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

OPERATIONS RESEARCH
Articles in Advance, pp. 1–15
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Ration Gaming and the Bullwhip Effect

Robert L. Bray,[a] Yuliang Yao,[b] Yongrui Duan,[c] Jiazhen Huo[c]

[a] Department of Operations Management, Kellogg School of Management, Northwestern University, Evanston, Illinois 60208; [b] Department of Management, College of Business and Economics, Lehigh University, Bethlehem, Pennsylvania 18015; [c] Department of Management Science and Engineering, School of Economics and Management, Tongji University, 200092 Shanghai, China

Contact: r-bray@kellogg.northwestern.edu, https://orcid.org/0000-0003-2773-0663 (RLB); yuy3@lehigh.edu (YY); yrduan@tongji.edu.cn (YD); huojiazhen@tongji.edu.cn (JH)

Received: December 29, 2016
Revised: December 22, 2017; April 14, 2018
Accepted: May 10, 2018
Published Online in Articles in Advance: ■■ ■■, 2018

Subject Classifications: economics: econometrics; inventory/production: applications; inventory/production: multi-item/echelon/stage
Area of Review: Operations and Supply Chains

https://doi.org/10.1287/opre.2018.1774

Abstract. We model a single-supplier, 73-store supply chain as a dynamic discrete choice problem. We estimate the model with transaction-level data, spanning 3,251 products and 1,370 days. We find two interrelated phenomena: the bullwhip effect and ration gaming. To establish the bullwhip effect, we show that shipments from suppliers are more variable than sales to customers. To establish ration gaming, we show that upstream scarcity triggers inventory runs, with stores simultaneously scrambling to amass private stocks in anticipation of impending shortages. These inventory runs increase our bullwhip measures by between 6% and 19%, which corroborates the long-standing hypothesis that ration gaming causes the bullwhip effect.

Funding: This work was supported by the National Natural Science Foundation of China [Grants 71371139, 71528007, 71532015, and 71771179].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/opre.2018.1774.

Keywords: bullwhip effect • ration gaming • (s, S) inventory policies • dynamic discrete choice • empirical supply chain management • structural estimation

1. Introduction

Demand fluctuations amplify up a supply chain like the crack of a whip. Lee et al. (1997a) called this phenomenon the bullwhip effect, which they attributed to (i) demand signal processing, (ii) order batching, (iii) cost shock fluctuations, and (iv) ration gaming. Ration gaming is the only one of these four bullwhip drivers without empirical evidence. Empiricists have established the amplification properties of demand signal processing (Metzler 1941, Lovell 1961, Kahn 1987), order batching (Blinder et al. 1981, Blinder and Maccini 1991), and cost shock fluctuations (Maccini and Rossana 1984, Miron and Zeldes 1988, Eichenbaum 1989). However, no empirical study has shown that ration gaming increases supply chain variability. In fact, no empirical study has shown that ration gaming even exists.

We present the first hard evidence of ration gaming.[1] The phenomenon manifests as inventory runs, the supply chain analog of bank runs. The context is a Chinese grocery supply chain that spans one upstream distribution center (DC) and 73 downstream stores. If the stores were self-sacrificing, they would curtail their orders when the DC's inventory runs low, scrimping for those in need. However, they are self-serving, and therefore, they accelerate their orders, stockpiling inventory to hedge against a potential upstream stock out. We estimate that these inventory runs account for about one-tenth of the bullwhip effect.

2. Stylized Ration Gaming Model

2.1. Positioning

Lee et al. (1997a) theorize that stores may game the means by which inventory is rationed. In addition to competing for customer demand, retailers must compete for vendor supply. Thus, stores will jockey for stock in times of scarcity—they will request excess inventory when they anticipate curtailed shipments, hoping to end up with the desired amount of product. These inflated orders amplify supply chain volatility, exacerbating the bullwhip effect.

Cachon and Lariviere (1999) show that a supplier can obviate this subterfuge by adopting a lexicographic allocation rule, ranking the stores and fulfilling their orders sequentially. This policy is truth inducing, because overordering under lexicographic allocation only earns a store inventory that it does not want. The DC that we observe follows a lexicographic allocation rule and thus, is immune to the "strategic manipulation" of Lee et al. (1997a) and Cachon and Lariviere (1999).[2]

However, the DC suffers another supply chain malady: inventory runs. The truth-inducing lexicographic allocation rule does not prevent inventory runs, because the impulse to hoard is not a lie—the stores are submitting inflated orders, because they want inflated shipments. These inventory runs represent a new version of ration gaming. The ration gaming of Lee et al. (1997a) pertains to an order's size; our ration gaming pertains to an order's timing. The ration gaming of Lee et al. (1997a) pertains to deceit (a store ordering two weeks' worth of supply in hopes of receiving one week's worth); our ration gaming pertains to hoarding (a store ordering two weeks' worth of supply in anticipation of next week's bare shelves).

2.2. Analysis

N stores sell a single product. The product is in demand until stopping time T, after which it becomes obsolete. The obsolescence time is exponentially distributed with mean τ. The stores observe τ but not T—they cannot anticipate when the product will go out of fashion. Although the product is in demand, customer arrivals follow independent Poisson processes, with arrival intensity α. Each customer demands one unit of inventory. The stores incur inventory underage cost μ for each unit of unfulfilled demand and inventory overage cost η for each unit of obsolete stock held at time T.

The stores order inventory from a common upstream DC. At time zero, there is one unit of inventory at each store, and there are u_0 units at the DC. The DC does not receive additional supply, and therefore, the average store can sell at most $u_0/N+1$ products. There is no shipping lead time from the DC to the stores, and the DC fulfills the orders that it can, promptly and in full. However, if the sum of orders in a given instant exceeds upstream supply, the DC dispenses stock according to a lexicographic allocation rule, fulfilling orders to the fullest extent possible in a random sequential manner. Each store observes its own inventories and the DC's inventories but not the other stores' inventories.

The following proposition characterizes a symmetric equilibrium with inventory runs that instantaneously liquidate upstream supply (proofs are shown in the online appendix).

Proposition 1. *If $\frac{\eta}{\eta+\mu} \le \left(\frac{\alpha\tau}{1+\alpha\tau}\right)^2$, then there exists $M \in \mathbb{N}$, such that, for all $N \ge M$, there is a Nash equilibrium in which each store orders its inventory up to $\rho(u)$ when the DC has inventory u, where*

$$\rho(u) = \begin{cases} 1 & \text{if } u > N(\omega-1) \\ \omega & \text{otherwise} \end{cases},$$

$$\text{and} \quad \omega = \operatorname{floor}\left(\frac{\ln\left(\frac{\eta}{\mu+\eta}\right)}{\ln\left(\frac{\alpha\tau}{1+\alpha\tau}\right)}\right) \ge 2.$$

The DC inventory level falls steadily with demand until it reaches threshold $N(\omega-1)$, at which point it vanishes as each store simultaneously cashes out $\omega-1$ units.[3] This coordinated raid on upstream supply is an inventory run.

The following propositions establish that inventory runs occur with positive probability and are costly.

Proposition 2. *Under Proposition 1's equilibrium, the probability of an inventory run is* $\min\left(1, \left(\frac{N\alpha\tau}{1+N\alpha\tau}\right)^{u_0-N\omega}\right) > 0.$

Proposition 3. *An inventory run increases each store's expected costs by at least* $\eta\omega - \eta(1+\alpha\tau)\left(1-\left(\frac{\alpha\tau}{1+\alpha\tau}\right)^{\omega}\right) > 0.$

Inventory runs are inefficient, because they sacrifice the flexibility of pooled inventory and move stock downstream, where it is more costly to store.

2.3. Empirical Framing

Our model indicates that storing inventory in the DC is like storing money in the bank: it is only viable when the institution is solvent. Stores withdraw their inventory when the DC's liabilities—its expected future orders—exceed its assets—its on-hand supply. However, the inventory runs in our sample are less extreme than our bang-bang equilibrium implies, and henceforth, we more permissively define an inventory run as any increase in the rate of downstream orders in response to low upstream inventories.

3. Empirical Setting

3.1. Overview

We study the sixth largest supermarket chain in China, with revenues of \$4.53, \$4.75, and \$4.55 billion in 2012, 2013, and 2014, respectively. By the end of 2014, it had 1,719 convenience stores, 2,415 supermarkets, and 157 "hypermarkets" (combination department stores and grocery stores). We focus on the hypermarkets, because the retailer operates them, whereas it franchises the smaller stores. Specifically, we study the 73 hypermarkets fulfilled by the Shanghai DC. The stores are located in Shanghai, Anhui, and Jiangsu.

3.2. Personnel

Each store has one manager who operates with autonomy. Most have bachelor's degrees and multiple years of experience at the company. Each manager receives a base salary and a performance bonus, which depends on the store's (i) total sales, (ii) company brand product sales, (iii) gross profits, (iv) net profits, (v) operating costs, (vi) price discounts, (vii) stock outs, and (viii) shrinkage.

Each store manager oversees a team of around 15 directors. Each director manages the inventories of one or two product categories. Most of the directors also have bachelor's degrees and several years of retail experience. Each director receives a base salary and a performance bonus, which depends on the product category's (i) total sales, (ii) gross profits, (iii) net profits, and (iv) stock-out rate.

A vice president of the company manages the DC. Beneath him are five executives who oversee (i) an information department, (ii) a logistics department, (iii) a convenience store department, (iv) a supermarket department, and (v) a hypermarket department. The hypermarket department head supervises a team of directors who manage the distribution of specific product categories to our stores. Each director receives a base salary and a performance bonus, which depends on (i) the fill rate of orders sent to the DC's supplier, (ii) the fill rate of orders received from the stores, (iii) the inventory turnover rate, and (iv) the shrinkage rate.

3.3. Order Process

Every store orders inventory every day. The process starts with the information technology (IT) system (which was last upgraded in 2010 before our sample was collected). Each morning, it generates a recommended order quantity for each product. It derives these order recommendations from 10 product-level operational variables, such as the average sales over the prior 35 days, the season, the ideal safety stock, and the batch size (the DC inventory level is not among these 10 operational variables). The formula mapping these operational variables to recommended orders is complex and ad hoc.

Each store's directors place their orders a few hours later. They do so via a communal computer terminal in the store's back room. Each director logs on to the computer and opens a spreadsheet that details by product (i) the current inventory level at the store, (ii) the current inventory level at the DC, (iii) the recent sales rate, (iv) the current retail price, (v) the current promotion level, and (vi) the recommended order quantity. The director specifies the actual order quantity in the table's seventh column. This column is initially populated with the IT system's recommended orders, but the director fine-tunes it, manually adjusting around 10% of orders.

