

Information Transmission and the Bullwhip Effect: An Empirical Investigation

Robert L. Bray, Haim Mendelson

Graduate School of Business, Stanford University, Stanford, California 94305
{rlbray@stanford.edu, haim@stanford.edu}

The bullwhip effect is the amplification of demand variability along a supply chain: a company bullwhips if it purchases from suppliers more variably than it sells to customers. Such bullwhips (amplifications of demand variability) can lead to mismatches between demand and production and hence to lower supply chain efficiency. We investigate the bullwhip effect in a sample of 4,689 public U.S. companies over 1974–2008. Overall, about two-thirds of firms bullwhip. The sample's mean and median bullwhips, both significantly positive, respectively measure 15.8% and 6.7% of total demand variability. Put another way, the mean quarterly standard deviation of upstream orders exceeds that of demand by \$20 million. We decompose the bullwhip by information transmission lead time. Estimating the bullwhip's information-lead-time components with a two-stage estimator, we find that demand signals firms observe with more than three-quarters' notice drive 30% of the bullwhip, and those firms observe with less than one-quarter's notice drive 51%. From 1974–1994 to 1995–2008, our sample's mean bullwhip dropped by a third.

Key words: bullwhip effect; martingale model of forecast evolution; production smoothing; bullwhip decomposition; demand uncertainty

History: Received June 11, 2010; accepted September 16, 2011, by Christian Terwiesch, operations management. Published online in *Articles in Advance* March 9, 2012.

1. Introduction

This paper studies the existence and structure of the bullwhip effect, one of supply chain management's most celebrated hypotheses. When Cachon et al. (2007, p. 457) seek the bullwhip effect in industry-level data, they find that "retail industries generally do not exhibit the effect, nor do most manufacturing industries." Like Cachon et al. (2007), we look at the bullwhip across the entire U.S. economy, but we study the effect at the firm rather than the industry level. In firm-level data, mean and median bullwhips are significantly positive; 65% of our sample's firms bullwhip.

A number of case studies illustrate the bullwhip: Hammond (1994), Lee et al. (1997), Fransoo and Wouters (2000), Lai (2005), and Wong et al. (2007), respectively, find it in pasta, soup, frozen dinner, toy, and grocery supply chains. However, two factors make drawing conclusions from single-firm studies difficult: First, a publication bias may favor positive results—after all, bullwhip case studies will feature companies that bullwhip. Second, as our own results show, companies exhibit substantial bullwhip heterogeneity: the bullwhip standard deviation is nearly three times larger than the bullwhip mean. In fact, we find that only 24% of firm bullwhips lie between half and twice the global average—case-study estimates

likely lie nowhere near the economy-wide mean. Moreover, 35% of our sample exhibits no bullwhip whatsoever.

Aware of these small-sample pitfalls, Cachon et al. (2007) search for the phenomenon in a wide panel of industries. They find mixed results, as seasonal smoothing—the attenuation of seasonal variation—dampens much of their effect: out of 75 industries, 61 exhibit a bullwhip when they remove seasonality, but only 39 do when they do not. However, Cachon et al. (2007, pp. 477–478) explain that "it is possible that firms exhibit the bullwhip effect but the industry does not" and hence conclude that "Now, attention should turn toward probing data from individual firms... so that we can deepen our understanding of this phenomenon." Accordingly, we study the bullwhip in a panel of U.S. companies. The bullwhip largely manifests itself in firm-level data: out of 31 industries, 30 exhibit positive mean bullwhips when we remove seasonality, and 26 when we do not. And the effect is economically meaningful: the mean quarterly standard deviation of upstream orders exceeds that of demand by \$20 million.

Methodologically, our study differs from the study of Cachon et al. (2007) in four noteworthy ways: First, our data—quarterly and firm level, rather than monthly and industry level—sacrifice temporal for

cross-sectional granularity. Second, rather than estimate the bullwhip in the fractional growth rate, by log differencing, we estimate it in the level, a measure that better aligns with the theoretical bullwhip literature. Third, we do not just test for the existence of the bullwhip—we also measure its prevalence: we estimate the entire distribution of company bullwhips rather than just their industry-level means. Fourth, and most importantly, we decompose the effect into an infinite number of flavors based on demand signal transmission lead times.

Following Aviv (2007) and Chen and Lee (2009, 2010),¹ our bullwhip measures distinguish between demand variability and demand uncertainty. And they further decompose demand uncertainty by information availability. We study the bullwhip effect in the context of a martingale model of forecast evolution (MMFE) demand process (Hausman 1969, Hausman and Peterson 1972), in which demand uncertainty resolves gradually through a series of “lead l ” demand signals, i.e., signals with l period transmission lead times, $l = 0, 1, 2, \dots$. Following the MMFE, we decompose the bullwhip into a series of lead l bullwhips, the variance amplifications of lead l demand signals. The lead $-l$ bullwhips provide a profile of information distortion—their patterns reflect demand-signal twisting.

The mean bullwhip in our sample measures 15.8% of the magnitude of demand variability when we incorporate seasonality, and 19.6% when we eliminate it. Both signals with short and long information lead times contribute to the bullwhip effect: the mean lead 0 bullwhip, attributable to signals with information lead times shorter than one quarter, measures 10.0% the magnitude of demand variability, and the mean lead 3+ bullwhip, attributable to signals with information lead times longer than three quarters, measures 5.8%. Thus, the beer-game impression of the bullwhip—a manager frantically amending orders, chasing a runaway demand—does not tell the entire story; managers can anticipate nearly a third of the signals driving the phenomenon nine months early.

Others have estimated firm production in response to dynamic demand forecasts. Cohen et al. (2003)

¹ Chen and Lee (2009, p. 795) write,

So far, most researchers, including Cachon et al. (2007), have been looking at order variability as the measure of the bullwhip effect. Maybe we need to develop a new measure of the harmful effects of the bullwhip, i.e., a measure that captures the order uncertainty and not just the order variability.

And Chen and Lee (2010, p. 18) explain that a “bullwhip measure should be properly discounted to account for the actual demand uncertainty faced by the upstream stage (which is a conditional variance as opposed to the total variability captured by the bullwhip measure).”

use production decisions to estimate semiconductor equipment manufacturing costs. Terwiesch et al. (2005) and Krishnan et al. (2007) study the relationship between customers placing orders and a producer satisfying them in the semiconductor industry. Both find gaming inefficiencies. Dong et al. (2011) study the effect of demand forecast sharing on supply chain performance. Finally, Serman (1989) and Croson and Donohue (2003, 2006) study the bullwhip effect in the laboratory.

In §2 we study the bullwhip effect theoretically. We first develop a model of firm production, a context in which to study the bullwhip. We then show that the bullwhip decomposes by information transmission lead time into an infinite set of lead l bullwhips. In §3, we construct a consistent estimator of the lead l bullwhip from differences in the variances of demand and order forecast errors. In §4, we present our bullwhip estimates. In §5, we provide robustness checks. In §6, we provide our concluding remarks.

2. Modeling the Bullwhip Effect

We begin our analysis with a model that extends the Graves et al. (1998) single-stage production problem. Our model, like the model of Chen and Lee (2009), pertains to a single firm that observes a demand described by the MMFE and replenishes with a generalized order-up-to policy (GOUTP). The MMFE generalizes most commonly used, exogenous demand models, and the GOUTP allows any order scheme that is stationary and affine in observed demand signals. Chen and Lee (2009) argue for such order policies, citing their parsimony and common usage (e.g., Graves et al. 1998, Balakrishnan et al. 2004, Aviv 2007). We model a single firm because our data do not contain buyer-seller relationships, and the bullwhip across a supply chain is roughly the sum (or product, if one measures the variance ratio, rather than the variance difference) of its contributing firm-level amplifications, as Fransoo and Wouters (2000, p. 87) explain:

The total bullwhip effect is the coefficient of variation of the production plan, divided by the coefficient of variation of consumer demand. Under specific conditions, this is the product of the measured effect at each echelon. Suppose Echelon 3 is the retail franchisee, Echelon 2 is the distribution center, and Echelon 1 is production, then

$$\frac{c_{out1}c_{out2}c_{out3}}{c_{in1}c_{in2}c_{in3}} = \frac{c_{out1}}{c_{in3}}$$

provided there is consistency between D_{inl} and D_{outl+1} so $c_{inl} = c_{outl+1}$.

Indeed, demand amplification across a *single firm* has become an almost universally accepted measure, in

both the theoretical² and empirical³ bullwhip literatures. Nevertheless, estimating the bullwhip effect across firms, rather than entire supply chains, limits our study.

2.1. Production Model

We consider a firm that produces a single output unit from a single input unit, ordered from a supplier that meets orders promptly (see Gavirneni et al. 1999, Lee et al. 2000, Chen and Lee 2009). The firm may freely return stock, so it can meet any desired order-up-to level (see Kahn 1987, Lee et al. 1997, Aviv 2003, Chen and Lee 2009). The supplier delivers orders with a lead time of $L \geq 0$. Without loss of generality, the firm’s production time is zero, so goods can be sold as soon as inputs arrive, and the firm only stores finished-good inventories.

The firm’s period t demand is

$$d_t \equiv \mu + \sum_{l=0}^{\infty} \epsilon_{t,l}, \tag{1}$$

where μ is a baseline mean, and $\epsilon_{t,l}$ is a demand signal with an l period information lead time; namely, the firm observes $\epsilon_{t,l}$ in period $(t - l)$. In period t , the firm observes signals $\epsilon_t \equiv [\epsilon_{t,0}, \epsilon_{t+1,1}, \epsilon_{t+2,2}, \dots]'$. The first component, $\epsilon_{t,0}$, gives the portion of period t demand unknown until period t . The remaining signals, with longer information lead times, reflect future demands. We model ϵ_t as independent and identically distributed (i.i.d.) mean-zero multivariate normal random variables, with covariance matrix Σ . We do not restrict Σ ’s top-left $(L + 1) \times (L + 1)$ submatrix, but beyond that, we make it diagonal—namely, $\epsilon_{t+1,1}$ and $\epsilon_{t+j,j}$ may be correlated as long as $l, j \leq L$.⁴

²See Lee et al. (1997, Theorem 1); Cachon and Lariviere (1999, Theorem 3); Graves (1999, Equation (12)); Chen et al. (2000, Theorem 2.2); Aviv (2007, Proposition 4); Chen and Lee (2009, Proposition 6); and Chen and Lee (2010, Proposition 1).

³Lai (2005, p. 3) considers “amplification at one party in the chain, so one way to qualify [his] paper is that it is about the contribution by a retailer to the bullwhip effect along the supply chain.” The primary bullwhip measure of Cachon et al. (2007, p. 464) is “the amount of volatility and industry contributes to the supply chain,” an industry-level analog to the firm bullwhip. And Fransoo and Wouters (2000, p. 88) explain that

The [bullwhip] measurement needs to be determined for each echelon separately, such that the benefits of partial solutions may be traded off against benefits of integral solutions. Each of the echelons may contribute to creating a bullwhip effect to a greater or smaller extent. Therefore, in order to make a proper trade-off, it is important to distinguish the contribution of each of the echelons in the supply chain.

