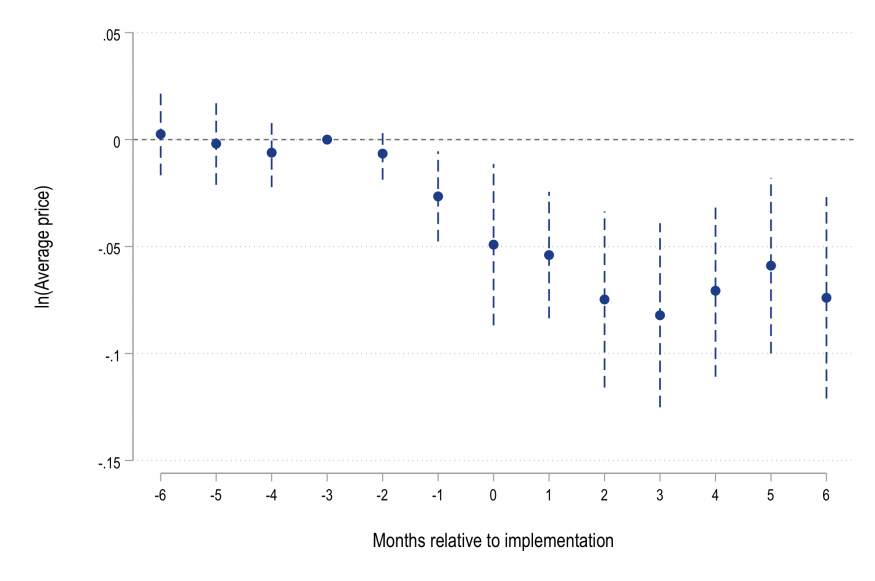


Figure 9: Treatment effect heterogeneity across retailers

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: $\ln(\text{Average price})_{ipt} = \alpha_{ip} + \alpha_t + \sum_k \beta_k (\textit{Treated}_i \times \textit{Post}_t \times \mathbb{1}[\textit{Retailer} = k]) + \mathbf{X}'_{i0}\gamma \times \textit{Post}_t + \epsilon_{ipt}$, 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 *Treated* × *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 in January 2022 and the six months following the InstructGPT API introduction. 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.

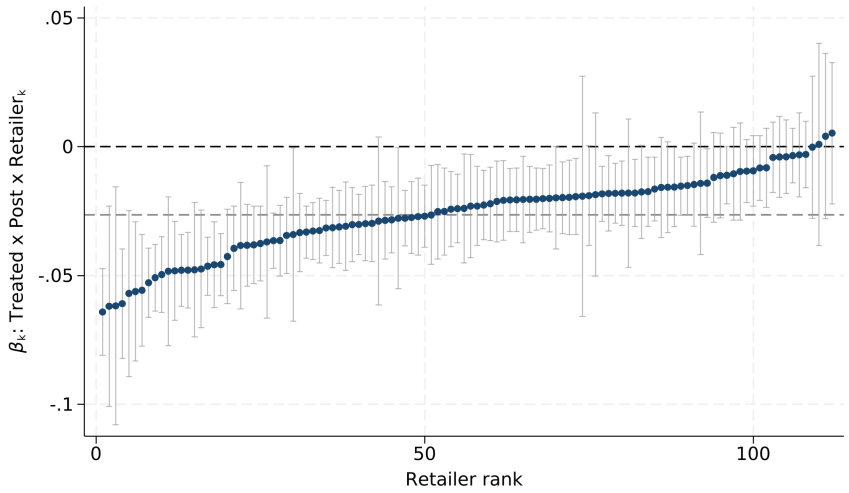
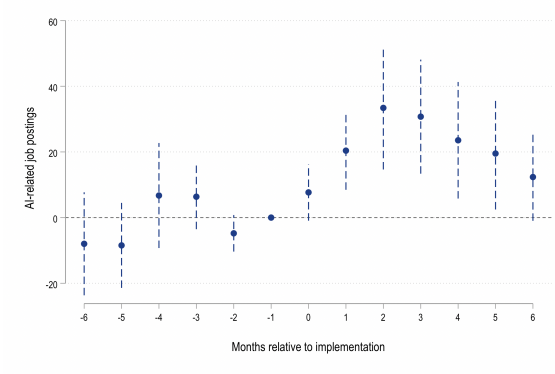


Figure 10: Evidence of AI adoption

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

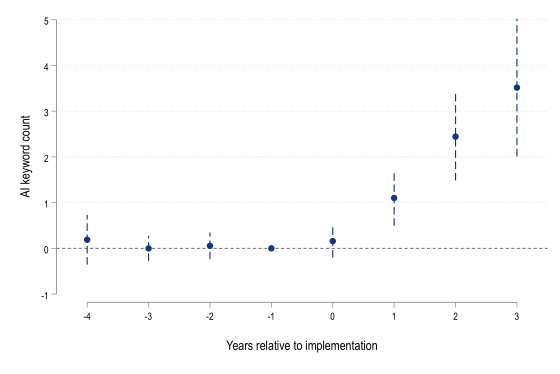


Table 2: Summary Statistics

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
Δln(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($\lvert\%\Delta P\rvert > 5\%$)	90,211	0.337	0.473	0.000	0.000	1.000
P($\lvert\%\Delta P\rvert > 10\%$)	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 3: 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 4: Pricing behavior around brand acquisitions

This table examines changes in pricing behavior around 68 brand acquisitions completed between 2012 and 2023. The sample includes 52 synergistic acquisitions, in which the acquirer already sold products in the acquired brand's Nielsen product modules, and 16 non-synergistic acquisitions, in which the acquirer's legacy products belonged to adjacent modules in the same broader product groups. The estimation window is ±12 months around each closing date, and the unit of observation is UPC × retailer × acquisition cohort × month. Following Bhattacharya, Illanes, and Stillerman (2026), we use pre-acquisition observations to estimate a counterfactual path for each outcome. We then estimate the average post-acquisition departure from the counterfactual path, weighting each acquisition equally. Columns (1) and (2) measure the absolute distance of cross-store price dispersion and regular repricing frequency from the acquired UPC's own pre-acquisition mean. Columns (3) and (4) measure the absolute distance of the same outcomes from the acquirer's legacy benchmark for synergistic acquisitions. Columns (5) and (6) report the corresponding product-group comparisons for non-synergistic acquisitions. *SD* is the within-retailer cross-store standard deviation of log prices. Repricing frequency is the share of stores with a change in the UPC's rolling-median reference price. Acquirer benchmarks vary by retailer and month and include only products owned by the acquirer before the transaction. Positive coefficients in columns (1) and (2) indicate departure from the acquired UPC's prior pricing behavior. Negative coefficients in columns (3) through (6) indicate convergence toward the acquirer's legacy behavior. All specifications include UPC × retailer × acquisition cohort fixed effects, fixed effects for January through December, and acquisition-specific linear event-time trends. Standard errors clustered at the deal 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.

	Distance from own pre-acquisition pricing		Distance from acquirer's legacy behavior			
			Synergistic acquisitions		Non-synergistic acquisitions	
	SD(ln price)	Repricing frequency	SD(ln price)	Repricing frequency	SD(ln price)	Repricing frequency
	(1)	(2)	(3)	(4)	(5)	(6)
Post acquisition	0.012***	0.034***	−0.013***	−0.013**	0.017	0.002
	(0.002)	(0.003)	(0.003)	(0.005)	(0.010)	(0.015)
Fixed effects						
UPC × retailer × stack	✓	✓	✓	✓	✓	✓
Deal × event-time trend	✓	✓	✓	✓	✓	✓
Month of year	✓	✓	✓	✓	✓	✓
Mean dep. var. (pre-period)	0.043	0.109	0.076	0.179	0.082	0.162
Deals	68	68	52	52	16	16
Observations	4,191,727	4,231,269	1,964,003	2,002,862	277,555	300,946

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 × retailer × county × 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 × UPC × 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^{***}	0.005^{***}	0.006^{***}
	(0.000)	(0.000)	(0.000)
ln(Units sold)	-0.004^{***}	0.001^{***}	0.001^{***}
	(0.000)	(0.000)	(0.000)
Fixed effects			
Month × Product module	✓		
Retailer	✓		
Retailer × Product module		✓	
Retailer × Month		✓	
Retailer × Product module × 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 × ln(Wage)	0.003	0.019^{***}	0.018^{***}
	(0.006)	(0.003)	(0.003)
Private label	-0.220^{***}	-0.387^{***}	-0.378^{***}
	(0.053)	(0.026)	(0.027)
ln(Wage)	0.362^{***}	−0.005	−0.004
	(0.015)	(0.004)	(0.004)
ln(Units sold)	-0.174^{***}	-0.158^{***}	-0.159^{***}
	(0.002)	(0.001)	(0.001)
Fixed effects			
Retailer × Product module × Month	✓		
Retailer × Product module × County		✓	✓
Calendar month		✓	
Retailer × Product module × 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 (4) is the standard deviation of log prices, and in columns (5) to (8) is the interquartile range of log prices. Columns (1) and (5) compute dispersion across all stores within a firm-product-month; columns (2) and (6) compute dispersion within counties; and columns (3) and (7) compute dispersion within retailers. Columns (4) and (8) measure dispersion across retailers. Within each county-week in which at least two retailers carry a UPC, we calculate the dispersion in average log prices across those retailers and then average this measure to the firm-product-month level. *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. 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 stores	Within county	Within retailer	Across retailers	All stores	Within county	Within retailer	Across retailers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treated × Post	0.025***	0.008**	0.009**	0.017***	0.020***	0.001*	0.007**	0.013***
	(0.006)	(0.003)	(0.004)	(0.004)	(0.006)	(0.001)	(0.003)	(0.003)
Ln(Avg. pre-period units) × Post	0.004***	0.002***	0.002***	0.000	0.005***	0.000*	0.001***	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.001)
Ln(Employees) × Post	−0.001	0.002	0.005	−0.001	0.008	0.001	0.006*	−0.001
	(0.006)	(0.003)	(0.005)	(0.004)	(0.007)	(0.001)	(0.004)	(0.003)
Ln(Assets) × Post	−0.000	−0.003	−0.007	−0.000	−0.012	−0.001**	−0.008**	0.000
	(0.006)	(0.003)	(0.005)	(0.004)	(0.007)	(0.001)	(0.004)	(0.003)
Tobin's q × Post	0.000	−0.002*	0.000	0.000	0.002	0.000	0.002	0.000
	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)	(0.000)	(0.001)	(0.001)
ROA × Post	−0.030	−0.054	0.011	0.028	0.208	0.036***	−0.006	0.027
	(0.160)	(0.064)	(0.105)	(0.062)	(0.157)	(0.012)	(0.098)	(0.047)
Cash/assets × Post	−0.004	−0.006	−0.042	0.010	−0.069*	−0.003	−0.040*	0.005
	(0.040)	(0.017)	(0.028)	(0.027)	(0.040)	(0.004)	(0.021)	(0.020)
Fixed assets/assets × Post	0.002	0.013	−0.033**	0.033*	−0.039	−0.005*	−0.067***	0.028**
	(0.019)	(0.010)	(0.015)	(0.017)	(0.024)	(0.003)	(0.015)	(0.013)
R&D/assets × Post	0.157	−0.037	0.256	0.017	−0.069	−0.080**	−0.021	−0.023
	(0.237)	(0.127)	(0.167)	(0.191)	(0.258)	(0.031)	(0.111)	(0.143)
Market-to-book × 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)	(0.000)	(0.000)	(0.000)	(0.000)
Ln(Age) × Post	0.002	0.003	0.003	0.002	0.000	0.000	0.003	0.002
	(0.005)	(0.003)	(0.004)	(0.004)	(0.005)	(0.001)	(0.003)	(0.003)
Fixed effects								
UPC	✓	✓	✓	✓	✓	✓	✓	✓
Calendar month	✓	✓	✓	✓	✓	✓	✓	✓
Observations	77,607	63,304	75,242	44,528	77,607	63,304	75,242	44,528
Adjusted R^2	0.46	0.46	0.46	0.50	0.50	0.78	0.45	0.53

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). *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. 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. 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.