3.4. Shipment Process

The DC directors check the stores' orders once a day. The directors also observe the stores' inventories, prices, and sales, but they generally ignore this information. The directors allocate inventory lexicographically, fulfilling orders in the sequence received.

After the DC directors decide what each store gets, the DC pickers execute the fulfillment. The pickers traverse the storage racks, filling their carts with inventory to be shipped. The pickers attach to each package a barcode sticker that specifies the intended recipient store. When their carts are full, the pickers offload their inventory to a conveyor belt, which reads the barcodes and sorts the inventory into piles based on the recipient stores' geographic districts. Each pile of inventory is loaded onto a truck that delivers to each store in a geographic district.

A team of backroom receivers greets a truck when it arrives at a store. First, the receivers find and unload the packages intended for their store as specified by the shipment manifest (the various stores' inventories are usually jumbled in the truck). Second, the receivers scan each box's barcode, logging its delivery. Third, the computer prints out a transaction receipt, and the truck driver and a store receiver inspect and sign it. Fourth, the receivers unpack and floor-ready the merchandise. Fifth, the receivers walk the inventory to the appropriate storage shelves. Thus, for each product that they receive, the store receivers must (i) find it in the truck, (ii) scan it into the computer, (iii) confirm that it was recorded in the transaction receipt, (iv) unpack and floor-ready it, and (v) transport it to the shelving units. In our model, we call this incremental cost of receiving an additional product a "shipping cost" (Section 6.2).

4. Data

Our sample describes the logistics of 3,251 products (store items) with daily frequency.[4] Tables 1 and 2 provide summary statistics. From April 1, 2011 to December 31, 2014, we observe (i) sales, (ii) wholesale prices, (iii) retail prices, (iv) price promotions, (v) start-

Table 1. Panel Dimensions

	Stores	Items	Dates	Stores × items	Stores × dates	Items × dates	Stores × items × dates
Detergent	67	38	1,370	1,011	85,989	47,086	1,074,011
Drinks	64	27	1,370	681	80,557	32,514	691,056
Oil/vinegar	64	14	1,370	359	83,071	18,677	416,673
Oral care	46	10	1,370	175	57,871	12,756	187,855
Shampoo	64	16	1,370	378	82,390	20,687	435,268
Tissues	49	6	1,370	142	62,154	8,100	166,911
Toilet paper	67	12	1,370	505	87,561	16,059	620,435
Total	73	123	1,370	3,251	94,144	155,879	3,592,209

Notes. This table provides the count of distinct store, item, and date combinations by product category. Our sample comprises 3,251 products and 3.6 million observations.

Table 2. Variable Overview

	Inventory		Order		Inbound		Outbound		Price	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Stores										
Detergent	74.8	28.0	46.3	20.0	46.4	18.0	5.3	2.0	16.4	12.8
Drinks	312.0	45.0	92.3	24.0	89.8	24.0	10.0	3.0	12.5	7.9
Oil/vinegar	44.7	28.0	30.9	20.0	29.9	20.0	3.6	2.0	3.8	2.9
Oral care	97.4	54.0	87.2	54.0	88.5	54.0	6.0	4.0	6.1	6.4
Shampoo	40.9	24.0	24.4	12.0	24.5	12.0	2.9	2.0	30.3	15.9
Tissues	26.1	18.0	20.1	12.0	19.7	12.0	2.4	2.0	15.6	12.6
Toilet paper	180.4	47.0	61.1	24.0	58.6	24.0	9.3	4.0	7.2	6.3
DC										
Detergent	1,765.6	740.0	1,309.8	480.0	1,408.9	560.0	172.0	66.0	16.1	11.3
Drinks	2,770.1	1,148.0	3,098.2	900.0	3,346.0	900.0	232.0	96.0	11.3	7.1
Oil/vinegar	1,390.8	810.0	1,112.2	520.0	1,014.1	480.0	131.1	90.0	3.2	2.5
Oral care	2,963.1	1,920.0	2,804.0	1,440.0	2,819.7	1,440.0	247.8	144.0	5.4	5.9
Shampoo	1,451.6	840.0	1,294.0	900.0	1,306.0	780.0	127.8	72.0	29.4	13.8
Tissues	629.4	496.0	590.1	474.0	534.5	360.0	57.8	45.0	11.8	9.3
Toilet paper	5,639.5	3,598.0	4,268.8	2,400.0	5,556.2	3,000.0	450.3	200.0	6.2	5.1

Notes. This table reports the means and medians of five store variables and five DC variables measured with daily frequency at the product level. It expresses the price variables in Chinese renminbi and the remaining variables in physical units. The inventory variables correspond to start-of-day stock levels. The order variables correspond to the nonzero store-to-DC and DC-to-vendor orders. The inbound variables correspond to the nonzero DC-to-store and vendor-to-DC shipments. The outbound variables correspond to the nonzero sales and aggregate DC-to-store shipments. Additionally, the price variables correspond to the retail and wholesale prices net promotions.

of-day inventories at the store, (vi) start-of-day inventories at the DC, (vii) store-to-DC orders, and (viii) DC-to-store shipments.[5] From our sales and store inventory series, we create a limited demand series. Specifically, we set demand equal to sales when the store's inventory is above its bottom decile, and we treat demand as unobserved otherwise. Thus, our demand variable is only 90% populated: it takes missing values when there is a credible risk of demand censoring caused by inventory stock out.

Figure 1 illustrates the data of a representative product: a five-pack of Guben brand 250 g whitening laundry detergent sold in a Shanghai store. The figure illustrates six prominent features of our sample.

1. The store orders in fixed lot sizes: the distinct order quantities are 6, 12, 18, 24, and zero (i.e., no order). Sample-wide, each product's six most common order quantities account for 99.3% of orders.

2. The store and DC inventories follow (s_t, S_t) policies with unstable S_t and s_t: the span between the largest and smallest reorder points is eight times the standard lot size, and the span between the largest and smallest stock-up-to levels is nine times the standard lot size. Sample-wide, 58% of reorder point ranges and 74% of stock-up-to level ranges exceed five times the standard lot size.

3. The shipping lead time is one day: of the 99 shipments to the store, 95 arrive the following day. Sample-wide, 95.2% of shipments arrive within one day.

4. The DC generally fulfills orders fully or not at all: of the 104 nonzero orders placed, only two are partially fulfilled, with $s \notin \{0, q\}$. Sample-wide, only 1.2% of orders are partially fulfilled.

5. Prices are stable: the retail and wholesale price coefficients of variation are 0.013 and 0.028, respectively. Sample-wide, the median retail and wholesale price coefficients of variation are 0.076 and 0.036, respectively.

6. Sales are interdependent: there is a 22% correlation between today's sales and tomorrow's sales. Sample-wide, 97% of store items exhibit significantly positive sales autocorrelation.

5. Reduced Form Results

5.1. Ration Gaming

Ration gaming has two aspects: rationing—the DC curtailing shipments when its supply runs short—and gaming—the stores selfishly manipulating the inventory allocation scheme. Figure 2 depicts both phenomena.

First, the DC rations inventory. It fulfills 95% of orders when its inventory level is above the first decile but only 36% when its inventory level is below the first decile (overall, 13% of orders go unfulfilled). Additionally, a stint of rationing can last awhile: if we define a "rationing spell" as a span of time when the estimated order fulfillment probability is less than one-half, then 50% of rationing spells last at least 5 days, 10% last at least 14 days, and 1% last at least 35 days. Sample-wide, these rationing spells make up 10% of our sample and account for 74% of unfulfilled orders. Moreover, these rationing spells are predictable, because the DC inventory declines at a steady pace (Figure 1). For example, with just an intercept and today's DC inventory level, we can predict tomorrow's DC inventory level (excluding inbound shipments) with a median R^2 of 0.96.

