

Ration Gaming and the Bullwhip Effect

Online Appendix

Robert L. Bray¹, Yuliang Yao², Yongrui Duan³, and Jiazhen Huo³

¹Kellogg School of Management, Northwestern University

²College of Business and Economics, Lehigh University

³School of Economics and Management, Tongji University

December 22, 2017

A Sample Selection

Each observation in our sample reports, for a given product on a given day, (i) the sales quantity, (ii) the wholesale price, (iii) the retail price, (iv) the current price discount, (v) the store's start-of-day inventories, (vi) the DC's start-of-day inventories, (vii) the store-to-DC orders, and (viii) the DC-to-store shipments. We keep a product in our sample if (i) it has at least 500 observations;¹ (ii) at least 4% of its observations have a positive store-to-DC order; (iii) at least 4% of its observations have a positive DC-to-store shipment; (iv) at least 2% of its observations have a positive vendor-to-DC shipment; (v) at least 8% of its observations have a store inventory level change; (vi) at least 6% of its observations have a DC inventory level change; (vii) at least 80% of its shipments arrive within a day; and (viii) it is stored at the DC rather than cross-docked. The first six conditions ensure we have enough variation in our state variables to estimate our state transition probabilities, and the last two ensure that the product fits our empirical model (recall that we specified a one-day shipping lead time to curtail the number of required state variables).

We chose the cutoff values in an ad hoc fashion before estimating our model. Filtering the data took several hours, so we didn't experiment too much with these thresholds. However, for

¹ We observe some products for fewer than 500 days because not all stores carry all items all the time. The category managers choose the stores' product assortments.

a robustness check, we reran the analysis after increasing each of these thresholds by 25%. The results were qualitatively the same.

B Proposition Proofs

Proposition 1: We show that the representative store finds it optimal to follow policy ρ when the $N - 1$ other stores also follow ρ , and N is sufficiently large. We restrict ourselves to the $u_0 > N(\alpha - 1)$ case; the $u_0 \leq N(\alpha - 1)$ case is similar, except its inventory run occurs at time zero. We also restrict ourselves to the $\frac{\eta}{\eta + \mu} \leq \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)^2$ case to ensure α is at least two, the minimum amount for there to be an inventory run.

First, we show that the representative store's optimal order-up-to level is a deterministic function of the DC inventory level. Following ρ , the other stores each hold (i) one unit of inventory when the DC has more than $N(\alpha - 1)$ units; (ii) α units of inventory when the DC has between one and $N(\alpha - 1)$ units; and (iii) an irrelevant amount of inventory when the DC has zero units (when the DC is stocked out, the other stores' inventories are moot). Hence, the other stores' inventory levels are either (i) extraneous (when the DC is stocked out) or (ii) a deterministic function of the DC inventory level (when the DC isn't stocked out). So the DC inventory level characterizes the state of the system, from the representative store's perspective.

Second, we define α . At any time, the representative store's aggregate future demand has a geometric distribution with mean $\lambda\tau$, probability mass function $\phi_{\lambda\tau}$, and inverse cumulative distribution function $\Phi_{\lambda\tau}^{-1}$. (Note, mixing a Poisson with an exponential yields a geometric random variable.) Thus, the store's expected cost when it has i units of inventory and the DC has none is

$$\begin{aligned} \pi(i) &= \eta \sum_{d=0}^i (i - d) \phi_{\lambda\tau}(d) + \mu \sum_{d=i}^{\infty} (d - i) \phi_{\lambda\tau}(d) \\ &= \eta(i - \lambda\tau) + (\mu + \eta)\lambda\tau \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)^i. \end{aligned}$$

The newsvendor solution implies that π is minimized at $\Phi_{\lambda\tau}^{-1}\left(\frac{\mu}{\mu + \eta}\right) = \text{floor}\left(\frac{\ln\left(\frac{\eta}{\mu + \eta}\right)}{\ln\left(\frac{\lambda\tau}{1 + \lambda\tau}\right)}\right) = \alpha$.