	$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$			$\mathbb{P}(\%\Delta P > 5\%)$			$\mathbb{P}(\%\Delta P < -5\%)$		
	All stores	Within county	Within retailer	All stores	Within county	Within retailer	All stores	Within county	Within retailer
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated × Post	0.064***	0.064***	0.033**	0.018	0.029**	0.016	0.046***	0.035***	0.017
	(0.018)	(0.017)	(0.013)	(0.012)	(0.012)	(0.011)	(0.012)	(0.013)	(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 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. Column (1) reports the baseline price-level treatment effect on the firm-product-month panel. 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. 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. Column (1) includes firm-product and calendar month fixed effects. 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. 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	−0.043***	−0.219**	−0.223**
	(0.011)	(0.098)	(0.100)
Treated × Post × ln(Wage)		0.023**	0.023**
		(0.011)	(0.012)
Treated × ln(Wage)		0.004	0.004
		(0.010)	(0.009)
ln(Wage) × Post		−0.020*	
		(0.011)	
ln(Wage)		−0.012	
		(0.013)	
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
County		✓	
Calendar month	✓	✓	
County × 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 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. 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)). 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. 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. 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			$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$			ln(Average price)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated × Post	0.011^{***}	0.007^{***}	0.005^{*}	0.027^{**}	0.033^{**}	0.028^{**}	-0.036^{***}	-0.033^{***}	-0.024^{***}
	(0.003)	(0.003)	(0.003)	(0.012)	(0.013)	(0.012)	(0.007)	(0.008)	(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 lifecycle and price localization

This table reports regressions of SD of ln(price) on lagged information stock during product lifecycle events. Information stock is the log of one plus cumulative store-weeks observed through month $t-3$, measured at the UPC level in column (1), the UPC×county level in column (2), the UPC×retailer level in column (3), and the UPC×retailer×county level in column (4). Column (1) follows new UPC launches. Column (2) follows established UPCs entering new counties. Column (3) follows established UPCs entering new retailers. Column (4) follows existing UPC-retailer relationships expanding into new counties. Panel A measures dispersion across all stores selling the product. Panel B restricts the sample to observations for which both dispersion measures can be constructed. Its first row re-estimates the Panel A specification on this restricted sample using the original store set, which can change over time. Its second row uses the identical sample and specification but recomputes dispersion using only the stores present in each cell's entry cohort. The difference between the two rows therefore isolates the effect of changes in the stores over which dispersion is measured. Coefficients and standard errors are multiplied by 100. Controls include log units sold, the log store count (log entry-cohort store count in the entry-cohort rows), and promotional share. The fixed effects used in each specification are listed at the bottom of each panel. Standard errors clustered at the unit of entry 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: Original store set, full sample

	SD of ln(price)			
	New UPC	New county	New retailer	New county within retailer
	(1)	(2)	(3)	(4)
Information stock$_{t-3}$	0.307***	0.132***	0.231***	0.038***
	(0.002)	(0.003)	(0.002)	(0.005)
Controls	✓	✓	✓	✓
Fixed effects				
UPC × retailer	✓		✓	
Module × month	✓			
UPC × county		✓		
UPC × month		✓	✓	
Module × county × month		✓		
Retailer × module × month			✓	
UPC × retailer × county				✓
UPC × retailer × month				✓
Observations	173,888,061	148,501,331	121,246,990	93,949,693
Adjusted R^2	0.48	0.58	0.61	0.58

Panel B: Entry-cohort store set, restricted sample

	SD of ln(price)			
	New UPC	New county	New retailer	New county within retailer
	(1)	(2)	(3)	(4)
Information stock$_{t-3}$, original store set	0.306***	0.081***	0.199***	0.024***
	(0.002)	(0.004)	(0.002)	(0.005)
Information stock$_{t-3}$, entry-cohort store set	0.145***	−0.004	0.054***	0.038***
	(0.002)	(0.004)	(0.002)	(0.005)
Controls	✓	✓	✓	✓
Fixed effects, as in Panel A	✓	✓	✓	✓
Observations	164,612,039	116,234,976	107,257,186	94,701,645
Adjusted R^2 (original)	0.48	0.59	0.61	0.58
Adjusted R^2 (entry cohort)	0.42	0.49	0.55	0.50

Table 11: Product-level heterogeneity

This table presents results from difference-in-differences regressions of prices around the introduction of the InstructGPT 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* × *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 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. All specifications include UPC and product group × 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 × Post	−0.003	−0.011	−0.008
	(0.009)	(0.010)	(0.011)
Treated × Post × Product indicator	−0.038***	−0.025**	−0.023**
	(0.013)	(0.011)	(0.011)
Product indicator × Post		0.015*	−0.003
		(0.008)	(0.008)
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
Product group × Calendar month	✓	✓	✓
Observations	89,063	89,050	89,038
Adjusted R^2	0.95	0.95	0.95

Table 12: 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 high- and low-data-assets firms, defined as firms with above- and below-median scores, respectively, on the 10-K Data Assets measure from Eisfeldt et al. (2025). The measure uses GPT 3.5 Turbo to assess whether business descriptions in firms' 10-Ks filed in 2022 indicate access to data valuable for large language model analysis. *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. 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.

Panel A: Pricing-task vs. non-pricing-task AI exposure

	SD of ln(price)	Pr($\lvert\Delta P\rvert > 5\%$)	ln(Average price)
	(1)	(2)	(3)
Pricing-task exposure × Post	0.395	2.619**	−1.414*
	(0.507)	(1.010)	(0.740)
Non-pricing-task exposure × Post	−0.039	0.073	0.370
	(0.179)	(0.396)	(0.294)
p(pricing = non-pricing)	0.499	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($\lvert\Delta P\rvert > 5\%$)	ln(Average price)
	(1)	(2)	(3)
Augmentation exposure × Post	0.017**	0.047**	−0.039***
	(0.007)	(0.020)	(0.010)
Automation exposure × Post	0.006	0.022	−0.014
	(0.010)	(0.022)	(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($\lvert\Delta P\rvert > 5\%$)	ln(Average price)
	(1)	(2)	(3)
High data assets × Post	0.025***	0.066***	−0.046***
	(0.006)	(0.016)	(0.012)
Low data assets × Post	0.005	0.091**	−0.004
	(0.008)	(0.044)	(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 13: 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**	4.448**	4.385**	16.014**	22.451**
	(3.381)	(1.771)	(1.773)	(7.247)	(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 × Post	1.516***	0.498**	0.708***	0.038
	(0.392)	(0.215)	(0.218)	(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 × Post	1.026***	0.405**	0.433***	0.079**	0.089**	0.002
	(0.283)	(0.174)	(0.135)	(0.037)	(0.035)	(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
Panel A: Dependent variables		
ln(Price)	The natural log of the store-level unit price for a UPC in a store-month.	NielsenIQ Retail Panel
ln(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
Δln(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(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(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
$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$	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
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

Variable	Definition	Source
Post	An indicator variable equal to one in the event month and the six months that follow.	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
Panel C: Information-channel variables		
Information stock	The natural log of one plus cumulative store-weeks observed through month $t-3$ for the relevant life-cycle unit: UPC, UPC×county, UPC×retailer, or UPC×retailer×county.	NielsenIQ Retail Panel
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
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
10-K data-assets score	A GPT 3.5 Turbo score from Eisfeldt et al. (2025) based on business descriptions in firms' 10-Ks filed in 2022. The score ranges from zero to three and captures whether the filing indicates access to data that would be valuable for large language model analysis.	SEC EDGAR; Eisfeldt et al. (2025)

Variable	Definition	Source
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(Units sold)	The natural log of total units sold for a firm-product in a given month.	NielsenIQ Retail Panel
ln(Avg. pre-period units) × 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(Employees)	The natural log of the number of employees, measured in the month before the event.	Revelio Labs
ln(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
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

Variable	Definition	Source
ln(Age)	The natural log of firm age.	CRSP
Panel E: Local conditions and robustness variables		
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
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)
Limited WFH capacity (DN)	The complement of the Dingel and Neiman (2020) indicator for whether an occupation can be performed entirely from home, aggregated to the firm level using employment shares measured before the event.	Dingel and Neiman (2020); Revelio Labs
Limited WFH capacity (ATUS)	The complement of the share of American Time Use Survey respondents within an occupation who report that they can work from home, aggregated to the firm level using employment shares 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
ln(County Covid deaths)	The natural log of one plus monthly county-level cumulative COVID-19 deaths.	NYT County-Level COVID-19 Database
Panel F: Product-level heterogeneity variables		
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
Panel G: AI adoption variables		

Variable	Definition	Source
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
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? Information and the Delegation of Pricing Authority

Alvin Chen, Leandro Sanz, and Michael Wittry

IA.A Data

IA.A.1 The manufacturer-retailer contract archive

Section 2 draws on a hand-collected archive of manufacturer-retailer contracts and pricing documents to characterize how pricing authority is allocated along the vertical chain. This section explains how we construct the archive, what it contains, and how we use it.

We assemble the archive from public records that state the pricing relationship between a manufacturer or brand principal and its downstream resellers in the parties' own words. The sources consist of material contract exhibits filed on SEC EDGAR, contracts and program agreements from litigation and regulatory dockets, franchise disclosure documents filed with state registries, and historical trade archives. We locate candidate documents through full-text searches and code the identity and role of the parties, the vertical layer governed by the document, and the pricing and payment instruments it contains. Every retained document is verified against its source. A core set of 52 instances received two rounds of quote verification at the clause level, and the remaining instances were confirmed through independent coding of the source text by multiple models. The resulting archive contains 95 verified document instances. We retain five automobile and marine dealer agreements separately because dealer franchise law governs their pricing terms, bringing the total to 100 when those documents are included.

The documents span roughly four decades, from the early 1980s to 2022, and cover more than a dozen consumer product markets. Table IA.A.2 reports their distribution across markets and identifies representative brand principals. The archive also retains 14 documents with weak or qualified evidence of upstream control rather than selecting only documents that closely fit the model. Table IA.A.1 reproduces representative language organized by model primitive. The clauses document price floors and retailer discretion above them, local sales reporting, promotional transfers, enforcement through chargebacks or termination, and brand protection as a rationale for limiting retailer discretion.

Two features determine how we interpret the archive. First, it is an opportunistic record of documents that entered the public domain, not a random sample of manufacturer-retailer rela-

tionships. We therefore use it to document the existence and form of pricing instruments, not to estimate their prevalence. Public filings are especially likely to understate minimum advertised price and unilateral resale price policies because these instruments generally operate through unilateral policies rather than bilateral contracts. We supplement the bilateral contract count with four additional sources of evidence. These cover pricing authority in 21 franchise systems, unilateral pricing policies and related litigation records, pricing terms used across retailers, and historical fair trade records documenting uniform pricing under laws that permitted resale price maintenance between 1911 and 1975. The policy and litigation records include the Johnson & Johnson Vision Care ACUVUE policy and *Leegin Creative Leather Products v. PSKS*. Second, the archive directly documents the contractual and policy instruments through which a brand owner constrains a retail price that it does not itself set.

IA.A.2 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

Table IA.A.1: Representative contract language by model primitive

This table reproduces representative clauses from the hand-collected manufacturer-retailer contract archive and organizes them by the model primitive each clause implements. Section 2 describes the institutional setting, and Table 1 maps the primitives to their corresponding contract instruments. Quotations are verbatim. "[sic]" marks drafting errors in the source, and minor spacing errors introduced by optical character recognition have been corrected. Ellipses denote text omitted within a clause. Document numbers refer to the archive manifest. The ACUVUE unilateral pricing policy comes from the litigation-sourced policy layer rather than the bilateral contract count. Johnson & Johnson Vision Care characterizes the policy, and the *In re Disposable Contact Lens Antitrust* consolidated complaint reproduces it (MDL No. 2626, Dkt. 395).