⁴Chen and Lee (2009, p. 12) explain that the bulk of signal variations lie in the general covariance region, as the scenario where “forecast information is not available beyond the lead time L is fairly common in practice.”

In response to observed demand signals, the firm follows a GOUTP (see Chen and Lee 2009), stocking up to

$$S_{t+L} \equiv m + L\mu + \sum_{l=0}^{\infty} \tilde{S}_{t+L,l}, \tag{2}$$

where

$$\tilde{S}_{t+L,l} \equiv \sum_{i=-L+1}^{\infty} w_{i,l} \epsilon_{t+L+i,l} + \sum_{i=L}^{\infty} (w_{-i,l} - 1) \epsilon_{t+L-i,l}.$$

The coefficients have a clear interpretation: m is the mean inventory level, and $w_{i,l}$ the cumulative fraction of $\epsilon_{t,l}$ that the firm produces i periods early, i.e., by period $(t - i)$. This policy preserves the MMFE structure, both in order quantity and inventory level:

$$\begin{aligned} o_t &= \mu + \sum_{l=0}^{\infty} \epsilon_{t,l}^o, & \epsilon_t^o &\equiv [\epsilon_{t,0}^o, \epsilon_{t+1,1}^o, \dots]' = A\epsilon_t, \\ i_t &= m + \sum_{l=0}^{\infty} \epsilon_{t,l}^i, & \epsilon_t^i &\equiv [\epsilon_{t,0}^i, \epsilon_{t+1,1}^i, \dots]' \\ & & &= C(D^L A - I)\epsilon_t, \end{aligned} \tag{3}$$

where o_t is the period t order quantity, with lead l signal $\epsilon_{t,l}^o$; i_t is the end-of-period t inventory level, with lead l signal $\epsilon_{t,l}^i$; C and D^L are square matrices with (i, j) th elements $C_{i,j} \equiv \mathbb{1}_{\{i \geq j\}}$ and $D_{i,j}^L \equiv \mathbb{1}_{\{i=j+L\}}$, respectively; and A is a square matrix with (i, j) th element, $A_{i,j} \equiv w_{j-i-L,j} - w_{j-i-L+1,j}$, so $A_{i,j}$ gives the fraction of lead j demand signals routed to lead i order signals.

Finally, to produce, the firm acquires a fixed amount of in-house production capacity, z , for which it pays $s > 0$ per unit per period to maintain. With this capacity, the firm produces the first z units in-house at unit cost \underline{c} , and outsources the rest at unit cost $\bar{c} + \bar{c} + s > \underline{c}$. Hence, the firm faces newsvendor production capacity costs of \bar{c} per unit per period of capacity shortage, and s per unit per period of capacity surplus (see Ernst and Pyke 1993, Balakrishnan et al. 2004). In addition, the firm faces newsvendor inventory costs of b per unit per period of backlogged demand, and h per unit per period of excess stock. Recapping, in period t the firm (1) observes ϵ_t (and thus d_t); (2) orders o_t ; (3) receives the period $(t - L)$ orders and finishes associated production; (4) adjusts inventory to i_t , satisfying the demand it can; and (5) pays newsvendor costs $C(o_t, i_t) \equiv h(i_t)^+ + b(-i_t)^+ + \bar{c}(o_t - z)^+ + s(z - o_t)^+ + (\bar{c} + \underline{c})o_t$.

Under the GOUTP, the firm

$$\begin{aligned} &\text{minimizes } E[C(o_t, i_t)], \\ &\quad \quad \quad w_{i,l,m,z} \\ &\text{subject to } w_{i,l} = 0 \quad \forall i > l - L. \end{aligned}$$

The constraints prevent the firm from conditioning on signals it has not yet observed. Because Equations (3)

Copyright: INFORMS holds copyright to this Article in Advance version, which is made available to subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

set inventory and order quantities to normal random variables, the optimal production capacity and mean inventory, with respect to the stock-up-to variables, are (see Porteus 2002, p. 13)

$$\begin{aligned} z(w_{i,l}) &= \Phi^{-1}\left(\frac{\bar{c}}{\bar{c}+s}\right)\sqrt{\text{Var}(o_t | w_{i,l})}, \quad \text{and} \\ m(w_{i,l}) &= \Phi^{-1}\left(\frac{b}{b+h}\right)\sqrt{\text{Var}(i_t | w_{i,l})}. \end{aligned} \quad (4)$$

Using (4), we recast the objective to depend only on $w_{i,l}$:

$$\begin{aligned} \min_{w_{i,l}} E[C(o_t, i_t) | w_{i,l}] \\ = k_i \sqrt{\text{Var}(i_t | w_{i,l})} + k_p \sqrt{\text{Var}(o_t | w_{i,l})}, \end{aligned} \quad (5)$$

where $k_i = (b+h)\phi(\Phi^{-1}(b/(b+h)))$ and $k_p = (\bar{c}+s)\phi(\Phi^{-1}(\bar{c}/(\bar{c}+s)))$. The following proposition characterizes the optimal order policy with respect to newsvendor-cost ratio $k \equiv k_i/k_p$ (consult the online appendix, available on Robert Bray's website, for proofs).

PROPOSITION 1. *The optimal order-up-to variables $w_{i,l}$ and transformation matrix A satisfy*

$$\begin{aligned} w_{i,l} &= \begin{cases} \frac{1-\lambda^{l-L+1-i}}{\lambda^i - \lambda^{2+2l-2L-i}} & 0 \geq l-L \geq i, \\ \frac{1-\lambda^{2+2l-2L-i}}{1+\lambda} & l-L > 0 \text{ and} \\ & l-L \geq i > 0, \\ 1 - \frac{1+\lambda^{2l-2L+1}}{1+\lambda} \lambda^{-i+1} & l-L > 0 \text{ and } i \leq 0, \\ 0 & l-L < i; \end{cases} \\ A_{\ell,l} &= \begin{cases} (1-\lambda)\lambda^\ell & l-L \leq 0, \\ \frac{1-\lambda}{1+\lambda} \lambda^\ell (\lambda^{l-L} + \lambda^{l-L+1}) & \ell \geq l-L > 0, \\ \frac{1-\lambda}{1+\lambda} \lambda^{l-L} (\lambda^{-\ell} + \lambda^{\ell+1}) & l-L > \ell; \end{cases} \\ \lambda &= \theta/2 + 1 - \sqrt{(\theta/2 + 1)^2 - 1}, \\ \theta &= k \sqrt{\frac{\text{Var}(o_t | w_{i,l})}{\text{Var}(i_t | w_{i,l})}}. \end{aligned} \quad (6)$$

One can solve this system by searching over θ —the ratio of the marginal costs of inventory variability to production variability, evaluated at the optimum—which entirely characterizes the solution.

Figure 1 depicts the solution. We find (1) the firm controls inventory more tightly as inventory-misalignment costs increase; (2) longer signal lead times allow the firm to shift production from peak periods to earlier periods; (3) because there is no preferred backlogging, the firm treats all delinquent

demand signals the same; and (4) because demand variations are negative as often as they are positive, consistently producing early costs as much as consistently producing late (i.e., $w_{i,\infty}$ is rotationally symmetric).

2.2. Bullwhip Effect

The following proposition characterizes the sign of the bullwhip effect, $\beta \equiv \text{Var}(o_t) - \text{Var}(d_t)$, with respect to the newsvendor ratio $k = k_i/k_p$.

PROPOSITION 2. *For some threshold T , $\beta > 0$ if and only if $k > T$ and $[\sum_{l=0}^L e'_l] \Sigma [\sum_{l=0}^L e_l] > \sum_{l=0}^L e'_l \Sigma e_l$.*

Following Lee et al. (1997), our model shows that the optimal policy can yield a bullwhip: sometimes it pays off to sacrifice the bullwhip to stabilize inventory levels. The newsvendor ratio k determines whether the firm bullwhips: when inventory costs are relatively high the best policy yields a bullwhip, but when production costs are relatively high it does not. Whether the bullwhip exists is thus an empirical question.

Proposition 2 suggests two ways to reduce the bullwhip effect. The first is to reduce the autocorrelation among signals with lead times no longer than the procurement lead time—i.e., reduce $[\sum_{l=0}^L e'_l] \Sigma [\sum_{l=0}^L e_l] - \sum_{l=0}^L e'_l \Sigma e_l$. These autocorrelations drive the “demand signal processing” underpinning the effect (see Lee et al. 1997). To reduce these autocorrelations, a firm can decrease its signal-exposure window L , or improve its demand forecasts, which, under the MMFE, is equivalent to increasing its signal transmission lead times. The second way is to decrease k , the costliness of inventory misalignments relative to the costliness of production-capacity misalignments. For example, our model illustrates that the firm can reduce the bullwhip effect by increasing product shelf life: a longer shelf life means a lower holding cost h , which means the firm carries a higher safety stock, which in turn means it reacts more calmly to demand spikes.

2.3. Bullwhip Decomposition

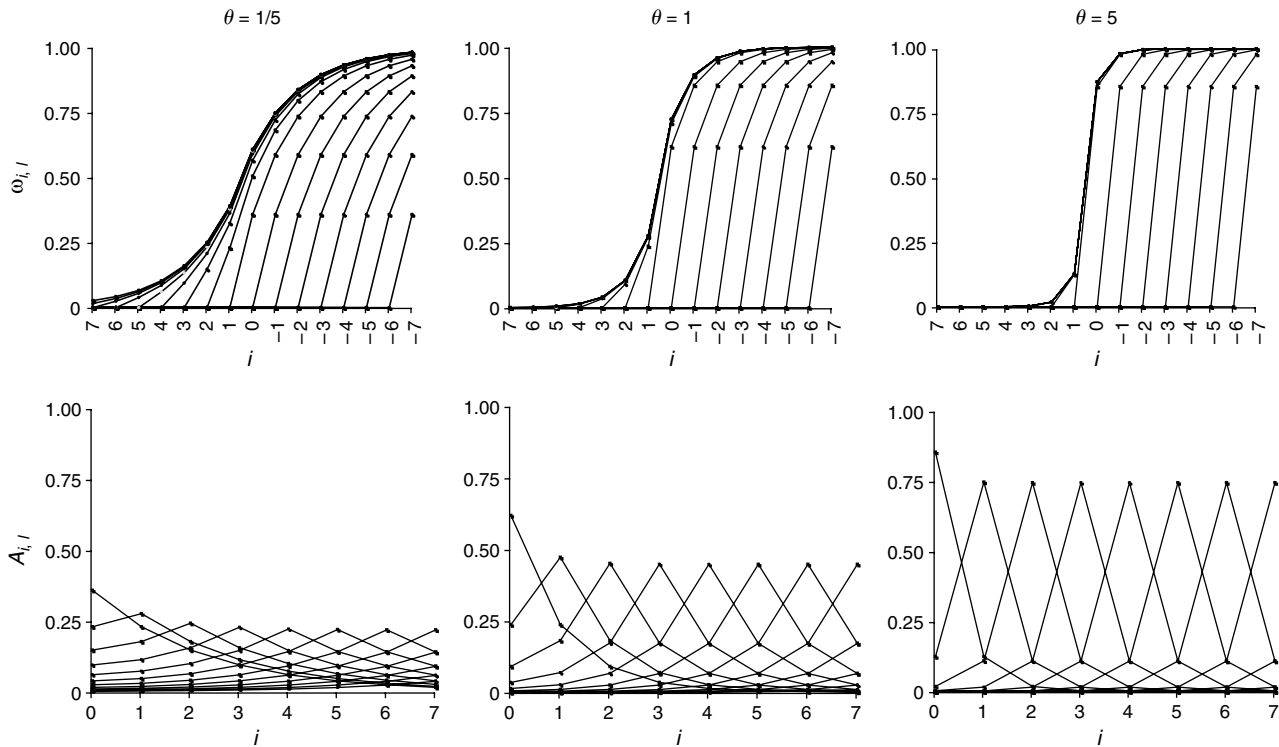
Within the framework of our model, we decompose the bullwhip, β , by information lead time. Defining e_l as a unit vector indicating the $(l+1)$ th position, we find

$$\begin{aligned} \beta &\equiv \text{Var}(o_t) - \text{Var}(d_t) = \text{Tr}[A \Sigma A'] - \text{Tr}[\Sigma] \\ &= \sum_{l=0}^{\infty} e'_l (A \Sigma A' - \Sigma) e_l = \sum_{l=0}^{\infty} \beta_l, \end{aligned} \quad (7)$$

where $\beta_l \equiv \text{Var}(o_{t,l}) - \text{Var}(d_{t,l}) = e'_l (A \Sigma A' - \Sigma) e_l$ is the lead l bullwhip, the variance amplification of lead l demand signals. Equation (7) provides an information distortion profile, a drill-down demonstrating the

