JOURNAL OF BUSINESS LOGISTICS
STRATEGIC SUPPLY CHAIN RESEARCH

Journal of Business Logistics, 2018, 1–12 doi: 10.1111/jbl.12174

Consistency and Recovery in Retail Supply Chains

Elliot Bendoly[1], Nathan Craig[1], and Nicole DeHoratius[2]

[1]*The Ohio State University*
[2]*University of Chicago*

Practitioners and researchers describe inventory service level with metrics that communicate the likelihood of demand fulfillment without considering the ongoing capabilities of the supplier, for example, in-stock and fill rate. We develop a method for measuring inventory service level that incorporates such supplier capabilities, namely consistency (the ability of a supplier to fulfill orders repeatedly) and recovery (the ability of a supplier to fulfill orders after a lapse in service). Using data from two retail supply chains, we illustrate our approach. To demonstrate the impact of consistency and recovery on supply chain performance, we model a retailer purchasing from competing suppliers with different levels of consistency and recovery. The model incorporates the retailer's uncertainty about demand and the retailer's uncertainty about its suppliers' service levels. We characterize how the retailer's orders and profitability change with a supplier's delivery performance through numerical experiments calibrated with field data. We find notable differences in market share across suppliers with similar traditional inventory service level metrics but differences in consistency and recovery. Further, we observe that a retailer can increase its profitability by determining orders via consistency and recovery in lieu of common metrics like in-stock. Given the influence of consistency and recovery on supply chain outcomes, we discuss implications for practice and future research.

Keywords: product availability; inventory service level; supplier reliability; supplier performance; B2B

INTRODUCTION

Supplier reliability affects numerous supply chain outcomes, including supplier market shares, inventory at different levels of a supply chain, and retailer cost and prices (Stank et al. 2003; Dada et al. 2007; Federgruen and Yang 2009; Liu et al. 2009; Davis-Sramek et al. 2010; Craig et al. 2016). Researchers typically model suppliers with imperfect product availability using the type 1 and type 2 inventory service level metrics. Type 1 service level, or in-stock, is the probability that a supplier will fill all demand in a given period. Type 2 service level, or fill rate, is the expected proportion of demand that a supplier will fill in a given period (Nahmias 2008). The literature on supplier reliability, however, offers a dynamic perspective on inventory service level not captured by traditional metrics. Researchers and practitioners often describe service level in terms of consistency, or predictability (Dana and Petruzzi 2001; Swait and Erdem 2002; Christopher 2005; Su and Zhang 2009; Solomon 2012). On the other hand, the literature also highlights the importance of a supply chain's ability to recover from service disruptions (Bakshi and Kleindorfer 2007; Craighead et al. 2007; Sheffi 2007; Turner 2011).

We extend prior research on retailers ordering from imperfect suppliers by incorporating aspects of supplier performance identified in the literature, namely consistency (the ability of a supplier to fulfill orders repeatedly) and recovery (the ability of a supplier to fulfill orders after a lapse in service). We propose a stylized model of supplier service level that captures both consistency and recovery, and we develop a method for estimating consistency and recovery using data commonly available within supply chains. We demonstrate this method using data from a supplier of consumer packaged goods (CPG) as well as a supplier of apparel. Our results demonstrate the model's ability to capture distinctions in each supplier's performance, specifically differences between the consistency and recovery rates for each supplier. These differences reinforce the need to distinguish between these service dimensions.

To further explore consistency and recovery, we model a retailer purchasing identical products from two suppliers. The model incorporates both the retailer's uncertainty about its demand and the retailer's uncertainty about its suppliers' service levels. The retailer observes each supplier's delivery history, updates its beliefs about each supplier's consistency and recovery, and then places orders that minimize the expected market mediation costs (i.e., the expected cost of overages and underages). The retailer's beliefs may be informal—as in a buyer's opinion of a particular supplier—or formal—as in supplier scorecards (Duffy 2004). The model identifies how supply chain outcomes—in particular, market share across suppliers and retailer cost—vary with supplier performance.