Model primitive	*Source (firm, date; doc.)*	*Verbatim contract language*
Price floor	J&J Vision Care, ACUVUE UPP, 2014 (policy layer)	"... the final net price for the product covered by the UPP remains at or above the UPP Minimum Retail Price, after discounts are applied."
	MACH 3 Tech. → Essentially Yours Ind., 2006 (doc. 85)	"all parties hereto will maintain a Uniform Retail Price structure set jointly by MACH 3 and DISTRIBUTOR ... Violation to sell the PRODUCTS by DISTRIBUTOR ... at discounts unauthorized by MACH 3 is prohibited."
Interval delegation	MACH 3 Tech. → Essentially Yours Ind., 2006 (doc. 85)	"DISTRIBUTOR ... shall, however, be allowed to market and sell product at prices that are more expensive than the Uniform Retail Price structure."
	Remote Knowledge → Sea Smart, 2003 (doc. 60)	"Distributor shall have discretion to distribute and price the units, provided, however, RK shall have the right to impose and maintain a minimum advertised price."
Enforcement of the floor	Sears/KCD → Sears Hometown & Outlet, 2012 (doc. 88)	"Seller may charge back to Buyer all price adjustments, charges, and penalties ... including failure to comply with Vendor policies"; and "Seller retains the right to offer minimum advertised price policies ('MAP') and unilateral pricing policies ('UPP')."
	ProdecoTech authorized-dealer agreement, 2015 (doc. 87)	"... 'Dealer's' whom [sic] do not follow MAP pricing will have their account terminated and be in default of the agreement."
Information capacity (ϕ)	R.J. Reynolds (RMSC), Total Nicotine Plan, 2022 (doc. 17)	"Retailer must purchase the Products only from wholesale supplier(s) that provide certain sales volume data, known as Shipment To Retail data ('STR data'), to RMSC through the Management Science Associates, Inc. ('MSAi') clearinghouse"; "'Local SOM' means the local trading area share of market for the Products, as determined by RMSC."
	Philip Morris USA, Retail Leaders Program, 2015 (doc. 14)	"authorizes PM USA and its affiliates to obtain, and Retailer's suppliers to release, any information requested by PM USA ... regarding Retailer's purchases and sales of tobacco products"; "Retailer must submit complete and accurate transaction-level PM USA Cigarette Scan Data ... for each STARS Week."
Lump-sum transfer	Philip Morris USA, Retail Leaders Program, 2015 (doc. 14)	"a PM USA Promotion pursuant to which a Store must reduce the Non-Promoted Price ... by no less than a specified amount in consideration of a PM USA Retail Price Promotional Allowance paid directly by PM USA to Retailer."
Rationale for the floor (γ)	Elevai Labs, Unilateral Price Policy, 2022 (doc. 30)	"In order to successfully compete in the marketplace and maintain its premium brand and product image, Company has adopted this Unilateral Price Policy for the mutual benefit of its Distributors, customers and shareholders."

Table IA.A.2: Composition of the manufacturer-retailer contract archive

This table reports the composition of the hand-collected archive of manufacturer-retailer contracts and pricing documents described in Section 2 and this appendix. Each row reports a product market and the number of verified contracts, amendments, program agreements, and filed documents in that market. Panel A groups markets in which direct price discipline and retail program agreements concentrate. Panel B reports wholesale, selective dealer, and adjacent documents that contain authorization and monitoring terms but little or no direct retail price discipline. Documents in which a retailer is the principal are excluded. Automobile and marine dealer agreements are reported separately because dealer franchise statutes, rather than the brand owner's contract, govern their pricing terms. Representative clauses from the archive, organized by model primitive, appear in Table IA.A.1.

Product market	*Documents*	*Representative brand principals*
Panel A: Product markets in which direct price-discipline dyads and retail program agreements concentrate		
Tobacco retail programs and retail pricing agreements	15	Philip Morris USA; R.J. Reynolds retail programs and everyday-low-price retail pricing agreements (1996, 1999); Brown & Williamson retail, commissary, and merchandising forms
Mattress and bedding	9	Sealy, Tempur-Pedic, Simmons, Serta, Purple, and Tempur-Sealy to Mattress Firm
Consumer electronics, mobile accessories, and vape	13	Vuzix; ZAGG; Skullcandy; Samsung; Bidi Vapor; TiVo to Best Buy; Palm; National Concessions Group
Beauty, skincare, personal care, and fragrance	8	Elevai; Jupiter Wellness; Neoteric; Church & Dwight; Obagi; Guess? to Parlux
Food, confectionery, coffee, and beverages	8	JJC to Canterbury Coffee; Rocky Mountain Chocolate; Dewmar (LEAN, ×2); Skinny Water to Canada Dry Bottling; Sow Good
Supplements and health products	5	Thorne Research (Pattern, Emerson Ecologics); Emerald Health; ViSalus
Dealer and distributor networks (non-marine durables, specialty consumables)	6	MACH 3 Technologies; ProdecoTech dealer form; Grote Molen; Ciralight
Brand licensing and direct-to-consumer marketing	3	Docklight (Marley licenses, ×2); Marius to MangoRx
Franchise and brand-separation retail	2	Lululemon (Oqoqo); Sears (Kenmore, Craftsman, DieHard) to Sears Hometown
Bread and bakery direct-store-delivery	2	CK Sales / Flowers Foods route agreements
Lawn and garden	1	Monsanto to Scotts (Roundup, 1998; explicit resale price maintenance)
Jewelry and accessories	1	Charles & Colvard to Rio Grande
Intimate and adult products	1	Tenga to OneUp Innovations (binding resale-price floor)
Panel B: Wholesale-tier, selective-dealer, and adjacent documents		
Marine and recreational durable dealer networks	9	Boston Whaler, Sea Ray, and Azimut to MarineMax
Beverage and alcohol distribution (exclusive territory)	6	Monster to Coca-Cola; Craft Brewers Alliance, Redhook, and Hansen to Anheuser-Busch; BrewDog
Wholesale and reseller agreements of sample brand owners	5	Hewlett-Packard to SYNNEX; Virgin Mobile to Sprint; HzO to ZAGG; Vayyar; Sonic Solutions
Wholesale-tier grocery supply	1	C&S Wholesale Grocers to Southeastern Grocers / BI-LO (2018)
Total	**95**	14 of the 95 documents are coded as negative or qualified variation
Optional automobile and marine dealer add-ons	5	Ford, GM, Honda, Toyota, Ferretti / Bertram (dealer-franchise law regimes; robustness only)
Total including add-ons	**100**	

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: $\leq$30 words explaining why AI helps humans perform this task*
- *Automation_reason: $\leq$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 Price variance decomposition

Table IA.B.1 reports the full decomposition. Panel A shows that the systematic component of price levels resides primarily at the product and manufacturer levels. UPC × month effects account for 99.0% of the explained variation in the product specification. In the direct comparison across vertical and geographic sources, manufacturer identity accounts for 91.4%, compared with 7.6% for retailers and 1.0% for counties. The bilateral specification divides the explained variation almost equally between UPC × month effects and retailer × UPC relationships, while the four-factor decomposition also assigns most of the explained variation to UPC identity.

Panels B and C show related patterns for price dispersion and price changes. In the manufacturer specification, manufacturers account for 56.0% of the explained variation in cross-store dispersion, compared with 36.5% for retailers and 7.5% for counties. Panel B, column (4), further attributes 21.7% of explained dispersion to manufacturer × retailer effects and only 4.0% to manufacturer × county effects, after accounting for UPC × month variation. For price changes, manufacturer identity accounts for 97.1% of the explained variation. Column (4) isolates who synchronizes price changes over time. Manufacturer × month effects account for 72.8% of the explained variation, compared with 21.7% for retailer × month and 5.5% for county × month. The low R^2 values for dispersion and price changes indicate that much of the underlying variation is not explained by these systematic components. Within the systematic component, however, manufacturers are the largest source, and price changes are synchronized primarily upstream.

Table IA.B.1: 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 a four-factor extension of the interactive fixed-effects Shapley decomposition in Couts, Gonçalves, and Loudis (2026), splitting 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 × Month	99.9%		92.5%	
Manufacturer		97.1%		
Manufacturer × Month				72.8%
Retailer × Month				21.7%
Retailer	0.1%	2.2%	0.1%	
Retailer × UPC			7.4%	
County	0.0%	0.6%	0.0%	
County × Month				5.5%
R^2	0.081	0.002	0.083	0.016
Observations (billion)	3.75	3.75	3.75	3.70

IA.B.2 Price levels around brand acquisitions

Section 6 measures how an acquired product's prices vary across stores and how often they change, not where the level settles. Applying the same two-stage design to the log of the CPI-deflated average price fills that gap. Prices rise after acquisition, and the increase is confined to synergistic deals: 0.022 where the acquirer already sells in the module, against −0.007 and insignificant where it does not, with the difference carrying a deal-clustered wild bootstrap p-value of 0.046. Figure IA.B.1 plots the corresponding event-time paths.

Table IA.B.2: Price levels around brand acquisitions

This table examines price levels around the brand acquisitions of Table 4. Estimation follows the same two-stage design: the first stage fits UPC × retailer × stack fixed effects, deal-specific linear event-time trends, and month-of-year effects on pre-acquisition months only, and the second stage regresses the departure of the outcome from the extrapolated pre-acquisition path on a post-acquisition indicator, weighting each deal equally. The outcome is the log of the CPI-deflated average price, so the coefficient is the signed post-acquisition price change relative to the extrapolated pre-acquisition path. Column (1) pools all acquisitions; columns (2) and (3) restrict to *synergistic* and *non-synergistic* acquisitions, respectively. Standard errors clustered at the deal level 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 deal-clustered standard errors is denoted by ***, **, and *, respectively.

	Ln(Real price)		
	All	Synergistic	Non-synergistic
	(1)	(2)	(3)
Post acquisition	0.015^{**}	0.022^{***}	-0.007
	(0.006)	(0.006)	(0.012)
Wild bootstrap p-value	0.010	0.002	0.546
Syn. = non-syn., wild p-value			0.046
Fixed effects			
UPC × retailer × stack	✓	✓	✓
Deal × event-time trend	✓	✓	✓
Month of year	✓	✓	✓
Mean dep. var. (pre-period)	1.723	1.684	1.852
Deals	68	52	16
Observations	4,231,684	3,605,823	625,861

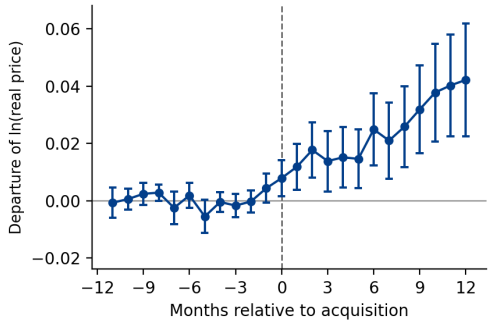
IA.B.3 Category-specific CPI deflation

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

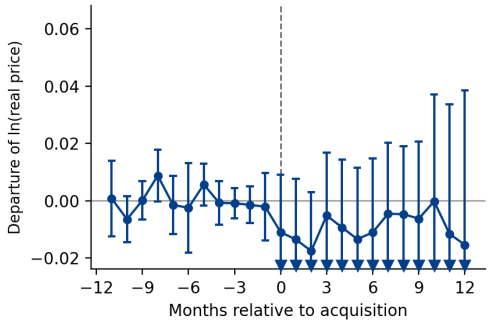
Figure IA.B.1: Price levels around brand acquisitions: event-time profiles

This figure plots monthly event-time profiles of the outcome in Table IA.B.2. Each point is the equal-weighted mean across deals of the departure of the log of the CPI-deflated average price from its extrapolated pre-acquisition path, computed from the first stage of the two-stage design (UPC×retailer×stack fixed effects, deal-specific linear event-time trends, and month-of-year effects, fit on pre-acquisition months only; UPC×retailer cells are weighted equally within each deal), with 90% pointwise confidence intervals computed from the cross-deal standard error. The first stage centers each deal's pre-acquisition residuals on zero, so pre-period levels and linear trends are absorbed by construction. The vertical dashed line marks the acquisition month ($k = 0$). Panel A plots *synergistic* acquisitions and Panel B *non-synergistic* acquisitions.

