Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Menu Costs and the Bullwhip Effect: Supply Chain Implications of Dynamic Pricing

Robert L. Bray, Ioannis Stamatopoulos

To cite this article:
Robert L. Bray, Ioannis Stamatopoulos (2022) Menu Costs and the Bullwhip Effect: Supply Chain Implications of Dynamic Pricing. Operations Research 70(2):748-765. https://doi.org/10.1287/opre.2021.2175

informs
http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 70, No. 2, March–April 2022, pp. 748–765
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Contextual Areas

Menu Costs and the Bullwhip Effect: Supply Chain Implications of Dynamic Pricing

Robert L. Bray,[a] Ioannis Stamatopoulos[b]

[a] Operations Department, Northwestern University, Evanston, Illinois 60208; [b] Information, Risk, and Operations Management Department, The University of Texas at Austin, Austin, Texas 78705

Contact: r-bray@kellogg.northwestern.edu, https://orcid.org/0000-0003-2773-0663 (RLB); yannis.stamos@mccombs.utexas.edu, https://orcid.org/0000-0002-3651-0788 (IS)

Received: December 29, 2019
Revised: August 25, 2020; May 21, 2021
Accepted: April 21, 2021
Published Online in Articles in Advance: December 23, 2021

Area of Review: Operations and Supply Chains

https://doi.org/10.1287/opre.2021.2175

Abstract: We study the supply chain implications of dynamic pricing. Specifically, we estimate how reducing menu costs—the operational burden of adjusting prices—would affect supply chain volatility. Fitting a structural econometric model to data from a large Chinese supermarket chain, we estimate that removing menu costs would (i) reduce the mean shipment coefficient of variation by 7.2 percentage points (pp), (ii) reduce the mean sales coefficient of variation by 4.3 pp, and (iii) reduce the mean bullwhip effect by 2.9 pp. These stabilizing changes are almost entirely attributable to an increase in the mean sales rate.

Supplemental Material: The online appendix and replication files are available at https://doi.org/10.1287/opre.2021.2175.

Keywords: dynamic pricing • menu costs • supply chain • bullwhip effect • empirical operations management • structural estimation

1. Introduction

The fields of revenue management and supply chain management have relatively little overlap, given their respective sizes.[1] Indeed, prices are stable at the mouth of most supply chains, as sophisticated revenue management techniques are not practical in most retail settings. This is because the operational costs of updating retail prices—known as "menu costs" after restaurateurs' menu-printing fees—overwhelm the benefits of dynamic pricing in most shops. However, new technologies, such as electronic shelf labels, which obviate the hassle of changing price stickers (Harrison et al. 2018, Barba 2019, Stamatopoulos et al. 2021), are mitigating these menu costs. And this means that supply chain managers must soon contend with dynamic pricing.

What will be the supply chain implications of dynamic pricing at retail outlets? To help answer this question, we create a structural econometric model of a large Chinese supermarket chain and use it to simulate the effect of removing menu costs on the chain's product flows. We estimate that removing menu costs would significantly stabilize the chain, decreasing its downstream sales volatility, its upstream shipment volatility, and its bullwhip effect, which is the difference of the two (i.e., the downstream-to-upstream volatility amplification). Specifically, we calculate that removing menu costs would increase the chain's sales rate, which—given that the inventory's standard deviation grows sublinearly in the sales rate—would make product flows relatively less lumpy. In short, we predict that dynamic pricing will help supply chains "outgrow" the bullwhip effect.

2. Related Literature

Chan et al. (2004, p. 2) asserted that pricing flexibility can make supply chains more efficient. For example, in their chapter on price and inventory management, they wrote,

> This *integration* of pricing, production and distribution decisions in retail or manufacturing environments is still in its early stages in many companies, but it has the potential to radically improve supply chain efficiencies in much the same way as revenue management has changed airline, hotel and car rental companies. Thus we are motivated in this paper to consider strategies which integrate pricing decisions with other aspects of the supply chain, particularly those related to production or inventory.

But this chapter provides no hard evidence that dynamic pricing increases or decreases supply chain efficiency. And neither do the surveys of Bitran and Caldentey (2003), Elmaghraby and Keskinocak (2003), Chen and Simchi-Levi (2012), Chen and Chen (2015), or Van den Boer (2015).

There is a good reason that the Chan et al. (2004) conjecture remains unverified: endogenizing prices requires explicitly modeling menu costs—the operational burden of changing prices (Netessine 2006)—which is difficult to do, both theoretically and empirically. The pure theory route is stymied because models that have both a fixed shipment cost and a fixed menu cost are largely intractable (Chen and Simchi-Levi 2012, p. 815). Accordingly, only a few theoretical works study inventory and revenue management with menu costs, and none of these study the upstream implications: Çelik et al. (2009) provided a heuristic pricing policy for perishable inventory under stochastic demand; Chen et al. (2011) provided two joint ordering and pricing policies, one with inventory carryovers but without menu costs and the other with menu costs but without inventory carryovers; and Chen and Hu (2012) provided a joint ordering and pricing policy for nonperishable inventory under deterministic demand.

The pure empirical route is also stymied by a dearth of exogenous menu-cost shocks. Accordingly, only a few empirical works study the operational effects of menu costs, and none of these study the upstream implications: Levy et al. (1997) and Owen and Trzepacz (2002) compared chains that are subject to a pricing law mandating item-specific price tags with chains that are not subject to this law to show that menu costs cause price stickiness; Anderson et al. (2015) compared single-variant products with multiple-variant products at a retailer that requires uniform pricing across variants to demonstrate the same; and Stamatopoulos et al. (2021) compare stores that installed electronic shelf labels with stores that did not to show that menu costs reduce profits.

To avoid the pitfalls of the pure theory and pure empirical routes, we use a blended approach that combines the strengths of both math and data. We develop and estimate a structural econometric ordering-and-pricing model and use it to compare the optimal supply chain policies both with and without menu costs. Our counterfactual study is numerical—and thus does not require a closed-form solution—but also hypothetical—and thus does not require an explicit menu cost experiment.

Our counterfactual is the mirror image of Aguirregabiria's (1999) counterfactual. Aguirregabiria estimated that shipping costs drive inventory volatility—stock levels rising and falling in saw-toothed fashion—which in turn drives price volatility—prices rising and falling countercyclically with inventory levels. Consequently, he estimated that eliminating the fixed shipping cost would steady the pricing process. In effect, we perform the inverse analysis, showing that eliminating the fixed price-change cost would steady the ordering process.

3. Theoretical Models

Removing menu costs can accentuate both demand-based pricing—setting prices to correlate positively with demand levels—and inventory-based pricing—setting prices to correlate negatively with inventory levels. We study these two pricing mechanisms with two models. Unfortunately, their combined effect on supply chain stability is unclear because one of the mechanisms appears to increase the bullwhip effect, whereas the other appears to decrease it. This theoretical limitation motivates the structural econometric approach we pursue in Section 6.

3.1. Demand-Based Pricing

Our first model illustrates that removing menu costs can stabilize the supply chain by undermining Lee et al.'s (1997) first bullwhip driver: demand signal processing. Specifically, the model suggests that the demand-based pricing enabled by eradicating menu costs would smooth the spikes and slumps in demand, which would lead to even more pronounced smoothing upstream, for a net reduction in the bullwhip effect.

Our model is the Lee et al. (1997) demand signal processing (DSP) model but with endogenous prices. The model describes a store that orders and prices a representative product, subject to stochastic and serially correlated demand, holding costs, backlogging costs, and free returns. We define and analyze our DSP model in Appendix A. We use it to anticipate how volatility along the supply chain would change as menu costs decrease from infinity to zero. We summarize our results in the following three propositions:[2]

Proposition 1. *Under the DSP model, the volatility of sales is smaller with dynamic pricing than with static pricing.*

Proposition 2. *Under the DSP model, the volatility of shipments is smaller with dynamic pricing than with static pricing.*

Proposition 3 *Under the DSP model, the bullwhip effect is smaller with dynamic pricing than with static pricing.*

Under a constant price, the store responds to positive demand signals by increasing order-up-to levels. But with a variable price, the store has another recourse: temper positive demand signals by increasing the price. And exploiting this second degree of freedom has two effects. First, it smooths sales because the profit-maximizing price always biases demand toward its mean. Second, it smooths shipments even more because (i) the downstream smoothing mechanically translates to upstream smoothing and (ii) the price reaction substitutes for the order reaction.

3.2. Inventory-Based Pricing

Our second model illustrates that removing menu costs can destabilize the supply chain by accentuating the Lee et al. (1997) third bullwhip driver: order batching. Specifically, the model suggests that the inventory-based pricing enabled by eradicating menu costs would decrease the cost of holding inventory relative to the cost of shipping it, which would compel stores to increase their shipping batch sizes, which, in turn, would increase their shipping volatilities.

Rather than develop our own model, we use the Stamatopoulos et al. (2019) generalized economic order quantity (EOQ) model, which is well-suited to our purposes. The model describes a store that orders and prices a representative product, subject to deterministic demand, holding costs, and shipping costs. We define and analyze this model in Appendix B. We use it to anticipate how the bullwhip effect changes as menu costs decrease from infinity to zero. We summarize our results with the following proposition:[3]

Proposition 4. *Under the EOQ model, the bullwhip effect is larger with dynamic pricing than with static pricing.*

Under a constant price, the inventory level follows a straight line down to zero. But with a variable price, the inventory level follows a convex curve, which decreases quickly at the beginning and slowly at the end. This convex curve has a smaller integral than the straight line for any given order quantity, which means that it comes with a smaller inventory-holding burden than the straight line. Accordingly, the economic order quantity increases as the menu cost decreases, which increases the bullwhip effect.

3.3. Reconciliation

Neither model tells the full story. The EOQ model has dynamics but not demand uncertainty, whereas the DSP model has demand uncertainty but not dynamics (because it has free returns). And their results are contradictory: the DSP model suggests that demand-based pricing leads to demand smoothing—stores tempering demand spikes with high prices—which mitigates demand signal processing, which reduces the bullwhip effect (Lee et al. 1997), whereas the EOQ model suggests that inventory-based pricing leads to inventory smoothing—stores burning off inventory gluts with low prices—which leads to lower inventory holding costs, which leads to larger shipping batch sizes, which increases the bullwhip effect (Lee et al. 1997).

Because no closed-form analytical model can cleanly capture both demand smoothing and inventory smoothing, we build a more general, structural econometric model in Section 6. Our structural model nests both the DSP and EOQ models as special cases. But before getting to that, let us discuss the data.

4. Empirical Setting and Data

4.1. Setting

We study 78 hypermarkets owned and operated by the sixth-largest supermarket chain in China. These hypermarkets are larger than the chain's normal grocery stores (which the chain generally franchises). They are located in Shanghai, Anhui, and Jiangsu and are fulfilled by a distribution center (DC) in Shanghai. Each store has one manager, who operates with autonomy and whose compensation depends on store sales and operating costs (including inventory costs).

Each store manager oversees a team of category managers, who make inventory and pricing decisions daily. These category managers decide whether and how much to order with the assistance of an IT system. Overall, the order placement and fulfillment process is fairly routine (see Bray et al. 2019 for a comprehensive description), so we focus on the price-updating process.

Local managers choose prices after consulting a global "guiding price," which the company headquarters sets in response to the supplier's price, the competitive environment, the target profit margin, and several other factors. All stores receive the same guiding price; but local managers can deviate from this guiding price, and they do so regularly. For example, on average, an item has 3.4 different prices across the 78 stores.[4]

Changing prices is an involved process: First, a category manager uses a computer terminal located in the store's back room to "apply" for a price update in the price maintenance system. Second, a price manager reviews and approves or denies the application. Third, if the price manager approves the application, the category manager prints the new price tag on a backroom printer. Fourth, a store clerk physically updates the price tag on the store floor, which, across all products, takes about one labor hour per day per store.

4.2. Data

We observe sales, wholesale prices, retail prices, store-to-DC orders, and start-of-the-day inventories from April 1, 2011 to December 31, 2014. We observe demand if and only if inventory is not stocked out. To avoid a censoring bias, we treat this variable as "unobserved" when the inventory level is below the lowest decile—that is, when there is a credible threat of stocking out; otherwise, we treat demand as observed and equal to sales.