Copyright: INFORMS holds copyright to this Article in Advance version, which is made available to subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

Figure 1 Optimal Order Policy



Notes. These plots characterize the optimal stock-up-to coefficients and routing matrix. The top panels plot $w_{i,l}$, the cumulative fraction of lead-time l signals produced i periods early (e.g., values at $i = 0$ give the fraction produced on time); the curves, from left to right, correspond to $l = \infty, l = L + 7, l = L + 6, \dots$, and $l = L - 7$. The bottom panels plot $A_{i,l}$, the fraction of lead l demand signals, $\epsilon_{t+l,l}$, routed to lead i production signals, $\epsilon_{t+i,i}^o$; the curves, from left to right, correspond to $l \leq L, l = L + 1, l = L + 2, \dots$, and $l = L + 7$. Inventory misalignments become relatively more costly as θ increases.

signals that drive the bullwhip. Naturally, bullwhips skewed toward short-lead-time distortions cost more, as short-notice order revisions require suppliers to produce hastily, in a helter-skelter fashion. For example, the supply chain scorecards of Graves et al. (1998) and Aviv (2007) more severely penalize short-lead-time order revisions. Also, a bullwhip's fix depends on its lead time. Suppliers need time to act (Aviv 2007), so information sharing better mitigates long-lead-time bullwhips, rather than short-lead-time ones. On the other hand, order fixing can address short-lead-time bullwhips, but not long-lead-time ones—a firm can commit orders for the next quarter, but not the next year (Balakrishnan et al. 2004).

The final proposition explains that information distortion requires an element of surprise—a testable implication of our model:

PROPOSITION 3. *The firm never bullwhips signals with arbitrarily long lead times: $\lim_{l \rightarrow \infty} \beta_l \leq 0$, where the inequality holds strictly when $\text{Var}(\epsilon_{t,l}) > 0$ and k is finite.*

This finding best speaks to the “seasonal bull whip”—the difference in the variances of the predictable seasonal components of demands and orders—because firms can fully anticipate these variations. Thus, Proposition 3 predicts a negative seasonal

bullwhip. For exposition purposes, we henceforth consider seasonal signals as having infinite, rather than “arbitrarily long,” lead times. That is, we let $d_t \equiv \mu + \epsilon_{t,\infty} + \sum_{l=0}^{\infty} \epsilon_{t,l}$, $o_t \equiv \mu + \epsilon_{t,\infty}^o + \sum_{l=0}^{\infty} \epsilon_{t,l}^o$, and $\beta = \beta_{\infty} + \sum_{l=0}^{\infty} \beta_l$, where $\epsilon_{t,\infty}$ and $\epsilon_{t,\infty}^o$ are demand's and order's respective seasonal components and $\beta_{\infty} \equiv \text{Var}(\epsilon_{t,\infty}^o) - \text{Var}(\epsilon_{t,\infty})$ is the seasonal bullwhip. In contrast, we call the β_l coefficients, for finite l , uncertainty bullwhips.

Bridging from theory to empirics, we conclude this section with a real-world example. Figure 2 displays the demand and order MMFE decompositions, and the bullwhip decomposition, for Teradyne Inc., a manufacturer of automatic test equipment for the telecommunications and electronics sectors. The company has no seasonal bullwhip, as it faces effectively no seasonality, but it has meaningful uncertainty bullwhips. The following section describes how we estimate the lead l bullwhips of Figure 2.

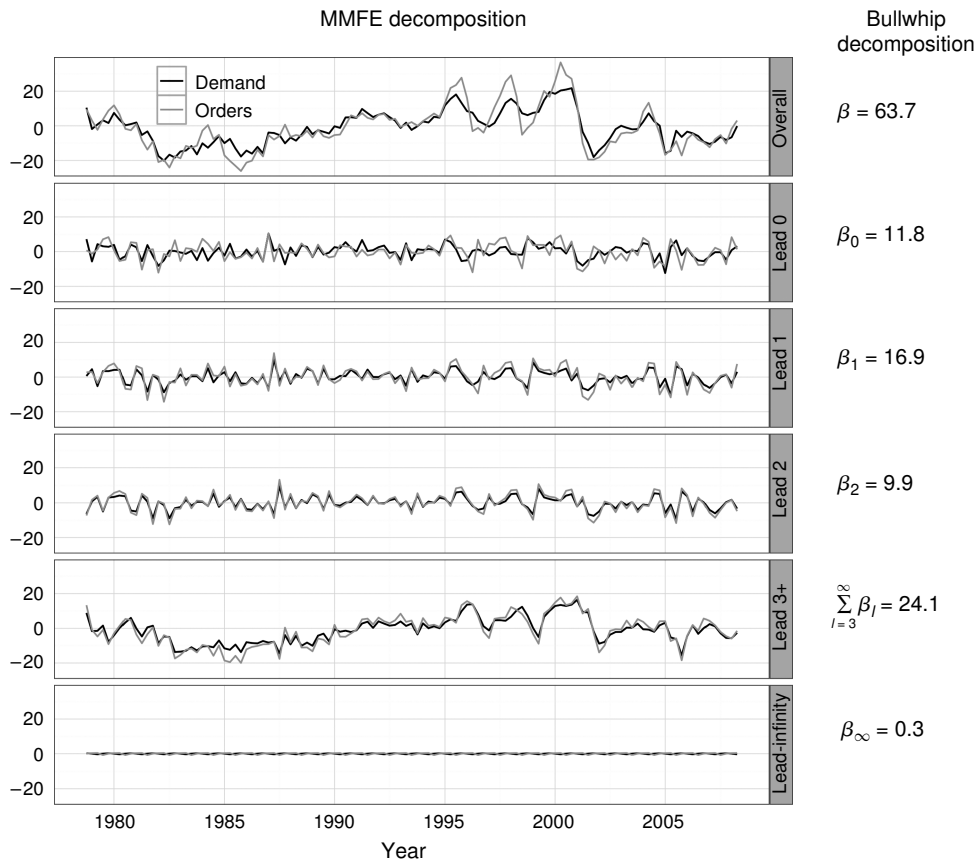
3. Estimation Procedure

3.1. Data

We use COMPUSTAT data, originating from quarterly financial statements of public U.S. companies,

Copyright: INFORMS holds copyright to this *Articles in Advance* version, which is made available to subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

Figure 2 Information-Lead-Time Decompositions



Notes. These figures decompose Teradyne Inc.'s demands, orders, and bullwhip by information lead time. The first plots overall detrended demands and orders, $\{d_t - \mu, o_t - \mu\}$, and the lower five plot their decomposed components, $\{\epsilon_{t,0}, \epsilon_{t,0}^o\}$, $\{\epsilon_{t,1}, \epsilon_{t,1}^o\}$, $\{\epsilon_{t,2}, \epsilon_{t,2}^o\}$, $\{\sum_{l=3}^{\infty} \epsilon_{t,l}, \sum_{l=3}^{\infty} \epsilon_{t,l}^o\}$, and $\{\epsilon_{t,\infty}, \epsilon_{t,\infty}^o\}$, in that order. To the right of each figure is a corresponding bullwhip. The unit of measure for both the plots and bullwhips is the percent of total demand variance. The lower five bullwhips sum to the top bullwhip, as the lower five plots sum to the top plot.

between 1974 and 2008 from the retailing, wholesaling, manufacturing, and resource extracting sectors (SIC 5200–5999, 5000–5199, 2000–3999, and 1000–1400, respectively). Lead l signals correspond to demands between l and $l + 1$ quarters hence, and lead ∞ signals to quarterly seasonal means. We proxy COGS for demand and production for orders (see Cachon et al. 2007, Lai 2005, Wong et al. 2007, Dong et al. 2011). (Recall, in our model sales equals demand and production equals orders.) We calculate production with the accounting identity $o_t = d_t + i_t - i_{t-1}$. Also, for consistency, we translate all COGS observations to LIFO form, adding the LIFO reserve to inventory, and subtracting its change from reported COGS.

We eliminate untrustworthy data, observations in which firms change their reporting schedule or fiscal calendar, or post total assets of less than a million dollars or nonpositive inventories or sales. Also, we allow companies to acquire others, but we remove companies from the sample after they have been acquired, or merge with another. Finally, we select each firm's longest series of uninterrupted data within a single industry, as long as the series

has at least 25 observations, the minimum necessary to estimate our time-series models. Our final sample comprises 187,901 observations from 4,297 firms. Table 1 reports summary statistics.

We transform each firm's demands and orders by (1) dividing by total assets, (2) detrending with linear and quadratic functions of t , (3) Winsorizing the

Table 1 Summary Statistics

	Sample	Retail	Wholesale	Manufacturing	Extraction
No. of firms	4,297	602	339	3,161	195
No. of obs.	187,901	27,118	13,964	139,369	7,450
$\text{Var}(\epsilon_{t,0})$	29.24	19.88	29.81	30.68	35.31
$\text{Var}(\epsilon_{t,1})$	6.61	4.70	7.18	6.95	5.95
$\text{Var}(\epsilon_{t,2})$	3.45	2.13	3.82	3.71	2.64
$\sum_{l=3}^{\infty} \text{Var}(\epsilon_{t,l})$	23.09	16.78	23.66	24.27	22.83
$\text{Var}(\epsilon_{t,\infty})$	15.72	34.79	17.13	12.30	7.65
Total assets	1.98	1.47	0.66	2.20	2.19
Inventory	0.23	0.30	0.33	0.22	0.05
Margin	0.19	0.29	0.21	0.16	0.38

Notes. Variable means by industry sector. Total assets are expressed in billions of 2008 dollars, and inventory as a fraction of total assets.