Panel A: Ln(Real price), synergistic



Panel B: Ln(Real price), non-synergistic



(1), and the wage-interaction coefficients remain positive in columns (2) and (3), though significant only at the 10 percent level.

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 InstructGPT 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 × Post	−0.030***	−0.153*	−0.158*
	(0.010)	(0.082)	(0.085)
Treated × Post × ln(Wage)		0.017*	0.017*
		(0.010)	(0.010)
Treated × ln(Wage)		0.011	0.010
		(0.010)	(0.010)
ln(Wage) × Post		−0.014	
		(0.009)	
ln(Wage)		−0.017	
		(0.013)	
Controls	✓	✓	✓
Fixed effects			
UPC	✓	✓	✓
County		✓	
Calendar month	✓	✓	
County × 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 event month and 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		$\Pr(\lvert\Delta P\rvert > 5\%)$		ln(Average price)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.024^{***}		0.056^{***}		-0.046^{***}	
	(0.005)		(0.016)		(0.012)	
Exposure × Post		1.324^{***}		2.919^{***}		-2.621^{***}
		(0.356)		(0.938)		(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

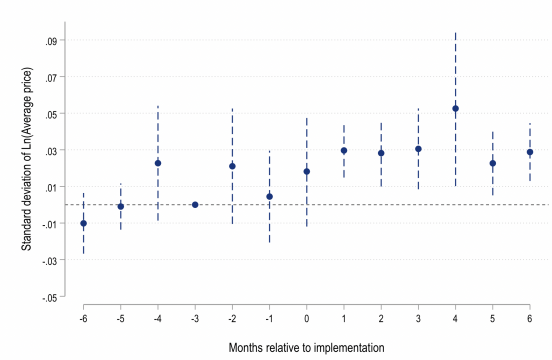
IA.B.5 Specifying the target and judge model

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.

Yin, Vu, and Persico (2026) show that this specification choice is consequential. Replicating a common exposure rubric with three frontier models on identical tasks, they find mean exposure differing by a factor of 3.6 and task-level agreement as low as 57 percent, with downstream difference-in-differences magnitudes varying by a factor of 2.4 and some county-level estimates changing sign. Their variation is across the model *doing* the scoring; the exercise here varies the model *being* scored, holding the scorer fixed. The two are complementary rather than substitutes, and we report the second because the object being proxied—what a brand owner's staff could actually accomplish with the technology available in early 2022—is a property of the target model. Table IA.B.5, Table IA.B.6, and Table IA.B.7 replicate the baseline price dispersion, price change, and price level analyses using the Davinci-01 exposure measure. The alternative measure yields the same three headline findings. Price dispersion increases, absolute repricing frequency rises, and average prices decline. The directional repricing estimates and the interaction with local wages generally have the same signs as in the baseline analysis, although they are less precisely estimated. The corresponding event studies (Figure IA.B.2, Figure IA.B.3, and Figure IA.B.4) show similar pre-event patterns and post-event dynamics.

Figure IA.B.2: 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.

Panel A: Standard deviation



Panel B: Interquartile range

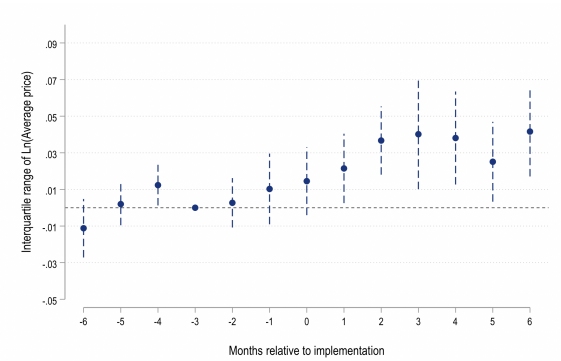
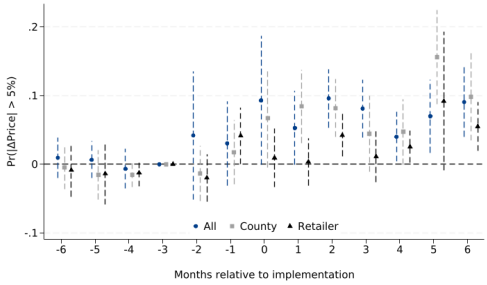
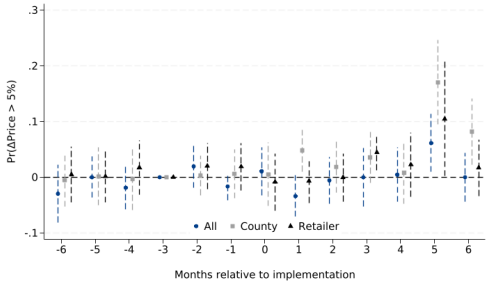


Figure IA.B.3: Price changes event study, using AI exposure based on Davinci-01
This figure plots the β_k coefficients and 90% confidence intervals from event-study regressions of the form: $\mathbb{P}(|\Delta \ln \bar{P}| > 5\%)_{ipt} = \alpha_{ip} + \alpha_t + \sum_{k \neq -3} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{ipt}$, 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.

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



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



Panel C: $\mathbb{P}(\Delta P < -5\%)$

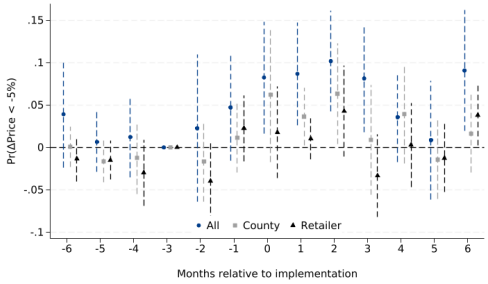


Figure IA.B.4: Price level event study, using AI exposure based on Davinci-01

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.

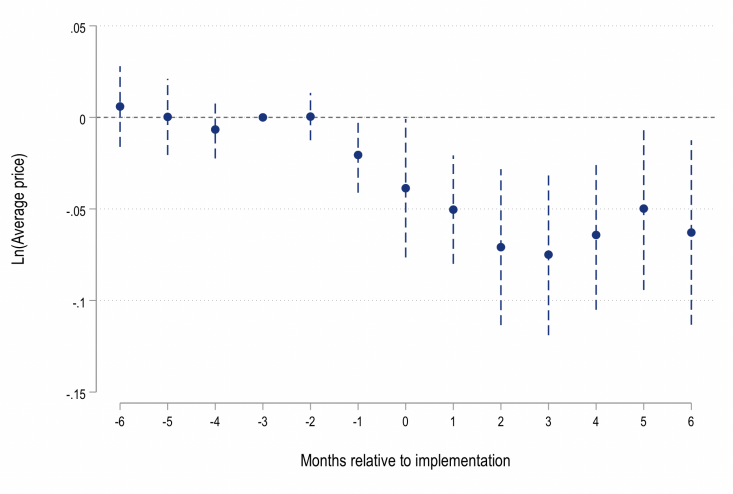


Table IA.B.5: Price dispersion, using AI exposure based on Davinci-01

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 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 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 stores	Within county	Within retailer	All stores	Within county	Within retailer
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.023***	0.010***	0.009**	0.022***	0.002**	0.007**
	(0.005)	(0.003)	(0.004)	(0.006)	(0.001)	(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 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 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×county and calendar month fixed effects; and the Retail specifications include UPC×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.

	$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$			$\mathbb{P}(\%\Delta P > 5\%)$			$\mathbb{P}(\%\Delta P < -5\%)$		
	All stores	Within county	Within retailer	All stores	Within county	Within retailer	All stores	Within county	Within retailer
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated × Post	0.061***	0.087***	0.035***	0.013	0.051***	0.013	0.049***	0.036	0.022*
	(0.015)	(0.028)	(0.013)	(0.010)	(0.012)	(0.011)	(0.013)	(0.024)	(0.012)
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Fixed effects									
UPC	✓			✓			✓		
UPC×County		✓			✓			✓	
UPC×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 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. 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***	−0.189*	−0.194*
	(0.011)	(0.102)	(0.104)
Treated × Post × ln(Wage)		0.019	0.020*
		(0.012)	(0.012)
Treated × ln(Wage)		0.006	0.006
		(0.011)	(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 limited capacity to work from home based on Dingel and Neiman (2020) and 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 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. 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. *Limited WFH capacity (DN)* is the complement of the Dingel and Neiman (2020) indicator for whether an occupation can be performed entirely from home. *Limited WFH capacity (ATUS)* is the complement of the share of American Time Use Survey respondents within an occupation who report that they can work from home. Both measures are aggregated to the firm level using employment shares measured before the event 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×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		$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$		ln(Average price)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.023**	0.027***	0.045**	0.060***	−0.051***	−0.047***
	(0.009)	(0.008)	(0.017)	(0.022)	(0.012)	(0.012)
Limited WFH capacity (DN) × Post	0.000		−0.141		−0.060	
	(0.044)		(0.085)		(0.067)	
Limited WFH capacity (ATUS) × Post		0.091		−0.047		−0.091
		(0.075)		(0.198)		(0.128)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC	✓	✓	✓	✓	✓	✓
UPC × 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		$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$		ln(Average price)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.011^{***}	0.011^{***}	0.064^{***}	0.064^{***}	-0.020^{***}	-0.020^{***}
	(0.003)	(0.003)	(0.017)	(0.017)	(0.005)	(0.005)
ln(County Covid cases)	−0.001		−0.006		-0.006^{**}	
	(0.003)		(0.008)		(0.003)	
ln(County Covid deaths)		0.001		0.001		0.003
		(0.003)		(0.004)		(0.002)
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC						
UPC × 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×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		$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$		ln(Average price)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.011***	0.011***	0.064***	0.064***	−0.020***	−0.020***
	(0.003)	(0.003)	(0.017)	(0.017)	(0.005)	(0.005)
County IR exposure × Post	−0.000		−0.001		−0.002***	
	(0.000)		(0.001)		(0.000)	
County IR exposure (ex-ante) × Post		−0.000		−0.001		−0.002***
		(0.000)		(0.001)		(0.000)
County IR exposure	−0.000		−0.000		0.001***	
	(0.000)		(0.001)		(0.000)	
Controls	✓	✓	✓	✓	✓	✓
Fixed effects						
UPC						
UPC × 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 Input cost exposure

The post-event window contains Russia's invasion of Ukraine in February 2022 and the associated commodity, energy, and freight shock. Treated firms are larger and hold a smaller share of their assets in property, plant, and equipment, and their input costs are a smaller share of revenue than those of control firms—gross margins average 0.50 against 0.38 in the pre-period. If that structural difference translated into a differential cost shock over the window, the relative price decline we estimate could reflect input costs rather than information.