To study the impact of consistency and recovery in practice, we construct numerical experiments calibrated with empirical data. The numerical experiments reveal the extent to which a supplier's orders from a retailer depend on the supplier's consistency and recovery. In particular, the numerical experiments show that equivalent performances as measured by a traditional inventory service level metric (in-stock) can result in materially different market shares across two otherwise identical suppliers. Further, the experiments show that a retailer can reduce its costs substantially by placing orders based on consistency and recovery rather than type 1 service level.

Our results suggest that metrics like consistency and recovery that capture information about supply chain dynamics can be useful to managers in retail supply chains. Suppliers that provide service levels similar to their competitors according to common metrics may find they lag the competition in market share due to differences in consistency and recovery. Further, retailers can reduce the cost of supply uncertainty by tracking the consistencies and recoveries of their suppliers.

Corresponding author:
Nathan Craig, Fisher College of Business, The Ohio State University, 644 Fisher Hall, 2100 Neil Avenue, Columbus, OH 43212, USA; E-mail: craig.186@osu.edu

LITERATURE REVIEW

Our work contributes to research in the domain of supplier reliability, particularly the business-to-business interaction between a retailer and multiple suppliers. Research in this area includes analytical studies as well as empirical studies that employ laboratory experiments and field-based data collection. Research on supplier reliability often examines imperfect, or unreliable, suppliers using models with random yields or stochastically proportional supply (Yano and Lee 1995). The downstream firm places an order for an item, and the supplier's delivery quantity is the downstream firm's order scaled by a random variable, where the random variable models the supplier's service level.

In the multiple-supplier context, a retailer may spread its orders across suppliers to reduce the risks associated with supply uncertainty. Gerchak and Parlar (1990) study a model with two suppliers, finding that, if the suppliers' costs are equal, the retailer's order quantity for a supplier increases with that supplier's service level. Anupindi and Akella (1993) find the same outcome for the two-supplier case with random demand in both single- and multiperiod contexts. In the case where the retailer qualifies its suppliers from a set of potential suppliers before placing orders, Dada et al. (2007), Burke et al. (2009), and Federgruen and Yang (2009) show that increased supplier service level leads to increased orders from a retailer among qualified suppliers.

Researchers have also used laboratory studies and field experiments to study the relationship between supplier inventory service level and orders from retailers. Gurnani et al. (2013) use laboratory experiments in a multisupplier context to determine how subjects spread orders across two suppliers, one that is perfectly reliable and more expensive, and one that is imperfect and less expensive. In these experiments, subjects increase their orders for the unreliable supplier as that supplier's service level increases. Craig et al. (2016) analyze a field experiment at an apparel supplier and find a marked increase in retailer orders for a subset of products after managers at the supplier increased the service level of those products.

Prior research highlights other aspects of supplier reliability beyond stochastically proportional supply, namely, consistency and recovery. Christopher (2005) argues that consistency, within industrial markets, may be more influential for winning orders than product or technical features and is a key driver of customer loyalty and retention. Swait and Erdem (2002) argue that "consistency in availability will increase utility because product unavailability on the shelf may force the consumer to reevaluate their commitment to the SKU" (p. 306). Davis-Sramek et al. (2010) demonstrate the importance of consistent supplier performance in shaping the perceptions and behavior of a retailer's salespeople, who, in turn, influence end consumers. Dana and Petruzzi (2001) find that higher inventory levels—and, hence, more consistent supply—attract customers. In the business-to-business context, Malmbak and Albaum (2007) determine, through a survey of retailers, that inconsistent product availability is one of the top 10 reasons that a retailer discontinues a supplier's brand.