Our original sample comprises 78 stores, 91 items, and 2,598 products. An item is a size-specific SKU—for example, item 103060 is a 700-g package of "Baimao laundry detergent for colorful clothes"—whereas a product is a store-item combination—for example, product 103060_1016 denotes item 103060 at store 1016. To facilitate our estimation procedure, we remove products with fewer than 200 daily observations, fewer than 20 orders, fewer than 25 days with an inventory change, or more than 20% of shipments arriving later than one day (because our model assumes next-day delivery).[5]

Because no individual product has enough data to support a full structural estimation, we aggregate similar products into clusters with the k-means algorithm.

Table 1. Sample-Composition Summary Statistics

	Stores	Items	Products	Clusters
Detergent	73	22	762	89
Drinks	72	28	748	87
Oil/vinegar	55	6	116	15
Oral care	52	10	183	23
Shampoo	66	11	318	38
Tissues	55	4	100	13
Toilet paper	68	10	371	43
Total	78	91	2,598	308

Notes. This table reports the number of distinct stores, items, products, and product clusters in our final sample. A product is a store-item combination, and a product cluster is a collection of products comprising one item and multiple stores.

We pool the product-level data by item rather than by store because the sample is more similar across stores than across items. Specifically, we divide each item's data into 10 product clusters, where each cluster represents the item at a collection of stores. For example, we divide item 13204's data into product clusters 13204_A, 13204_B, 13204_C, 13204_D, 13204_E, 13204_F, 13204_G, 13204_H, 13204_I, and 13204_J, where cluster 13204_A represents item 13204 at stores 1005, 1021, 1023, 1037, 3772, 3880, 4125, 4517, 7906, and 8442; cluster 13204_B represents item 13204 at stores 1003, 1009, 1053, 3775, 7600, and 8415, and so on. Each item has its own store grouping: for example, clusters 13204_A and 108004_A correspond to different sets of stores. We create our item-store groupings based on similarity in average inventories, average order quantities, and average demands. For example, the average inventory is 33.5 units in cluster 13204_A and 74.4 units in cluster 13204_G. We remove product clusters with fewer than 100 retail price changes, fewer than 200 orders, or fewer than two distinct order quantities with at least 10 corresponding orders (Bray et al. 2019 proved that our model is empirically identified when there are at least two nonzero order quantities).[6]

Our final sample comprises 308 product clusters spanning 91 items, 78 stores, and 2,598 products across seven categories: detergents, drinks, oil/vinegar, oral care, shampoo, tissues, and toilet paper (Table 1). Most of our items are stable sundries, but a few of the beverages are perishable.

5. Reduced Form Evidence

The bullwhip effect, menu costs, demand smoothing, and inventory smoothing are all apparent in our sample.

5.1. Bullwhip Effect

We measure the bullwhip effect with the DC-to-store shipments coefficient of variation minus the store-to-customer sales coefficient of variation. This measure is positive when upstream volatility exceeds downstream volatility, that is, when the supply chain exhibits the bullwhip effect.

We report our measure both at the product level and the item level. The product-level measure uses the shipments and sales of a given item at a specific store; the item-level measure uses the shipments and sales of a given item aggregated across all stores. Roughly speaking, the product-level metric measures the bullwhip from the stores' perspective (store managers care about the flow through their individual facilities), whereas the item-level metric measures the bullwhip from the DC's perspective (the DC manager cares about the total flow through all facilities).

We find a robust bullwhip effect at both levels of analysis: the shipment coefficient of variation exceeds the sales coefficient of variation in 91 out of 91 items and 2,596 out of 2,598 products (Figure 1). Indeed, the upstream volatilities are generally more than three times the downstream volatilities: at the product level, the median shipment coefficient of variation is 5.60, whereas the median sales coefficient of variation is only 1.37; at the item level, the median shipment coefficient of variation is 2.57, whereas the median sales coefficient of variation is only 0.84.

5.2. Menu Costs

The prices in our sample indicate menu costs. For example, the average price reoptimization changes prices by ¥3.50, or 16% (Table 2). And the "lumpiness"

Table 2. Price-Change Summary Statistics

	Magnitude	Duration
Detergent	0.17	26.44
Drinks	0.12	25.87
Oil/vinegar	0.16	45.34
Oral care	0.16	33.04
Shampoo	0.27	34.98
Tissues	0.14	46.75
Toilet paper	0.12	31.68
Total	0.16	30.23

Notes. This table reports the average price change magnitude, measured as a fraction of the prior price, and the average price duration, measured in days. For example, the average price change shifts the price by 16% and the average price lasts 30.23 days.

Figure 1. Reduced-Form Evidence for the Bullwhip Effect

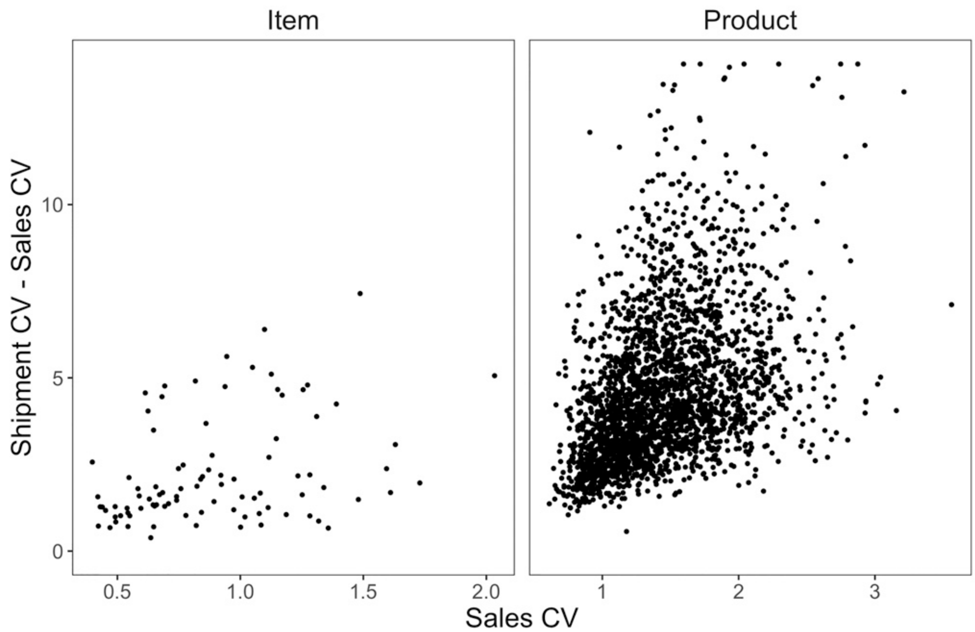

Notes. To create these plots, we calculate the coefficients of variation (CV) of sales and shipments by item and product and plot the difference between the shipment and sales measures as a function of the sales measure. This difference is positive in every case, so all items and products exhibit the bullwhip effect.

holds in both directions: the new-price-to-old-price ratio has quantiles 1.05, 1.11, and 1.27 following price increases and has quantiles 0.78, 0.89, and 0.95 following price decreases. Such coarse corrections are difficult to explain without menu costs.

More convincing are the profit patterns: profits follow a downward drifting process, punctuated by upward jumps when prices change (Figure 2). This, of course, is the expected pattern—the freshest prices should yield the highest profits. But what is striking is the magnitude of the profit discontinuity. For example, there are 16,420 price changes for which the price held constant in the 40 days leading up to the change and in the 40 days following the change. Focusing on this sample of price changes and denoting the number of days since the price change with t, we find that profits are on average 16.3% higher for $t \in \{20, \cdots, 40\}$ than for $t \in \{-40, \cdots, -20\}$ (with this difference being statistically significant at the $p < 10^{-12}$ level). Also, in the same sample, we find that 57% of products have higher average profits for $t \in \{20, \cdots, 40\}$ than for $t \in \{-40, \cdots, -20\}$ (with this probability being statistically greater than 0.5 at the $p < 10^{-12}$ level). The profit deterioration that follows price changes is also significant. For example, limiting attention to prices that remained constant for at least 180 days and defining t as we did above, we estimate that the average daily profit is as follows:

$$
\begin{aligned}
&¥3.41 \text{ when } t \in \{0, \cdots, 9\} \text{ days,}\\
&¥3.16 \text{ when } t \in \{10, \cdots, 19\} \text{ days,}\\
&¥3.10 \text{ when } t \in \{20, \cdots, 29\} \text{ days,}\\
&¥3.08 \text{ when } t \in \{30, \cdots, 39\} \text{ days,}\\
&¥2.98 \text{ when } t \in \{60, \cdots, 69\} \text{ days,}\\
&¥2.98 \text{ when } t \in \{70, \cdots, 79\} \text{ days,}\\
&¥2.95 \text{ when } t \in \{80, \cdots, 89\} \text{ days,}\\
&¥2.89 \text{ when } t \in \{90, \cdots, 99\} \text{ days,}\\
&¥2.85 \text{ when } t \in \{100, \cdots, 109\} \text{ days,}\\
&¥2.80 \text{ when } t \in \{110, \cdots, 119\} \text{ days,}\\
&¥2.79 \text{ when } t \in \{120, \cdots, 129\} \text{ days,}\\
&¥2.77 \text{ when } t \in \{130, \cdots, 139\} \text{ days,}\\
&¥2.76 \text{ when } t \in \{140, \cdots, 149\} \text{ days,}\\
&¥2.73 \text{ when } t \in \{150, \cdots, 159\} \text{ days,}\\
&¥2.71 \text{ when } t \in \{160, \cdots, 169\} \text{ days, and}\\
&¥2.66 \text{ when } t \in \{170, \cdots, 179\} \text{ days.}
\end{aligned}
$$

The persistence of the profit decline suggests that it stems from a gradual mismatch between an evolving demand and a static price, rather than from some other factor that coincides with the price change.

5.3. Demand and Inventory Smoothing

Our theoretical models explain that stores could adjust prices to smooth both demand and inventory: the DSP model suggests that stores increase prices to temper demand spikes and decrease prices to spur slow sales, and the EOQ model suggests that stores increase prices to preserve scarce supplies and decrease prices to liquidate excess supplies. We find both forms of smoothing in the data: the probability that a price change increases the price (i) increases with sales over the previous seven days and (ii) decreases with the current inventory level (Figure 3). The former result suggests demand smoothing and the latter suggests inventory smoothing.

6. Structural Econometric Model

6.1. Overview

To reconcile Section 3's theory with Section 4's data, we create a dynamic discrete choice model of the supply chain. We build on Aguirregabiria's (1999) model, making four substantive changes. First, we incorporate a Markov modulated demand process (Chen and Song 2001), whereas Aguirregabiria treated demand as independent and identically distributed (i.i.d.).[7] Second, we allow the DC to not fulfill the store's order quantity. Third, we model prices and order quantities as discrete decision variables, whereas Aguirregabiria treated them as continuous. This change makes our model more conducive to estimation. Specifically, our empirical likelihood function captures both the decision to place an order and the order size, whereas Aguirregabiria's