We test this directly. Table IA.B.10 regresses quarterly Compustat measures of costs and margins on the interaction of treatment with the post-event indicator, with firm and calendar quarter fixed effects, over 2019Q1–2022Q3. Treated and control firms show no differential change in gross margins, in the ratio of cost of goods sold to sales, or in the levels of costs or sales. The margin estimate is the informative one: at 0.004 with a standard error of 0.010, a differential margin movement larger than roughly two percentage points can be ruled out. The cost and sales levels are estimated less precisely and we read them as supporting rather than decisive. Together with the product-group-by-month specifications reported above, which absorb any shock common to a product category in a month, this evidence indicates that the pricing results are not a manifestation of the 2022 input-cost environment.

Table IA.B.10: Input costs and margins around the event

This table reports firm-quarter difference-in-differences regressions of cost and margin measures from Compustat over 2019Q1–2022Q3. *Treated* equals one for firms with above-median AI exposure, measured in January 2020, and *Post* equals one from 2022Q1. Gross margin is sales less cost of goods sold, divided by sales; ratios are winsorized at the first and ninety-ninth percentiles. All specifications include firm and calendar quarter fixed effects, and standard errors are clustered at the firm level.

	Gross margin	COGS / sales	ln(COGS)	ln(Sales)
	(1)	(2)	(3)	(4)
Treated × Post	0.0036	−0.0036	−0.0510	−0.0347
	(0.0097)	(0.0097)	(0.0430)	(0.0348)
Firm fixed effects	✓	✓	✓	✓
Quarter fixed effects	✓	✓	✓	✓
Observations	1,409	1,409	1,439	1,439
Firms	103	103	104	104

IA.B.9 Where the dispersion increase comes from

The cross-retailer dispersion result establishes that chains carrying the same product move apart after the event, but it does not by itself say why. A brand that relaxes a binding floor and a brand whose common wholesale terms are passed through unevenly both generate rising dispersion across chains. The two differ in where in the price distribution the movement occurs. Releasing a floor is a one-sided intervention: it permits prices below a level that was previously enforced and does nothing to prices above it. Uneven pass-through of a common change is not one-sided.

Table IA.B.11 splits the dispersion increase around a reference price, the median chain price for each product and county over the second half of 2021, and decomposes the movement of the mean into the mass falling below that reference and the mass rising above it. The increase is entirely below the reference: semi-dispersion below rises by 0.029 on a pre-period mean of 0.110, while semi-dispersion above is flat and statistically indistinguishable from zero. The mass decomposition tells the same story: growing below-reference mass of -0.023 against an essentially unchanged above-reference component of -0.005 carries a mean deviation of -0.028, so roughly five-sixths of the movement in the mean comes from mass opening below the reference rather than from the distribution shifting as a whole. That mean deviation has the same sign and order of magnitude as the -0.043 relative price decline in Table 8, though the two are estimated on different samples and we do not read one as a decomposition of the other. In event time, the below-reference series shows no pre-trend and a sustained post-event increase, while the above-reference series is flat throughout.

Table IA.B.12 asks whether a floor is visible in the pre-period price distribution at all. It is: roughly twenty-three percent of chain-product-county prices sit within one percent of the lowest price observed for that product in that county before the event, and the share is the same for treated and control firms, so floors are a general feature of branded pricing rather than something specific to exposed firms. After the event, treated products open new mass strictly below the earlier floor, an increase of 2.2 percentage points that ramps in over the first months and attenuates by the sixth. The mass sitting exactly at the floor does not thin. Prices that move do not inch under the old floor; they jump well below it, and floor-sitting persists elsewhere. That pattern fits a floor released in some markets and not others, rather than a floor lowered uniformly for everyone.

Table IA.B.11: Where the dispersion increase falls in the price distribution

This table decomposes the increase in within-county cross-retailer price dispersion around a reference price, the median chain price for each product and county over July to December 2021. Columns (1) and (2) report semi-dispersion below and above the reference; columns (3) and (4) report the mass of deviations below and above it; and column (5) reports the mean deviation, which is their sum. The sample is products with at least two chains in a county. The sample includes 75 firms. All specifications include product and month fixed effects and the controls of Equation 30, with standard errors clustered by firm.

	Semi-dispersion		Mass decomposition		Mean
	Below ref.	Above ref.	Below ref.	Above ref.	deviation
	(1)	(2)	(3)	(4)	(5)
Treated × Post	0.029***	−0.005	−0.023***	−0.005	−0.028***
	(0.007)	(0.005)	(0.006)	(0.004)	(0.008)
$\bar{y}$ (pre-period)	0.110	0.067	-0.080	0.049	-0.032
Controls	✓	✓	✓	✓	✓
UPC fixed effects	✓	✓	✓	✓	✓
Month fixed effects	✓	✓	✓	✓	✓
Observations	44,289	44,289	44,289	44,289	44,289
Adjusted R^2	0.32	0.61	0.31	0.59	0.31

Table IA.B.12: Price floors and their release

This table reports the share of chain-product-county prices at and below a pre-period floor, defined as the lowest price observed for that product in that county before the event, using a monthly-median definition robust to temporary promotions in columns (1) and (2) and a weekly definition in column (3). A price is "at" the floor if it lies within one percent of it. The sample includes 75 firms in columns (1) and (2) and 74 firms in column (3). All specifications include product and month fixed effects and the controls of Equation 30, with standard errors clustered by firm.

	Monthly-median floor		Weekly floor
	At floor	Below floor	Below floor
	(1)	(2)	(3)
Treated × Post	0.004	0.022**	0.013*
	(0.015)	(0.010)	(0.008)
Pre-period share at floor, treated	0.223		0.114
Pre-period share at floor, control	0.230		0.128
Controls	✓	✓	✓
UPC fixed effects	✓	✓	✓
Month fixed effects	✓	✓	✓
Observations	46,432	46,432	43,652
Adjusted R^2	0.56	0.34	0.28

IA.B.10 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. We test whether similar differences between high- and low-exposure firms emerge before the InstructGPT API introduction. We re-estimate the baseline specifications using January 2019 and January 2021 as placebo event dates, 36 and 12 months before the actual event, respectively. Across all three outcomes and both placebo dates, the coefficients are small, statistically insignificant, and vary in sign (Table IA.B.13). These results make pre-existing differential changes between high- and low-exposure firms less likely to explain the baseline estimates.

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.5 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.13: Placebo tests using January 2019 and January 2021 as event dates

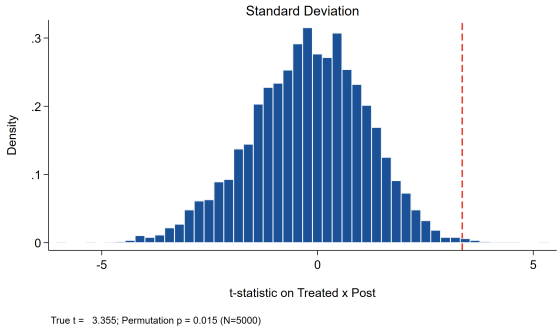
This table presents results from placebo tests that move the event date away from the actual InstructGPT API introduction in January 2022. Columns (1) to (3) use January 2019 (36 months before) and columns (4) to (6) use January 2021 (12 months before) as placebo event dates. Within each placebo date, 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	$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$	ln(Average price)	Standard deviation	$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$	ln(Average price)
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.010	−0.007	−0.008	0.002	0.015	−0.015
	(0.007)	(0.012)	(0.011)	(0.009)	(0.014)	(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

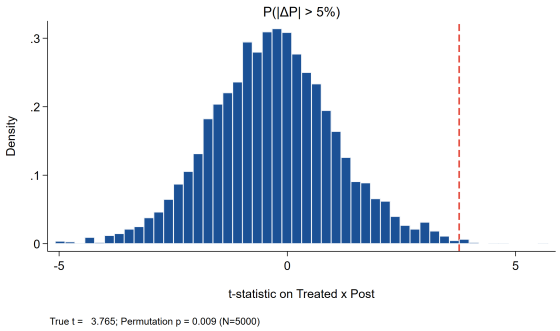
Figure IA.B.5: Placebo tests randomizing treatment status

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.

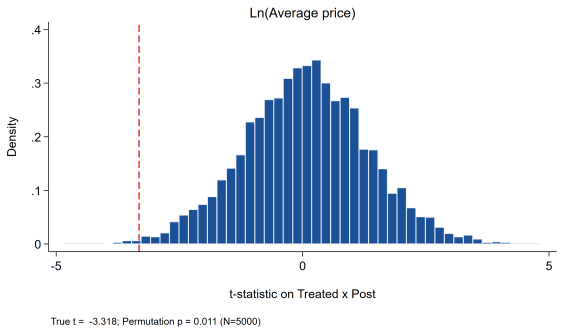
Panel A: Standard deviation



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



Panel C: ln(Average price)



IA.B.11 Alternative augmentation vs. automation decomposition

As a complement to the augmentation vs. automation decomposition in Table 12, 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.14 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.14: 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 in January 2022 and the six months following the InstructGPT API introduction. 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	$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$	ln(Average price)	Standard deviation	$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$	ln(Average price)
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.020***	0.055***	−0.042**	0.014	0.040*	−0.014
	(0.006)	(0.019)	(0.017)	(0.010)	(0.021)	(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.12 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.15 reports PPML regressions of AI keyword counts on the *Treated* × *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.6 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.6: 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$_{it}$ $= \alpha_i + \alpha_t + \sum_{k \neq -1} \beta_k \mathbb{1}[\text{Event time} = k] \times \text{Treated}_i + \epsilon_{it}$, 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.

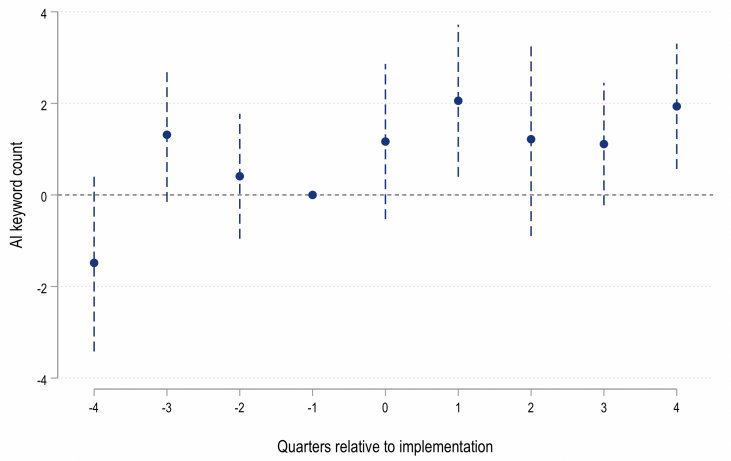


Table IA.B.15: AI keyword counts in earnings conference calls

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.

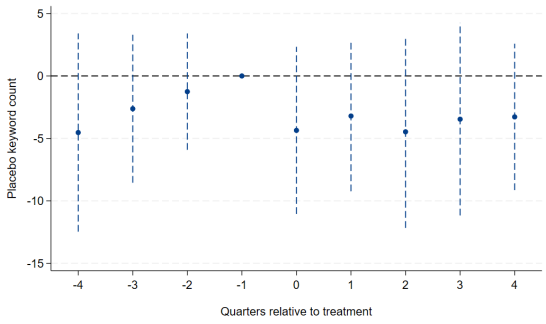
	AI keyword count			
	Total	Prepared remarks	Management Q&A	Analyst Q&A
	(1)	(2)	(3)	(4)
Treated × Post	1.800***	1.629**	2.265***	0.348
	(0.609)	(0.746)	(0.740)	(2.092)
Fixed effects				
Firm	✓	✓	✓	✓
Quarter	✓	✓	✓	✓
Observations	297	238	155	30

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.7 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.