Figure 1. Raw Data of Representative Product

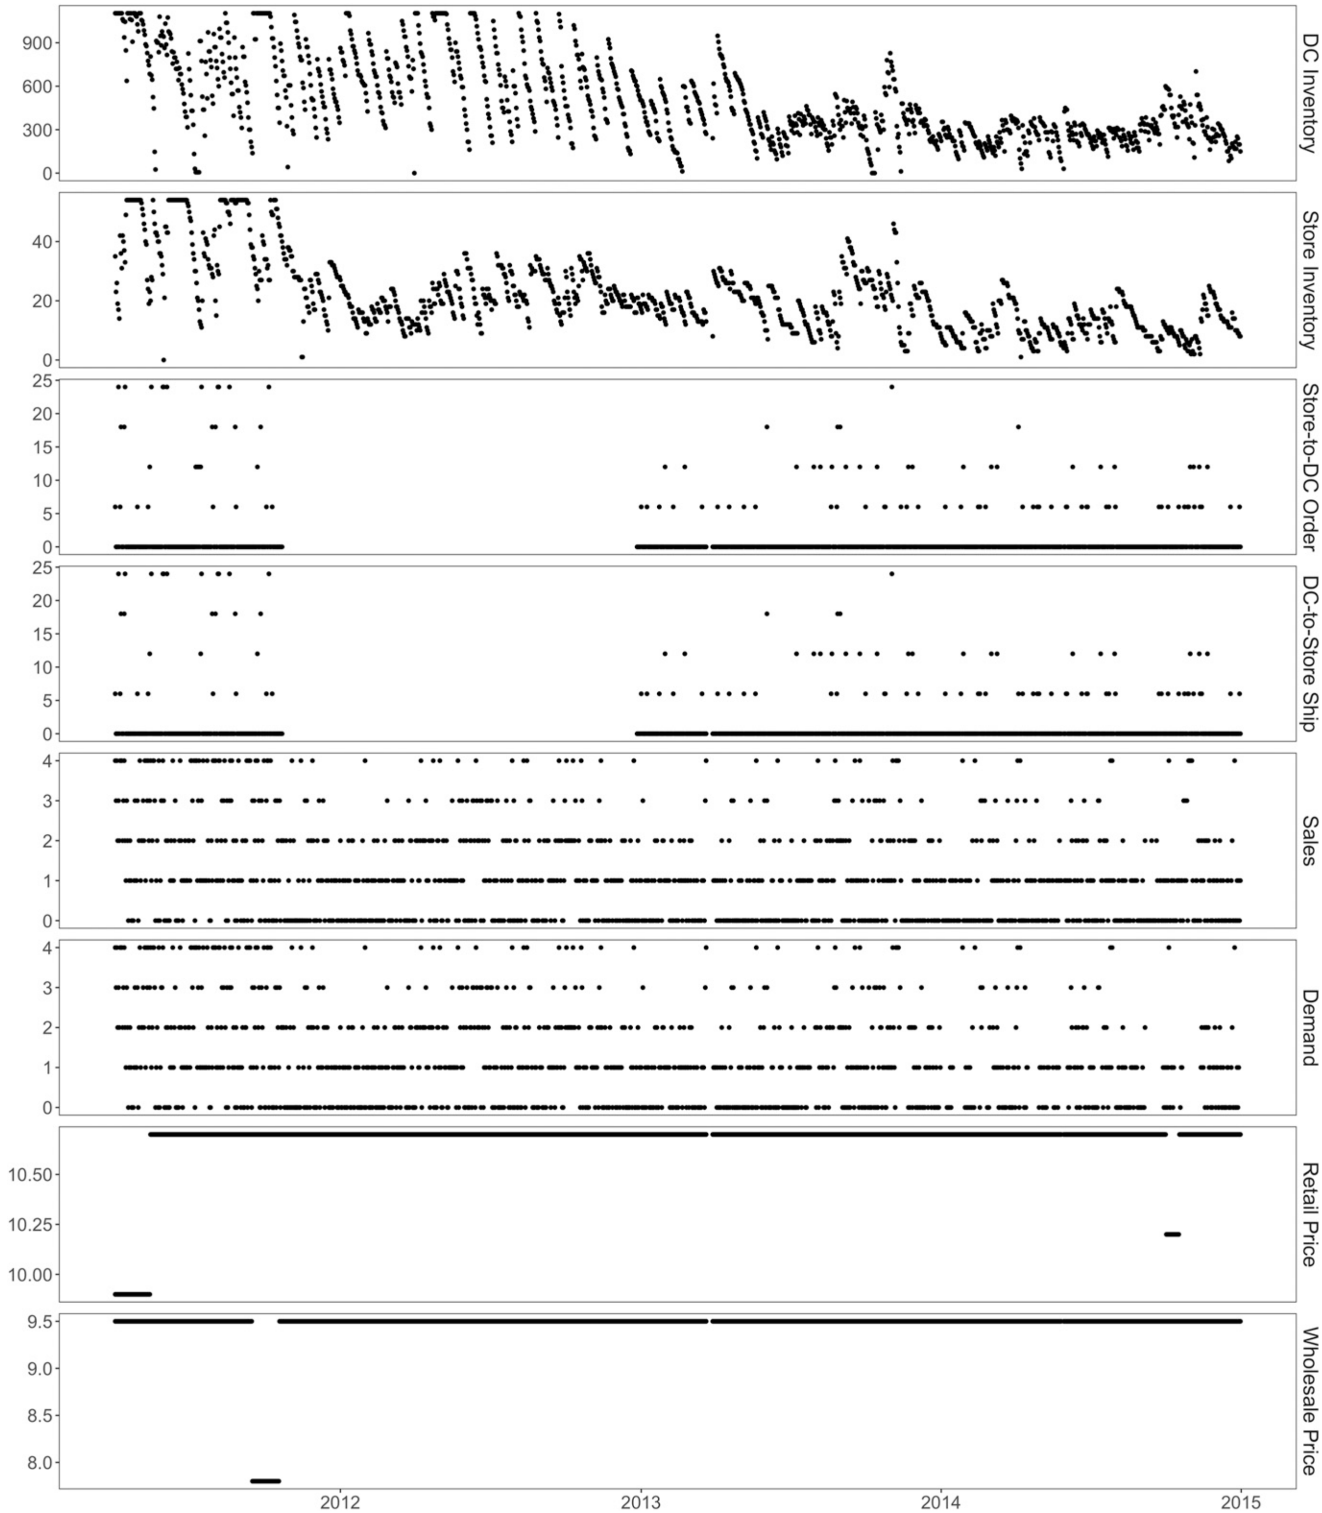

Notes. These scatter plots depict the raw data of a representative product: a five-pack of 250 g Guben brand whitening laundry detergent sold in a Shanghai store. The graphs denote the price variables in Chinese renminbi and the remaining variables in physical units. They depict the time series with daily frequency. The 2012 orders and shipments data were lost.

Second, the stores game the inventory rationing scheme. They have a 9.6% probability of ordering inventory when the DC inventory is above the first decile and a 12.8% probability of ordering inventory when the DC inventory is below the first decile (an increase of (12.8% − 9.6%)/9.6% = 31%). This phenomenon is

Figure 2. Signatures of Rationing and Gaming

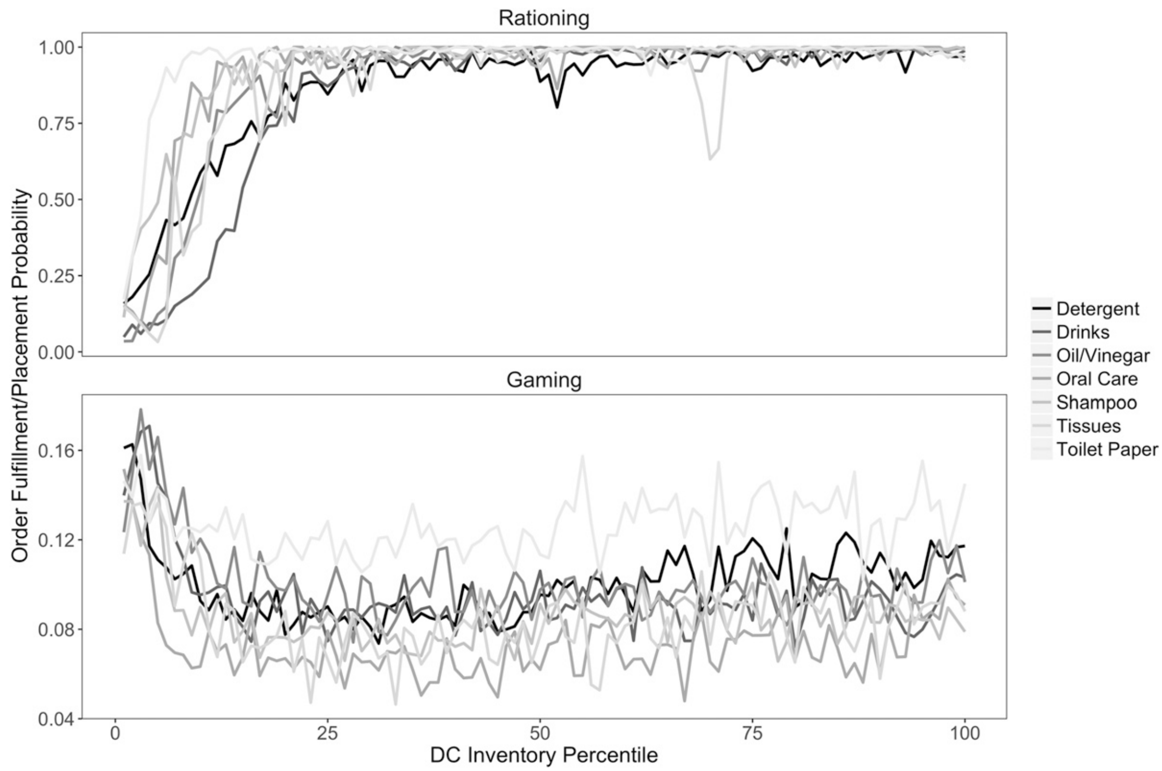

Notes. These line plots depict the degree of rationing and gaming by product category. The rationing plot graphs the order fulfillment probability (the fraction of orders that the DC fulfills) as a function of the DC inventory level. Additionally, the gaming plot graphs the order placement probability (the fraction of observations with a positive order quantity) as a function of the DC inventory level. We measure DC inventories in percentages with the empirical cumulative distribution. The probability of the DC fulfilling an order is lower when the DC inventory level is in the lowest decile, which implies rationing. In contrast, the probability of the store placing an order is higher when the DC inventory level is in the lowest decile, which implies gaming.

broad: 60 of the 61 stores that have at least 10,000 observations in our sample order more frequently when the DC inventory is in its lowest decile.

We establish the statistical significance of this result with 16 ordinary least squares (OLS) regressions. Each regression's dependent variable is a dummy that indicates whether a given store ordered a given item on a given day.[6] Each regression's primary independent variables are nine dummies that indicate the decile of the DC inventory level (we reserve the lowest inventory decile as our reference value). Additionally, each regression's control variables are an intercept, the store's current inventory level, and one of the $2^4 = 16$ combinations of the following four sets of variables.

1. Fixed effects: 72 store dummies, 122 item dummies, and 31 month dummies.

2. Sales: today's sales of the given item at the given store, the given item across all stores, and the given store across all items plus the average of these three variables across the prior week.

3. Price: the retail price, wholesale price, and discount rate of the given item at the given store plus the average of these three variables across the prior week.

4. Future: the average of the six contemporaneous sales and price variables across the subsequent week.

The future variables are proxies for their forecasted values. A potential source of endogeneity, next week's sales projection, will influence both the DC's inventory level and the store's order propensity. Incorporating the future (and past) sales should control for the variation in DC inventories caused by the variation in demand (and its expectation).

Table 3 reports the DC inventory coefficient estimates. The figures represent the store order probabilities when the DC inventory is in the nine highest deciles minus the store order probabilities when the DC inventory is in the lowest decile. All nine estimates across all 16 specifications are significantly negative: order probabilities are significantly higher when the DC inventory level is in the bottom decile.