Figure 2. Reduced-Form Evidence of Menu Costs

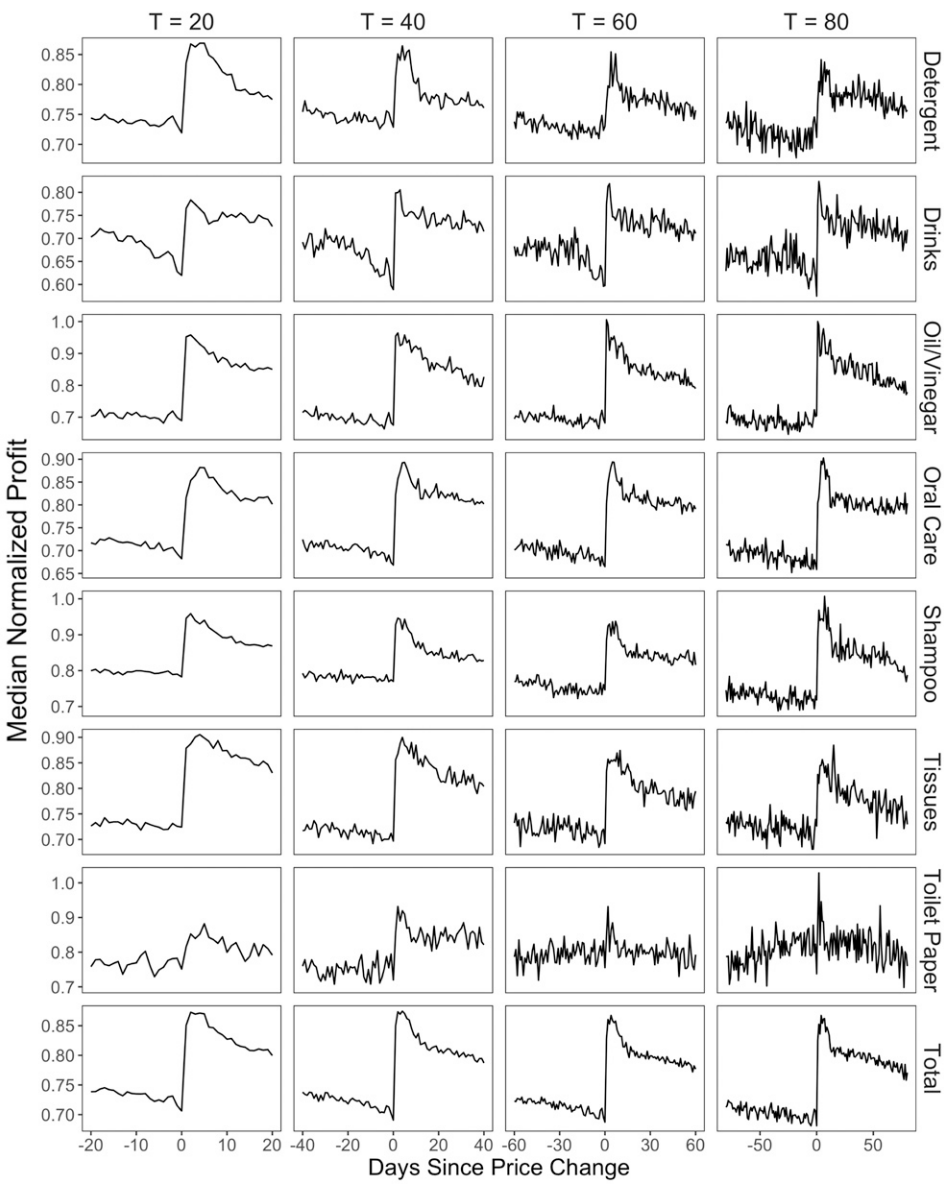

Notes. Producing these figures takes several steps. First, we isolate all instances in which a product's price held constant in the T days leading up to a price change and the T days following a price change. Second, we normalize each product's daily average profits to one in the $(2T+1)$–day window surrounding the price change. Third, we calculate the median normalized profit for each day $t \in \{-T, \cdots, T\}$, where $t = 0$ represents the day of the price change. Fourth, we plot these median profits, by product category, for $T \in \{20, 40, 60, 80\}$.

Figure 3. Reduced-Form Evidence for Demand-Based and Inventory-Based Pricing

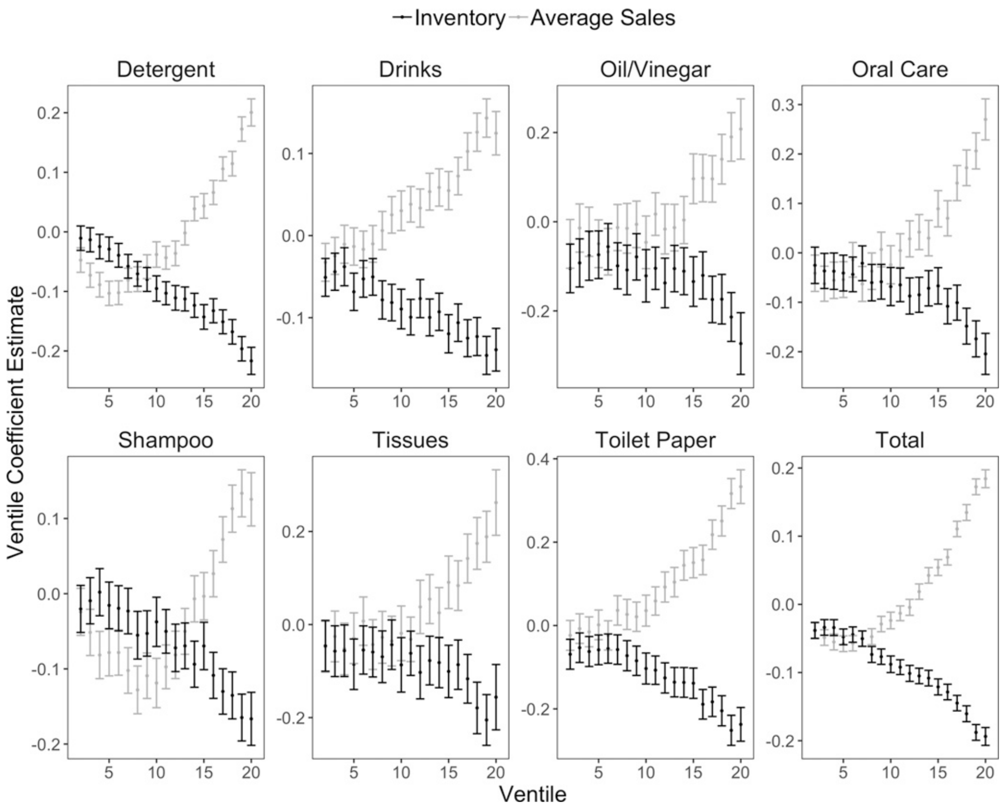

Notes. Producing these figures takes several steps. First, we filter the sample to the set of observations that have a price change. Second, we calculate the average sales over the previous seven days for each observation. Third, we group each product's inventory and average sales values into ventile buckets. (Ventiles are like deciles, except they divide a variable into 20 groups instead of 10.) Fourth, we regress a dummy variable that specifies that the price change was a price increase on dummy variables that specify (i) the inventory-level ventile, (ii) the average-sales ventile, (iii) the store, (iv) the item, and (v) the month. Fifth, we plot our 19 inventory-level-ventile coefficient estimates and our 19 average-sales-ventile coefficient estimates with points and plot their corresponding 95% confidence intervals with vertical bars. We find that the inventory estimates decrease by ventile, which suggests that stores are more likely to *decrease* prices when inventories are high. For example, the likelihood that a detergent price change increases prices is roughly 20 percentage points lower when the inventory level is in the top ventile than when it is in the bottom ventile. We also find that the average sales estimates increase by ventile, which suggests that stores are more likely to *increase* prices when sales are high. For example, we find that the likelihood that a detergent price change increases the price is roughly 20 percentage points higher when the average sales over the previous seven days is in the top ventile than when it is in the bottom ventile.

captures in only the former. Likewise, our likelihood function captures both the decision to change a price and the magnitude of the price change, whereas Aguirregabiria's captures only the former. Fourth, we incorporate Hendel and Nevo's (2006) two-step decision structure, specifying the store to first choose order quantities (conditional on nested-logit inclusive values) and then choose prices (conditional on order quantities), whereas Aguirregabiria modeled the store's decision as a single-shot, discrete-choice problem over price-quantity pairs. Aguirregabiria's framing is less economically sound because it requires the idiosyncratic value of two price-quantity pairs to be independent, even if they share the same price or the same quantity. His framing also is more restrictive because it requires the retail price statistical error to be exactly as noisy as the order quantity statistical error. In contrast, our specification incorporates an additional parameter, σ, which gives these distinct errors distinct variances. This additional parameter significantly improves our model fit.

6.2. Model Details

Our model describes a store that chooses order quantities and prices to maximize the expected discounted utility it receives from a representative product under per-period discount factor $\beta \in [0, 1)$. Each period lasts one day. A representative day proceeds as follows:

1. The day begins with inventory $i \in \mathbb{i} = \{0, \cdots, \bar{i}\}$, reference retail price $r \in \mathbb{p} = \{p_1, \cdots, p_{\bar{p}}\}$, and demand state $s \in \mathbb{s} = \{s_1, \cdots, s_{\bar{s}}\}$. More specifically, r is the previous day's retail price and s is a statistic that characterizes the distribution of demand.

2. Order shock vector $e = (e_1, \cdots, e_{\bar{q}})' \in \mathbb{R}^{\bar{q}}$ resolves. The elements of e are i.i.d., mean-zero Gumbel random

variables that correspond one-to-one with the set of feasible order quantities $\mathbb{q} = \{q_1, \cdots q_{\bar{q}}\}$, where $0 \in \mathbb{q}$. The store observes e, but we do not. We treat e as a statistical error term, as it captures all unobserved factors that influence the store's order quantity decision. Informally, one can interpret the elements of e as random fluctuations in shipping costs (e.g., it is more burdensome to place an order when the truck is almost full).

3. The store orders $q \in \mathbb{q}$ units of inventory from the DC.

4. Price shock vector $u = (u_1, \cdots, u_{\bar{p}})' \in \mathbb{R}^{\bar{p}}$ resolves. The elements of u are i.i.d., mean-zero Gumbel random variables that correspond one-to-one with the set of feasible retail prices $\mathbb{p} = \{p_1, \cdots, p_{\bar{p}}\}$. The store observes u, but we do not. We treat u as a statistical error term, as it captures all unobserved factors that influence the store's pricing decision. Informally, one can interpret the elements of u as random fluctuations in menu costs (e.g., the availability of sticker changers) combined with a bit of store manager caprice (e.g., the "gut feeling" associated with each new price).

5. The store chooses retail price $p \in \mathbb{p}$.[8]

6. Demand $d \in \mathbb{N}$ resolves from distribution $\delta_d(d|s,p)$.

7. The store sells $\min(i,d)$ units of inventory.

8. The store receives revenue $p\min(i,d)$.

9. Fulfillment indicator variable $f \in \{0,1\}$ resolves from distribution $\delta_f(f)$.

10. The store receives fq units of inventory from the DC.

11. The store receives utility

$$\kappa p \min(i,d) - \lambda \mathbb{1}(fq \neq 0) - \mu \mathbb{1}(p \neq r) - \eta \max(i-d,0) + e_q + \sigma u_p,$$

where κ is the benefit of one additional renminbi (the currency in use) in revenue, λ is the fixed shipping cost, μ is the menu cost, η is the cost of holding one unit of inventory for one day, and σ is a scalar that parameterizes the volatility of the pricing error term. In vector notation, the store's utility is $\theta'\gamma(i,r,q,p,e,u,d,f)$, where $\theta = (\kappa, \lambda, \mu, \eta, 1, \sigma)'$ and $\gamma(i,r,q,p,e,u,d,f) = (p\min(i,d), -\mathbb{1}(fq \neq 0), -\mathbb{1}(p \neq r), -\max(i-d,0), e_q, u_p)'$. We normalize the fifth element of θ to one to be consistent with Aguirregabiria's (1999) and the Bray et al. (2019) utility function scaling. Thus, we implicitly measure all utilities relative to the variance of the order quantity error term.

12. Tomorrow's inventory transitions to $i' = \max(i-d,0) + fq$.

13. Tomorrow's reference price is set to $r' = p$.

14. Tomorrow's demand state resolves from distribution $\delta_s(s'|s)$.[9]

This sequence of events characterizes a Markov decision process with action variables $\{p,q\}$ and state variables $\{i,r,s,e,u\}$. We estimate our model separately for each product cluster with Rust's (1987) nested fixed-point algorithm (see Appendix C for details).

6.3. Empirical Identification

The Bray et al. (2019) fourth theorem formally established that θ is empirically identified up to scale. But informally, it is easy to see that our utility parameters are identified.

First, the newsvendor model indicates that the service level should increase with κ/η, which identifies this ratio. For example, if the store almost never stocks out, then the inventory underage cost must be large relative to the inventory overage cost.

Second, inverting the classic EOQ solution, we find that the square of the average order quantity divided by twice the average demand should increase with λ/η, which identifies this ratio. For example, if the order quantity is large relative to the average demand, then the shipping cost must be large relative to the inventory holding cost.

Third, the classic menu cost models suggest that the frequency of price changes should decrease with μ/η, which identifies this ratio. For example, if the store drops prices when inventories are high, then the cost of holding excess stock must exceed the cost of changing prices.

Fourth, the classic logit logic suggests that the predictability of price changes decreases with σ, which identifies this parameter. For example, if prices are almost perfectly predictable given x and q, then σ must be near zero.