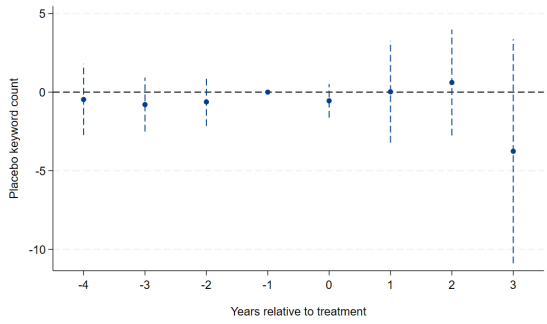
Figure IA.B.7: Placebo keyword event studies

This figure plots the β_k coefficients and 90% confidence intervals from event-study regressions of the form: 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 InstructGPT 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.

Panel A: Earnings conference calls



Panel B: 10-K filings



IA.B.13 Sample composition and inference

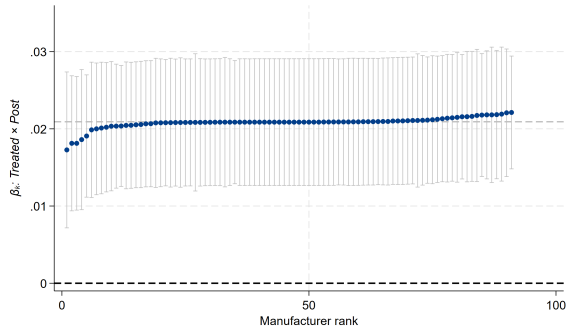
Leave-one-out by manufacturer and category. The matched sample contains 91 manufacturers and roughly 100 product categories, raising the concern that the estimates could be driven by a small number of firms or categories. We address this by re-estimating the baseline specification 91 times, dropping each manufacturer in turn, and a further 100 times, dropping each product category in turn. Figure IA.B.8 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.9 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.16 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.

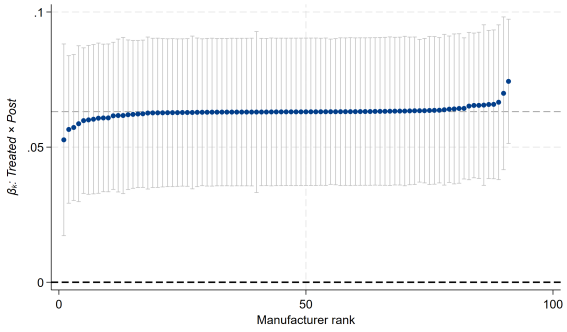
Figure IA.B.8: Leave-one-out manufacturer analysis

This figure presents the estimated difference-in-differences coefficient on *Treated* × *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.

Panel A: Standard deviation



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



Panel C: ln(Average price)

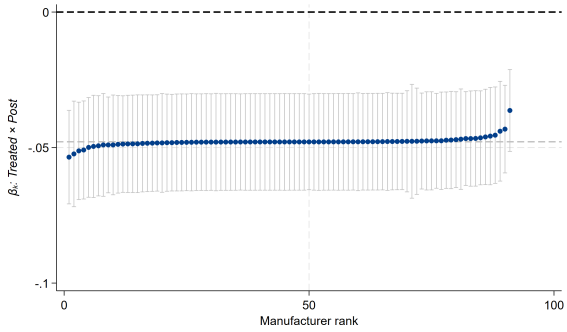
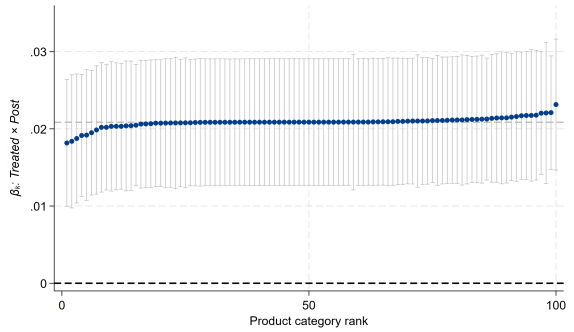


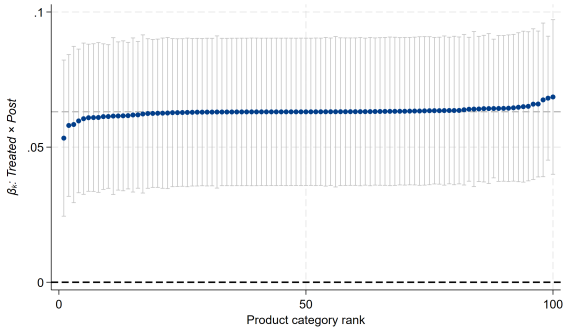
Figure IA.B.9: Leave-one-out product category analysis

This figure presents the estimated difference-in-differences coefficient on *Treated* × *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.

Panel A: Standard deviation



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



Panel C: ln(Average price)

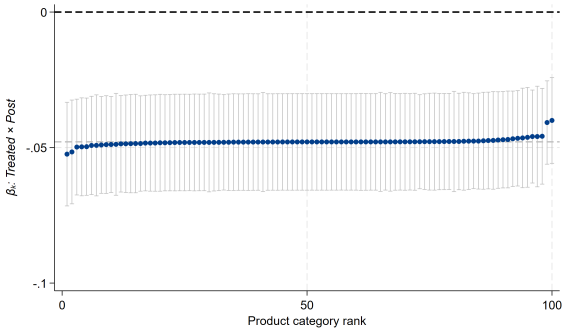


Table IA.B.16: Wild cluster bootstrap inference

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(*Average price*), respectively. Each column reports the coefficient on *Treated* × *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	$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$	ln(Average price)
	(1)	(2)	(3)
Treated × Post	0.025***	0.071***	−0.043***
	(0.006)	(0.018)	(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.14 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.17 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.17: 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(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	$\mathbb{P}(\lvert\%\Delta P\rvert > 5\%)$	ln(Average price)
	(1)	(2)	(3)
Tech-task intensity × Post	0.060	2.305**	0.210
	(0.439)	(1.117)	(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.15 Manufacturer-retailer contract disclosures

To complement the contract archive, we audit the 10-K filings of the sample brand owners for disclosures about their commercial relationships with retailers. The audit captures what firms report publicly rather than the full set of contract terms, which need not be publicly disclosed. We use each firm's last fiscal year ending before January 2021, so the evidence predates the treatment window.

We classify the disclosures into four categories covering trade promotions and allowances, pricing policies and guidance, category management and shelf arrangements, and rebates and performance payments. As a benchmark, we separately identify firms that mention wholesale or contract pricing but none of these four categories. Panel A of Table IA.B.18 reports the prevalence of each disclosure type.

Most brand owners disclose financial arrangements tied to retailer activity. Approximately 77% disclose at least one of the four categories, with rebates and performance payments reported by

73% and trade promotions by 44%. Nearly 29% report a dollar amount of consideration paid to customers. Although 28% mention wholesale or contract pricing, only 1.6% do so without also reporting one of the four categories. These disclosures show that brand owner payoffs are commonly tied to retailer sales volume and promotional activity.

Direct disclosures of pricing authority are rare. Only 1.1% of firms discuss pricing policies or guidance, and only 4.2% discuss category management or shelf arrangements. Just 4.2% disclose three or more categories. These low reporting rates should be interpreted in light of the disclosure regime described in Section 2. Minimum advertised price and resale price policies often operate through unilateral policy statements and private contracts that firms are not required to file. Panel B reproduces illustrative excerpts and shows the language captured by the audit, ranging from promotional allowances and performance rebates to pricing guidance and shelf arrangements.

Table IA.B.18: Manufacturer-retailer contract disclosures

This table summarizes manufacturer-retailer disclosures in the 10-K filings of 190 sample brand owners with machine-readable annual filings. We use each firm's last fiscal year ending before January 2021, so the audit predates the treatment window. We exclude 33 sample firms whose filings identify them as retailers rather than brand owners. Panel A reports the share of firms disclosing trade promotions and allowances, pricing policies and guidance, category management and shelf arrangements, and rebates and performance payments. "Wholesale only" identifies firms that mention wholesale or contract pricing but none of the four categories. "Broad contract disclosure" identifies firms disclosing three or more categories. "Consideration to customers disclosed" identifies firms reporting a dollar amount of consideration paid to customers. A large language model classifies the filing text and must return a verbatim supporting quotation for each positive classification. Every quotation is checked against the source filing. A second model independently rescored a stratified sample of 32 firms without observing the first model's classifications. Agreement across the four categories ranges from 94 to 97%, with Cohen's κ between 0.65 and 0.93. Panel B reproduces illustrative excerpts spanning all four categories. Ellipses denote omitted text, and brackets identify editorial additions.

Panel A: Contract disclosure rates

	N	Share (%)
Trade promotions & allowances	84	44.2
Pricing policies & guidance	2	1.1
Category management & shelf	8	4.2
Rebates & performance payments	138	72.6
Wholesale or contract pricing (control)	53	27.9
Any contract disclosure (categories 1–4)	147	77.4
Consideration to customers disclosed	55	28.9
Wholesale only (no categories 1–4)	3	1.6
Broad contract disclosure (three or more categories)	8	4.2
Mean contract categories (0–4)	1.22 (SD = 0.84)	
Median contract categories	1	
Firms in sample	190	

Appendix B: Proofs

Proof of Lemma 1. An uninformed brand assigns probability θ to an adverse shock, so the expected demand intercept is $\mathbb{E}[\alpha_{mj}] = \alpha - \theta\mu$.

If the floor binds every retailer, all charge a common price p and expected channel profit is $\Pi(p) = p(\alpha - \theta\mu - (\beta - \gamma)p)$, which is strictly concave because $\beta > \gamma$. The first-order condition gives

$$\underline{p}_m^0 = \frac{\alpha - \theta\mu}{2(\beta - \gamma)} = p_m^B(\alpha - \theta\mu) \tag{IA.B.1}$$

Panel B: Illustrative excerpts from 10-K filings

Firm	Category	Excerpt
Procter & Gamble	Trade promotion	"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 promotion	"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 promotion	"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 management	"[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 promotion	"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."

and

$$\Pi^0 = \frac{(\alpha - \theta\mu)^2}{4(\beta - \gamma)}. \tag{IA.B.2}$$

The uninformed price floor is the brand-preferred common price evaluated at the prior mean intercept. Floors above $\underline{p}_m^0$ also bind every retailer and are dominated by concavity.

Facing the price floor $\underline{p}_m^0$, retailer j's unconstrained best response is $(\alpha_{mj} + \gamma \underline{p}_m^0)/(2\beta)$, which is increasing in α_{mj}. The floor therefore binds every retailer in every state if and only if it binds

the highest intercept, $\alpha_{mj} = \alpha$:

$$\alpha \leq (2\beta - \gamma)\underline{p}_m^0, \tag{IA.B.3}$$

which corresponds to

$$\theta \leq \frac{1}{\mu}\frac{\gamma\alpha}{2\beta - \gamma} := \theta_{\text{binding}}. \tag{IA.B.4}$$

Because α is the highest possible demand intercept, any floor that binds retailers in markets without a demand shock also binds retailers in markets with a demand shock. Note that $\underline{p}_m^0$ maximizes total channel profits in any regime that features a binding price floor for all retailers in all demand states. Every remaining candidate, therefore, satisfies $\underline{p}_m < \frac{\alpha}{2\beta-\gamma}$ and is slack in markets without a demand shock ($z_m = 0$), resulting in the unconstrained-retailer equilibrium profit

$$\Pi^R(\alpha) = \frac{\beta\alpha^2}{(2\beta - \gamma)^2}. \tag{IA.B.5}$$

In markets with a demand shock, total channel profit under any common floor is bounded above by the profit the brand would obtain if it could directly set retailer-specific prices after observing the shock and each retailer's exposure. In that relaxed problem, the brand solves

$$\max_{p_L, p_H} \quad \tfrac{1}{2}p_L\left(\alpha - \mu - \sigma - \beta p_L + \gamma\frac{p_L + p_H}{2}\right) + \tfrac{1}{2}p_H\left(\alpha - \mu + \sigma - \beta p_H + \gamma\frac{p_L + p_H}{2}\right), \tag{IA.B.6}$$

where p_L is the price for high-exposure retailers and p_H is the price charged for low-exposure retailers. The optimal profit in this case is

$$\Pi^{Max} = \frac{(\alpha - \mu)^2}{4(\beta - \gamma)} + \frac{\sigma^2}{4\beta}. \tag{IA.B.7}$$

Therefore, the optimal price floor binds all retailers in all demand states whenever

$$\frac{(\alpha - \theta\mu)^2}{4(\beta - \gamma)} \geq (1 - \theta)\frac{\beta\alpha^2}{(2\beta - \gamma)^2} + \theta\left[\frac{(\alpha - \mu)^2}{4(\beta - \gamma)} + \frac{\sigma^2}{4\beta}\right]. \tag{IA.B.8}$$

Note that this condition holds strictly when $\theta = 0$. Because the condition is continuous in θ, let θ_{profit} denote the smallest $\theta \in (0, 1]$ at which the condition holds with equality, and set $\theta_{profit} = 1$

if it holds for all $\theta \in (0,1)$.