Third, we establish that the representative store holds one unit of inventory when the DC has more than $N(\alpha - 1)$ units. The $\frac{\eta}{\eta + \mu} \leq \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)^2$ regularity condition ensures that the store always carries at least one unit of inventory. And the store doesn't carry more than one unit because it's

guaranteed access to α units—the newsvendor-optimal amount—as soon as the DC inventory no longer exceeds $N(\alpha - 1)$ units (i.e., at the moment of the inventory run). Thus, hoarding inventory when the DC has more than $N(\alpha - 1)$ units increases expected inventory-overflow costs without decreasing expected inventory-underage costs.

Fourth, we establish that the representative store joins the inventory run when N is large, ordering up to α as soon as the DC inventory dips below $N(\alpha - 1) + 1$ units. If the representative store doesn't join the run, then immediately after the run, (i) the representative store has one unit of inventory; (ii) the DC has $\alpha - 1$ units of inventory; and (iii) the other stores each have α units of inventory. And the other stores thereafter synchronize their orders with their sales, so that they deplete the DC's stock one unit at a time. Since the DC's inventory decrements gradually, the representative store will not hold more than one unit of inventory until it is ready to inherit the entire DC supply: as long as the DC offers inventory, the other stores will consume it, so the representative store must claim every unit to claim the marginal unit. Thus, the representative store holds one unit of inventory until the upstream inventory level reaches some threshold $u \leq \alpha - 1$, at which point it orders u . In this case, the representative store's expected cost is

$$\pi_N(u) = \left(\frac{N\lambda\tau}{1 + N\lambda\tau} \right)^{\alpha - u - 1} \pi(u + 1) + \left(1 - \left(\frac{N\lambda\tau}{1 + N\lambda\tau} \right)^{\alpha - u - 1} \right) \eta,$$

where $\left(\frac{N\lambda\tau}{1 + N\lambda\tau} \right)^{\alpha - u - 1}$ is the probability of the representative store depleting the DC inventory before the product goes obsolete, $\pi(u + 1)$ is the expected cost of holding $u + 1$ units of inventory when the DC stocks out, and η is the cost of holding one unit of inventory when the product becomes obsolete. Note, $\pi_N(u)$ converges with N to $\pi(u + 1)$, which takes its minimum value at $u = \alpha - 1$. So ordering $\alpha - 1$ at the moment of the inventory run is optimal for sufficiently large N .

Proposition 2: When $u_0 \leq N\alpha$, there is an inventory run at time zero. When $u_0 > N\alpha$, there is an inventory run when the stores sell at least $u_0 - N\alpha$ units in aggregate. The probability of at least $u_0 - N\alpha \geq 0$ demand realizations before the product goes obsolete is $\left(\frac{N\lambda\tau}{1 + N\lambda\tau} \right)^{u_0 - N\alpha}$.

Proposition 3: Each store holds α units of inventory immediately after the inventory run. Thus, each store has expected cost $\pi(\alpha)$, conditional on there being an inventory run. In contrast, if the stores continued to follow the globally optimal policy of carrying one unit of inventory while

supplies last, they would each have expected cost $\left(\frac{N\lambda\tau}{1+N\lambda\tau}\right)^{N(\alpha-1)}\pi(1) + \left(1 - \left(\frac{N\lambda\tau}{1+N\lambda\tau}\right)^{N(\alpha-1)}\right)\eta$, where $\left(\frac{N\lambda\tau}{1+N\lambda\tau}\right)^{N(\alpha-1)}$ is the probability of the DC stocking out before the part goes obsolete, $\left(1 - \left(\frac{N\lambda\tau}{1+N\lambda\tau}\right)^{N(\alpha-1)}\right)$ is the the probability of the part going obsolete before the DC stocks out, $\pi(1)$ is the expected cost when the DC stocks out before the part goes obsolete, and η is the expected cost when the part goes obsolete before the DC stocks out. Thus, an inventory run increases each store's expected cost by the following amount:²