Fifth, the classic logit logic suggests that the predictability of orders increases with the magnitudes of κ, λ, μ, and η, which identifies the scale of these parameters. For example, if order quantities were almost perfectly predictable given x, then the magnitude of these parameters must be high.

7. Structural Estimates

We calculate $\widehat{\kappa}$, $\widehat{\lambda}$, $\widehat{\mu}$, $\widehat{\eta}$, and $\widehat{\sigma}$ for each of our 232 product clusters and bootstrap a standard error for each of these $5 \cdot 232 = 1{,}160$ parameter estimates (Figure 4). Across product clusters, our structural parameter estimates have the following quantiles:

	$\widehat{\kappa}$	$\widehat{\lambda}$	$\widehat{\mu}$	$\widehat{\eta}$	$\widehat{\sigma}$
Q1	0.0075	3.4	8.7	0.00046	0.8233
Q2	0.016	4.3	230	0.0012	4.38
Q3	0.066	5.2	850	0.0034	156

Note that we implicitly express these estimates relative to the variance of e_q, which we normalize to $\pi^2/6$ (the variance of a standard Gumbel random variable).

Our fixed shipping cost and inventory holding cost estimates are similar to that of Bray et al. (2019): their mean and median $\widehat{\lambda}$ estimates were 3.83 and 3.48, which lie within our $\widehat{\lambda}$'s interquartile range; and their mean and median $\widehat{\eta}$ estimates were 0.0030 and 0.0019, which lie within our $\widehat{\eta}$'s interquartile range. The $\widehat{\kappa}$ estimates are more difficult to benchmark because Bray et al. (2019) set the cost of a lost sale to a constant, whereas we set it to a constant multiplied by the

Figure 4. Distribution of Structural Utility Parameter Estimates and Corresponding t-Statistics

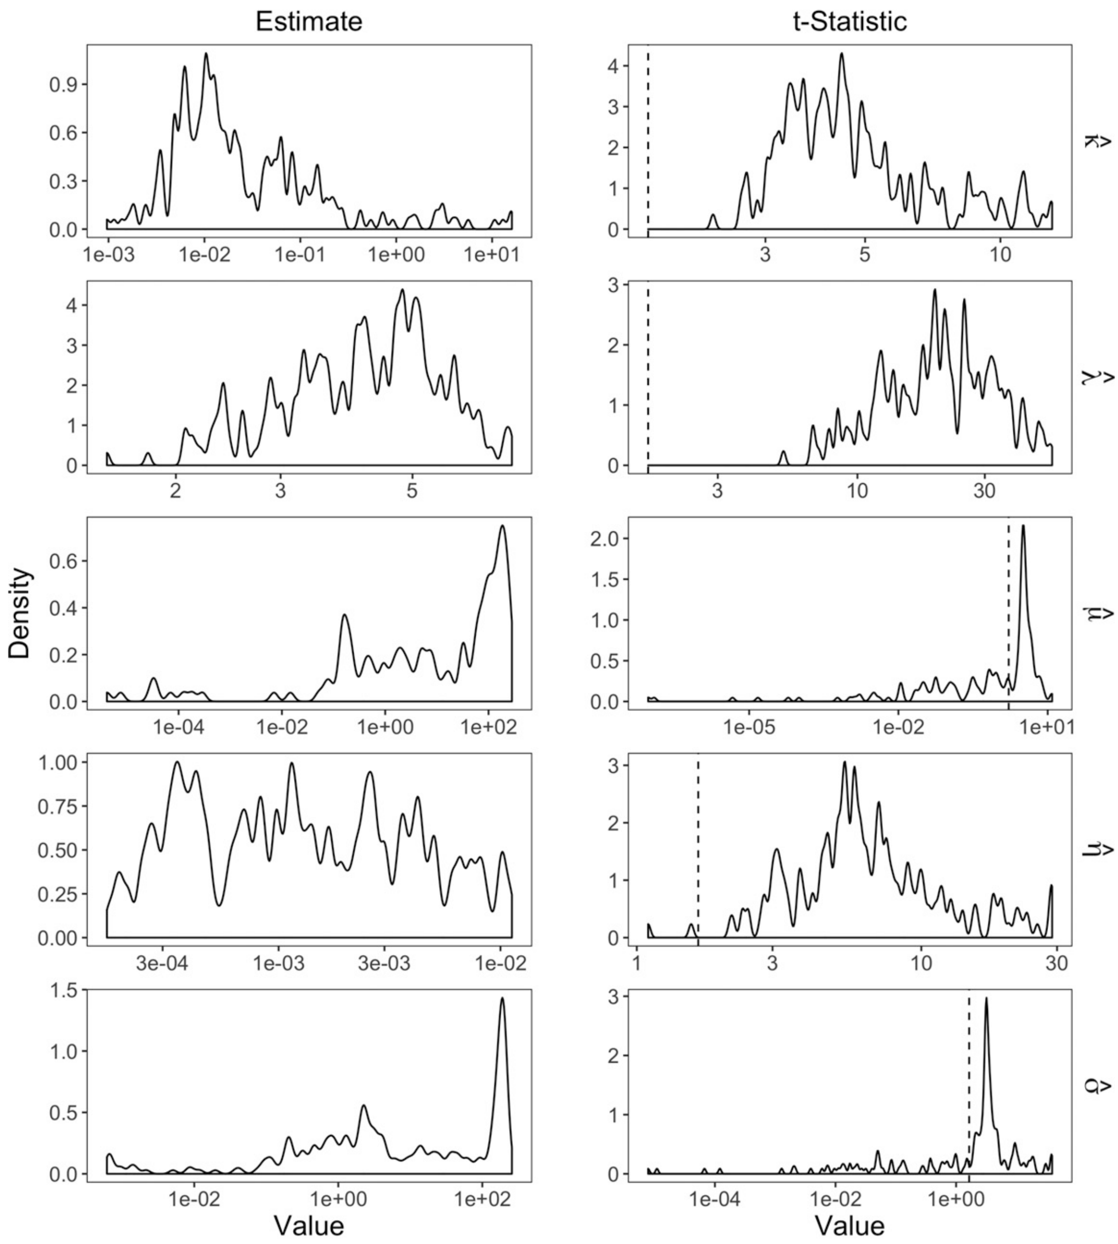

Notes. To create these plots, we estimate these distributions with a kernel density estimator and graph them on a log scale. The dashed lines in the t-statistic plots mark the $p = 0.05$ statistical significance threshold; anything to the right of these lines is significantly greater than zero.

prevailing price. But we get comparable estimates under the steady-state average price: the Bray et al. (2019) median lost sale cost was 0.080, just outside the interquartile range of our average-price lost sale cost, which ranges from 0.082–0.785. Bray et al. (2019) did not report menu costs but Aguirregabiria (1999) did, and his estimates were relatively close to ours: Aguirregabiria's (1999) $\widehat{\mu}$ estimates are between 50 and 139, whereas our median $\widehat{\mu}$ estimate is 230 (which is not significantly larger than 139 at the $p = 0.05$ level). Finally, neither Aguirregabiria (1999) nor Bray et al. (2019) incorporated a parameter analogous to σ. But we find that this parameter significantly improves our model fit. Specifically, we reject the $\sigma = 1$ null hypothesis at the $p = 0.05$ level in 82% of product clusters. Overall, 72% of our $\widehat{\sigma}$ estimates exceed one: prices are generally more erratic than orders.

As predicted in Section 6, we observe that (i) $\widehat{\kappa}/\widehat{\eta}$ increases with the service level; (ii) $\widehat{\lambda}/\widehat{\eta}$ increases with the inverted EOQ (i.e., the average order quantity squared, divided by twice the average demand); and (iii) $\widehat{\mu}/\widehat{\eta}$ decreases with the frequency of price changes (Table 3).

Table 3. Explanation of the Variation in Our Structural Estimates

	$\widehat{\kappa}/\widehat{\eta}$	$\widehat{\lambda}/\widehat{\eta}$	$\widehat{\mu}/\widehat{\eta}$
Intercept	0.309*	−0.002	0.276*
	(0.038)	(0.027)	(0.043)
Mean inventory	−0.047	0.254	−0.150
	(0.147)	(0.105)	(0.164)
Mean sales	−0.191	0.081	0.174
	(0.136)	(0.098)	(0.152)
Service level	0.324*	0.097	0.174
	(0.065)	(0.047)	(0.073)
Inverted EOQ	0.265*	0.516*	0.240*
	(0.082)	(0.059)	(0.092)
Price change frequency	−0.236*	0.064	−0.237*
	(0.077)	(0.055)	(0.086)

Notes. Producing this table takes several steps. First, we calculate five operational statistics for each product cluster: average inventory, average sales, service level, "inverted EOQ," and price change frequency. The service level is the fraction of days without a stockout. The inverted EOQ is the average order quantity squared divided by twice the average demand (i.e., what the λ/η ratio would be in the EOQ model). And the price change frequency is the fraction of days with a price change. Second, we difference these five operational statistics across product clusters that share an item (keeping the within-item variation and discarding the between-items variation). Third, we similarly difference estimate ratios $\widehat{\kappa}/\widehat{\eta}, \lambda/\widehat{\eta}$, and $\widehat{\mu}/\widehat{\eta}$. Fourth, we regress the estimate ratio differences on the operational statistic differences. Fifth, we tabulate the coefficient estimates and corresponding standard errors.

$^*p = 0.01$.

8. Counterfactuals

8.1. Specification

We next use our structural model to anticipate the supply chain implications of removing menu costs.[10] The most natural counterfactual would be a ceteris paribus change in the menu cost, from $\widehat{\mu}$ to zero. Unfortunately, this specification yields nonsensical counterfactual policies: without menu costs keeping them in check, the pricing error terms run amok, compelling the stores to set prices helter-skelter in a manic pursuit of the largest Gumbel shocks. So, we consider two alternative specifications.

The first compares the supply chain without price errors (i.e., with parameter vector $(\widehat{\kappa},\widehat{\lambda},\widehat{\mu},\widehat{\eta},1,0)'$) with the supply chain without menu costs or price errors (i.e., with parameter vector $(\widehat{\kappa},\widehat{\lambda},0,\widehat{\eta},1,0)'$). This specification implicitly assumes that prices will respond to demand and inventory in a controlled fashion after menu costs are removed, without a reckless degree of volatility. The second specification compares the current supply chain (i.e., with parameter vector $(\widehat{\kappa},\widehat{\lambda},\widehat{\mu},\widehat{\eta},1,\widehat{\sigma})'$) with the supply chain without menu costs or price errors (i.e., with parameter vector $(\widehat{\kappa},\widehat{\lambda},0,\widehat{\eta},1,0)'$). This specification implicitly interprets the price errors as arising from menu costs so that they go to zero together. Our two counterfactuals yield nearly the same findings, so we present the results of the former here and relegate the results of the latter to the online appendix. We henceforth refer to the case with parameter vector $(\widehat{\kappa},\widehat{\lambda},0,\widehat{\eta},1,0)'$ as the "counterfactual scenario" and to the case with parameter vector $(\widehat{\kappa},\widehat{\lambda},\widehat{\mu},\widehat{\eta},1,0)'$ as the "current scenario." (This, of course, is a slight abuse of language, because the true current scenario has $\sigma = \widehat{\sigma}$.)

8.2. Results

We find that removing menu costs intensifies both demand smoothing and inventory smoothing (Figure 5). Specifically, we estimate that removing menu costs amplifies the positive correlation between prices and demand states from an average of 2.54% to an average of 33.2% and amplifies the negative correlation between prices and inventory levels from an average of –3.6% to an average of –28.8%. Further, we estimate that the hypothetical change increases the average sales rate by $0.13/3.49 = 3.7\%$ and decreases the average inventory flow time by $0.51/23.99 = 2.1\%$ (Table 4). In other words, without menu costs, more inventory flows more quickly through the supply chain.

Relatedly, we estimate that removing menu costs stabilizes the supply chain (Table 5). The mean shipment coefficient of variation falls from 4.14 to 4.07 (a change that is significant at the $p < 10^{-6}$ level); the mean sales coefficient of variation falls from 1.26 in the current scenario to 1.21 in the counterfactual scenario (a change that is significant at the $p < 10^{-10}$ level); and the mean bullwhip falls from 2.89 to 2.86 (a change that is significant at the $p < 10^{-3}$ level).