Define

$$\bar{\theta} := \min\{\theta_{binding}, \theta_{profit}\} > 0. \tag{IA.B.9}$$

Then, imposing $\theta < \bar{\theta}$ implies that the optimal uninformed price floor binds all retailers in all states. □

Proof of Lemma 2. This proof follows the logic of the Proof of Lemma 1 closely, substituting θ_{ms} for θ and noting that $\theta_{ms} < \theta$ implies that

$$\theta_{ms}\mu < \theta\mu \leq \frac{\gamma\alpha}{2\beta - \gamma}. \tag{IA.B.10}$$

□

Proof of Lemma 3. When the brand observes $s_m = 1$, it knows that an adverse demand event occurred in market m. Hence $z_m = 1$. Retailer j's specific negative demand shock is $\mu \pm \sigma$ with equal probability. Because, conditional on an adverse market demand shock, the retailer-specific exposures (d_{mj}) are i.i.d. and there is a continuum of retailers in each market, half of retailers have demand intercept $\alpha - \mu + \sigma$ and half have $\alpha - \mu - \sigma$.

First consider the no-delegation price floor p that binds for all retailers. In this case, total channel profit in market m is

$$\Pi^N(p) = p(\alpha - \mu - \beta p + \gamma p), \tag{IA.B.11}$$

which is strictly concave because $\beta > \gamma$. The first-order condition is

$$\alpha - \mu - 2(\beta - \gamma)p = 0. \tag{IA.B.12}$$

Thus the optimal floor that binds all retailers is

$$\underline{p}^N = \frac{\alpha - \mu}{2(\beta - \gamma)}. \tag{IA.B.13}$$

The associated total channel profit is

$$\Pi^{N*} = \frac{(\alpha - \mu)^2}{4(\beta - \gamma)}. \tag{IA.B.14}$$

Next, consider a partial-delegation price floor that binds only retailers with a high exposure to the demand shock (i.e., $d_{mj} = \mu + \sigma$). Let p denote the price floor charged by those retailers, and let h denote the price charged by retailers with a low exposure to the demand shock (i,e,, $d_{mj} = \mu - \sigma$). Because half of retailers are of each type, the average price in market m is

$$\bar{p} = \tfrac{1}{2}(p + h). \tag{IA.B.15}$$

The retailers with demand loss $d_{mj} = \mu - \sigma$ are unconstrained. Their first-order condition is

$$\alpha - \mu + \sigma - 2\beta h + \gamma \frac{p + h}{2} = 0. \tag{IA.B.16}$$

The best response function h is

$$h(p) = \frac{2(\alpha - \mu + \sigma) + \gamma p}{4\beta - \gamma}. \tag{IA.B.17}$$

Given $h(p)$, total channel profit under a partial-delegation price floor p is

$$\Pi^P(p) = \frac{1}{2} p \left(\alpha - \mu - \sigma - \beta p + \gamma \tfrac{p + h(p)}{2} \right) + \frac{1}{2} h(p) \left(\alpha - \mu + \sigma - \beta h(p) + \gamma \tfrac{p + h(p)}{2} \right). \tag{IA.B.18}$$

Differentiating with respect to p yields

$$\frac{d\Pi^P(p)}{dp} = \frac{8\beta^2(\alpha - \mu) - (2\beta p + \sigma)(8\beta^2 - 8\beta\gamma + \gamma^2)}{(4\beta - \gamma)^2}. \tag{IA.B.19}$$

The second derivative is

$$\frac{d^2\Pi^P(p)}{dp^2} = -\frac{2\beta(8\beta^2 - 8\beta\gamma + \gamma^2)}{(4\beta - \gamma)^2} < 0, \tag{IA.B.20}$$

implying the objective is strictly concave. Hence, the first-order condition identifies the optimal

partial-delegation price floor:

$$\underline{p}^P = \frac{4\beta(\alpha-\mu)}{8\beta^2-8\beta\gamma+\gamma^2} - \frac{\sigma}{2\beta}. \tag{IA.B.21}$$

The corresponding price charged by unconstrained retailers is

$$h(\underline{p}^P) = h^P = \frac{2(\alpha-\mu)(2\beta-\gamma)}{8\beta^2-8\beta\gamma+\gamma^2} + \frac{\sigma}{2\beta}. \tag{IA.B.22}$$

Substituting $\underline{p}^P$ and h^P into channel profit gives

$$\Pi^{P*} = \frac{2\beta(\alpha-\mu)^2}{8\beta^2-8\beta\gamma+\gamma^2} + \frac{\sigma^2}{4\beta}. \tag{IA.B.23}$$

Note that the brand prefers partial delegation if and only if $\Pi^{P*} > \Pi^{N*}$, which corresponds to

$$\frac{2\beta(\alpha-\mu)^2}{8\beta^2-8\beta\gamma+\gamma^2} + \frac{\sigma^2}{4\beta} > \frac{(\alpha-\mu)^2}{4(\beta-\gamma)}. \tag{IA.B.24}$$

Rearranging the inequality gives

$$\sigma > (\alpha-\mu)\gamma\sqrt{\frac{\beta}{(\beta-\gamma)(8\beta^2-8\beta\gamma+\gamma^2)}} := \bar{\sigma}. \tag{IA.B.25}$$

Therefore, if $\sigma \leq \bar{\sigma}$, the brand weakly prefers the no-delegation price floor ($\underline{p}^N$). If $\sigma > \bar{\sigma}$, the brand strictly prefers the partial-delegation price floor ($\underline{p}^P$).

We also verify the binding patterns of the two price floors. If $\sigma \leq \bar{\sigma}$, then

$$\sigma \leq \bar{\sigma} < \frac{(\alpha-\mu)\gamma}{2(\beta-\gamma)} \tag{IA.B.26}$$

because

$$8\beta^2-8\beta\gamma+\gamma^2 > 4\beta(\beta-\gamma). \tag{IA.B.27}$$

Thus,

$$\frac{\alpha-\mu}{2(\beta-\gamma)} \geq \frac{\alpha-\mu+\sigma+\gamma\frac{\alpha-\mu}{2(\beta-\gamma)}}{2\beta}, \tag{IA.B.28}$$

meaning that the no-delegation price floor binds for even the retailers with the highest realized demand intercept ($\alpha - \mu + \sigma$). Therefore, it binds for all retailers.

If $\sigma > \bar{\sigma}$. Under the partial-delegation floor $\underline{p}^P$, the price difference between unconstrained and constrained retailers is

$$h^P - p^P = \frac{\sigma}{\beta} - \frac{2(\alpha - \mu)\gamma}{8\beta^2 - 8\beta\gamma + \gamma^2}. \tag{IA.B.29}$$

Because

$$\bar{\sigma} > \frac{2\beta(\alpha - \mu)\gamma}{8\beta^2 - 8\beta\gamma + \gamma^2}, \tag{IA.B.30}$$

the condition $\sigma > \bar{\sigma}$ implies that $h^P > p^P$, meaning that retailers with a realized demand intercept of $\alpha - \mu + \sigma$ are not constrained. In addition, consider retailers with a realized demand intercept of $\alpha - \mu - \sigma$. Given the average price $(p^P + h^P)/2$, their unconstrained best response is

$$\frac{\alpha - \mu - \sigma + \gamma\frac{p^P + h^P}{2}}{2\beta}, \tag{IA.B.31}$$

which is strictly below the partial-delegation price floor:

$$\underline{p}^P - \frac{\alpha - \mu - \sigma + \gamma\frac{p^P + h^P}{2}}{2\beta} = \frac{2(\alpha - \mu)\gamma}{8\beta^2 - 8\beta\gamma + \gamma^2} > 0. \tag{IA.B.32}$$

□

Proof of Lemma 4. Let Π^0 denote the brand's expected channel profit in a market when it does not acquire information. By Lemma 1, the optimal uninformed price floor binds all retailers and is

$$\underline{p}^0_m = \frac{\alpha - \theta\mu}{2(\beta - \gamma)}. \tag{IA.B.33}$$

Hence

$$\Pi^0 = \frac{(\alpha - \theta\mu)^2}{4(\beta - \gamma)}. \tag{IA.B.34}$$

Now suppose the brand acquires information. With probability $1 - \theta\phi$, the brand observes $s_m = 0$. Conditional on this signal, the posterior probability that an adverse demand shock

occurred is

$$\theta_0 = \Pr(z_m = 1 \mid s_m = 0, a_m = 1) = \frac{\theta(1-\phi)}{1-\theta\phi}. \tag{IA.B.35}$$

By Lemma 2, the optimal price floor is

$$\underline{p}_m(0)^* = \frac{\alpha - \theta_0\mu}{2(\beta-\gamma)} \tag{IA.B.36}$$

and binds all retailers. Hence, total channel profit is

$$\Pi_0(\phi) = \frac{(\alpha-\theta_0\mu)^2}{4(\beta-\gamma)}. \tag{IA.B.37}$$

With probability $\theta\phi$, the brand observes $s_m = 1$ and verifies that an adverse demand shock occurred (i.e., $\theta_1 = \Pr(z_m = 1 \mid s_m = 1, a_m = 1) = 1$). By Lemma 3, total channel profit is

$$\Pi_1(\sigma) = \begin{cases} \frac{(\alpha-\mu)^2}{4(\beta-\gamma)}, & \sigma \le \bar{\sigma}, \\ \frac{2\beta(\alpha-\mu)^2}{8\beta^2-8\beta\gamma+\gamma^2} + \frac{\sigma^2}{4\beta}, & \sigma > \bar{\sigma}. \end{cases} \tag{IA.B.38}$$

Hence, the expected total channel profit given information acquisition is

$$\bar{\Pi}(\phi) = \theta\phi\Pi_1(\sigma) + (1-\theta\phi)\Pi_0(\phi). \tag{IA.B.39}$$

The brand acquires information if and only if $\Phi(\phi) = (\bar{\Pi}(\phi) - I) - \Pi^0 \ge 0$. Define $\bar{I}$ as

$$\bar{I} = \bar{\Pi}(1) - \Pi^0. \tag{IA.B.40}$$

Then, $I < \bar{I}$ implies that $\Phi(1) > 0$. Note that when $\phi = 0$, $\bar{\Pi}(0) = \Pi^0$, implying that $\Phi(0) = -I < 0$. Finally, note that Φ strictly increases with ϕ, implying that there exists a $\bar{\phi} \in (0,1)$ such that $\Phi \ge 0$ if and only if $\phi \ge \bar{\phi}$. □

Proof of Proposition 1. When $\phi_0 < \bar{\phi}$, Lemma 4 implies that the brand does not acquire information. By Lemma 1, the brand sets an uninformed floor that binds all retailers in all demand states, implying that both dimensions of price dispersion are 0: $\mathbb{E}_j[Var(p^0_{mj})] = \mathbb{E}_m[Var(p^0_{mj})] = 0$.