In Table 4, we correlate the stores' gaming with the DC's rationing. We divide our sample into four equal subsamples based on the degree of inventory rationing. We control for the product category, and therefore, each subsample has roughly the same mix of items (e.g., each has four shampoo items and three toilet paper items). We then run Table 3's final regression—the one without fixed effects or the sales, price, and future control variables—across each subsample. We find that gaming worsens when rationing worsens. In fact, Table 4's

Table 3. Gaming Regression Estimates

	D2	D3	D4	D5	D6	D7	D8	D9	D10
F.E.									
Sales									
Price									
Future	−3.21 (0.23)	−3.59 (0.18)	−3.66 (0.23)	−3.33 (0.18)	−3.04 (0.18)	−2.96 (0.18)	−2.90 (0.20)	−2.82 (0.21)	−2.61 (0.26)
No future	−3.54 (0.22)	−3.92 (0.24)	−4.07 (0.22)	−3.69 (0.19)	−3.44 (0.24)	−3.29 (0.22)	−3.24 (0.23)	−3.11 (0.20)	−2.81 (0.24)
No price									
Future	−3.22 (0.16)	−3.59 (0.17)	−3.67 (0.18)	−3.34 (0.20)	−3.05 (0.22)	−2.97 (0.18)	−2.90 (0.21)	−2.82 (0.25)	−2.62 (0.19)
No future	−3.66 (0.19)	−3.98 (0.21)	−4.09 (0.22)	−3.69 (0.22)	−3.37 (0.21)	−3.23 (0.23)	−3.19 (0.22)	−3.00 (0.22)	−2.71 (0.23)
No sales									
Price									
Future	−3.23 (0.19)	−3.64 (0.20)	−3.72 (0.22)	−3.36 (0.26)	−3.11 (0.23)	−3.00 (0.23)	−2.96 (0.20)	−2.86 (0.18)	−2.67 (0.24)
No future	−3.21 (0.14)	−3.64 (0.12)	−3.70 (0.13)	−3.49 (0.14)	−3.23 (0.12)	−3.01 (0.15)	−2.94 (0.13)	−2.98 (0.16)	−2.57 (0.15)
No price									
Future	−3.06 (0.16)	−3.49 (0.21)	−3.62 (0.14)	−3.24 (0.18)	−3.00 (0.19)	−2.85 (0.19)	−2.85 (0.22)	−2.67 (0.19)	−2.48 (0.18)
No future	−3.21 (0.12)	−3.58 (0.13)	−3.60 (0.13)	−3.35 (0.13)	−3.05 (0.12)	−2.78 (0.12)	−2.72 (0.13)	−2.74 (0.14)	−2.28 (0.17)
No F.E.									
Sales									
Price									
Future	−3.24 (0.20)	−3.56 (0.26)	−3.55 (0.20)	−3.18 (0.23)	−3.01 (0.21)	−2.96 (0.24)	−2.93 (0.26)	−2.86 (0.24)	−2.72 (0.29)
No future	−3.57 (0.20)	−3.87 (0.20)	−3.95 (0.20)	−3.52 (0.20)	−3.44 (0.24)	−3.27 (0.23)	−3.26 (0.23)	−3.14 (0.23)	−2.91 (0.28)
No price									
Future	−3.26 (0.20)	−3.57 (0.16)	−3.57 (0.18)	−3.20 (0.19)	−3.01 (0.18)	−2.96 (0.18)	−2.94 (0.21)	−2.86 (0.20)	−2.73 (0.18)
No future	−3.68 (0.22)	−3.92 (0.19)	−3.98 (0.25)	−3.52 (0.19)	−3.38 (0.27)	−3.22 (0.26)	−3.24 (0.21)	−3.03 (0.28)	−2.81 (0.27)
No sales									
Price									
Future	−3.29 (0.20)	−3.64 (0.22)	−3.65 (0.22)	−3.26 (0.21)	−3.12 (0.21)	−3.03 (0.22)	−3.03 (0.18)	−2.93 (0.21)	−2.84 (0.26)
No future	−3.38 (0.12)	−3.83 (0.11)	−3.89 (0.13)	−3.66 (0.11)	−3.44 (0.15)	−3.17 (0.12)	−3.07 (0.14)	−3.08 (0.15)	−2.64 (0.14)
No price									
Future	−3.07 (0.22)	−3.39 (0.19)	−3.50 (0.22)	−3.01 (0.17)	−2.88 (0.22)	−2.74 (0.19)	−2.76 (0.24)	−2.63 (0.22)	−2.43 (0.18)
No future	−3.35 (0.14)	−3.72 (0.13)	−3.76 (0.14)	−3.49 (0.12)	−3.23 (0.12)	−2.93 (0.13)	−2.85 (0.14)	−2.84 (0.17)	−2.33 (0.16)

Notes. This table presents the coefficient estimates of 16 OLS regressions. For each, the dependent variable is a dummy that indicates whether a given store ordered a given item on a given day, and the primary independent variables are nine dummies that indicate the DC's inventory decile of the given item on the given day (the lowest inventory decile does not have a dummy, because it serves as the benchmark). Columns D2–D10 report the coefficient estimates of the second through 10th inventory decile dummies. These estimates report the probability of a store placing an order when the DC's inventory is in the second through 10th decile minus the probability of the store placing an order when the DC's inventory is in the first decile. For example, the top left estimate suggests that stores are, on average, 3.21% less likely to order when the DC's inventory is in the second lowest decile than when the DC's inventory is in the lowest decile. The different rows correspond to different regression specifications. Each includes an intercept and the store's inventory level as control variables. Additionally, the "F.E." regressions include store, item, and month fixed effects. The "Sales" regressions include the sales of the given item at the given store, the sales of the given item across all stores, the sales at the given store across all items, and the average of these variables over the previous week. The "Price" regressions include the retail price, the wholesale price, the discount level, and the average of these variables over the previous week. Additionally, the "Future" regressions include the following week's average retail price, wholesale price, discount level, sales of the given item at the given store, sales of the given item across all stores, and sales at the given store across all items. We calculate the standard errors (in parentheses) with the bootstrap, sampling by product cluster. Each estimate is significantly negative, which indicates that order probabilities are significantly higher when the DC inventory level is in the lowest decile.

columns are perfectly ordered: the degree of gaming moves in lockstep with the degree of rationing.

To corroborate these empirical results, we interview the managers of two Shanghai stores. Store 1 is in a famous mall in Puxi, and Store 2 is at the intersection of two busy subway lines in Pudong. Store 1 is the 67th-largest store in our sample, and Store 2 is the 39th largest. We interviewed Store 1's manager by phone on August 26, 2017, and September 1, 2017, and Store 2's manager in person on October 15, 2017. Both managers admitted that their directors place orders in response to low DC inventory levels. They were forthcoming about their directors' strategic stockpiling, because they did not view it as antisocial. Instead, they considered assurance of supply a critical aspect of the job; according to them, not monitoring upstream availability would be negligent. Both managers speculated that the other stores also time orders to avoid upstream stock outs, and neither manager could think of another reason for order quantities to correlate with DC inventory levels after conditioning on demand.

5.2. Bullwhip Effect

In addition to ration gaming, our supply chain exhibits the bullwhip effect. Following Cachon et al. (2007), Bray and Mendelson (2012), Chen and Lee (2012), and Bray and Mendelson (2015), we measure the "material flow" bullwhip effect with the SD of DC-to-store shipments divided by the SD of sales, and we measure the "information flow" bullwhip effect with the SD of store-to-DC

Table 4. Gaming Regression Estimates Moderated by Degree of Rationing

	D2	D3	D4	D5	D6	D7	D8	D9	D10
Q1	−1.27 (0.15)	−1.38 (0.15)	−1.28 (0.13)	−1.36 (0.15)	−1.27 (0.14)	−1.26 (0.16)	−0.74 (0.17)	−0.49 (0.17)	−0.02 (0.18)
Q2	−2.15 (0.22)	−2.43 (0.29)	−2.28 (0.27)	−2.09 (0.23)	−2.17 (0.24)	−1.92 (0.25)	−1.68 (0.25)	−1.93 (0.33)	−0.92 (0.42)
Q3	−4.34 (0.40)	−4.14 (0.37)	−4.02 (0.34)	−3.27 (0.30)	−3.04 (0.32)	−2.72 (0.31)	−2.76 (0.33)	−2.92 (0.38)	−2.78 (0.42)
Q4	−5.71 (0.47)	−6.91 (0.54)	−7.34 (0.56)	−7.15 (0.58)	−6.22 (0.52)	−5.53 (0.47)	−5.89 (0.51)	−5.76 (0.51)	−4.96 (0.46)

Notes. This table presents the coefficient estimates of one regression model applied to four subsamples, Q1 to Q4. We create the subsamples by dividing each product category's items into quartiles based on the degree of inventory rationing, where subsample Q1 comprises the least-rationed items and subsample Q4 comprises the most-rationed items. We measure rationing with the fraction of orders that go unfulfilled. We run Table 3's simplest regression specification—without fixed effects and without sales, price, or future control variables—across each subsample. We calculate standard errors with the bootstrap, in the fashion of Table 3. More-negative estimates indicate more-extreme gaming, so the degree of gaming decreases from subsample Q4 to subsample Q1. Thus, we find a positive association between rationing and gaming.

orders divided by the SD of demand. A bullwhip is present when these ratios exceed one. Figure 3 illustrates a robust bullwhip effect: 97% of our material flow ratios and 98% of our information flow ratios exceed one.[7] At the product level, the median shipment SD is 4.6 times the median sales SD, and the median order quantity SD is 4.9 times the median demand SD.

Order batching accounts for most of this effect. For example, Figure 4 illustrates that a strawberry milkshake's three most common sales quantities are 1 (133 observations), 0 (131 observations), and 2 (111 observations), whereas its three most common shipment quantities are 0 (509 observations), 9 (30 observations), and 45 (12 observations). Although material, however, order batching's contribution to the bullwhip effect is moot, because it is unavoidable: there will always be order batching as long as there are shipping costs.

More interesting is the bullwhip caused by ration gaming, because that problem is solvable. To confirm the hypothesis of Lee et al. (1997a) that ration gaming contributes to the bullwhip effect, we calculate the magnitude of the phenomenon without inventory runs. Additionally, to perform this counterfactual analysis, we construct a structural econometric model of the supply chain.

6. Structural Econometric Model

6.1. Positioning

Our empirical inventory model belongs to the (s_t, S_t) class: its inventories exhibit the saw-toothed pattern characteristic of (s, S) policies, but its reorder point and order-up-to level vary dynamically. We create the fourth microeconometric (s_t, S_t) inventory model, building on those of Aguirregabiria (1999), Erdem et al. (2003), and Hendel and Nevo (2006).[8] We extend their specifications in five ways.