8.3. Mechanism

The change in the sales rate largely explains the change in supply chain stability: regressing the change in the sales coefficient of variation on the change in the sales rate yields an R^2 of 0.88; regressing the change in the shipment coefficient of variation on the change in the sales rate yields an R^2 of 0.96; and regressing the change in the bullwhip effect on the change in the sales rate yields an R^2 of 0.91 (Figure 6). Specifically, each volatility measure decreases with the throughput rate:

Figure 5. How the Pricing Process Changes from the Current Scenario to the Counterfactual Scenario

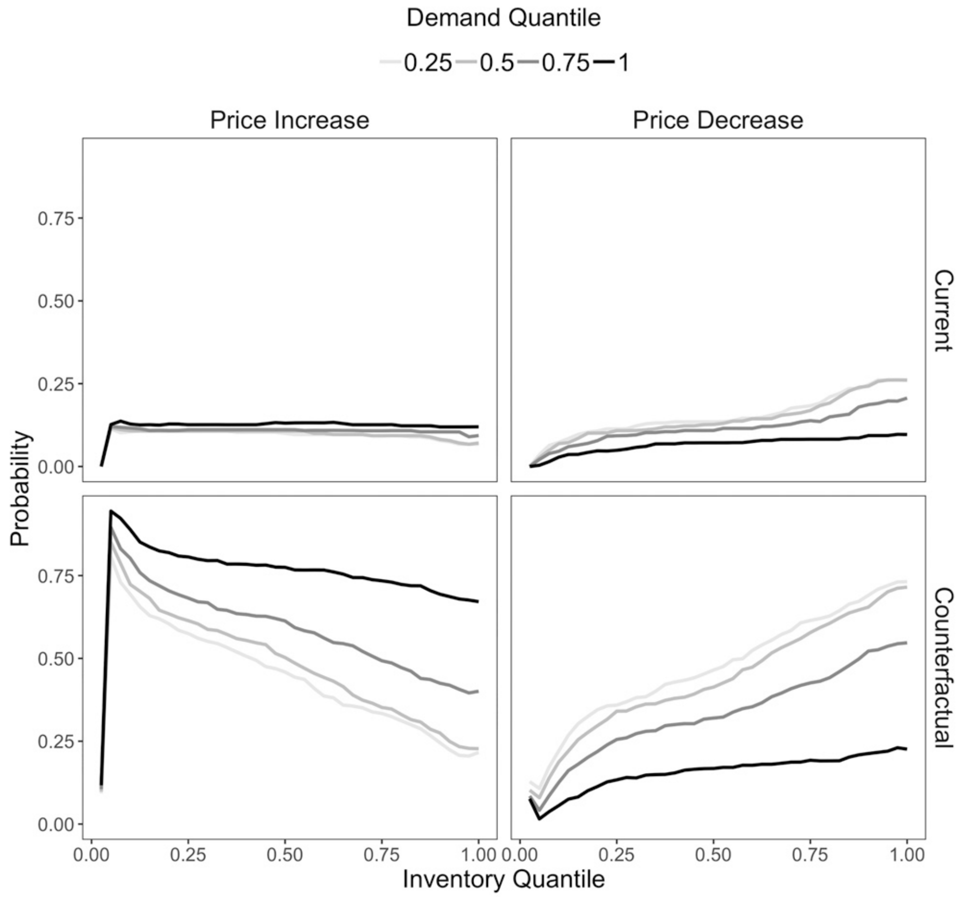

Notes. Producing these plots takes several steps. First, we calculate each product cluster's optimal policy under the current scenario, with parameter vector $(\widehat{\kappa}, \widehat{\lambda}, \widehat{\mu}, \widehat{\eta}, 1, 0)'$, and under the counterfactual scenario, with parameter vector $(\widehat{\kappa}, \widehat{\lambda}, 0, \widehat{\eta}, 1, 0)'$. Second, we map all state variables to their quantiles to compare across product clusters. For example, we map each product cluster's median inventory level to 0.5. Third, we calculate the average price-increase and price-decrease probabilities under the median price and under the various quantiles for the demand and inventory state variables. Fourth, we plot these price change probabilities by the quantile of the demand state variable (which is specified by line color) and the quantile of the inventory state variable (which is specified by the horizontal axis). In general, we find that price increases are more likely when the inventory state variable is low and the demand state variable is high, and vice versa for price decreases. These effects are more pronounced without menu costs, as prices are more sensitive to the inventory and demand state variables in this case.

the supply chain is more stable at higher sales rates, like a bicycle is more stable at higher speeds. This principle holds both upstream and downstream but for different reasons.

The upstream flows are governed by EOQ dynamics. Specifically, the EOQ model suggests that the shipping batch size should increase with the square root of sales. And because the shipping standard deviation scales

Table 4. Downstream Operations from the Current Scenario to the Counterfactual Scenario

	Current	Counterfactual	Difference
Price	16.529	16.949	0.419
Demand	3.904	3.921	0.017
Sales	3.487	3.616	0.129
Lost sales	0.417	0.305	−0.112
Inventory	23.986	23.473	−0.513
Shipment prob.	0.082	0.085	0.003
Price change prob.	0.004	0.106	0.102

Notes. To produce this table, for each product cluster, we calculate the Markov chain's stationary distribution under the current parameter vector, $(\widehat{\kappa}, \widehat{\lambda}, \widehat{\mu}, \widehat{\eta}, 1, 0)'$, and under the counterfactual parameter vector, $(\widehat{\kappa}, \widehat{\lambda}, 0, \widehat{\eta}, 1, 0)'$. Under each steady-state distribution, we calculate the average price, demand, sales, lost sales, inventory, shipment probability, and price change probability. We tabulate the average value of these statistics across products cluster for both scenarios and tabulate their differences from one scenario to the next. Prices are measured in renminbi (the local currency); demand, sales, and lost sales are measured in units per day; and inventory is measured in days' worth of supply.

Table 5. Change in Supply Chain Stability from the Current Scenario to the Counterfactual Scenario

	Shipment CV	Sales CV	Bullwhip
Detergent	−0.036	−0.030*	−0.006
	(0.025)	(0.010)	(0.016)
Drinks	−0.089*	−0.051*	−0.038
	(0.027)	(0.011)	(0.017)
Oil/vinegar	−0.203*	−0.078*	−0.125*
	(0.058)	(0.022)	(0.036)
Oral care	−0.015	−0.024*	0.009
	(0.027)	(0.008)	(0.019)
Shampoo	−0.102	−0.054*	−0.049
	(0.049)	(0.018)	(0.031)
Tissues	−0.095	−0.040	−0.055
	(0.086)	(0.035)	(0.052)
Toilet paper	−0.066*	−0.043*	−0.023
	(0.018)	(0.007)	(0.011)
Total	−0.072*	−0.043*	−0.029*
	(0.014)	(0.005)	(0.009)

Notes. To produce this table, for each product cluster, we calculate the Markov chain's stationary distribution under the current parameter vector, $(\widehat{\kappa},\widehat{\lambda},\widehat{\mu},\widehat{\eta},1,0)'$, and under the counterfactual parameter vector, $(\widehat{\kappa},\widehat{\lambda},0,\widehat{\eta},1,0)'$. Under each steady-state distribution, we calculate the shipment coefficient of variation, the sales coefficient of variation, and the difference between the two, which serves as our bullwhip measure. Finally, we present the mean within-product-cluster changes in these three volatility measures, from the current scenario to the counterfactual scenario, by category.

*$p = 0.01$.

linearly with the shipping batch size, the shipping coefficient of variation should thus decrease with the square root of sales. In short, increasing the sales rate decreases the shipping batch size, in relative terms.[11]

We find a similar pattern downstream: the sales standard deviations grow more slowly than the sales means, for a net reduction in the sales coefficient of variation. This, however, stems from classic demand

Figure 6. Change in the Sales Rate, from the Current to the Counterfactual Scenario, Explains the Changes in the Shipment Coefficient of Variation, in the Sales Coefficient of Variation, and in the Bullwhip Effect

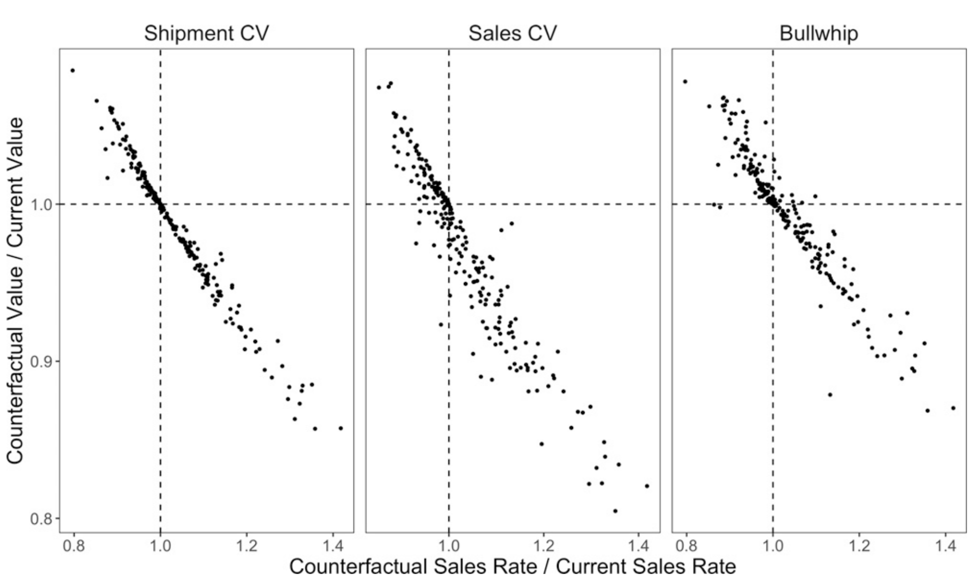

Notes. To produce these plots, for each product cluster, we calculate the shipment coefficient of variation, the sales coefficient of variation, and the bullwhip effect, under both the current and counterfactual scenarios, in the fashion of Table 5. We likewise calculate the mean sales rates—measured in items per day per store—under the two scenarios. We then scatter plot the fractional change in our volatility measures as a function of the fractional change in the sales rates. The dashed lines denote the value of one, where there is no change. These plots depict the change in supply chain stability from the current scenario to the counterfactual scenario. For each product cluster, we calculate the Markov chain's stationary distribution under the current parameter vector, $(\widehat{\kappa},\widehat{\lambda},\widehat{\mu},\widehat{\eta},1,0)'$, and under the counterfactual parameter vector, $(\widehat{\kappa},\widehat{\lambda},0,\widehat{\eta},1,0)'$. Under each steady-state distribution, we calculate the shipment coefficient of variation, the sales coefficient of variation, and the difference between the two, which serves as our bullwhip measure. Finally, we plot the distributions of the within-product-cluster changes in these three volatility measures, from the current scenario to the counterfactual scenario, by category.

pooling: the standard deviation of demand generally scales with the square root of the mean demand. For example, our demands are negative binomials, which are similar to Poison random variables; a Poison's standard deviation equals the square root of its mean. (Note that this demand pooling effect differs from the demand smoothing effect that we anticipated with our DSP model.)

The bullwhip effect also diminishes as the supply chain scales. In fact, this "scaling effect" by far dominates the demand smoothing or inventory smoothing effects we theorized in Section 3.

Finally, the increase in the throughput rate largely stems from a decrease in the stockout rate. Specifically, the reduction in stockouts explains $0.112/0.129 = 87\%$ of the increase in sales (Table 4). Intuitively, reducing menu costs decreases stockouts because the stores can increase prices when inventories run low. More precisely, stockouts decrease because the marginal unit is priced higher, which makes stocking out more costly (the average price increases by $0.42/16.53 = 2.5\%$ in the counterfactual).[12]

8.4. Strategic Consumers

Does strategic consumer behavior threaten the robustness of our results? Not if the returns to strategic behavior given the stores' optimal policies are low. The logic goes like this: (i) stores devise their pricing and ordering policies ignoring strategic consumer behavior; (ii) given the stores' policies, the returns to strategic consumer behavior are low; (iii) given the low returns to strategic behavior, consumers choose not pay the cost of being strategic (i.e., not to monitor prices, inventory, and demand or try to learn the store's policy), so no strategic behavior occurs; (iv) stores are correct to ignore strategic behavior when devising their pricing and ordering policies.