Copyright: INFORMS holds copyright to this Article in Advance version, which is made available to subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

top and bottom 1%, and (4) normalizing the demand variances to one. The first and second transformations stabilize our series' first two moments (Granger and Newbold 1974); the third dampens the effect of significant outliers; and the fourth allows us to express bullwhips as a percent of demand variance, a concrete, unitless measure.

3.2. Estimator Construction

We estimate lead l bullwhips with a sequential method of moments estimator. Our estimation procedure exploits two well-known features of the MMFE. The first is that $F_{t,l}$, the mean-square-error-minimizing forecast of period t demand from period $(t - l)$, is the true demand net unobserved signals: $F_{t,l} = \mu + \epsilon_{t,\infty} + \sum_{i=l}^{\infty} \epsilon_{t,i}$. (We define $F_{t,0} \equiv d_t$ and $F_{t,\infty} \equiv \mu + \epsilon_{t,\infty}$.) The second is that the signals contributing to period t demand are uncorrelated, so $\text{Var}(\sum_{i=0}^l \epsilon_{t,i}) = \text{Var}(\sum_{i=0}^{l-1} \epsilon_{t,i}) + \text{Var}(\epsilon_{t,l})$. Combining these features we find $\text{Var}(\epsilon_{t,l}) = \text{Var}(d_t - F_{t,l+1}) - \text{Var}(d_t - F_{t,l})$. Accordingly, we define the following lead l bullwhip estimator:

$$\begin{aligned} \hat{\beta}_l &\equiv \widehat{\text{Var}}(\epsilon_{t,l}^o) - \widehat{\text{Var}}(\epsilon_{t,l}) \\ &= [\widehat{\text{Var}}(o_t - F_{t,l+1}^o) - \widehat{\text{Var}}(o_t - F_{t,l}^o)] \\ &\quad - [\widehat{\text{Var}}(d_t - F_{t,l+1}) - \widehat{\text{Var}}(d_t - F_{t,l})], \end{aligned} \tag{8}$$

where $F_{t,l}^o$ is an equivalent order forecast. Our estimation procedure follows three steps: (1) estimate the demand and order forecasts, $\hat{F}_{t,l}$ and $\hat{F}_{t,l}^o$; (2) estimate the forecast error variances, $\widehat{\text{Var}}(d_t - F_{t,l})$ and $\widehat{\text{Var}}(o_t - F_{t,l}^o)$; and (3) calculate $\hat{\beta}_l$ from (8).

To estimate forecasts, we specify that demands and orders follow deterministic seasonal shifts combined with linear functions of an underlying vector autoregressive process. In this case, fitted values of regressions of future demands and orders on contemporaneous explanatory variables and quarter dummies consistently estimate $F_{t,l}$ and $F_{t,l}^o$ (see Lütkepohl 2005). For explanatory variables we use current inventory levels, and demands and orders from the current and prior four quarters.⁵

Next, we estimate forecast error variances with their sample moments:⁶

$$\begin{aligned} \widehat{\text{Var}}(d_t - F_{t,l}) &\equiv \sum_{t=1}^T (d_t - \hat{F}_{t,l})^2 / T, \\ \widehat{\text{Var}}(o_t - F_{t,l}^o) &\equiv \sum_{t=1}^T (o_t - \hat{F}_{t,l}^o)^2 / T. \end{aligned} \tag{9}$$

⁵ We consider alternate specifications in §5.

⁶ Alternatively, you can observe Bessel's correction, and divide the sum of the square residuals by $T - 1$ instead of by T . Both denominators are valid, however (see Davidson and MacKinnon 2004, Equations 3.46, 3.49).

Plugging the relevant variance estimates into (8) yields $\hat{\beta}_l$. (We similarly define $\hat{\beta}_\infty = [\widehat{\text{Var}}(o_t) - \widehat{\text{Var}}(o_t - F_{t,\infty}^o)] - [\widehat{\text{Var}}(d_t) - \widehat{\text{Var}}(d_t - F_{t,\infty})]$ and $\sum_{i=l}^{\infty} \hat{\beta}_i = [\widehat{\text{Var}}(o_t - F_{t,\infty}^o) - \widehat{\text{Var}}(o_t - F_{t,l}^o)] - [\widehat{\text{Var}}(d_t - F_{t,\infty}) - \widehat{\text{Var}}(d_t - F_{t,l})]$.)

To estimate a bullwhip's mean across a collection of companies, we estimate each firm's forecasts individually, and then estimate the unconditional variances, (9), jointly across the relevant firms' forecast errors. To account for temporal and cross-sectional correlations, as well as heteroskedasticity, we use two-way cluster robust standard errors (Petersen 2009, Gow et al. 2010, Cameron et al. 2011). Because the moment conditions across our estimator's two stages—the first estimating the forecasts, and the second estimating their error variances—are asymptotically uncorrelated, we can use second-stage standard errors directly, without having to correct for first-stage misestimation (Newey 1984). We translate the two-way cluster-robust estimator covariance matrix of $[\widehat{\text{Var}}(o_t - F_{t,l+1}^o), \widehat{\text{Var}}(o_t - F_{t,l}^o), \widehat{\text{Var}}(d_t - F_{t,l+1}), \widehat{\text{Var}}(d_t - F_{t,l})]$ into $\hat{\beta}_l$ standard errors with the Delta method (Cameron and Trivedi 2005).

3.3. Estimator Properties

Our forecast-error variance estimators, belonging to the sequential m-estimator class characterized in §6.6 of Cameron and Trivedi (2005), are root- n consistent and asymptotically normal.⁷ Our bullwhip estimates, linear combinations of these forecast-error variance estimates, are thus also root- n consistent and asymptotically normal.

Our estimates are robust to measurement error. Suppose we observe $\tilde{d}_t = d_t + \eta_t$, and $\tilde{o}_t = \tilde{d}_t + i_t - i_{t-1}$ (recall, we calculate production from demand and inventory changes), where η_t is a measurement error term uncorrelated with demands, orders, and our forecast variables. Despite measurement error, the lead l bullwhip estimate remains consistent:

$$\begin{aligned} \hat{\beta}_l &= [\widehat{\text{Var}}(\tilde{o}_t - F_{t,l+1}^o) - \widehat{\text{Var}}(\tilde{o}_t - F_{t,l}^o)] \\ &\quad - [\widehat{\text{Var}}(\tilde{d}_t - F_{t,l+1}) - \widehat{\text{Var}}(\tilde{d}_t - F_{t,l})] \\ &= \left[\sum_{t=1}^T (\tilde{o}_t - \hat{F}_{t,l+1}^o)^2 / T - \sum_{t=1}^T (\tilde{o}_t - \hat{F}_{t,l}^o)^2 / T \right] \\ &\quad - \left[\sum_{t=1}^T (\tilde{d}_t - \hat{F}_{t,l+1})^2 / T - \sum_{t=1}^T (\tilde{d}_t - \hat{F}_{t,l})^2 / T \right] \\ &\xrightarrow{p\text{-lim}} [\text{Var}(o_t - F_{t,l+1}^o) + \text{Var}(\eta_t) - \text{Var}(o_t - F_{t,l}^o) - \text{Var}(\eta_t)] \\ &\quad - [\text{Var}(d_t - F_{t,l+1}) + \text{Var}(\eta_t) \\ &\quad - \text{Var}(d_t - F_{t,l}) - \text{Var}(\eta_t)] = \beta_l. \end{aligned}$$

⁷ The moment conditions defining $\widehat{\text{Var}}(d_t - F_{t,l})$ are $E[(d_t - \hat{F}_{t,l}(W_{t-1}, \theta))W_{t-1}] = 0$ and $E[\widehat{\text{Var}}(d_t - F_{t,l}) - (d_t - F_{t,l}(W_{t-1}, \theta))^2] = 0$, where W_{t-1} are forecast variables and θ forecast parameters.

Copyright: INFORMS holds copyright to this Article in Advance version, which is made available to subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

4. Results

4.1. Existence and Prevalence of the Bullwhip Effect

First, we estimate mean firm-level bullwhips across industries, sectors, and the entire sample, listing the results in Table 2.⁸ As predicted, the seasonal bullwhip is largely negative: out of 31 industries, 26 have negative seasonal mean bullwhips. These negative seasonal bullwhips induce a drop in seasonality across sectors, from retailing to resource extraction (see Table 1). However, out of 31 industries, 30, 29, 25, and 30 have positive lead 0, lead 1, lead 2, and lead 3+ mean bullwhips, respectively. Additionally, 26 industries exhibit positive overall bullwhip means, so the positive uncertainty bullwhips generally outweigh their negative seasonal counterparts, a finding that differs from the Cachon et al. (2007) conclusion that seasonal smoothing generally outweighs uncertainty amplification, and hence that most industries exhibit no bullwhip.

Industry aggregation, which overweighs the negative seasonal bullwhip, can explain this discrepancy. Seasonal signals correlate more highly across companies than do firm-specific shocks. Thus, industry aggregation attenuates uncertainty bullwhips more than it does seasonal bullwhips, as stochastic variations largely cancel out upon aggregation, whereas seasonal variations do not. To demonstrate, we explore the effect of industry aggregation ourselves, measuring at the firm-level, the four-digit SIC, the three-digit SIC, and the two-digit SIC the relative mean seasonal bullwhip, $|\hat{\beta}_\infty|/(|\hat{\beta}_\infty| + |\sum_{l=0}^{\infty} \hat{\beta}_l|)$, and the mean overall bullwhip, $\hat{\beta}$. As the level of aggregation increases, $|\hat{\beta}_\infty|/(|\hat{\beta}_\infty| + |\sum_{l=0}^{\infty} \hat{\beta}_l|)$ indeed increases, from 19% to 33%, to 46%, to 53%. In turn, $\hat{\beta}$ converges to the Cachon et al. (2007) zero bullwhips, going from 15.8% the magnitude of underlying demand variability to 5.5%, to 1.2%, to -1.7% .

Next we consider our decomposition, which partitions the bullwhip into economically meaningful components: the sample's mean uncertainty bullwhips—all significantly positive, yet diminishing with information lead time as signals become less informative—decompose into those with

short lead times (<1 quarter)	51%,
midrange lead times (1 – 3 quarters)	19%,
long lead times (>3 quarters)	30%.

The bullwhip effect boasts a long tail: signals arriving with more than nine months' notice drive nearly a third of the effect.

⁸ In Table 2, $\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2, \sum_{i=3}^{\infty} \hat{\beta}_i$, and $\hat{\beta}_\infty$ do not quite sum to $\hat{\beta}$, as Winsorizing the data slightly rattles our estimates.

However, means tell only part of the story, so we now consider the entire distribution of firm-level bullwhips. Figure 3 characterizes the bullwhips' marginal distributions. The boxplots illustrate a striking degree of heterogeneity: the coefficients of variation are all larger than two, and each interquartile range spans both positive and negative values—the bullwhip is by no means universal. Also, the boxplots depict skewed distributions. Because of these skews, the median bullwhips fall short of the means: the across-sample medians $\hat{\beta}, \hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2, \widehat{\sum_{i=3}^{\infty} \beta_i}$, and $\hat{\beta}_\infty$ measure 6.7, 4.2, 1.0, 0.2, 2.7, and -1.2 , respectively; these figures each differ from zero significantly at $p = 0.01$. We block bootstrap to calculate the median estimators' standard errors (see Hahn 1995, Hall et al. 1995).