1. Our model describes a two-tier supply chain, whereas the other models describe a single stock of inventory.

Figure 3. Bullwhip Effect Estimates

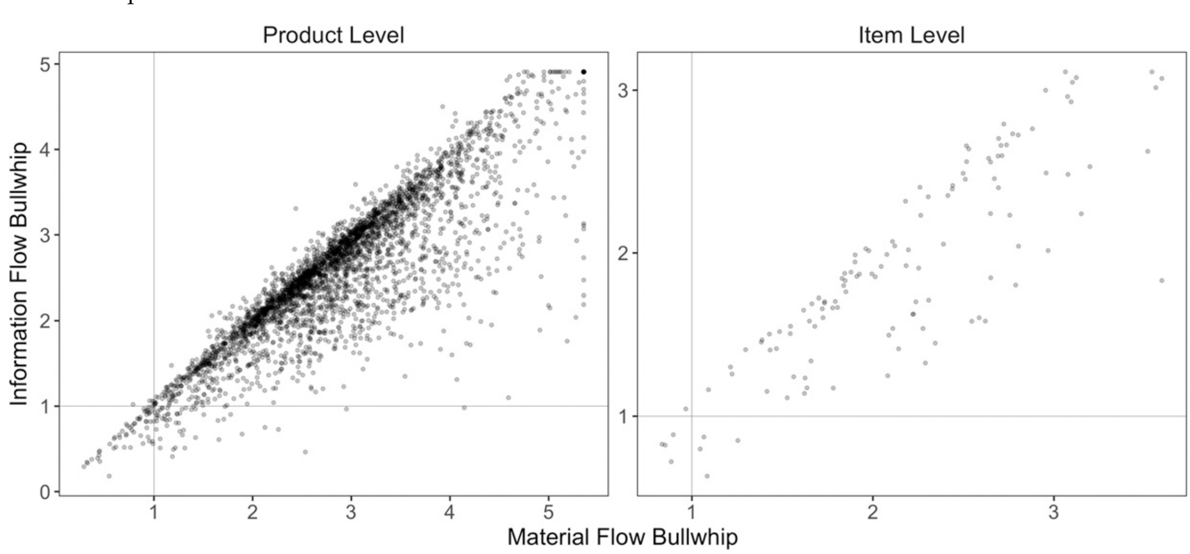

Notes. These scatter plots depict our bullwhip effect estimates both disaggregated by product and aggregated across stores by item. The material flow bullwhip is the SD of shipments divided by the SD of sales, and the information flow bullwhip is the SD of orders divided by the SD of demand. Overall, 97% of these ratios exceed one, which indicates a strong bullwhip effect.

Figure 4. Bullwhip Effect from Order Batching

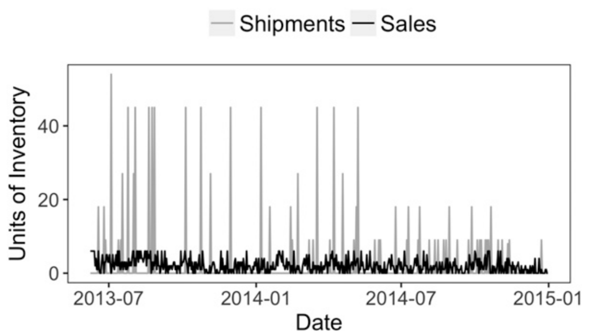

Notes. This line plot illustrates the sales and shipments of a representative product: a four-pack of 125mL Wangwang brand strawberry milkshake sold in a store in Yancheng, Jiangsu. The SD of shipments is 4.36 times that of sales, and therefore, the product exhibits the bullwhip effect. However, most of this bullwhip stems from order batching (e.g., the distinct sales quantities are {0, 1, 2, 3, 4, 5, 6}, and the distinct shipment quantities are {0, 9, 18, 27, 45, 54}).

2. Our estimator factors both reorder point s_t and order-up-to level S_t, whereas the other estimators disregard the latter. Erdem et al. (2003) and Hendel and Nevo (2006) do not observe S_t, and Aguirregabiria (1999) does not incorporate it in his likelihood function, because doing so would make his statistical model degenerate. Aguirregabiria (1999, p. 293) "exploit[s] moment conditions associated with the optimal discrete choice [of whether to order (i.e., s_t)], but not moment conditions associated with the marginal conditions [of how much to order (i.e., $S_t - s_t$)]," because his order quantity—a continuous variable observed without error—is too informative (theoretically, he could determine his model primitives from a handful of order quantities). To avoid this model degeneracy, we exploit order batching. Each order in our data is an integer multiple of some standard lot size. Thus, we do not observe ideal order-up-to level—we observe ideal order-up-to level rounded to the nearest batch. This rounding error prevents our empirical likelihood function from becoming overdetermined.

3. Our model incorporates the general Markov-modulated demand process of Chen and Song (2001), whereas the other models assume demand to be *i.i.d.*

4. We prove that our model is empirically identified, but the other authors do not. In fact, Hendel and Nevo (2006, p. 1653) concede that they "have no reason to believe that costs and preferences are identified nonparametrically or even that flexible functional forms can be estimated," and Erdem et al. (2003, pp. 52–53) explain that their "model is too complex for [them] to provide analytic results on identification" and find with simulations that they "cannot separately identify" the linear cost of holding inventory and the transport cost.

5. We describe our data with 246 dynamic programs, whereas Erdem et al. (2003) use only four and Aguirregabiria (1999) and Hendel and Nevo (2006) use only one. We further customize our model, because we have a richer sample with more products (N) and time periods (T):

Sample	N	T
The sample of Aguirregabiria (1999)	534	29
The sample of Hendel and Nevo (2006)	218	104
The sample of Erdem et al. (2003)	838	123
Our sample	3,251	1,370

6.2. Overview

We now present our supply chain model. The supply chain comprises a single DC and multiple stores. An external vendor ships a nonperishable product to the DC, the DC supplies the stores, and the stores satisfy local demands. The stores compete with one another for access to the DC's inventory, hoarding stock when they anticipate a shortage.

We frame our model from the perspective of a representative store. The store faces newsvendor-style inventory costs—unsatisfied demands are lost, and unsold stocks are stored at a cost—and an economic order quantity (EOQ)-style shipping cost—each inventory delivery has a fixed fee.[9] Comprising a Markov decision process, the store's order quantity $q \in \mathbb{q} = \{q_0, \cdots, q_{\bar{q}}\}$ is a deterministic function of four state variables: store inventory $i \in \mathbb{i} = \{0, \cdots, \bar{i}\}$, DC inventory $u \in \mathbb{u} = \{0, \cdots, \bar{u}\}$, demand state $m \in \mathbb{m} = \{m_0, \cdots, m_{\bar{m}}\}$, and shipping cost shock $e = \{e(q) \mid q \in \mathbb{q}\} \in \mathbb{R}^{|\mathbb{q}|}$. Variable m is a sufficient statistic for the distribution of future demands (Chen and Song 2001), and variable e is a vector of *i.i.d.* mean-zero Gumbel random variables that shifts the cost associated with each potential shipment size. The store observes e, but we do not—it is our statistical error term. However, we observe the other state variables, which we house in vector $x = [i, u, m]'$. Thus, the observable portion of the state space is $\mathbb{x} = \mathbb{i} \times \mathbb{u} \times \mathbb{m}$. Henceforth, tomorrow's variables wear a prime symbol (e.g., m'), and today's stand bare (e.g., m).

6.3. Sequence of Events

Today's events proceed as follows.

1. The day begins in demand state $m \in \mathbb{m}$, with downstream inventory $i \in \mathbb{i}$ at the store and upstream inventory $u \in \mathbb{u}$ at the DC.

2. Shipping cost shifter $e \in \mathbb{R}^{|\mathbb{q}|}$ resolves independent of the other model variables.

3. The store orders q units of inventory from the DC in response to information set $\{x, e\}$. (Note that setting $q = 0 \in \mathbb{q}$ is equivalent to not placing an order.)

4. Demand $d \in \mathbb{N}$ resolves from probability mass function (PMF) $\delta^d(d \mid m)$, and the store sells $\min(i, d)$ units of inventory.

5. The store incurs newsvendor cost $\mu \max(d - i, 0) + \eta \max(i - d, 0)$.

6. Boolean $b \in \{0, 1\}$ resolves from PMF $\delta^b(b \mid u)$, and the DC ships $s = bq$ units to the store: the DC fulfills orders fully or not at all.

7. The store incurs shipping cost $\lambda \mathbb{1}(s>0)-e(s)$, where λ is the average cost of receiving a shipment and $-e(s)$ is a shock to the cost of receiving a shipment of size s.[10]
8. The store's inventory transitions to $i'=i-\min(i,d)+s$; in other words, $i' \in \mathbb{i}$ resolves from PMF $\delta^i(i'|d,i,s)=\mathbb{1}(i'=i-\min(i,d)+s)$.
9. The DC's upstream inventory $u' \in \mathbb{u}$ resolves from PMF $\delta^u(u'|u,s)$.
10. Demand state variable $m' \in \mathbb{m}$ resolves from PMF $\delta^m(m'|d,m)$.

6.4. Primitives

The following state transition function characterizes the distribution of tomorrow's observable state conditional on today's state and order quantity:

$$\delta(x'|x,q) = \sum_{b\in\{0,1\}} \sum_{d\in\mathbb{N}} \delta^b(b|u)\delta^d(d|m)\delta^i(i'|d,i,bq)\delta^u(u'|u,bq) \cdot \delta^m(m'|d,m).$$

Additionally, the following cost function specifies today's expected cost conditional on today's state and order quantity:

$$\begin{aligned} \phi(q|x,e) &= \pi(q|i,m) - \varepsilon(q|e), \\ \text{where} \quad \pi(q|x) &= \lambda\delta^b(1|u)\mathbb{1}(q>0) \\ &\quad + \sum_{d\in\mathbb{N}} \delta^d(d|m)(\mu\max(d-i,0) \\ &\quad + \eta\max(i-d,0)) \\ \text{and} \quad \varepsilon(q|e) &= \sum_{b\in\{0,1\}} \delta^b(b|u)e(qb). \end{aligned}$$

6.5. Value Function

The following equations characterize the store's Bellman equation (Aguirregabiria and Mira 2010, Arcidiacono and Ellickson 2011):

$$\begin{aligned} v(x) &= \mathrm{E}\left(\min_{q\in\mathbb{q}} \phi(q\,|\,x,e) + \beta\sum_{x'\in\mathbb{x}} \delta(x'\,|\,x,q)v(x')\middle|x\right) \\ &= \sum_{q\in\mathbb{q}} \rho(q\,|\,x)(\gamma(q\,|\,x)-\xi(q\,|\,x)), \end{aligned}$$

$$\begin{aligned} \text{where } \gamma(q\,|\,x) &= \pi(q\,|\,x) + \beta\sum_{x'\in\mathbb{x}} \delta(x'\,|\,x,q)v(x'), \\ \rho(q\,|\,x) &= \mathrm{E}\left(\mathbb{1}(q = \arg\min_{q_j\in\mathbb{q}} \gamma(q_j\,|\,x)-\xi(q_j\,|\,x))\middle|x\right) \\ &= \frac{\exp(-\gamma(q\,|\,x)/\delta^b(1\,|\,u))}{\sum_{q_j\in\mathbb{q}}\exp(-\gamma(q_j\,|\,x)/\delta^b(1\,|\,u))}, \\ \text{and} \quad \xi(q\,|\,x) &= \mathrm{E}\left(\varepsilon(q|e)|x,q\right) \\ &= -\delta^b(1\,|\,u)\log(\rho(q\,|\,x)). \end{aligned}$$

In the expressions above, v characterizes the store's expected discounted costs conditional on the current observable state, γ characterizes the store's expected discounted costs net the error term and conditional on the current observable state and order quantity, ρ specifies the probability of a given order quantity conditional on the current observable state, ξ represents the expected error term conditional on the current observable state and order quantity, and $\beta = 0.9997$ is the daily discount factor.