We estimate the returns to strategic behavior with a variation of Hendel and Nevo's (2013) method. Specifically, we assume that consumers (i) visit the store every nth day for $n \in \{1, 2, 4, 8, 16\}$ and (ii) know the expectation of the price n days later. Hence, the "effective price" that a strategic consumer faces on day t is $y_t = p_{t-n}\mathbb{1}(p_{t-n} \le \mathrm{E}_{t-n}(p_t)) + p_t\mathbb{1}(p_{t-n} > \mathrm{E}_{t-n}(p_t))$, where $\mathrm{E}_{t-n}(\cdot)$ represents expectation with respect to the day-$(t-n)$ information set. We calculate the expected value of a strategic consumer's price under the Markov chain's stationary distribution, $\mathrm{E}(y_t)$, and the expected value of a myopic consumer's price under this distribution, $\mathrm{E}(p_t)$, for both the current and counterfactual scenarios. Then, we calculate the strategic consumer's savings, $(\mathrm{E}(p_t) - \mathrm{E}(y_t))/\mathrm{E}(p_t)$. These savings estimates should be *generous* upper bounds to the real savings, as they assume that (i) consumers know the current state variables; (ii) consumers know the store's policy—and hence the dynamic program's transition probabilities; and (iii) consumers are able to calculate the distribution of the state variables after n transitions. Our simulations deliver savings estimates of roughly 1% in the current scenario and roughly 2% in the counterfactual scenario.[13] In summary, our analysis establishes that the money consumers leave on the table by acting myopically is negligible—both in the current and counterfactual settings. (See the online appendix for details.)

9. Conclusion

Our theoretical models did not fare particularly well at predicting our empirical results. These models identified some of the second-order effects but missed the primary effect, which was an uptick in sales rates that helped steady the inventory flows.

Our EOQ model suggests that removing menu costs would reduce the ratio between the average inventory level and the shipping quantity—as high stocks burn faster and low stocks linger longer—which would decrease the effective inventory holding cost and thus increase the order quantity. An increase in the order quantity should in turn increase the shipping volatility, which should increase the bullwhip effect. We find evidence of increased inventory smoothing, but its effect on supply chain stability is swamped by the scaling effect. The problem with our EOQ model is that it does not factor in demand uncertainty and therefore does not factor in the lost sales due to stockouts, the primary driver of the change in sales rates.

Our DSP model suggests that removing menu costs would help stores adjust prices to smooth demand shocks and that this demand smoothing would lead to even more pronounced order smoothing, for a net reduction in the bullwhip effect. We find empirical evidence for the downstream effect but not for the upstream effect. The problem with our DSP model is that it does not factor in the fixed shipping costs or the fixed batch sizes that inventories must move in, both of which attenuate the upstream effects of demand smoothing. For example, shipping costs can compel a store to order two weeks' worth of supply, and a single day's smoothing makes a small dent in a 14-day order. Indeed, these two factors—the fixed shipping cost and the fixed batch size—quash almost all demand signals, so that the change in the baseline demand rate almost entirely mediates the upstream effect of removing menu costs. And like our EOQ model, our DSP model did not anticipate the change in sales rates.

This exercise highlights the strength of structural estimation for the study of supply chains. Supply chains are too complex to solve outright with equations, but they are structured enough to model numerically. And although theoretical models, such as our DSP or

EOQ specifications, can articulate potential effects, they are not powerful enough to entertain multiple different mechanisms and determine which will be first order and which will be second order. But a structural econometric model, primed with real data from a real operation, can weigh the various vying forces and impartially crown a victor.

Appendix A. Demand Signal Processing Model

We illustrate the demand-based pricing effect of removing menu costs with an extension of the Lee et al. (1997) demand signal processing model. At the beginning of period t, the store chooses replenishment quantity q_t and price p_t. Next, demand $d_t = s_t - mp_t$ resolves, where $m > 0$ is a scalar and s_t a random state variable that follows an autoregressive process: $s_t = a + bs_{t-1} + \sigma e_t$, where $a > 0, \sigma > 0$, and $b \in (0,1)$ are scalars, and $e_t \sim N(0,1)$ is a random shock that resolves after q_t and p_t are chosen. Next, the inventory level transitions to $i_t = i_{t-1} + q_t - d_t$. Finally, the store receives revenue $p_t d_t$ and incurs ordering cost wq_t, menu cost $\mu \mathbb{1}(p_t \neq p_{t-1})$, holding cost $\eta \max(i_t, 0)$, and backlog cost $\kappa \max(-i_t, 0)$. Excess inventory can be returned at no cost; the wholesale price satisfies $w < a/(m(1-b))$, which ensures that the expected demand would be positive if the firm priced at the marginal cost. The store's objective is to maximize the long-run discounted profit under discount factor $\beta \in [0,1)$.[14]

Dynamic Pricing

When $\mu = 0$, the store chooses price p_t and order quantity q_t to maximize

$$\mathrm{E}\left(\sum_{t\geq 0} \beta^t (p_t d_t - wq_t - \eta \max(i_t, 0) - \kappa \max(-i_t, 0))\right). \quad \text{(A.1)}$$

If j_t is the order-up-to level on day t, then

$$\sum_{t\geq 0} \beta^t q_t = \sum_{t\geq 0} \beta^t (d_{t-1} + j_t - j_{t-1}) = d_{-1} - j_{-1} + \sum_{t\geq 0} \beta^t (\beta d_t + (1-\beta) j_t).$$

With this, we express the store's objective in terms of p_t and j_t:

$$\mathrm{E}\left(\sum_{t\geq 0} \beta^t ((p_t - \beta w) d_t - (1-\beta) w j_t - \eta \max(j_t - d_t, 0) - \kappa \max(d_t - j_t, 0))\right).$$

The corresponding first-order conditions are

$$p_t = \frac{a + bs_{t-1} + mw}{2m} \quad \text{and}$$
$$j_t = (a + bs_{t-1} - mw)/2 + \sigma \Phi^{-1}\left(\frac{\kappa - (1-\beta)w}{\kappa + \eta}\right).$$

Price p_t implies sales

$$d_t = a + bs_{t-1} + \sigma e_t - mp_t = (a + bs_{t-1} - mw)/2 + \sigma e_t,$$

which has mean and variance

$$\mathrm{E}(d_t) = (a + b\mathrm{E}(s_{t-1}) - mw)/2 = (a + b\frac{a}{1-b} - mw)/2 = \frac{a}{2(1-b)} - mw/2 \quad \text{and}$$
$$\mathrm{Var}(d_t) = \frac{b^2}{4}\mathrm{Var}(s_{t-1}) + \sigma^2 = \frac{b^2\sigma^2}{4(1-b^2)} + \sigma^2 \quad \text{(A.2)}$$

Next, order-up-to level j_t and sales d_t imply order quantity

$$q_t = j_t - j_{t-1} + d_{t-1} = b(s_{t-1} - s_{t-2})/2 + d_{t-1} = (bs_{t-1} + a - wm)/2 + \sigma e_{t-1},$$

which has mean and variance

$$\mathrm{E}(q_t) = \mathrm{E}(d_t) \quad \text{and}$$
$$\mathrm{Var}(q_t) = \frac{b^2\sigma^2}{4(1-b^2)} + (1+b)\sigma^2 = \mathrm{Var}(d_t) + b\sigma^2. \quad \text{(A.3)}$$

Static Pricing

When $\mu = \infty$, the store maximizes the same objective but with the additional constraint that $p_t = p_{t-1}$. Adding this constraint changes the problem's first-order conditions to[15]

$$j_t = a + bs_{t-1} - \frac{a}{2(1-b)} - mw/2 + \sigma \Phi^{-1}\left(\frac{\kappa - (1-\beta)w}{\kappa + \eta}\right) \quad \text{and}$$
$$p_0 = \frac{a}{2m(1-b)} + w/2.$$

Price p_0 implies sales

$$d_t = a + bs_{t-1} + \sigma e_t - mp_0 = a + bs_{t-1} + \sigma e_t - \frac{a}{2(1-b)} - mw/2,$$

which has mean and variance

$$\mathrm{E}(d_t) = \mathrm{E}\left(a + bs_{t-1} + \sigma e_t - \frac{a}{2(1-b)} - mw/2\right) = \frac{a}{2(1-b)} - mw/2$$
$$\text{and } \mathrm{Var}(d_t) = \mathrm{Var}\left(a + bs_{t-1} + \sigma e_t - \frac{a}{2(1-b)} - mw/2\right)$$
$$= b^2\mathrm{Var}(s_{t-1}) + \sigma^2 = \frac{\sigma^2}{1-b^2}. \quad \text{(A.4)}$$

Next, order-up-to level j_t and sales d_t imply order quantity

$$q_t = j_t - j_{t-1} + d_{t-1} = b(s_{t-1} - s_{t-2}) + d_{t-1} = bs_{t-1} + \sigma e_{t-1} + a - \frac{a}{2(1-b)} - wm/2,$$

which has mean and variance

$$\mathrm{E}(q_t) = \mathrm{E}(d_t) \quad \text{and}$$
$$\mathrm{Var}(q_t) = \frac{b^2\sigma^2}{1-b^2} + (1+2b)\sigma^2 = \mathrm{Var}(d_t) + 2b\sigma^2. \quad \text{(A.5)}$$

Comparison

Expressions (A.2) and (A.4) establish that decreasing the menu cost from $\mu = \infty$ to $\mu = 0$ would decrease the variance of sales by

$$\frac{\sigma^2}{1-b^2} - \frac{b^2\sigma^2}{4(1-b^2)} - \sigma^2 = \frac{3\sigma^2 b^2}{4(1-b^2)} > 0.$$

Because the mean sales are the same in both cases, the sales coefficient of variation would likewise decrease.

Expressions (A.3) and (A.5) establish that decreasing the menu cost from $\mu = \infty$ to $\mu = 0$ would decrease the variance of inbound shipments by

$$\frac{b^2\sigma^2}{1-b^2} + (1+2b)\sigma^2 - \frac{b^2\sigma^2}{4(1-b^2)} - (1+b)\sigma^2 = \frac{3b^2\sigma^2}{4(1-b^2)} + b\sigma^2 > 0.$$

Because the mean shipment quantities are the same in both cases, the shipment coefficient of variation would likewise decrease.

Expressions (A.2)–(A.5) establish that decreasing the menu cost from $\mu = \infty$ to $\mu = 0$ would decrease the bullwhip effect by

$$\frac{\sqrt{\frac{b^2\sigma^2}{1-b^2} + (1+2b)\sigma^2} - \sqrt{\frac{\sigma^2}{1-b^2}}}{\frac{a}{2(1-b)} - mw/2} - \frac{\sqrt{\frac{b^2\sigma^2}{4(1-b^2)} + (1+b)\sigma^2} - \sqrt{\frac{b^2\sigma^2}{4(1-b^2)} + \sigma^2}}{\frac{a}{2(1-b)} - mw/2}$$

$$= \frac{\sqrt{\frac{\sigma^2}{1-b^2}}}{\frac{a}{2(1-b)} - mw/2}\Big(\sqrt{1+2b(1-b^2)} - 1 - \sqrt{1-3b^2/4 + b(1-b^2)} + \sqrt{1-3b^2/4}\Big).$$

It is easy to check that the expression inside the parenthesis is zero for $b = 0$, strictly increasing until it peaks at $b \simeq 0.502$ and then strictly decreasing until it hits zero again at $b = 1$.

Appendix B. Economic Order Quantity Model

We illustrate the inventory-based pricing effect of removing menu costs with the Stamatopoulos et al. (2019) generalized EOQ model. The model is the continuous-time limit of a discrete-time specification. Specifically, we chop each day into n periods and let $n \to \infty$. In period t, the store chooses price p_t and sells $d_t = (a - mp_t)/n$ units, where $a, m > 0$ are fixed demand constants. To satisfy this demand, the store orders $q = \sum_{t=1}^{n\tau} d_t$ units of inventory every τ days. In other words, each replenishment cycle spans $n\tau$ periods. Note that replenishment cycles are identical because demand is deterministic.