Now consider $\phi_1 = \phi_0 + \Delta_\phi \geq \bar{\phi}$. Lemma 4 implies that the brand acquires information. With probability $1 - \theta\phi_1$, the brand observes $s_m = 0$. In this case, let $\theta_0 = \frac{\theta(1-\phi_1)}{1-\theta\phi_1}$ denote the posterior probability that an adverse demand shock occurred. By Lemma 2, all retailers in such markets charge $p_m^0 = \frac{\alpha - \theta_0\mu}{2(\beta-\gamma)}$.

With probability $\theta\phi_1$, the brand observes $s_m = 1$ and verifies that an adverse market demand shock occurred ($z_m = 1$). There are two cases to consider: $\sigma \leq \bar{\sigma}$ and $\sigma > \bar{\sigma}$.

In the first case ($\sigma \leq \bar{\sigma}$), Lemma 3 implies that all retailers in such markets charge $p_m^1 = \frac{\alpha-\mu}{2(\beta-\gamma)}$. Since $\theta_0 < 1$, we have $p_m^0 > p_m^1$. Thus, prices differ across markets, so $\Delta_M > 0$.

In the second case ($\sigma > \bar{\sigma}$), Lemma 3 implies that verified adverse-shock markets feature partial delegation. Retailers with high exposure to the adverse shock ($d_{mi} = \mu + \sigma$) charge a lower price than retailers with low exposure ($d_{mk} = \mu - \sigma$): $p_{mi}^* < p_{mk}^*$. Therefore, verified adverse-shock markets exhibit within-market cross-retailer price dispersion: $\Delta_R > 0$. Moreover, because the two possible adverse-shock prices differ from each other, they cannot both equal p_m^0. Hence, prices also vary across markets for the same retailer: $\Delta_M > 0$. □

Proof of Proposition 2. When $\phi_0 < \bar{\phi}$, Lemma 4 implies that the brand does not acquire information. By Lemma 1, the brand sets an uninformed floor that binds all retailers in all demand states. Thus, all retailers charge the same price regardless of z_m and d_{mj}, implying that $\Xi_z = 0$ and $\Xi_d = 0$.

Now consider $\phi_1 = \phi_0 + \Delta_\phi \geq \bar{\phi}$. Lemma 4 implies that the brand acquires information. If $z_m = 0$, the brand necessarily observes $s_m = 0$. Let $\theta_0 = \frac{\theta(1-\phi_1)}{1-\theta\phi_1}$ denote the posterior probability of an adverse event after observing $s_m = 0$. By Lemma 2, all retailers in such markets charge

$$p_0 = \frac{\alpha - \theta_0\mu}{2(\beta - \gamma)}, \tag{IA.B.41}$$

implying that

$$\mathbb{E}[p_{mj}^* \mid z_m = 0] = p_0. \tag{IA.B.42}$$

If $z_m = 1$, the brand observes $s_m = 1$ with probability ϕ_1 and observes $s_m = 0$ with probability

$1-\phi_1$. Hence

$$\mathbb{E}[p^*_{mj} \mid z_m = 1] = (1-\phi_1)p_0 + \phi_1 \bar{p}_1, \tag{IA.B.43}$$

where $\bar{p}_1$ is the average price in a verified adverse-event market. Specifically, if $\sigma \leq \bar{\sigma}$, Lemma 3 implies that

$$\bar{p}_1 = \frac{\alpha - \mu}{2(\beta - \gamma)}. \tag{IA.B.44}$$

If $\sigma > \bar{\sigma}$, Lemma 3 implies

$$\bar{p}_1 = \frac{(\alpha - \mu)(4\beta - \gamma)}{8\beta^2 - 8\beta\gamma + \gamma^2}. \tag{IA.B.45}$$

In either case, $\bar{p}_1 < p_0$. For $\sigma \leq \bar{\sigma}$, this follows from $\theta_0 < 1$. For $\sigma > \bar{\sigma}$, note that

$$\frac{(\alpha - \mu)(4\beta - \gamma)}{8\beta^2 - 8\beta\gamma + \gamma^2} < \frac{\alpha - \mu}{2(\beta - \gamma)} < \frac{\alpha - \theta_0\mu}{2(\beta - \gamma)} = p_0. \tag{IA.B.46}$$

Therefore,

$$\Xi_z = p_0 - \big((1-\phi_1)p_0 + \phi_1\bar{p}_1\big) = \phi_1(p_0 - \bar{p}_1) > 0. \tag{IA.B.47}$$

Finally, consider retailer-specific exposure. If $\sigma \leq \bar{\sigma}$, then even after a verified adverse event, all retailers charge the same price, so retailer-specific exposure does not affect prices. If $\sigma > \bar{\sigma}$, then verified adverse-event markets feature partial delegation. Retailers with low exposure, $d_{mj} = \mu - \sigma$, charge h^P, while retailers with high exposure, $d_{mj} = \mu + \sigma$, charge p^P, with $h^P > p^P$. Because verification occurs with probability ϕ_1 conditional on $z_m = 1$, we have

$$\Xi_d = \phi_1(h^P - p^P) > 0. \tag{IA.B.48}$$

□

Proof of Proposition 3. When $\phi_0 < \bar{\phi}$, Lemma 4 implies that the brand does not acquire information. By Lemma 1, the brand sets the uninformed floor

$$\underline{p}^0_m = \frac{\alpha - \theta\mu}{2(\beta - \gamma)}, \tag{IA.B.49}$$

which binds all retailers in all demand states. Hence, the average price level is

$$\bar{p}^0 = \frac{\alpha - \theta\mu}{2(\beta - \gamma)}. \tag{IA.B.50}$$

Now consider $\phi_1 = \phi_0 + \Delta_\phi \geq \bar{\phi}$. Lemma 4 implies that the brand acquires information. With probability $1 - \theta\phi_1$, the brand observes $s_m = 0$. Let $\theta_0 = \frac{\theta(1-\phi_1)}{1-\theta\phi_1}$ denote the posterior probability that an adverse event occurred after observing $s_m = 0$. By Lemma 2, all retailers in such markets charge

$$p_{m0} = \frac{\alpha - \theta_0\mu}{2(\beta - \gamma)}. \tag{IA.B.51}$$

With probability $\theta\phi_1$, the brand observes $s_m = 1$ and verifies that an adverse demand event occurred. Because $\sigma > \bar{\sigma}$, Lemma 3 implies that the brand uses a partial-delegation floor. The average price in verified adverse-event markets is

$$p_{m1} = \frac{1}{2}\left[\frac{4\beta(\alpha-\mu)}{8\beta^2 - 8\beta\gamma + \gamma^2} - \frac{\sigma}{2\beta}\right] + \frac{1}{2}\left[\frac{2(\alpha-\mu)(2\beta-\gamma)}{8\beta^2 - 8\beta\gamma + \gamma^2} + \frac{\sigma}{2\beta}\right] = \frac{(\alpha-\mu)(4\beta-\gamma)}{8\beta^2 - 8\beta\gamma + \gamma^2}. \tag{IA.B.52}$$

Therefore, the average price level after information acquisition is

$$\bar{p}^1 = (1 - \theta\phi_1)p_{m0} + \theta\phi_1 p_{m1}. \tag{IA.B.53}$$

Substituting in the expression for p_{m0} and p_{m1} yields

$$\bar{p}^1 = (1 - \theta\phi_1)\frac{\alpha - \theta_0\mu}{2(\beta - \gamma)} + \theta\phi_1\frac{(\alpha-\mu)(4\beta-\gamma)}{8\beta^2 - 8\beta\gamma + \gamma^2}. \tag{IA.B.54}$$

Substituting in the expression for θ_0 yields

$$\bar{p}^1 = \frac{\alpha - \theta\mu}{2(\beta - \gamma)} + \theta\phi_1\left[\frac{(\alpha-\mu)(4\beta-\gamma)}{8\beta^2 - 8\beta\gamma + \gamma^2} - \frac{\alpha-\mu}{2(\beta-\gamma)}\right]. \tag{IA.B.55}$$

The first term is exactly the no-information average price $\bar{p}^0$. Hence,

$$\bar{p}^1 - \bar{p}^0 = \theta\phi_1(\alpha-\mu)\left[\frac{4\beta-\gamma}{8\beta^2 - 8\beta\gamma + \gamma^2} - \frac{1}{2(\beta-\gamma)}\right] < 0 \tag{IA.B.56}$$

because $2(\beta-\gamma)(4\beta-\gamma) < 8\beta^2 - 8\beta\gamma + \gamma^2$. □

Proof of Proposition 4. Fix a market m and retailer j, and note that $\theta\phi$ is the per-period probability that the brand's signal is affirmative ($s_{mt} = 1$): $\Pr(s_{mt} = 1) = \theta\phi$ and $\Pr(s_{mt} = 0) = 1 - \theta\phi$. Because z_{mt} is drawn independently across periods and the acquisition decision is constant over time, $s_{mt} \overset{i.i.d.}{\sim} Bernouilli(\theta\phi)$, and the retailer exposures d_{mjt} are drawn independently across periods conditional on $z_{mt} = 1$.

If $a_{mt} = 0$, Lemma 1 gives the floor $\underline{p}^0 = \frac{\alpha-\theta\mu}{2(\beta-\gamma)}$, which depends only on model primitives and binds every retailer in every demand state. Hence, $p^*_{mjt} = \underline{p}^0$ for all t and $R = 0$.

If $a_{mt} = 1$ and $\sigma \leq \bar{\sigma}$, by Lemmas 2 and 3, the optimal informed price floor binds all retailers under either signal realization, so $p^*_{mjt} = \underline{p}_m(s_{mt})$, with

$$\underline{p}_m(0) = \frac{\alpha - \theta_0\mu}{2(\beta-\gamma)} \tag{IA.B.57}$$

and

$$\underline{p}_m(1) = \frac{\alpha - \mu}{2(\beta-\gamma)}. \tag{IA.B.58}$$

Because $\theta_0 = \frac{\theta(1-\phi)}{1-\theta\phi} < 1$, we have $\underline{p}_m(0) > \underline{p}_m(1)$, implying that prices change across periods if and only if the signal changes. By the independence across periods,

$$R = \Pr(s_{mt} \neq s_{m,t-1}) = 2\theta\phi(1-\theta\phi) > 0. \tag{IA.B.59}$$

If $a_{mt} = 1$ and $\sigma > \bar{\sigma}$, the optimal floor binds all retailers when $s_{mt} = 0$: $p^*_{mjt} = \underline{p}_m(0)$. When $s_{mt} = 1$, Lemma 3 gives

$$p^*_{mjt}(d_{mjt}) = \begin{cases} \frac{4\beta(\alpha-\mu)}{8\beta^2-8\beta\gamma+\gamma^2} - \frac{\sigma}{2\beta}, & d_{mjt} = \mu + \sigma, \\ \frac{2(\alpha-\mu)(2\beta-\gamma)}{8\beta^2-8\beta\gamma+\gamma^2} + \frac{\sigma}{2\beta}, & d_{mjt} = \mu - \sigma, \end{cases} \tag{IA.B.60}$$

each occurring with probability $\frac{1}{2}$, independently across periods. Hence, the probability that prices

change from period to period is

$$R = 1 - (1 - \theta\phi)^2 - \tfrac{1}{2}(\theta\phi)^2 = 2\theta\phi(1 - \theta\phi) + \tfrac{1}{2}(\theta\phi)^2 > 0. \tag{IA.B.61}$$

Finally, by Lemma 4 the brand does not acquire information for $\phi_0 < \bar{\phi}$ and does acquire for $\phi_0 + \Delta_\phi \geq \bar{\phi}$. Repricing frequency therefore rises from $R = 0$ to a strictly positive value.

□