6.6. Empirical Identification

Our identification argument is straightforward. First, the stock-out rate identifies μ relative to η: the newsvendor model suggests that the service level should be around $\mu/(\mu+\eta)$. Second, the magnitude of orders identifies λ relative to η: the EOQ model suggests that the average nonzero order should be around $\sqrt{2\mathrm{E}(d)\lambda/\eta}$. Third, our model's predictive power identifies η relative to the SD of the error term (which is normalized to one): McFadden's R^2 decreases with the relative magnitude of the error term.

The following proposition formalizes these intuitions.

Proposition 4. *If $|\mathbb{q}| \geq 3$, then cost parameters $\theta = \{\lambda,\mu,\eta\}$ are empirically identified from functions δ and ξ, which we can estimate nonparametrically, and from scalar β, which we take as given.*

This identification argument does not extend to the specifications of Aguirregabiria (1999), Erdem et al. (2003), and Hendel and Nevo (2006), which treat order-up-to level S_t as unobserved.

7. Estimation Procedure

7.1. Overview

We estimate our dynamic discrete choice model with refinement by Bray (2018) of the nested pseudolikelihood estimator of Aguirregabiria and Mira (2002). We estimate across each item separately, and therefore, we define our estimator in terms of one representative item. We pool across our 73 stores with the expectation–maximization (EM) algorithm method of Arcidiacono and Miller (2011), which specifies a finite mixture of store types.

7.2. State and Action Spaces

Before we can define state space $\mathbb{x}$, we must define demand state variable m. We set m to the product's expected next day demand, which we estimate with the fitted value of product-level OLS regressions of next day demand on (i) monthly dummy variables, (ii) the sales of the given item at the given store, (iii) the sales of the given item across all stores, (iv) the sales across all items at the given store, (v) the listed wholesale price, (vi) the listed retail price, and (vii) the posted price discount.

After calculating m, we set state space $\mathbb{x}$ to a $20\cdot 15\cdot 10 = 3{,}000$-element grid spanning 20 values of i, 15 values of u, and 10 values of m. We round each state in our sample to the nearest grid element. We set the grid's breakpoints to the variables' empirical quantiles, and therefore, the data are evenly dispersed.

Next, we set action space $\mathbb{q}$ to the five most common order quantities and the median value of the remaining orders. For example, the Colgate Protective Toothbrush Family Pack has seven distinct order quantities:

Order quantity	0	24	48	72	120	240	744
Observation count	12,506	1,053	87	27	19	4	1

Therefore, for this item, we set $\mathbb{q} = \{0, 24, 48, 72, 120, 240\}$ and round the stray 744-unit order down to 240. We round fewer than 1% of orders in this fashion.

7.3. Unobserved Heterogeneity

We observe a given item in upward of 73 stores. We pool these stores into $\bar{k}$ latent classes.[11] The stores in class $k \in \mathbb{k} = \{1, \ldots, \bar{k}\}$ have cost parameters $\theta_k = \{\lambda_k, \mu_k, \eta_k\}$ and state transition function δ_k. We do not observe the class assignments, and therefore, the structural parameters have a finite mixture distribution from our perspective: a priori, a given store has probability p_k of belonging to type k.

Theoretically, we could estimate primitives $\theta = \{\theta_k\}$, $\delta = \{\delta_k\}$, and $p = \{p_k\}$ with the following maximum likelihood:

$$\{\widehat{\theta}, \widehat{\delta}, \widehat{p}\} = \underset{\{\theta,\delta,p\}}{\arg\max} \sum_{n\in\mathbb{n}} \log\Bigg(\sum_{k\in\mathbb{k}} \prod_{t\in\mathbb{t}} p_k \rho(q_{nt} \,|\, x_{nt}, \theta_k, \delta_k) \cdot \delta_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1})\Bigg),$$

where $\mathbb{n}$ is the set of stores and $\mathbb{t}$ is the set of days. However, this maximization problem is intractable, and therefore, we translate it to a simpler form with a technique developed by Arcidiacono and Jones (2003) and Arcidiacono and Miller (2011).

1. We derive the following system of equations from the optimization problem's first-order conditions:

$$\{\tilde{\theta}, \tilde{\delta}\} = \underset{\theta,\delta}{\arg\max} \sum_{n\in\mathbb{n}} \sum_{t\in\mathbb{t}} \sum_{k\in\mathbb{k}} \tilde{w}_{nk}(\log(\rho(q_{nt} \,|\, x_{nt}, \theta_k, \delta_k)) + \log(\delta_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1})))$$

and $$\tilde{p}_k = \sum_{n\in\mathbb{n}} \frac{\tilde{w}_{nk}}{|\mathbb{n}|},$$

where

$$\tilde{w}_{nk} = \frac{\prod_{t\in\mathbb{t}} \tilde{p}_k \rho(q_{nt} \,|\, x_{nt}, \tilde{\theta}_k, \tilde{\delta}_k) \tilde{\delta}_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1})}{\sum_{j\in\mathbb{k}} \prod_{t\in\mathbb{t}} \tilde{p}_j \rho(q_{nt} \,|\, x_{nt}, \tilde{\theta}_j, \tilde{\delta}_j) \tilde{\delta}_j(x_{nt} \,|\, x_{nt-1}, q_{nt-1})}.$$

Note, $\tilde{w}_{nk}$ is the ex post probability of store n belonging to type k.

2. We extricate the maximization problem from the fixed point by replacing structural estimates $\rho(q_{nt} \,|\, x_{nt}, \tilde{\theta}_k, \tilde{\delta}_k)$ and $\tilde{\delta}_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1})$ with reduced form analogs $\breve{\rho}_k(q_{nt} \,|\, x_{nt})$ and $\breve{\delta}_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1})$:

$$\{\breve{\theta}, \breve{\delta}\} = \underset{\theta,\delta}{\arg\max} \sum_{n\in\mathbb{n}} \sum_{t\in\mathbb{t}} \sum_{k\in\mathbb{k}} \breve{w}_{nk}(\log(\rho(q_{nt} \,|\, x_{nt}, \theta_k, \delta_k)) + \log(\delta_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1}))),$$

where $$\breve{w}_{nk} = \frac{\prod_{t\in\mathbb{t}} \breve{p}_k \breve{\rho}_k(q_{nt} \,|\, x_{nt}) \breve{\delta}_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1})}{\sum_{j\in\mathbb{k}} \prod_{t\in\mathbb{t}} \breve{p}_j \breve{\rho}_j(q_{nt} \,|\, x_{nt}) \breve{\delta}_j(x_{nt} \,|\, x_{nt-1}, q_{nt-1})},$$

$$\breve{p}_k = \sum_{n\in\mathbb{n}} \frac{\breve{w}_{nk}}{|\mathbb{n}|},$$

$$\breve{\delta}_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1}) = \frac{\sum_{m\in\mathbb{n}} \sum_{s\in\mathbb{t}} \breve{w}_{mk} \mathbb{1}(x_{ms} = x_{nt} \cap x_{ms-1} = x_{nt-1} \cap q_{ms-1} = q_{nt-1})}{\sum_{n\in\mathbb{m}} \sum_{s\in\mathbb{t}} \breve{w}_{mk} \mathbb{1}(x_{ms-1} = x_{nt-1} \cap q_{ms-1} = q_{nt-1})},$$

and $$\breve{\rho}_k(q_{nt} \,|\, x_{nt}) = \frac{\sum_{m\in\mathbb{n}} \sum_{s\in\mathbb{t}} \breve{w}_{mk} \mathbb{1}(x_{ms} = x_{nt} \cap q_{ms} = q_{nt})}{\sum_{n\in\mathbb{m}} \sum_{s\in\mathbb{t}} \breve{w}_{mk} \mathbb{1}(x_{ms} = x_{nt})}.$$

3. We calculate $\breve{w}_{nk}$, $\breve{p}_k$, $\breve{\delta}_k$, and $\breve{\rho}_k$ by iterating their respective equations. (This step is equivalent to the EM algorithm.)

4. We split the remaining likelihood function into $2\bar{k}$ constituent parts in the fashion of (Rust 1994, p. 3108)

$$\widehat{\delta}_k = \underset{\delta_k}{\arg\max} \sum_{n\in\mathbb{n}} \sum_{t\in\mathbb{t}} \breve{w}_{nk} \log(\delta_k(x_{nt} \,|\, x_{nt-1}, q_{nt-1}))$$

and $$\widehat{\theta}_k = \underset{\theta_k}{\arg\max} \sum_{n\in\mathbb{n}} \sum_{t\in\mathbb{t}} \breve{w}_{nk} \log(\rho(q_{nt} \,|\, x_{nt}, \theta_k, \widehat{\delta}_k)).$$

We consider the former optimization problem in Section 7.4 and the latter in Section 7.5.