To study the bullwhip in this setting, we first calculate shipment and sales volatility—both with respect to time, as there is no uncertainty. Note that the shipment quantity is q every $n\tau$ periods and is zero otherwise. Hence, it resembles q times a Bernoulli random variable with mean $1/n\tau$. Accordingly, the shipment quantity has mean $q/n\tau$, variance $q^2(n\tau - 1)/n^2\tau^2$, and coefficient of variation $\sqrt{n\tau - 1}$. When it comes to sales, note that $d_t \in [0, a/n]$, where the lower limit corresponds to price $p_t = \infty$ and the upper limit to price $p_t = 0$. Hence, Popoviciu's inequality establishes that demand's variance is upper bounded by $(a/n)^2/4$ and its coefficient of variation is upper bounded by $a\tau/2q$. And, of course, these terms are lower bounded by zero.

These bounds imply that if BW_n is the bullwhip effect when days are divided into n periods, $\sqrt{n\tau - 1} - a\tau/2q \le \mathrm{BW}_n \le \sqrt{n\tau - 1}$. This, in turn, implies that

$$\lim_{n\to\infty} \mathrm{BW}_n/\sqrt{n} = \sqrt{\tau}.$$

So, from an EOQ perspective, the square root of the time between orders—equivalently, the square root of the days of inventory in each order—measures the bullwhip effect. Intuitively, increasing this time is the same as increasing the order batch size, the Lee et al. (1997) third bullwhip driver.

And the Stamatopoulos et al. (2019) analysis tells us how removing menu costs affects τ. Specifically, they found that the optimal inventory holding cost per cycle will be the same under $\mu = 0$ and $\mu = \infty$. However, they also showed that for any given τ, the per-cycle inventory holding cost will be strictly less under $\mu = 0$ than under $\mu = \infty$ (given our linear demand model). These two facts imply that τ must be strictly greater under $\mu = 0$ than under $\mu = \infty$ (see Stamatopoulos et al. 2019, appendix C, for details).

Appendix C. Estimation Procedure

Value Function and Policy Function

Before developing the estimator, we first define our dynamic program's value function and policy function. These functions integrate over unobserved error terms e and u in the standard fashion (see Aguirregabiria and Mira 2010). We define these functions in six steps:

1. The Gumbel distribution has two useful properties: if $\{\epsilon_i\}_{i=1}^N$ are i.i.d., mean-zero Gumbel random variables and $\{y_i\}_{i=1}^N$ are fixed scalars, then

$$\mathrm{E}(\epsilon_j | j = \arg\max_i (y_i + \epsilon_i)) = -\ln(\mathrm{P}(j = \arg\max_i (y_i + \epsilon_i))), \text{ and} \tag{C.6}$$

$$\mathrm{P}(j = \arg\max_i (y_i + \epsilon_i)) = \frac{\exp(y_j)}{\sum_{i=1}^n \exp(y_i)}. \tag{C.7}$$

2. Define $\mathbb{x} = \mathbb{i} \times \mathbb{p} \times \mathbb{s}$ as the collection of observable state variables. And define $\phi(q|x)$ as the probability that the optimally behaving store chooses order quantity q in state $x = \{i, r, s\}$ and $\psi(p|q, x)$ as the probability that the optimally behaving store chooses price p in state x given order q. Note that from our perspective, functions ϕ and ψ fully characterize the store's policy. We characterize these functions in step 6 but until then take them as given.

3. Line (6) establishes that the expected value of e_q conditional on the agent choosing order q in state x is $-\ln(\phi(q|x))$ and that the expected value of u_p conditional on the agent choosing price p and order q in state x is $-\ln(\psi(p|q, x))$. Using these observations, we express the store's expected state-x utility in terms of ϕ, ψ, δ_d, and δ_f:

$$\begin{aligned} &\mathrm{E}(\theta'\gamma(i, r, q, p, e, u, d, f) \mid x) = \theta'\pi(x),\\ &\text{where } \pi(x) = \mathrm{E}(\gamma(i, r, q, p, e, u, d, f) \mid x)\\ &= \sum_{q\in\mathbb{q}} \sum_{p\in\mathbb{P}} \phi(q|x)\psi(p|q, x)\omega(p, q, x) \end{aligned} \tag{C.8}$$

$$\begin{aligned} &\text{and } \omega(p, q, x) = \mathrm{E}(\gamma(i, r, q, p, e, u, d, f) \mid p, q, x)\\ &= \sum_{d\in\mathbb{N}} \sum_{f\in\{0,1\}} \delta_d(d|s, p)\delta_f(f)(p\min(i, d), -\ 1(fq \ne 0),\\ &-\ \mathbb{1}(p \ne r), -\max(i - d, 0), -\ln(\phi(q|x)), -\ln(\psi(p|q, x)))'. \end{aligned} \tag{C.9}$$

4. We likewise express the Markov chain's transition probabilities in terms of ϕ, ψ, δ_d, δ_f, and δ_s: the probability that the state jumps from x to x' is

$$\delta(x'|x) = \sum_{q\in\mathbb{q}} \sum_{p\in\mathbb{p}} \phi(q|x)\psi(p|q, x)\zeta(x'|p, q, x), \tag{C.10}$$

where $$\zeta(x'|p,q,x) = \sum_{d\in\mathbb{N}}\sum_{f\in\{0,1\}} \delta_d(d|s,p)\delta_f(f)\delta_s(s'|s)\,\mathbb{1}(r'=p\cap i'=\max(i-d,0)+fq). \quad \text{(C.11)}$$

5. Next, we express the dynamic program's value function in terms of utility function π and transition function δ. The expected discounted value of entering state x is $\theta'\nu(x)$, where vector $\nu(x)$ is the unique solution of the following Bellman equation:

$$\nu(x) = \pi(x) + \beta\sum_{x'\in\mathbb{x}} \delta(x'|x)\nu(x'). \quad \text{(C.12)}$$

6. Finally, we characterize policy functions ϕ and ψ. Expression (C.7) establishes that the optimal choice probabilities have the following multinomial logistic forms:

$$\phi(q|x) = \frac{\exp\left(\sum_{p\in\mathbb{p}}\psi(p|q,x)(\overline{\theta}'\omega(p,q,x)+\beta\sum_{x'\in\mathbb{x}}\zeta(x'|p,q,x)\theta'\nu(x'))\right)}{\sum_{j\in\mathbb{q}}\exp\left(\sum_{p\in\mathbb{p}}\psi(p|j,x)(\overline{\theta}'\omega(p,j,x)+\beta\sum_{x'\in\mathbb{x}}\zeta(x'|x,p,j)\theta'\nu(x'))\right)} \quad \text{(C.13)}$$

$$\text{and } \psi(p|q,x) = \frac{\exp\left((\underline{\theta}'\omega(p,q,x)+\beta\sum_{x'\in\mathbb{x}}\zeta(x'|p,q,x)\theta'\nu(x'))/\sigma\right)}{\sum_{j\in\mathbb{p}}\exp\left((\underline{\theta}'\omega(j,q,x)+\beta\sum_{x'\in\mathbb{x}}\zeta(x'|x,j,q)\theta'\nu(x'))/\sigma\right)}, \quad \text{(C.14)}$$

$$\text{where } \overline{\theta} = (\kappa,\lambda,\mu,\eta,0,\sigma)'$$
$$\text{and } \underline{\theta} = (\kappa,\lambda,\mu,\eta,0,0)'.$$

Expressions (C.8)–(C.14) characterize the optimal policy implicitly. This system of equations has a unique solution, which we can identify with the policy iteration algorithm, iterating the equations to convergence.

Estimation Methodology

Our estimation procedure has 10 steps. We first discuss these steps at a high level and then discuss them in detail.

In step 1, we estimate the causal effect of prices on demand with an instrumental variables (IV) regression. This step represents the most notable difference between our estimation procedure and the procedure from Bray et al. (2019), who treated demand as exogenous. Our IV regression includes several exogenous factors that influence demand, such as month dummies and product dummies. Ideally, we would incorporate each of these demand shifters as a state variable in our dynamic program, but the curse of dimensionality prevents us from doing so. So in step 2, we concentrate these various demand shifters into the single demand state variable s. Specifically, we set s to the fitted value of the IV regression when prices are normalized to zero—that is, to the intercept of the demand curve, conditional on the current exogenous demand shifters.

After calculating s, we have all of our state and action variables. We limit the number of possible values each of these variables can take in step 3. Specifically, we restrict the state space to a grid of $40\cdot4\cdot4=640$ points and the action space to a grid of $4\cdot4=16$ points.[16]

Next, we estimate the transition probabilities. In step 4, we estimate δ_d with the negative binomial distribution that best rationalizes the observed demands given the expected demands we get from our IV regression. In steps 5 and 6, we estimate δ_f and δ_s with their empirical distributions. And in step 7, we combine these estimates to create an estimate of ζ, the action-specific transition probability function.

We define our dynamic program solver in step 8. The solver uses the policy iteration algorithm, iterating equations (C.8)–(C.14) until convergence, to calculate the optimal policies that correspond to a given parameter vector, θ.

And, finally, we nest this dynamic programming problem within a maximum likelihood problem in steps 9 and 10, setting $\widehat{\theta}$ to the parameter vector that maximizes the empirical likelihood of the observed actions.

In more detail, our estimation process proceeds as follows:

1. We run an IV regression of tomorrow's demand on (i) store dummy variables; (ii) month dummy variables; (iii) today's demand; (iv) today's sales of the given item across all stores; (v) today's sales of the given store across all items, and (vi) tomorrow's retail price, using the wholesale price from today and the retail price from four weeks before as instruments for tomorrow's retail price. These instruments are valid because they directly influence prices but not demand. The current retail price should clearly respond to both the current wholesale price (because the stores prefer stable margins) and last month's retail price (because menu costs make prices sticky). But, otherwise, wholesale prices and month-old retail prices should not appreciably influence the customer's purchase decision.

We test the strength of our instruments with an F-test of the null hypothesis that the coefficients corresponding to our two instrumental variables are zero in the first stage of the IV regression. These tests strongly reject the hypothesis that the wholesale price from today and the retail price from four weeks before do not influence tomorrow's retail price. Specifically, of our 308 product cluster tests, 298 reject the null at the $p=0.05$ level, 292 reject it at the $p=0.01$ level, and 284 reject it at the $p=0.001$ level.

2. We test the necessity of our instruments with a Hausman test for the endogeneity of retail prices. Specifically, we test whether the coefficient of the fitted retail price (i.e., the projection of the retail price onto the instruments) goes to zero when we add the actual retail price into the second stage of the IV regression. Overall, we find it prudent to use instruments: we reject the null hypothesis that retail prices are exogenous in 99 out of 308 product clusters, at the $p=0.05$ level. Demand-based pricing makes retail prices endogenous.[17]

We define demand state s as the fitted value of our IV regression when the retail price is set to zero.

3. For numerical tractability, we restrict the state space to a grid comprising 40 inventory levels, four demand states, and four retail prices; and we set the action space to a grid comprising the same four retail prices and four order quantities. We set the order quantity grid points to the four most frequent orders. We set the retail price grid points to the centroid points of four retail price clusters (created with the k-means clustering algorithm) and set the demand state grid points analogously. Finally, we arrange the inventory grid points to lie with logarithmic density, because the value function is steeper for smaller inventory levels.

4. We define $\widehat{\mathrm{E}}(d|s,p)$ as the fitted value of our IV regression and define $\widehat{\delta}_d$ as the negative binomial distribution that best fits the realized demands, given these demand expectations. That is, we suppose that a given demand is drawn from a negative binomial distribution with mean $\widehat{\mathrm{E}}(d|s,p)$ and estimate its scale parameter with maximum likelihood.

5. We define $\widehat{\delta}_f$ as the empirical distribution of the fulfillment indicator. That is, we set $\widehat{\delta}_f(1)$ to the number of shipments divided by the number of orders.

6. We define $\widehat{\delta}_s$ as the empirical distribution of tomorrow's demand state given today's demand state. That is, we round

all demand states in our sample to one of the four demand state grid points and define $\widehat{\delta}_s(s'|s)$ as the number of times the demand state jumped from grid point s to grid point s' divided by the number of grid-point-s observations in the sample.

7. We create a transition probability function estimate $\widehat{\zeta}$ by plugging $\widehat{\delta}_d$, $\widehat{\delta}_f$, and $\widehat{\delta}_s$ into (C.11) in such a way as to assign all of the probability mass to the 40 inventory grid points. To do so, we reassign mass that falls between consecutive grid points i_n and i_{n+1} to these grid points, in inverse proportion to their distance. For example, if the last two grid points were 90 and 100, then we would reroute $(100-92)/(100-90) = 80\%$ of the mass at point 92 to point 90 and reroute $(92-90)/(100-90) = 20\%$ to point 100.