We present the firm-level bullwhip probability density functions, which we estimate nonparametrically, in Figures 3 and 4. The former plot depicts modes of zero: every sector has a handful of companies that precisely peg production to demand, which yields zero bullwhip. It also demonstrates that retailers, because of their strong proclivity to smooth seasonality, are the only sector without an average bullwhip. Figure 4 presents the joint distributions of β and β_0 . Although each sector has its mode at the origin, retailers and wholesalers have secondary production-smoothing peaks. In our sample, 65% of firms exhibit a positive overall bullwhip, 72% a positive lead 0 bullwhip, and 56% exhibit both.

4.2. Has the Bullwhip Changed over Time?

Chen et al. (2005, pp. 1015, 1024) found that “inventories were significantly reduced” over the 1981–2000 time span as the “manufacturing firms [they studied] improved their interactions with suppliers and their own internal operations.” Moreover, Kahn et al. (2002, p. 183) and Davis and Kahn (2008, p. 155) argue that “changes in inventory behavior stemming from improvements in information technology (IT) have played a direct role in reducing real output volatility,” causing a “striking decline in volatility of aggregate economic activity since the early 1980s.” These changes suggest a drop in the bullwhip effect. Indeed, citing the “significant improvements in information technology and supply chain management” Cachon et al. (2007, pp. 467, 476) hypothesize such a drop, yet find their industry bullwhips “mostly stable over [their] sample period.”

Conversely, our firm-level data indicate that the bullwhips drop dramatically from before 1995 to after—see Table 3. We choose the 1995 breakpoint because (Jorgenson 2001, Jorgenson et al. 2003, Basu et al. 2003) and others deem it the first year of the “information age,” as “a substantial acceleration in the IT price decline occurred in 1995, triggered by a much sharper acceleration in the price

Table 2 Mean Bullwhips

	$\hat{\beta}$	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\beta}_2$	$\widehat{\sum_{i=3}^{\infty} \beta_i}$	$\hat{\beta}_{\infty}$
Retail						
Hardware and garden	15.57** (6.94)	16.30*** (2.87)	0.85 (1.13)	0.34 (0.36)	4.62 (14.10)	-10.96*** (0.51)
General merchandise	-26.55*** (6.30)	3.65*** (0.39)	1.30*** (0.22)	0.26 (0.22)	1.85*** (0.59)	-32.70*** (6.71)
Food	2.37*** (0.57)	0.25** (0.10)	0.42*** (0.07)	0.11 (0.11)	0.88*** (0.17)	0.80* (0.48)
Apparel and accessory	1.78 (4.28)	7.27*** (1.26)	0.84* (0.47)	0.38 (0.79)	3.66** (1.51)	-15.07*** (5.02)
Furniture and home furnishings	3.44 (6.00)	8.63*** (1.07)	1.83*** (0.28)	-0.16 (0.61)	6.60*** (1.33)	-14.22*** (3.51)
Eating and drinking places	1.02*** (0.06)	0.43*** (0.01)	-0.05** (0.02)	-0.27*** (0.01)	0.54*** (0.02)	0.24*** (0.01)
Miscellaneous	-5.91 (4.30)	6.67*** (1.25)	1.51*** (0.50)	0.33 (0.43)	3.98*** (0.70)	-19.26*** (4.45)
Segment mean	-3.23 (3.19)	4.62*** (1.21)	0.89*** (0.27)	0.12 (0.19)	2.69*** (0.63)	-12.54*** (4.21)
Wholesale						
Durable goods	24.07*** (4.77)	10.56*** (1.76)	4.69*** (1.09)	2.56*** (0.65)	9.19*** (1.25)	-4.40** (1.98)
Nondurable goods	5.98*** (1.97)	3.24* (1.76)	1.38*** (0.49)	0.00 (0.27)	2.68*** (0.67)	-1.94 (1.35)
Segment mean	17.81*** (4.06)	8.03*** (1.53)	3.55*** (0.86)	1.67*** (0.54)	6.94*** (1.11)	-3.55*** (1.35)
Manufacturing						
Food	18.30*** (5.31)	6.90*** (1.50)	0.86** (0.35)	0.51* (0.31)	3.11*** (0.72)	2.76 (2.14)
Textile mill	5.99 (6.06)	6.48*** (2.29)	2.17** (0.98)	0.11 (0.57)	5.98*** (1.19)	-9.20** (3.59)
Apparel	-6.04** (2.90)	5.53** (2.40)	3.49*** (0.91)	1.22 (0.77)	5.83*** (1.28)	-22.41*** (2.89)
Lumber and wood	27.57** (11.59)	9.37** (3.79)	2.28*** (0.77)	0.60 (0.97)	6.23*** (1.83)	4.74 (4.92)
Furniture and fixtures	11.96*** (3.44)	8.66*** (1.59)	3.53*** (0.95)	1.44 (1.07)	2.65* (1.36)	-6.04*** (2.07)
Paper	11.48*** (2.67)	5.43** (2.24)	2.13*** (0.50)	0.60 (0.75)	2.90*** (1.05)	0.09 (0.38)
Printing and publishing	1.82 (2.55)	5.34*** (1.35)	0.34 (0.51)	-0.01 (0.43)	1.74*** (0.45)	-5.92** (2.86)
Chemicals	15.37*** (2.62)	10.39*** (1.34)	2.24*** (0.25)	0.47* (0.28)	5.00*** (0.64)	-3.24*** (0.98)
Petroleum and coal	4.28*** (0.90)	6.30*** (0.63)	0.53*** (0.13)	-0.01 (0.14)	-0.17 (0.20)	-2.35*** (0.23)
Rubber and plastics	15.03*** (3.98)	9.73*** (1.24)	2.16*** (0.71)	0.82 (0.50)	4.33** (1.78)	-3.49* (1.99)
Leather goods	9.54* (5.54)	12.14*** (0.32)	3.29*** (0.44)	1.91* (1.08)	4.32*** (0.74)	-13.28*** (3.77)
Stone and glass	-2.89 (3.68)	3.05* (1.71)	1.07 (0.72)	1.10** (0.49)	1.91*** (0.35)	-9.98*** (2.33)
Primary metal	29.50*** (3.74)	12.19*** (1.69)	4.72*** (1.13)	2.20*** (0.54)	8.83*** (1.12)	-0.78 (1.01)
Fabricated metal	17.41*** (5.33)	10.49*** (2.10)	2.58*** (0.57)	1.30* (0.70)	7.11*** (1.91)	-4.24*** (1.63)
Industrial machinery	27.96*** (2.80)	14.93*** (1.39)	4.64*** (0.45)	2.43*** (0.35)	9.21*** (0.82)	-4.48*** (0.94)
Electronic equipment	28.89*** (2.98)	15.37*** (1.62)	4.18*** (0.40)	1.47*** (0.36)	8.53*** (1.06)	-2.27*** (0.83)

Table 2 (Continued)

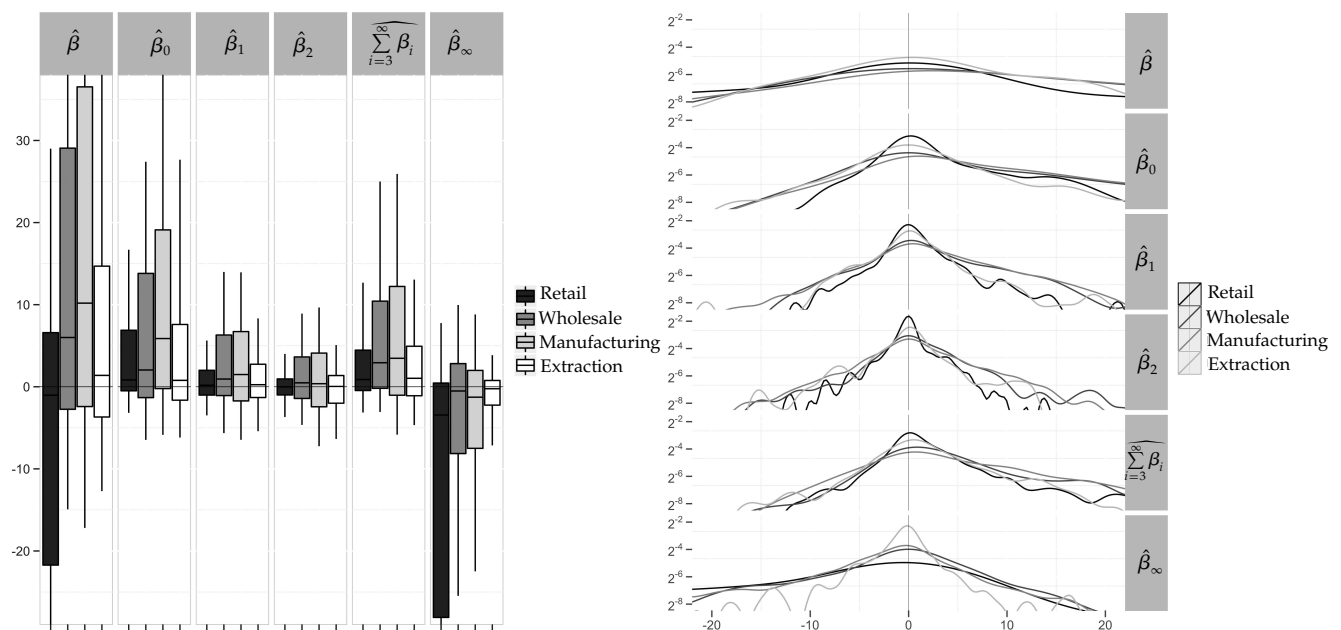
	$\hat{\beta}$	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\beta}_2$	$\sum_{i=3}^{\infty} \hat{\beta}_i$	$\hat{\beta}_{\infty}$
Manufacturing						
Transportation equipment	14.42*** (2.62)	7.43*** (1.24)	3.15*** (0.58)	1.32*** (0.31)	5.16*** (0.78)	-3.24*** (1.12)
Instruments and related	29.73*** (3.14)	15.91*** (1.50)	3.41*** (0.58)	1.52*** (0.37)	9.08*** (1.19)	-1.62** (0.80)
Miscellaneous	6.08 (5.29)	15.11*** (3.42)	3.88*** (1.19)	0.14 (0.50)	6.04*** (1.36)	-20.78*** (1.80)
Segment mean	19.62*** (1.47)	11.40*** (0.61)	3.10*** (0.21)	1.24*** (0.15)	6.43*** (0.46)	-3.87*** (0.56)
Extraction						
Metal	3.56 (3.01)	5.24*** (1.17)	1.49** (0.75)	0.32 (0.30)	0.78 (0.85)	-4.43* (2.42)
Coal	-3.91 (4.85)	-0.87 (3.93)	-2.00*** (0.44)	-1.05*** (0.27)	1.71*** (0.52)	-1.54* (0.84)
Oil and gas	14.64*** (1.34)	7.86*** (0.95)	1.46*** (0.29)	0.12 (0.22)	4.76*** (0.49)	-0.39* (0.21)
Segment mean	10.03*** (2.57)	6.56*** (0.89)	1.31*** (0.30)	0.13 (0.18)	3.28*** (0.96)	-1.80* (0.97)
Sample mean	15.81*** (1.51)	9.98*** (0.62)	2.74*** (0.21)	1.07*** (0.14)	5.80*** (0.42)	-5.02*** (0.71)