On the other hand, firms can mitigate the adverse effects of a stockout by quickly recovering after a service disruption (Bakshi and Kleindorfer 2007). Craighead et al. (2007) identify a supplier's recovery capability as critical to performance and find that a firm's recovery capability is negatively associated with the severity of service disruptions. In other words, suppliers that have the capability to return to service after a stockout will not necessarily suffer all negative effects of stockouts. Moreover, firms that have the ability to recover can win additional business from competitors (Christopher 2005; Sheffi 2007). In the third-party logistics context, Stank et al. (2003) find that a provider's relational performance—factors including responsiveness—drives customer satisfaction and market share.

This study extends prior research in several ways. First, we operationalize the concepts of consistency and recovery by developing metrics and illustrating the use of these metrics in practice. Second, we model a retail supply chain in which a supplier's current service level depends on its prior state, capturing both the probability and the persistence of stockouts. Prior researchers have examined suppliers with state-dependent behaviors, including production processes that go "out of control" (Porteus 1986), state-dependent lead times (Song and Zipkin 1996), state-dependent costs (Ozekici and Parlar 1999), and product availability in a continuous-time setting with many suppliers, wherein all suppliers are available or unavailable at the same time (Parlar et al. 1995). In contrast, we study a discrete-time setting that models the periodic inventory cycles found in many retail supply chains.

Finally, most prior research assumes that the retailer knows its suppliers' service levels with certainty. To the best of our knowledge, Tomlin (2009) is the first to develop a model wherein a retailer does not know its suppliers' service levels but must instead measure and build beliefs about service levels. Tomlin (2009) and Chen et al. (2010) study the case in which a firm strategically places orders with a supplier to learn about the supplier's capabilities and to determine whether to engage the supplier further. In contrast, we examine the case in which the retailer and its suppliers execute ongoing supply chain relationships. The retailer does not know its suppliers' service levels with certainty and therefore reacts to changes in a supplier's performance over repeated measurements.

CONSISTENCY AND RECOVERY

In this section, we develop a method for operationalizing the concepts of consistency and recovery using data readily available among supply chain partners. We calculate consistency and recovery using data provided by suppliers to retailers in distinct industries. To the best of our knowledge, we are the first to propose a quantitative method for measuring these two constructs.

Defining consistency and recovery

We employ a stylized model of supplier service level to develop a general method for calculating consistency and recovery in practice. In this model, the supplier delivers according to a binary Markov chain (Parlar et al. 1995; Song and Zipkin 1996). State 1 represents full product availability: The downstream party's orders are filled in full when the supplier is in this state. State 0 represents a stockout: The downstream party's orders are not filled by a supplier in this state. We define the probability of transitioning out of state $s \in \{0, 1\}$ as $u_s > 0$. Therefore, *consistency*, or the probability of repeatedly filling an order, is $1 - u_1$, whereas *recovery*, or the probability of filling an order after a stockout, is u_0. Figure 1 depicts this model of product availability.

Different arrangements of consistency and recovery may lead to identical service levels according to traditional metrics. As the Markov chain is ergodic, it has a unique steady-state distribution. Let Π_{sv}^{m} be the probability that a supplier in state $s \in \{0, 1\}$ moves to state v after m transitions. The long-run probability that the supplier is in state v is then $\lim_{m\to\infty} \Pi_{sv}^{m}$. Therefore, a supplier's type 1 service level is the long-run probability that the supplier is in state 1. We denote type 1 service level by θ, which satisfies $\theta = u_0/u_0 + u_1$ (Ross 2007). Consider two suppliers with differing values of consistency and recovery but identical type 1 service levels. The first supplier has consistency equal to 0.7 with recovery equal to 0.9. The second supplier has a higher consistency of 0.9 paired with a lower recovery equal to 0.3. The type 1 service level for both suppliers is 75%. However, it is reasonable to expect that supply chain outcomes, for example, supplier market share and retailer profitability, would differ across the two supplier behaviors.

Estimating consistency and recovery

We collected data from suppliers to retailers in two separate industries. We disguise the names of these suppliers and refer to them as the CPG supplier and the apparel supplier. We calculate consistency and recovery for each supplier using a maximum-likelihood estimate approach. If the observed values of consistency and recovery are substantially different from each other, we suggest that retailers respond to these dimensions, as argued by Turner (2011) and Solomon (2012), among others. On the other hand, if the observed values of consistency and recovery are similar, then metrics like type 1 service level should describe service effectively.