8. Next, we create a function that identifies the optimal policy associated with input parameter vector θ.

a. We start with initial policy function guesses $\widehat{\phi}$ and $\widehat{\psi}$.

b. We plug $\widehat{\phi}$ and $\widehat{\psi}$, along with $\widehat{\delta}_d$, $\widehat{\delta}_f$, and $\widehat{\zeta}$, into (C.8), (C.9), and (C.10) to create estimates $\widehat{\pi}$, $\widehat{\omega}$, and $\widehat{\delta}$.

c. We plug $\widehat{\pi}$ and $\widehat{\delta}$ into (C.12) and solve for the corresponding fixed point solution, $\widehat{\nu}$.

d. We plug $\widehat{\psi}$, $\widehat{\omega}$, $\widehat{\zeta}$, and $\widehat{\nu}$ into the right-hand sides of (C.13) and (C.14) to revise estimates $\widehat{\phi}$ and $\widehat{\psi}$.

e. We iterate steps 8(b)–8(d) until the difference between the consecutive values of $\widehat{\nu}$ is smaller than 10^{-10}.

f. We return $\widehat{\phi}_\theta = \widehat{\phi}$ and $\widehat{\psi}_\theta = \widehat{\psi}$.

9. We define $\mathcal{L}(\theta)$ as the sample average value of $\ln(\widehat{\phi}_\theta(q|x)) + \ln(\widehat{\psi}_\theta(p|q,x))$.

10. We set $\widehat{\theta}$ to the parameter vector that maximizes $\mathcal{L}(\theta)$. We solve this maximization problem with the Berndt–Hall–Hall–Hausman algorithm.

Endnotes

[1] For example, among all articles published in INFORMS journals since January 1, 1990, we count 581 that mention the phrase "supply chain" in their title, abstract, or keywords; 439 that mention either "revenue management" or "dynamic pricing;" and 6 that mention both "supply chain" and either "revenue management" or "dynamic pricing."

[2] In Propositions 1–3, we measure sales and shipment volatilities with the coefficient of variation and we measure the bullwhip effect with the shipment coefficient of variation minus the sales coefficient of variation.

[3] In Proposition 4, we measure the bullwhip effect with the squared root of the days of inventory that comprise each shipment. This measure is the continuous-time limit of the shipment coefficient of variation minus the sales coefficient of variation. See Appendix B for details.

[4] Two comments on guiding prices: First, because local demand variation generally exceeds global demand variation, one would expect guiding prices to be stickier than retail prices. Indeed, the modal across-stores price—a rough proxy for the guiding price—lasts between 6 and 76 days more than the average individual-store price, across product categories. Second, both the corporate planners and the local store managers significantly influence store prices. Specifically, regressing prices on date dummies, item-by-item, we get a median R^2 of 62.6%. Hence, roughly 62.6% of the within-item variation in prices is common across stores—which suggests a sizeable corporate planner effect—and the remaining 37.4% of price variation is attributable to between-store heterogeneity—which suggests a sizeable local manager effect.

[5] These four filters reduce the sample by 20%, 82%, 0.001%, and 1.8%, respectively. However, our results are not sensitive to the restrictiveness of these filters; we get basically the same results when we double the four threshold values (i.e., when we drop products with fewer than 400 daily observations, fewer than 40 orders, fewer than 50 days with an inventory change, or more than 40% of shipments arriving later than one day).

[6] These three filters reduce the sample by 34%, 4.8%, and 9.2%, respectively. However, our results are not sensitive to the restrictiveness of these filters; we get basically the same results when we double the three threshold values (i.e., when we drop product clusters with fewer than 200 retail price changes, fewer than 400 orders, or fewer than 2 distinct order quantities with at least 20 corresponding orders). See the online appendix for details.

[7] Aguirregabiria (1999, p. 298) reported lagged demand estimates in his fifth table but only as a robustness check; he did not incorporate these lagged demands as state variables in his dynamic program.

[8] As we discuss in Section 4, retail prices are influenced by both corporate and local managers. In this work, we view the entire process as our subject of inquiry and set the level of abstraction to the store and the prices it posts. But an interesting avenue for future research would be to open the black box and explore the microlevel workings of the price selection process.

[9] Note that we assume that the demand state variable resolves independently of the current price. Thus, we implicitly rule out strategic customer behaviors, supposing that customers do not factor future prices into their purchasing decisions. Fortunately, we can test the reasonableness of this assumption empirically, because we know the trajectory of prices over time. We discuss this point further in Section 7.

[10] Before rigorously testing the supply chain implications of dynamic pricing with our counterfactual comparisons, we also check our theoretical model's predictions with a simple reduced-form analysis. In summary, the analysis suggests that when product prices at the Chinese supermarkets are updated more frequently than usual, (i) revenues are more stable than usual, (ii) sales are higher than usual, and (iii) orders are (weakly) higher than usual. Although these correlations do not have a causal interpretation—because the underlying comparisons are not ceteris paribus—they are worth reporting, for completeness. See the online appendix for details.

[11] This trend is exacerbated by the fact that every shipment must be an integer multiple of some standard lot size, a restriction that limits the stores' ordering flexibility. For example, the standard lot size of item 11251—a 250 mL carton of "fresh, orange-flavor orange juice"—is 24 units. So when the average sales rate increases from 1.62–1.95 units per day, following a removal of menu costs, the stores cannot commensurately increase their order quantities, because they must swallow an additional 24 units to receive any more.

[12] You may wonder whether there is a simple way to predict the product clusters that see increasing sales in the counterfactual. Unfortunately, we cannot find one. See the online appendix for details.

[13] The key difference between Hendel and Nevo's (2013) exercise (which detects value to strategic behavior) and ours (which does not detect value to strategic behavior) is that Hendel and Nevo assume *perfect foresight*: consumers in their model know exactly when a product's price will change next and what the new price level will be after the change. This assumption is crucial. To see why, suppose that the price of a product is \$1 for exactly one randomly chosen second every week and is \$100 for the remaining 604,799 seconds. With perfect foresight, consumers stand to gain 99% on the product's price by being strategic. But without perfect foresight, the value of being strategic is essentially nil—there is only a $1/604{,}800$ chance of arriving in the correct second—despite a 99% price reduction every week.

[14] We focus on the case when lead time $L = 1$ and informally interpret a *period* as the time between two orders (approximately 15 days in our data). We do so because we want a model with at most one batch of pipeline inventory at any given time. A model with $L = 15$

allows for 15 distinct batches of pipeline inventory, which is not the case in our context. So it is more accurate to set $L = 1$ and interpret the length of *a period* to be around 15 days. In fact, the original DSP model, developed by Kahn (1987), has $L = 1$, with each period representing one month. (Lee et al. 1997 then extended their specification to permit multiple shipping lead times.)

[15] We derive this expression with the identity $\mathrm{E}(\sum_{t\geq 0}\beta^t \cdot s_t | s_{-1}) = a/(1-b)(1-\beta)$.

[16] We have made our dynamic program relatively small because we have to solve it so many times. Specifically, each estimation requires us to solve around 25 dynamic programs; we have to estimate our model for 308 product clusters, and we have to bootstrap each estimate 30 times. All told, that is around $25 \cdot 308 \cdot 30 = 231{,}000$ dynamic programs that need solving. We solve these dynamic programs in parallel on an AWS computer with 64 cores and 256 GiB of RAM.

[17] It can happen—for example, because the instruments happen to be weak for a product cluster—that some of the IV price coefficients are positive. To correct this issue, we run a set of "back-up" ordinary least squares regressions (which include the same set of controls as the IV regressions) and, whenever an IV price coefficient is negative, we replace it with the corresponding ordinary least squares price coefficient. We adjust intercepts accordingly. After applying this correction, all product clusters have downward-sloping demand curves. See the online appendix for details.

References

Aguirregabiria V (1999) The dynamics of markups and inventories in retailing firms. *Rev. Econom. Stud.* 66(2):275–308.

Aguirregabiria V, Mira P (2010) Dynamic discrete choice structural models: A survey. *J. Econometrics* 156(1):38–67.

Anderson E, Jaimovich N, Simester D (2015) Price stickiness: Empirical evidence of the menu cost channel. *Rev. Econom. Statist.* 97(4):813–826.

Barba J (2019) Best Buy bets on digital price labels. *Shopper Marketing* https://shoppermarketingmag.com/best-buy-bets-digital-price-labels.

Bitran G, Caldentey R (2003) An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* 5(3):203–229.

Bray R, Yao YO, Duan Y, Huo J (2019) Ration gaming and the bullwhip effect. *Oper. Res.* 67(2):453–467.

Çelik S, Muharremoglu A, Savin S (2009) Revenue management with costly price adjustments. *Oper. Res.* 57(5):1206–1219.

Chan LMA, Shen ZJM, Simchi-Levi D, Swann JL (2004) Coordination of pricing and inventory decisions: A survey and classification. Simchi-Levi D, Wu SD, Shen ZJ, eds. *Handbook of Quantitative Supply Chain Analysis*, International Series in Operations Research & Management Science, vol. 74 (Springer, Boston), 335–392.

Chen M, Chen Z-L (2015) Recent developments in dynamic pricing research: Multiple products, competition, and limited demand information. *Production Oper. Management* 24(5):704–731.

Chen X, Hu P (2012) Joint pricing and inventory management with deterministic demand and costly price adjustment. *Oper. Res. Lett.* 40(5):385–389.

Chen X, Simchi-Levi D (2012) Pricing and inventory management. *The Oxford Handbook of Pricing Management*, vol. 1 (Citeseer), 784–824.

Chen F, Song J-S (2001) Optimal policies for multiechelon inventory problems with Markov-modulated demand. *Oper. Res.* 49(2): 226–234.

Chen X, Zhou SX, Chen Y (2011) Integration of inventory and pricing decisions with costly price adjustments. *Oper. Res.* 59(5): 1144–1158.

Elmaghraby W, Keskinocak P (2003) Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Sci.* 49(10): 1287–1309.

Harrison N, Faigen G, Brewer D (2018) Why Amazon's grocery store may not be the future of retail. *Harvard Bus. Rev.* 1–5.

Hendel I, Nevo A (2006) Measuring the implications of sales and consumer inventory behavior. *Econometrica* 74(6):1637–1673.

Hendel I, Nevo A (2013) Intertemporal price discrimination in storable goods markets. *Amer. Econom. Rev.* 103(7):2722–2751.

Kahn JA (1987) Inventories and the volatility of production. *Amer. Econom. Rev.* 77(4):667–679.

Lee HL, Padmanabhan V, Whang S (1997) Information distortion in a supply chain: The bullwhip effect. *Management Sci.* 43(4): 546–558.

Levy D, Bergen M, Dutta S, Venable R (1997) The magnitude of menu costs: Direct evidence from large U.S. supermarket chains. *Quart. J. Econom.* 112(3):791–825.

Netessine S (2006) Dynamic pricing of inventory/capacity with infrequent price changes. *Eur. J. Oper. Res.*174(1):553–580.

Owen A, Trzepacz D (2002) Menu costs, firm strategy, and price rigidity. *Econom. Lett.* 76(3):345–349.

Rust J (1987) Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55(5):999–1033.

Stamatopoulos I, Chehrazi N, Bassamboo A (2019) Welfare implications of inventory-driven dynamic pricing. *Management Sci.* 65 (12):5741–5765.

Stamatopoulos I, Moreno A, Bassamboo A (2021) The effects of menu costs on retail performance: Evidence from adoption of the electronic shelf technology. *Management Sci.* 67(1):242–256.

Van den Boer A (2015) Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys Oper. Res. Management Sci.* 20(1):1–18.

Robert L. Bray is a professor in operations management at Northwestern University's Kellogg School of Management. He studies supply chain management and econometrics. He has a PhD in operations management from the Stanford Graduate School of Business and a BS in industrial engineering from the University of California, Berkeley.

Ioannis Stamatopoulos is a professor of operations management at the McCombs School of Business, The University of Texas at Austin. He studies technology and operations management. He holds a PhD in operations management from Northwestern's Kellogg School of Management and an MS in economics from the same institution, as well as a BS in mathematics from the University of Athens (Athens, Greece).