Notes. This table shows mean firm-level bullwhips, aggregated by industry, sector, and the entire economy, measured as a percent of total demand variance (e.g., a bullwhip of 10 means orders are 10% more variable than demands are). The numbers in parentheses report two-way cluster robust standard errors.
 * $p = 0.05$; ** $p = 0.01$.

decline of semiconductors in 1994" (Jorgenson 2001, p. 1). Comparing the sample-wide pre- and post-1995 bullwhips (not shown), we find the magnitudes of mean estimates $\hat{\beta}$, $\hat{\beta}_0$, $\hat{\beta}_1$, $\hat{\beta}_2$, $\sum_{i=3}^{\infty} \hat{\beta}_i$, and $\hat{\beta}_{\infty}$,

respectively, decline by 33, 41, 39, 53, 25, and 34%. The manufacturing sector significantly reduced all uncertainty bullwhips, the retail sector its lead 0 and lead 3+ bullwhips, and the extracting sector

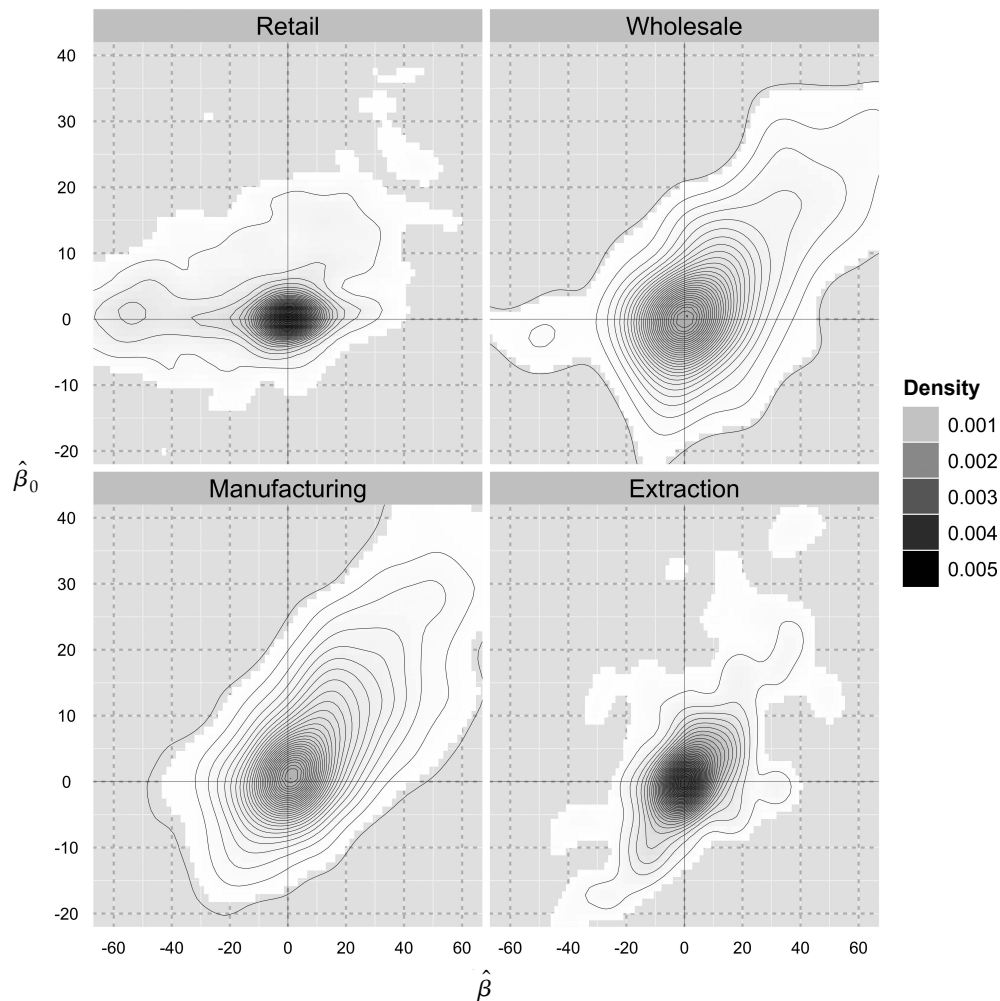
Figure 3 Marginal Bullwhip Distributions



Notes. The boxplots on the left identify bullwhip medians across firms, as well as their interquartile and interdecile ranges. The probability density functions on the right describe the entire marginal distribution of firm-level bullwhips on a log scale. We estimate these densities with kernel regressions.

Copyright: INFORMS holds copyright to this Article in Advance version, which is made available to subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

Figure 4 Joint Bullwhip Distributions



Notes. Contour plots of the joint probability density functions of $\hat{\beta}$ and $\hat{\beta}_0$. We estimate these densities with two-dimensional kernel regressions.

its lead 3+ bullwhips. Moreover, the retailing and wholesaling sectors significantly reduced the magnitudes of their seasonal bullwhips (i.e., lessened their seasonal smoothing), as, over time, these segments more tightly controlled inventories, and underlying demand seasonality dropped.

5. Robustness Checks

In this section we study four potential sources of bias in our estimations: product aggregation, temporal aggregation, forecast misspecification, and demand censoring. Table 4 summarizes the results. Although they certainly are not definitive, we cannot reject the hypothesis that there are no meaningful biases. That is, although the checks cannot disprove the existence of these biases, they increase our confidence in our results, as they do not suggest that any of them exist. Additional data and empirical methodologies may further illuminate these issues, as we discuss in the concluding remarks.

Product Aggregation. Our data aggregate across firm product offerings, which could bias bullwhip estimates (Chen and Lee 2010). In theory, this bias should work against our results, attenuating the bullwhip estimates (aggregating across products should have a similar effect as aggregating across firms, which §4 demonstrates dampens bullwhip estimates).⁹ Nevertheless, for completeness, we empirically explore the effect of product aggregation by measuring the change in our estimates attributable to further aggregation (see the online appendix for additional product-aggregation robustness checks). To create a higher degree of aggregation, we merge similar companies, fusing them into couplets by summing their sales and order quantities. This aggregation scheme simulates aggregating across a firm's products: we pool two

⁹ The dampening effect should be drastically smaller in this context, because aggregating across companies combines fewer and more similar products than does aggregating across industries.

Table 3 Bullwhip Trends

	$\hat{\beta}$	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\beta}_2$	$\widehat{\sum_{i=3}^{\infty} \beta_i}$	$\hat{\beta}_{\infty}$
Retail						
Hardware and garden	-14.65 (47.14)	-11.49*** (3.85)	2.35 (6.02)	-2.58 (6.38)	-2.57 (4.73)	8.09*** (0.89)
General merchandise	8.87 (12.21)	-3.48 (3.61)	-1.70* (0.96)	1.36 (4.20)	-1.65 (1.67)	12.38 (9.90)
Food	-4.59 (2.98)	-3.29 (2.37)	0.42 (0.51)	-0.14 (0.44)	-1.05 (1.17)	-0.71 (1.17)
Apparel and accessory	31.48*** (8.57)	-7.52 (6.43)	-1.64 (1.28)	-1.24 (4.22)	-4.56 (3.90)	41.74** (16.30)
Furniture and home furnishings	5.25 (12.62)	4.28 (5.65)	8.82* (5.05)	-5.36** (2.09)	-0.39 (1.82)	-1.85 (6.95)
Eating and drinking places	1.22*** (0.34)	0.80*** (0.23)	-1.78*** (0.08)	-0.46*** (0.07)	-0.28* (0.16)	2.64*** (0.18)
Miscellaneous	9.48 (17.94)	-5.48 (3.36)	0.90 (1.92)	-0.68 (1.01)	-4.69 (3.73)	16.95 (15.09)
Segment mean	6.73 (7.52)	-3.13* (1.87)	0.25 (0.98)	-0.82 (0.87)	-2.04* (1.08)	11.46* (6.34)
Wholesale						
Durable goods	17.99 (12.56)	-1.40 (4.54)	4.27 (7.88)	2.62 (4.60)	3.34 (4.83)	7.84* (4.52)
Nondurable goods	16.02 (11.67)	-1.23 (3.22)	2.73 (2.07)	-2.82 (1.95)	-0.01 (3.26)	14.85 (10.34)
Segment mean	17.48* (9.56)	-1.36 (3.49)	3.87 (5.92)	1.22 (3.68)	2.48 (3.86)	9.64** (4.64)
Manufacturing						
Food	-15.99* (8.20)	-11.70*** (3.44)	-1.62 (1.71)	-1.87 (1.57)	0.76 (1.37)	-2.35 (6.12)
Textile mill	-35.25 (21.76)	-14.85 (12.43)	-4.87 (3.93)	4.79 (4.22)	-3.90 (3.64)	-21.34 (19.39)
Apparel	17.83 (18.76)	-0.70 (7.78)	1.36 (1.80)	-6.56 (19.90)	0.05 (6.28)	21.35 (14.65)
Lumber and wood	0.06 (23.63)	-2.26 (9.05)	-5.28 (3.69)	-1.40 (3.11)	12.10** (5.99)	10.69 (17.72)
Furniture and fixtures	-0.44 (5.49)	-3.10 (4.78)	3.40 (4.70)	-3.27 (2.03)	0.14 (3.53)	1.54 (4.94)
Paper	-8.37 (7.36)	-4.79** (2.37)	0.61 (1.84)	-3.28** (1.59)	4.03*** (0.75)	-4.36 (4.46)
Printing and publishing	-6.48 (5.86)	-3.87 (3.09)	-1.25 (2.90)	-2.41 (1.54)	-0.69 (1.29)	3.43 (2.11)
Chemicals	-2.10 (7.00)	-5.64 (4.71)	-3.18* (1.68)	-1.76 (1.68)	0.86 (3.09)	9.21** (3.57)
Petroleum and coal	-0.77 (4.72)	-4.60 (3.58)	0.26 (0.73)	0.79 (1.68)	1.74 (2.51)	0.82 (0.56)
Rubber and plastics	11.57 (9.26)	1.73 (7.54)	-6.49 (4.97)	-1.05 (3.42)	-10.14 (6.50)	22.15*** (6.32)
Leather goods	-14.11*** (1.73)	6.85*** (0.64)	-0.43 (0.46)	-3.84*** (0.48)	-5.50*** (0.66)	-11.28*** (1.25)
Stone and glass	9.49* (5.21)	8.56*** (2.07)	-2.16 (1.81)	-0.90 (3.04)	-4.84*** (1.49)	8.32*** (2.78)
Primary metal	-13.36 (9.19)	-10.17** (4.37)	-1.59 (5.08)	2.68 (1.81)	-9.50*** (2.80)	3.19 (2.16)
Fabricated metal	-14.28 (20.47)	-10.41 (7.60)	-0.89 (2.46)	-6.23* (3.67)	-2.17 (9.20)	8.08 (6.96)
Industrial machinery	-27.73*** (8.60)	-11.88*** (2.70)	-4.49** (2.22)	-3.49*** (1.13)	-8.05** (3.22)	0.37 (3.61)
Electronic equipment	-6.75 (10.49)	-2.38 (2.84)	0.34 (2.36)	1.64 (1.54)	-2.09 (3.30)	-5.17* (3.01)

Copyright: INFORMS holds copyright to this Article in Advance version, which is made available to subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.