7.4. State Transitions

Maximizing the log likelihood of the observed state transitions under weights $\breve{w}_{nk}$ yields

$$\widehat{\delta}_k(x' \,|\, x, q) = \sum_{b\in\{0,1\}} \sum_{d\in\mathbb{N}} \widehat{\delta}_k^b(b \,|\, u) \widehat{\delta}_k^d(d \,|\, m) \widehat{\delta}_k^i(i' \,|\, d, i, bq) \cdot \widehat{\delta}_k^u(u' \,|\, u, bq) \widehat{\delta}_k^m(m' \,|\, d, m),$$

where $\widehat{\delta}_k^b$, $\widehat{\delta}_k^d$, and $\widehat{\delta}_k^i$ return the fitted value of $\breve{w}_{nk}$-weighted empirical frequency estimators, $\widehat{\delta}_k^u$ returns the fitted value of a $\breve{w}_{nk}$-weighted ordered logistic regression of u' on $u - s$, and $\widehat{\delta}_k^m$ returns the fitted value of a $\breve{w}_{nk}$-weighted ordered logistic regression of m' on d and m.[12]

7.5. Cost Function

Following Aguirregabiria and Mira (2002), we estimate type k's cost parameters by iterating

$$\widehat{\theta}_k^l = \underset{\theta}{\arg\max} \sum_{n\in\mathbb{n}} \sum_{t\in\mathbb{t}} \breve{w}_{nk} \log((\psi_k(\theta) \widehat{\rho}_k^l)(q_{nt} \,|\, x_{nt}))$$

and $$\widehat{\rho}_k^{l+1} = \psi_k(\widehat{\theta}_k^l) \widehat{\rho}_k^l$$

Figure 5. Cost Parameter Estimates

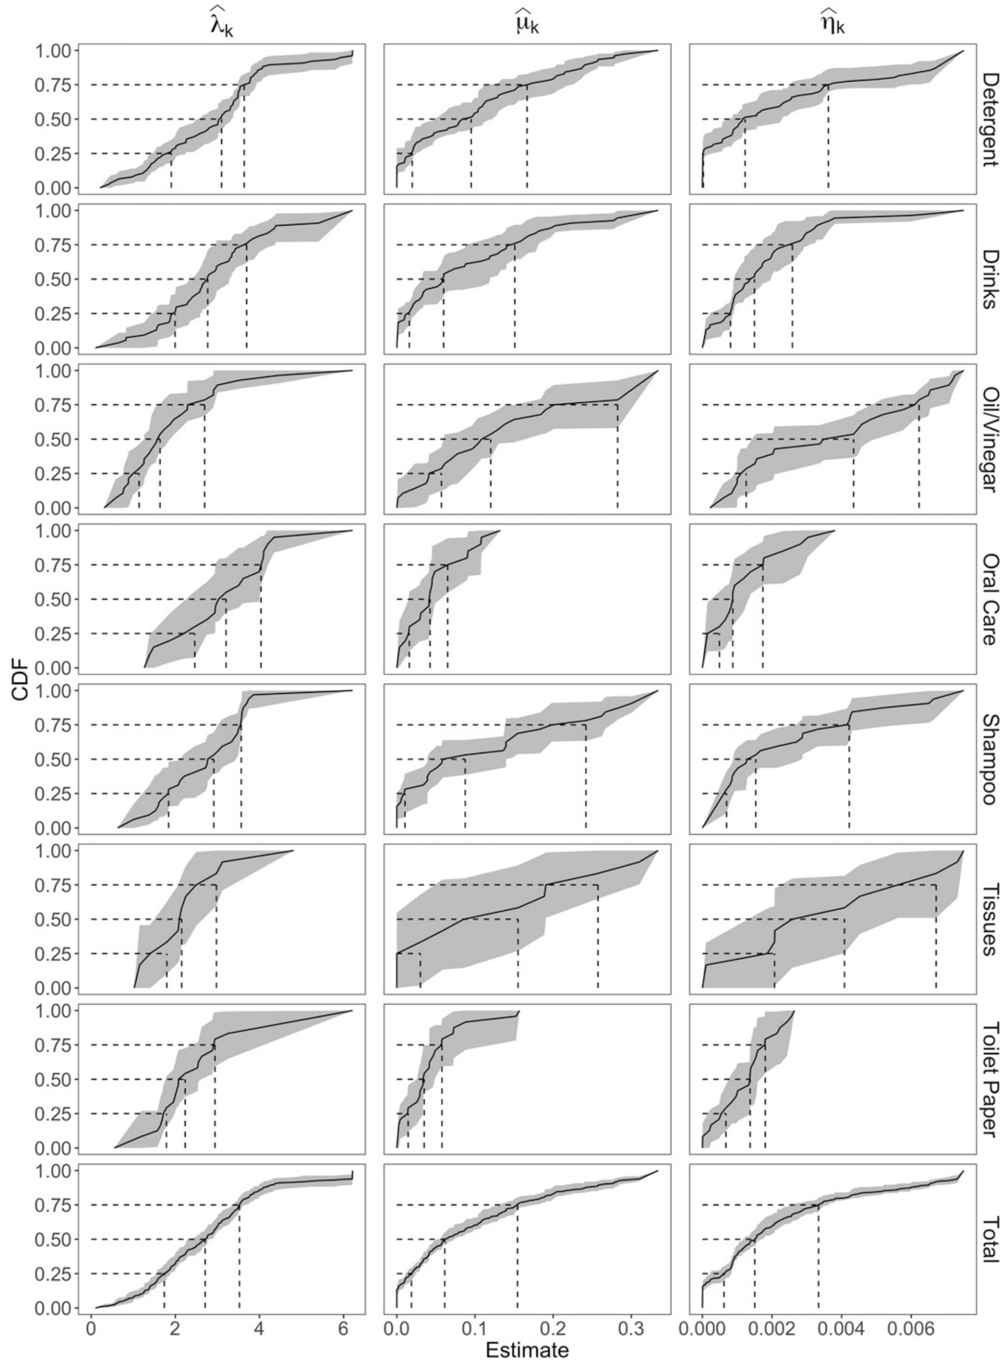

Notes. These line plots depict the distribution of our $\widehat{\theta}_k = \{\widehat{\lambda}_k, \widehat{\mu}_k, \widehat{\eta}_k\}$ cost estimates relative to the variance of the error term, which we have normalized to one. The curves are empirical cumulative distribution functions (CDFs), and the grey bands are their 99% confidence intervals (calculated with the bootstrap). The dashed lines highlight the estimates' quartiles.

to convergence, where $\psi_k(\theta)$ is the policy iteration operator evaluated under structural parameters $\widehat{\delta}_k$ and θ.

8. Structural Econometric Results

8.1. Structural Estimates

Figure 5 depicts the distribution of our cost parameter estimates. We express all costs relative to the error term's SD, which we normalized to one; $\widehat{\lambda}_k$, $\widehat{\mu}_k$, and $\widehat{\eta}_k$ have means 3.83, 0.14, and 0.0030, respectively, and medians 3.48, 0.080, and 0.0019, respectively. These estimates are significantly larger than zero, and therefore, shipping costs, stock-out costs, and holding costs are all relevant. We estimate that not satisfying a unit of demand is roughly as costly as storing a unit of inventory

Figure 6. Counterfactual Analysis

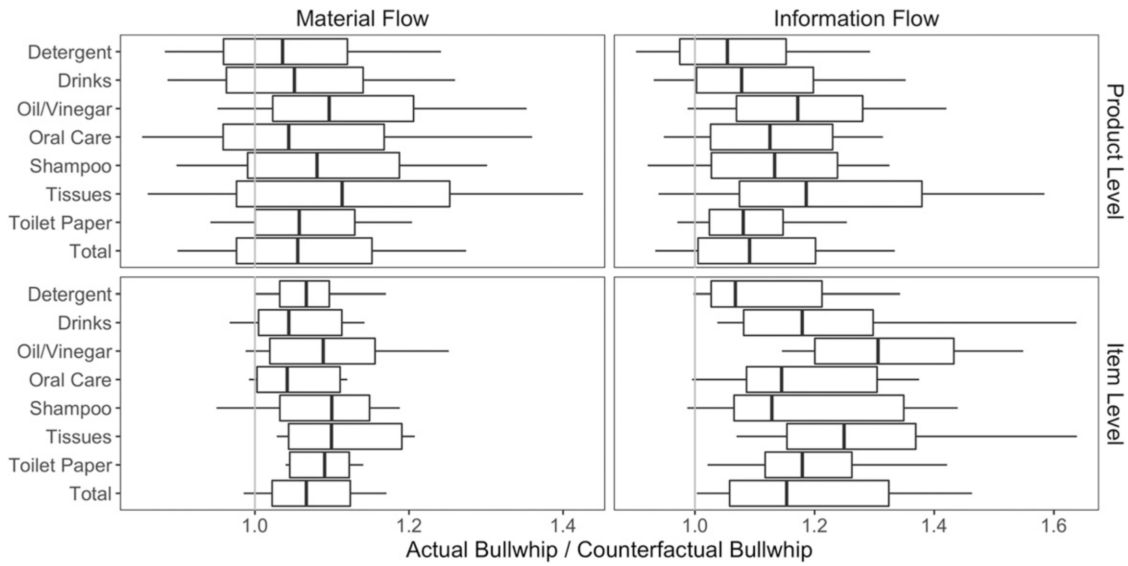

Notes. These box plots depict the medians, interquartile ranges, and interdecile ranges of the ratio of our actual bullwhip estimates to our counterfactual bullwhip estimates. For example, we estimate that ration gaming increases 77% of the product-level information flow bullwhips and 91% of the item-level information flow bullwhips.

for $0.14/0.0030 = 47.7$ days and that receiving a shipment is roughly as costly as not satisfying $3.83/0.14 = 27.4$ units of demand. This latter figure might seem high, but the cost of stocking out is mitigated by the fact that most products have multiple close substitutes (e.g., the median store has eight types of toilet paper).

8.2. Counterfactual Analysis

To determine the causal effect of ration gaming, we simulate the supply chain in its absence, speculating how the stores would have ordered if they did not strategically avoid upstream stock outs. We eliminate ration gaming from the supply chain by concealing the DC inventory level from the stores—the stores cannot muster inventory runs if they cannot foresee upstream shortages. Rather than condition on the DC's inventory level, the stores in our simulation believe that the DC's order fulfillment probability is fixed at the long-run average fulfillment rate. For example, the DC fulfilled 1,647 of the 2,068 orders that it received for the 115mL bottle of Li Jin brand sesame oil, and therefore, in our counterfactual scenario, the stores presume that this sesame oil has a constant fulfillment probability of $1{,}647/2{,}068 = 0.80$.

We use our primitive estimates to solve the stores' counterfactual optimal policies. We then use these counterfactual optimal policies to simulate a counterfactual sample of data. Additionally, we then use this counterfactual sample to estimate the counterfactual bullwhip effect without ration gaming.

Confirming the hypothesis of Lee et al. (1997a), we find that ration gaming underlies a meaningful portion of the bullwhip effect (Figure 6). Specifically, we estimate that ration gaming increases the geometric mean of the material flow bullwhip by 6.23% at the product level and 7.48% at the item level and increases the geometric mean of the information flow bullwhip by 10.58% at the product level and 19.4% at the item level. Ration gaming affects the information flow bullwhip more than the material flow bullwhip, because orders exceed shipments when inventory is scarce. Additionally, ration gaming affects the item-level bullwhips more than the product-level bullwhips, because the inventory runs not only make the store orders more variable but also, make the store orders more correlated. The common inventory signal coordinates the stores' orders.

9. Conclusion

In this article, we do seven things.