$$\begin{aligned}
& \pi(\alpha) - \left(\frac{N\lambda\tau}{1+N\lambda\tau}\right)^{N(\alpha-1)}\pi(1) - \left(1 - \left(\frac{N\lambda\tau}{1+N\lambda\tau}\right)^{N(\alpha-1)}\right)\eta \\
&= \eta(\alpha - 1 - \lambda\tau) + (\mu + \eta)\lambda\tau \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)^\alpha + \left(\frac{N\lambda\tau}{1+N\lambda\tau}\right)^{N(\alpha-1)} \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)(\eta - \mu\lambda\tau) \\
&\geq \eta(\alpha - 1 - \lambda\tau) + (\mu + \eta)\lambda\tau \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)^\alpha + \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)^{(\alpha-1)} \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)(\eta - \mu\lambda\tau) \\
&= \eta\alpha - \eta(1 + \lambda\tau) \left(1 - \left(\frac{\lambda\tau}{1 + \lambda\tau}\right)^\alpha\right).
\end{aligned}$$

This last expression is strictly positive for $\alpha \geq 2$ because (i) it is zero when $\alpha = 1$; (ii) its derivative with respect to α is positive when $\alpha = 1$; and (iii) its second derivative with respect to α is always positive.

Proposition 4: To simplify the notation, I assume that (i) the DC fulfills all orders promptly, (ii) demand never exceeds D , and (iii) there exists order quantity q_0 such that $q_0 \in \mathfrak{Q}$, $2q_0 \in \mathfrak{Q}$, and $3q_0 \in \mathfrak{Q}$. Relaxing these assumptions is cumbersome but straightforward.

Arcidiacono and Miller's (2011) first theorem implies

$$\nu(x) = \pi(q|x) - \xi(q|x) + \beta \sum_{x' \in \mathfrak{x}} \delta(x'|x, q) \left(\pi(q'|x') - \xi(q'|x') \right) + \beta \sum_{x'' \in \mathfrak{x}} \delta(x''|x', q') \nu(x'')$$

²Note, our $\frac{\eta}{\eta+\mu} \leq \left(\frac{\lambda\tau}{1+\lambda\tau}\right)^2$ regularity condition guarantees that $\eta - \mu\lambda\tau$ is negative.

for all $\{q, q'\} \in \mathfrak{Q}^2$. This implies

$$\begin{aligned} & \pi(q_1|x) - \xi(q_1|x) + \beta \sum_{x' \in \mathfrak{X}} \delta(x'|x, q_1) \left(\pi(q'_1|x') - \xi(q'_1|x') + \beta \sum_{x'' \in \mathfrak{X}} \delta(x''|x', q'_1) \nu(x'') \right) \\ & = \pi(q_2|x) - \xi(q_2|x) + \beta \sum_{x' \in \mathfrak{X}} \delta(x'|x, q_2) \left(\pi(q'_2|x') - \xi(q'_2|x') + \beta \sum_{x'' \in \mathfrak{X}} \delta(x''|x', q'_2) \nu(x'') \right) \end{aligned}$$

for all $\{q_1, q'_1, q_2, q'_2\} \in \mathfrak{Q}^4$. This equation simplifies to

$$\lambda + \xi(0|x) - \xi(q_0|x) + \beta \eta q_0 + \beta \sum_{x' \in \mathfrak{X}} (\delta(x'|x, 0) \xi(2q_0|x') - \delta(x'|x, q_0) \xi(q_0|x')) = 0$$

when $i = D$, $q_1 = q_0$, $q'_1 = q_0$, $q_2 = 0$, and $q'_2 = 2q_0$, and to

$$\lambda + \xi(0|x) - \xi(2q_0|x) + \beta \eta 2q_0 + \beta \sum_{x' \in \mathfrak{X}} (\delta(x'|x, 0) \xi(3q_0|x') - \delta(x'|x, q_0) \xi(q_0|x')) = 0$$

when $i = D$, $q_1 = 2q_0$, $q'_1 = q_0$, $q_2 = 0$, and $q'_2 = 3q_0$. These two equations identify λ and η . Once we've pinned down these parameters, we can use a similar trick to identify μ .

References

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