Table 3 (Continued)

	$\hat{\beta}$	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\beta}_2$	$\sum_{i=3}^{\infty} \hat{\beta}_i$	$\hat{\beta}_{\infty}$
Manufacturing						
Transportation equipment	-22.06** (9.06)	-11.20*** (3.80)	-5.70** (2.40)	2.70 (1.73)	-10.70*** (3.73)	2.76 (2.87)
Instruments and related	0.67 (15.23)	-2.36 (5.39)	-0.67 (3.84)	-5.27* (3.17)	2.36 (8.13)	2.36 (3.59)
Miscellaneous	-60.81*** (19.33)	-24.14*** (9.04)	-6.34 (4.37)	0.73 (3.68)	-12.52** (6.15)	-16.23 (10.26)
Segment mean	-11.59*** (3.68)	-6.67*** (1.25)	-2.04** (0.81)	-1.47** (0.59)	-2.88* (1.57)	1.15 (1.55)
Extraction						
Metal	-29.90 (19.82)	-6.54 (6.23)	1.98 (3.70)	-4.21 (4.10)	-0.91 (1.98)	-20.84 (14.46)
Oil and gas	3.38 (3.05)	1.75 (3.25)	2.61* (1.36)	1.07 (0.87)	-3.68* (2.04)	2.88 (2.15)
Nonmetallic minerals	-18.59*** (1.10)	-4.79*** (0.56)	-4.37*** (0.33)	-3.93*** (0.40)	-5.63*** (0.68)	0.13 (0.63)
Segment mean	-6.37 (7.78)	-0.80 (3.08)	1.71 (1.57)	-0.64 (1.57)	-3.28** (1.56)	-2.66 (4.77)
Sample mean	-6.86** (3.26)	-5.65*** (1.07)	-1.24 (0.78)	-1.19** (0.53)	-2.44* (1.30)	3.18** (1.55)

Notes. We estimate bullwhip changes between 1974–1994 and 1995–2008 by regressing the squared forecast errors on a constant and post-1995 indicator variable, reporting the latter's coefficients. We allow forecast processes and variances to change after 1995. We include only the 492 firms that have at least 25 clean, consecutive observations both before and after 1995. We include an equal number of observations in each subperiod, for a fair comparison.

companies because couplet-level aggregation is the next closest to firm-level aggregation, and we attempt to pool firms that sell similar products. To pair companies, we match them by four-digit SIC and the mean inventory-to-sales ratio.¹⁰ As Table 4 demonstrates, running our analysis across couplets yields nearly the same results as those in Table 2. So we do not find evidence of a meaningful product aggregation bias.

Temporal Aggregation. Our quarterly data are temporally aggregated. According to Chen and Lee (2010), temporal aggregation should attenuate bullwhip estimates: a positive “bullwhip ratio tends to decrease as the aggregation period increases” (Chen and Lee 2010, p. 13). Thus, like product aggregation, we have no reason to believe this feature of our data inflates our estimates. Nevertheless, we study its effect with the monthly, industry-level Census Bureau data analyzed by Cachon et al. (2007). We measure the effect of temporal aggregation, increasing the level of aggregation from one month, to two, to three.¹¹ Table 4 shows that the bullwhip estimates remain qualitatively unchanged as the level of temporal aggregation varies from one to three months.

¹⁰ Matching on other variables yields similar results.

¹¹ The two-month aggregation combines January and February, March and April, etc. And the three-month aggregation combines annual quarters. (Naturally, different aggregation schemes will yield different results.)

Forecast Misspecification. Misspecifying the demand and order forecasts can bias our estimates—but only to an extent, because $\hat{\beta}$ and $\hat{\beta}_{\infty}$ do not rely on these forecasts, and thus neither does the sum of the uncertainty-bullwhip estimates, $\sum_{i=0}^{\infty} \hat{\beta}_i = \hat{\beta} - \hat{\beta}_{\infty}$. What is sensitive to our forecast specification is the allocation of uncertainty bullwhips to information lead times. That is, forecast misspecification can lead us to attribute part of β_i to $\hat{\beta}_i$, but it cannot create any additional uncertainty bullwhip, as that quantity is fixed.

We measure our results' sensitivity to forecast specification by repeating our analysis with three alternative sets of explanatory variables: The first uses eight quarters of lagged demands and orders rather than four. The second uses four quarters of lagged demands and orders, but includes gross domestic product, total industrial production index, average three-month commercial paper interest rate, aggregate sales and production of the firm's two-digit SIC, and the change in firm store counts, if it is a retailer (see Gaur et al. 2005). The third includes these variables and uses eight quarters of lagged demands and orders.¹² Table 4 demonstrates that the coefficients'

¹² To accommodate additional forecast variables, we increase our firm-length cutoff to 30, 35, and 42 quarters, for the first, second, and third specifications, respectively.

Table 4 Robustness Checks

	$\hat{\beta}$	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\beta}_2$	$\sum_{i=3}^{\infty} \hat{\beta}_i$	$\hat{\beta}_{\infty}$
Baseline	15.81*** (1.51)	9.98*** (0.62)	2.74*** (0.21)	1.07*** (0.14)	5.80*** (0.42)	−5.02*** (0.71)
Couplet	17.54*** (2.16)	9.80*** (0.75)	2.59*** (0.27)	0.91*** (0.20)	8.68*** (0.76)	−6.04*** (1.02)
Census						
1 month	11.10*** (2.21)	13.12*** (1.79)	−0.20 (0.67)	0.56 (0.49)	2.74** (1.31)	−6.63*** (1.06)
2 months	11.48*** (4.39)	9.92*** (3.51)	0.95 (1.04)	0.87 (0.96)	6.66*** (1.96)	−8.20*** (1.89)
3 months	12.57** (5.93)	8.09** (3.49)	1.46* (0.80)	1.16 (1.23)	4.83** (2.40)	−3.21* (1.81)
Alternative forecasts						
Specification 1	15.35*** (1.52)	7.76*** (0.50)	1.99*** (0.19)	0.91*** (0.17)	8.65*** (0.67)	−5.01*** (0.72)
Specification 2	14.66*** (1.47)	7.09*** (0.46)	2.05*** (0.18)	0.88*** (0.11)	8.42*** (0.64)	−4.99*** (0.71)
Specification 3	14.16*** (1.48)	5.73*** (0.39)	1.72*** (0.19)	0.71*** (0.13)	10.08*** (0.76)	−4.98*** (0.72)
Inventory quartile						
Q1	13.34*** (1.44)	9.72*** (0.66)	2.29*** (0.28)	0.85*** (0.22)	4.36*** (0.55)	−4.55*** (0.68)
Q2	11.58*** (1.40)	7.73*** (0.57)	2.24*** (0.25)	1.13*** (0.20)	4.41*** (0.43)	−4.83*** (0.72)
Q3	14.08*** (1.59)	9.02*** (0.62)	2.46*** (0.29)	1.16*** (0.20)	5.44*** (0.50)	−5.16*** (0.79)
Q4	24.27*** (2.16)	13.70*** (0.89)	3.84*** (0.34)	1.00*** (0.24)	9.12*** (0.72)	−5.33*** (0.84)

Notes. This table summarizes the results of §4’s four robustness checks. All estimates correspond to sample-wide mean bullwhips. The first row repeats our main results from Table 2. The second lists the company-couplet bullwhips. The estimates under the “Census” heading report the bullwhips in the Cachon et al. (2007) Census Bureau data, temporally aggregated at one, two, and three months. Those under the “Alternative forecasts” heading present the results under different forecast specifications. Specification 1 extends the number of lagged demands and orders from four to eight; specification 2 includes gross domestic product, total industrial production index, average three-month commercial paper interest rate, aggregate sales and production of the firm’s two-digit SIC, and the change in firm store counts, if it is a retailer (see Gaur et al. 2005); and specification 3 includes these variables and extends the number of lagged demands and orders to eight. And the bottommost estimates report the mean bullwhips of our inventory-quartile subsamples, with Q1 indicating the lowest-inventory subsample, and Q4 indicating the highest.

* $p \leq 0.1$; ** $p \leq 0.05$; *** $p \leq 0.01$.

signs and significances hold under the alternative forecast specifications.

Censoring Bias. Sales, the minimum of demand and inventory availability, is a censored variable. Inventory censoring can inflate bullwhip estimates by truncating demand, making it appear less variable. To gauge whether a censoring bias drives our results, we seek to determine whether stockouts relate to our bullwhip measure. Because we cannot observe stockouts, we use period-start inventory levels as a proxy—according to the newsvendor model, the two should strongly negatively correlate, as higher inventories generally mean fewer stockouts. Hence, if a censoring bias drove our results, we would expect an inverse relationship between the amount of on-hand inventory at period start and the measured bullwhip effect. To test this relationship, we divide our sample, by

period start inventory levels, into four subsamples and compare the mean bullwhips of each. (We used the same approach in §4.2, but there we classified observations by date, rather than by inventory level.) To control for firm and seasonal characteristics, we allocate each firm-quarter evenly between subsamples; thus, we ultimately divide our sample by the inventory quartiles of each firm in each calendar quarter.

We do not find a censoring bias signature: the mean bullwhip does not decrease across the subsamples, as inventories increase. More importantly, the bullwhip effect is strongest in the highest-inventory subsample, when stockouts, and hence demand censoring, should be least likely. However, although suggestive, this robustness check is not definitive, as it hinges on an assumed negative relationship between period-start inventory levels and stockouts.

Table 5 Bullwhip Summary

	$\hat{\beta}$	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\beta}_2$	$\widehat{\sum_{i=3}^{\infty} \beta_i}$	$\hat{\beta}_{\infty}$
Mean	15.81	9.98	2.74	1.07	5.80	-5.02
Median	6.67	4.21	0.99	0.19	2.69	-1.23
Standard deviation	43.91	20.17	9.03	7.40	14.64	19.44
Fraction of firms	0.65	0.71	0.61	0.53	0.68	0.37
Fraction of industries	0.84	0.97	0.94	0.81	0.97	0.16

Notes. This table shows sample-wide bullwhip statistics. The mean and median estimates, measured as a percent of total demand variance, are all significant at $p = 0.01$. The last lines report the fraction of firms with positive bullwhips, and the fraction of industries with positive mean-firm-level bullwhips.

6. Concluding Remarks

This paper studies the bullwhip effect in firm-level data. Table 5 summarizes our findings. Overall, we find evidence for the effect—our sample's mean and median bullwhips are significantly positive. Yet, rather than universal, we find the effect idiosyncratic, as the bullwhip varies greatly across firms. The phenomenon results from a tug-of-war between two opposing forces: uncertainty amplification and production smoothing. Our bullwhip decomposition makes these forces apparent: firms generally amplify last-minute shocks—the mean lead 0 bullwhips are positive in 97% of the industries we consider—but smooth seasonal variations—the mean seasonal bullwhips are negative in 84% of the industries. Our model predicts such seasonal smoothing.