1. We perform the first structural estimation of a multiechelon supply chain.
2. We develop the first (s_t, S_t) inventory model estimator that factors both reorder point s_t and order-up-to level S_t.
3. We bridge the theoretical inventory management literature and empirical dynamic discrete choice literature. We show that the dynamic discrete choice paradigm can accommodate such operational features as order batch size constraints (Veinott 1965, Chen and Zheng 1994), fixed ordering costs (Yang et al. 2014), and Markov-modulated demand (Chen and Song 2001). We show that the paradigm can capture detailed supply chain dynamics (e.g., our state transitions depend on (i) the order from the store to the DC, (ii) the shipment from the DC to the store, (iii) the current period's actual demand, (iv) the next period's expected demand, and (v) the DC's inventory level).

4. We show that, like a bank can suffer a run on cash, a supply chain can suffer a run on inventory. This is a new version of ration gaming. Whereas the gaming of Lee et al. (1997a, p. 552) "is triggered only at an upswing of demand," our gaming is triggered by a dip in supply (a low upstream inventory level). Additionally, whereas their gaming stems from a conspiracy to deceive, our gaming stems from an honest tragedy of the commons.

5. We conduct the first empirical study that establishes the existence of ration gaming. Our stores engage in moderate inventory runs, ordering 31% more often when the upstream inventory level is in the bottom decile.

6. We confirm the two-decades-old hypothesis of Lee et al. (1997a) that ration gaming contributes to the bullwhip effect. We estimate that inventory runs lead to median increases of between 5.6% and 15.3% in our bullwhip measures.

7. We show that supply chain visibility can be harmful. The supply chain literature treats information sharing as an unalloyed good (surveys are in Chen 2003 and Kumar and Pugazhendhi 2012). However, more informed decisions are not necessarily better decisions: managers can use information selfishly. This is what we observed. The stores use the DC's inventory information to pass the stock-out risk upstream. Fortunately, there is a simple fix: inventory blinding.

Acknowledgments

The authors thank Victor Aguirregabiria, Thomas Bray, Martin Lariviere, Juan Serpa, Tunay Tunca, and Weiming Zhu for insightful comments.

Endnotes

[1] Lee et al. (1997b) cite potential gaming for computer memory, but they provide no empirical evidence; instead, they cite Li (1992), who also provides no empirical evidence. Lee et al. (1997b) also explain that there may have been gaming in Hewlett Packard's and IBM's supply chains, but they are not sure, because (i) "HP managers could not discern whether the orders genuinely reflected real market demands or were simply phantom orders from resellers trying to get better allocation product" and (ii) it was "unclear to IBM how much of the increase in orders was genuine market demand and how much was due to resellers placing phantom orders when IBM had to ration the product." Finally, Lee et al. (1997b) claim that telephone companies gamed the supply of Motorola cell phones in 1994, but their only evidence is an article from Kelly (1995), who only suggests a possibility of gaming. Next, Armony and Plambeck (2005) allude to gaming in Cisco's supply chain, but their only evidence is an article from Thurm (2001), whose only evidence is the following quote from Cisco's Chief Strategy Officer: "We knew there were multiple orders. We just didn't know the magnitude." Lai (2005) sought ration gaming in a Spanish supermarket but concluded that "gaming is unlikely to be significant in this retail case." Fransoo and Wouters (2000, p. 87) concluded that "shortage gaming did occur and this was a major problem" in a supply chain that they studied, but they did not quantify or illustrate this shortage gaming in any way, because they "encounter[ed] difficulties in measuring this, in particular in filtering the effect of order batching and shortage gaming." Finally, Sterman and Dogan (2015, p. 19) studied inventory hoarding with a laboratory experiment, but their specification explicitly barred "shortage gaming (because there is no horizontal competition)."

[2] The stores do not have an incentive to order more than they desire, because only 1.2% of orders are partially fulfilled.

[3] The sequence in which the stores' requests are fulfilled is irrelevant, because every store manages to stock up to ω during the inventory run.

[4] The online appendix outlines our sample selection procedure.

[5] We do not observe orders or shipments from October 23, 2011 to December 31, 2012, because of a lost Excel file.

[6] To avoid double counting standing orders, we remove observations in which the store ordered the previous day. However, including those observations yields similar results.

[7] The rightward skew in the data confirms the theory of Chen et al. (2016) that the material flow bullwhip generally exceeds the information flow bullwhip.

[8] The model of Hall and Rust (2000) is purely theoretical, and therefore, we exclude it from the list.

[9] The shipping cost is not the total expense of deploying a truck from the DC to the store, but rather, the incremental expense of receiving an additional product when the truck arrives (Section 3.4).

[10] We add the minus sign, because (i) the store minimizes costs instead of maximizing profits and (ii) the Gumbel's special properties correspond to its right tail.

[11] We set $\bar{k} = 2$ for our primary analysis, but setting $\bar{k} = 3$ yields similar results.

[12] Variable i' seems probabilistic given d, i, and s, because it must be rounded to a particular grid point; we assign probability mass to the grid points via linear interpolation.

References

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

Aguirregabiria V, Mira P (2002) Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica* 70(4):1519–1543.

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

Arcidiacono P, Ellickson PB (2011) Practical methods for estimation of dynamic discrete choice models. *Annual Rev. Econom.* 3(1):363–394.

Arcidiacono P, Jones JB (2003) Finite mixture distributions, sequential likelihood and the EM algorithm. *Econometrica* 71(3):933–946.

Arcidiacono P, Miller R (2011) Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* 79(6):1823–1867.

Armony M, Plambeck EL (2005) The impact of duplicate orders on demand estimation and capacity investment. *Management Sci.* 51(10):1505–1518.

Blinder AS, Maccini LJ (1991) Taking stock: A critical assessment of recent research on inventories. *J. Econom. Perspect.* 5(1):73–96.

Blinder AS, Lovell MC, Summers LH (1981) Retail inventory behavior and business fluctuations. *Brookings Papers Econom. Activity* 1981 (2):443–520.

Bray RL (2018) Strong convergence and dynamic economic models. *Quant. Econom.* Forthcoming.

Bray RL, Mendelson H (2012) Information transmission and the bullwhip effect: An empirical investigation. *Management Sci.* 58(5):860–875.

Bray RL, Mendelson H (2015) Production smoothing and the bullwhip effect. *Manufacturing Service Oper. Management* 17(2):208–220.

Cachon GP, Lariviere MA (1999) Capacity choice and allocation: Strategic behavior and supply chain performance. *Management Sci.* 45(8):1091–1108.

Cachon GP, Randall T, Schmidt GM (2007) In search of the bullwhip effect. *Manufacturing Service Oper. Management* 9(4):457–479.

Chen F (2003) Information sharing and supply chain coordination. Graves SC, de Kok AG, eds. *Handbooks in Operations Research and Management Science*, vol. 11 (Elsevier, Amsterdam), 341–421.

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

Chen F, Zheng Y-S (1994) Evaluating echelon stock (R, nQ) policies in serial production/inventory systems with stochastic demand. *Management Sci.* 40(10):1262–1275.

Chen L, Lee HL (2012) Bullwhip effect measurement and its implications. *Oper. Res.* 60(4):771–784.

Chen L, Luo W, Shang K (2016) Measuring the bullwhip effect: Discrepancy and alignment between information and material flows. *Manufacturing Service Oper. Management* 19(1):36–51.

Eichenbaum M (1989) Some empirical evidence on the production level and production cost smoothing models of inventory investment. *Amer. Econom. Rev.* 79(4):853–864.

Erdem T, Imai S, Keane MP (2003) Brand and quantity choice dynamics under price uncertainty. *Quant. Marketing Econom.* 1(1):5–64.

Fransoo JC, Wouters MJF (2000) Measuring the bullwhip effect in the supply chain. *Supply Chain Management Internat. J.* 5(2):78–89.

Hall G, Rust J (2000) An empirical model of inventory investment by durable commodity intermediaries. *Carnegie-Rochester Conf. Ser. Public Policy* 52(1):171–214.

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

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

Kelly K (1995) Burned by busy signals: Why Motorola ramped up production way past demand. *Business Week* (March 6), 36.

Kumar RS, Pugazhendhi S (2012) Information sharing in supply chains: An overview. *Procedia Engrg.* 38(C):2147–2154.

Lai R (2005) Bullwhip in a Spanish shop. Harvard NOM Working Paper no. 06-06, Harvard Business School Negotiation, Organizations and Markets Unit, Harvard Business School, Boston.

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

Lee HL, Padmanabhan V, Whang S (1997b) The bullwhip effect in supply chains. *Sloan Management Rev.* 38(3):93–102.

Li L (1992) The role of inventory in delivery-time competition. *Management Sci.* 38(2):182–197.

Lovell M (1961) Manufacturers' inventories, sales expectations, and the acceleration principle. *Econometrica* 29(3):293–314.

Maccini L, Rossana R (1984) Joint production, quasi-fixed factors of production, and investment in finished goods inventories. *J. Money Credit Banking* 16(2):218–236.

Metzler LA (1941) The nature and stability of inventory cycles. *Rev. Econom. Statist.* 23(3):113–129.

Miron JA, Zeldes SP (1988) Seasonality, cost shocks, and the production smoothing model of inventories. *Econometrica* 56(4):877–908.

Rust J (1994) Structural estimation of Markov decision processes. Engle RF, McFadden DL, eds. *Handbook of Econometrics*, vol. 4 (North Holland, Amsterdam), 3081–3143.

Sterman JD, Dogan D (2015) I'm not hoarding, I'm just stocking up before the hoarders get here. Behavioral causes of phantom ordering in supply chains. *J. Oper. Management* 39–40: 6–22.

Thurm S (2001) Missed signals: Behind Cisco's woes are some wounds of its own making. *Wall Street Journal* (April 18), A1.

Veinott A (1965) The optimal inventory policy for batch ordering. *Oper. Res.* 13(3):424–432.

Yang Y, Yuan Q, Xue W, Zhou Y (2014) Analysis of batch ordering inventory models with setup cost and capacity constraint. *Internat. J. Production Econom.* 155:340–350.

Robert L. Bray is an associate professor of operations management at the Kellogg School of Management, Northwestern University. He is primarily interested in empirical supply chain management.

Yuliang Yao is the George N. Beckwith '32 Professor of Information Systems in the Department of Management, College of Business and Economics, Lehigh University. His research interests are in the interdisciplinary fields of information systems and supply chain management, including business value of information-technology-enabled supply chains; economics of electronic commerce; social networks; and VMI, CPFR, and CRM.

Yongrui Duan is a professor at the School of Economics and Management, Tongji University. Her research interests include supply chain management, service operations management, system evaluation, and optimization.

Jiazhen Huo is a professor at the School of Economics and Management, Tongji University. He is the Chair Professor of DHL and Chair Professor of Bosch-Tongji University Global Supply Chain Management. His research interests include logistics, supply chain management, and service operations management.