Our estimates, however, come with several caveats: (1) we estimate bullwhips across firms, rather than across supply chains; (2) we proxy COGS for demand and production for orders, which could introduce a censoring bias; (3) we do not observe true forecasts; and (4) we use data aggregated temporally at the quarter, and cross-sectionally at the company. As a result, the bullwhip effect warrants further study. Developing a full understanding of the bullwhip effect will require comprehensive efforts by multiple researchers, as an ideal bullwhip sample—a multi-firm collection of separable supply chains, with high-frequency, product-level demand and order data—is unlikely to surface soon. Addressing any of the caveats listed above would substantially improve our perspective on the phenomenon.

The bullwhip resolves gradually over time as information about demands and order quantities is unveiled in the periods leading up to their final realizations. From this insight, we construct a decomposition of the bullwhip based on information-transmission lead times, which clarifies and enriches the bullwhip, providing an information distortion profile; rather than lump all demand variations, it demonstrates which variations firms amplify. Our decomposition identifies several bullwhip flavors: signals arriving with more than three-quarters' notice

drive 30% of the mean bullwhip, and those arriving with less than one-quarter's notice drive 51%. These bullwhip flavors have different supply chain effects—short-lead-time bullwhips, providing suppliers the least reaction time, presumably cause the most havoc. Perhaps worse than a big bullwhip is a late bullwhip.

Addressing the different bullwhip flavors requires different operational fixes. For example, Caterpillar Inc. waged a multi pronged attack on its various bullwhip components. Since 2000, Caterpillar has been engaged in a “supply chain makeover” (Songini 2000), to address “concerns about the potential disruptions that could come with a inventory bullwhip” (Aeppel 2010). The company dealt with long-lead-time bullwhips by sharing order forecasts (see Aviv 2007): since 2000 the company has been engaged in “high-speed sharing of key sales and business data throughout Caterpillar and between its product design department and the suppliers” (Songini 2000). The company addressed midrange-lead-time bullwhips by ensuring supply chain agility (see Lee 2004): Caterpillar required “a detailed written plan from its suppliers for each part they produce, explaining how the supplier will respond to the bullwhip.” Finally, it mitigated short-lead-time bullwhips by fixing orders (see Balakrishnan et al. 2004): “the company has promised to stick by ‘freeze periods’ as it transitions to growth: For a three-month span after it places an order, it promises not to change it” (Aeppel 2010). These efforts earned Caterpillar “a spot in 2010 on Gartner Inc.’s top 10 list of industrial supply chains” (Katz 2011). Perhaps more impressively, from before 2000 to after 2000, the company reduced its bullwhip profile, $\{\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2, \widehat{\sum_{i=3}^{\infty} \beta_i}\}$, from $\{18.9, 17.8, 15.1, 28.3\}$ to $\{-1.9, 2.9, -0.7, 2.8\}$.

Acknowledgments

The authors thank the associate and departmental editors, three anonymous reviewers, and the participants of the Wharton Empirical Operations Research Workshop, in particular Karen Donohue, for their helpful comments and suggestions.

References

- Aeppel, T. 2010. Bullwhip hits firms as growth snaps back. *Wall Street Journal* (January 27), <http://online.wsj.com/article/SB10001424052748704509704575019392199662672.html?mod=WSJ-hps-LEFTWhatsNews>.
- Aviv, Y. 2003. A time-series framework for supply-chain inventory management. *Oper. Res.* **51**(2) 210–227.
- Aviv, Y. 2007. On the benefits of collaborative forecasting partnerships between retailers and manufacturers. *Management Sci.* **53**(5) 777–794.
- Balakrishnan, A., J. Geunes, M. S. Pangburn. 2004. Coordinating supply chains by controlling upstream variability propagation. *Manufacturing Service Oper. Management* **6**(2) 163–183.
- Basu, S., J. G. Fernald, N. Oulton, S. Srinivasan. 2003. The case of the missing productivity growth, or does information technology explain why productivity accelerated in the United States but not in the United Kingdom? *NBER Macroeconomics Annual* (2004) 9–63.
- Cachon, G. P., M. A. Lariviere. 1999. Capacity choice and allocation: Strategic behavior and supply chain performance. *Management Sci.* **45**(8) 1091–1108.
- Cachon, G. P., T. Randall, G. M. Schmidt. 2007. Search of the bullwhip effect. *Manufacturing Service Oper. Management* **9**(4) 457–479.
- Cameron, A. C., P. K. Trivedi. 2005. *Microeconometrics: Methods and Applications*. Cambridge University Press, New York.
- Cameron, A. C., J. B. Gelbach, D. L. Miller. 2011. Robust inference with multiway clustering. *J. Bus. Econom. Statist.* **29**(2) 238–249.
- Chen, F., Z. Drezner, J. K. Ryan, D. Simchi-Levi. 2000. Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Sci.* **46**(3) 436–443.
- Chen, H., M. Z. Frank, O. Q. Wu. 2005. What actually happened to the inventories of American companies between 1981 and 2000? *Management Sci.* **51**(7) 1015–1031.
- Chen, L., H. L. Lee. 2009. Information sharing and order variability control under a generalized demand model. *Management Sci.* **55**(5) 781–797.
- Chen, L., H. L. Lee. 2010. Bullwhip effect measurement and its implications. Working paper, Fuqua School of Business, Duke University, Durham, NC.
- Cohen, M. A., T. H. Ho, Z. J. Ren, C. Terwiesch. 2003. Measuring imputed cost in the semiconductor equipment supply chain. *Management Sci.* **49**(12) 1653–1670.
- Crosron, R., K. Donohue. 2003. Impact of POS data sharing on supply chain management: An experimental study. *Production Oper. Management* **12**(1) 1–11.
- Crosron, R., K. Donohue. 2006. Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Sci.* **52**(3) 323–336.
- Davidson, R., J. G. MacKinnon. 2004. *Econometric Theory and Methods*, vol. 21. Oxford University Press, New York.
- Davis, S. J., J. A. Kahn. 2008. Interpreting the great moderation: Changes in the volatility of economic activity at the macro and micro levels. *J. Econom. Perspect.* **22**(4) 155–180.
- Dong, Y., M. Dresner, O. Yao. 2011. An empirical analysis of the value of managing information asymmetry in vendor managed inventory. Working paper, Robert H. Smith School of Business, University of Maryland, College Park.
- Ernst, R., D. F. Pyke. 1993. Optimal base stocks policies and truck capacity in a two-echelon system. *Naval Res. Logist.* **40**(7) 879–903.
- Fransoo, J. C., M. J. F. Wouters. 2000. Measuring the bullwhip effect in the supply chain. *Supply Chain Management: An Internat. J.* **5**(2) 78–89.
- Gaur, V., M. L. Fisher, A. Raman. 2005. An econometric analysis of inventory turnover performance in retail services. *Management Sci.* **51**(2) 181–194.
- Gavirneni, S., R. Kapuscinski, S. Tayur. 1999. Value of information in capacitated supply chains. *Management Sci.* **45**(1) 16–24.
- Gow, I. D., G. Ormazabal, D. Taylor. 2010. Correcting for cross-sectional and time-series dependence in accounting research. *Accounting Rev.* **85**(2) 483–512.
- Granger, C. W. J., P. Newbold. 1974. Spurious regressions in econometrics. *J. Econometrics* **2**(2) 111–120.
- Graves, S. C. 1999. A single-item inventory model for a nonstationary demand process. *Manufacturing Service Oper. Management* **1**(1) 50–61.
- Graves, S. C., D. B. Kletter, W. B. Hetzel. 1998. A dynamic model for requirements planning with application to supply chain optimization. *Oper. Res.* **46**(3-Supplement-3) S35–S49.
- Hahn, J. 1995. Bootstrapping quantile regression estimators. *Econometric Theory* **11**(1) 105–121.
- Hall, P., J. L. Horowitz, B. Y. Jing. 1995. On blocking rules for the bootstrap with dependent data. *Biometrika* **82**(3) 561–574.
- Hammond, J. H. 1994. Barilla SpA (A). HBS Case 9-694-046, Harvard Business School, Boston.
- Hausman, W. H. 1969. Sequential decision problems: A model to exploit existing forecasters. *Management Sci.* **16**(2) 93–111.
- Hausman, W. H., R. Peterson. 1972. Multiproduct production scheduling for style goods with limited capacity, forecast revisions, and terminal delivery. *Management Sci.* **18**(7) 370–383.
- Jorgenson, D. W. 2001. Information technology and the U.S. economy. *Amer. Econom. Rev.* **91**(1) 1–32.
- Jorgenson, D. W., M. S. Ho, K. J. Stiroh. 2003. Lessons for Europe from the U.S. growth resurgence. *CESifo Econom. Stud.* **49**(1) 27–47.
- Kahn, J. A. 1987. Inventories and the volatility of production. *Amer. Econom. Rev.* **77**(4) 667–679.
- Kahn, J. A., M. M. McConnell, P.-Q. Gabriel. 2002. On the causes of the increased stability of the U. S. economy. *Federal Reserve Bank of New York Econom. Policy Rev.* **1**(May) 183–202.
- Katz, J. 2011. Creating high-value supply chains. *Indust. Week* (February 16), http://www.industryweek.com/articles/creating_high-value_supply_chains_23884.aspx?SectionID=2.
- Krishnan, J., P. R. Kleindorfer, A. Heching. 2007. Demand distortions and capacity allocation policies. Working paper, International Monetary Fund, Washington, DC. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1021950.
- Lai, R. K. 2005. Bullwhip in a Spanish shop. Harvard NOM Research Paper 06-06, Harvard Business School, Boston.
- Lee, H. L. 2004. The triple-A supply chain. *Harvard Business Review* **82**(10) 102–113.
- Lee, H. L., V. Padmanabhan, S. Whang. 1997. Information distortion in a supply chain: The bullwhip effect. *Management Sci.* **43**(4) 546–558.
- Lee, H. L., K. C. So, C. S. Tang. 2000. The value of information sharing in a two-level supply chain. *Management Sci.* **46**(5) 626–643.
- Lütkepohl, H. 2005. *New Introduction to Multiple Time Series Analysis*. Springer, Berlin.
- Newey, W. K. 1984. A method of moments interpretation of sequential estimators. *Econom. Lett.* **14**(2–3) 201–206.
- Petersen, M. A. 2009. Estimating standard errors in finance panel data sets: Comparing approaches. *Rev. Financial Stud.* **22**(1) 435–480.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford Business Books, Stanford, CA.
- Songini, M. L. 2000. Caterpillar moves to revamp supply-chain operations via the Web. *Computer World* (October 11), http://www.computerworld.com/s/article/print/52290/Caterpillar_moves_to_revamp_supply_chain_operations_via_the_Web.
- Sterman, J. D. 1989. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision-making experiment. *Management Sci.* **35**(3) 321–339.
- Terwiesch, C., Z. J. Ren, T. H. Ho, M. A. Cohen. 2005. An empirical analysis of forecast sharing in the semiconductor equipment supply chain. *Management Sci.* **51**(2) 208–220.
- Wong, C. Y., M. M. El-Beheiry, J. Johansen, H. H. Hvolby. 2007. The implications of information sharing on bullwhip effects in a toy supply chain. *Internat. J. Risk Assessment Management* **7**(1) 4–18.