We characterize each supplier's service to retailers over a period of time across an entire category of stock keeping units (SKUs). Our empirical data include the SKU-level sum of orders from all of the suppliers' retailer customers—hundreds in each case—in the relevant region on a weekly basis. For both suppliers, the orders represent demand from retailers, not just sales to retailers, as the retailers did not have access to the suppliers' inventories and therefore would place orders without knowing whether the orders would be filled. We also obtained data that allow us to determine whether the supplier filled all orders from its retailer customers (i.e., whether the supplier was in state 1 or 0). Both suppliers operated on a weekly inventory cycle wherein orders they did not fill were dropped. The CPG supplier retained data regarding the fulfillment of each customer's orders. The apparel supplier tracked whether its warehouse inventory at the end of weekly cycle was zero. We assume that the apparel supplier was in state 0 during weeks in which it ran out of inventory. As the apparel supplier fulfilled all retailer orders from the warehouse, the assumption will be violated only in the case where weekly orders exactly matched inventory, which is unlikely.

Figure 1: Binary Markov inventory service level.

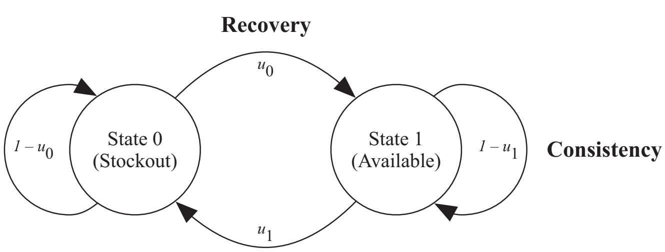

Table 1 summarizes these data, describing the number of SKUs and weeks we examined at each supplier, the order volume for each supplier, and each supplier's type 1 service level. For the CPG supplier, some of the SKUs entered and exited the category during our observation window, for example, due to promotions. The mean number of weeks the CPG supplier offered a SKU was 55 with a minimum of 25 and a maximum of 128. The apparel supplier offered each of its SKUs over the entire 140-week observation window.

With these data, we estimate values of consistency and recovery for the CPG and apparel suppliers. The maximum-likelihood estimates of consistency and recovery are as follows for a given supplier's data. Let m_{svt} be the count of transitions from state s to state v through week t of a supplier's data. We aggregate these transitions across SKUs. Let $\hat{u}_{st}$ be the maximum-likelihood estimate of a supplier's Markov chain transition probabilities through week t (Anderson and Goodman 1957). The maximum-likelihood estimates of consistency and recovery through week t for a supplier are:

$$\text{Consistency} = 1 - \hat{u}_{1t} = \frac{m_{11t}}{m_{10t} + m_{11t}} \text{ and}$$
$$\text{Recovery} \ = \hat{u}_{0t} = \frac{m_{01t}}{m_{00t} + m_{01t}}$$

Table 2 summarizes the transition counts as well as the estimates of consistency and recovery for the two suppliers. We observe substantial differences within each supplier in the values of consistency and recovery. Whereas consistency is relatively high for both suppliers, recovery is far lower. Given the magnitude of the differences in consistency and recovery within each supplier (e.g., 0.952 vs. 0.285 for the CPG supplier), it is reasonable to assume that these two aspects of supplier performance convey useful information to parties within a supply chain, and, further, that retailers should recognize and react to these two aspects of supplier performance.

The model and numerical analyses that follow study a retailer buying from two suppliers that differ only in consistency and recovery. Our objective was to determine whether measures of consistency and recovery across two similar suppliers to a single retailer (e.g., P&G and Unilever) influence supplier market shares and retailer profitability. We note that, while the empirical data show that consistency and recovery are not equal for a single supplier, these data do not necessarily show that consistency and recovery differ across competitors within an industry. Nonetheless, prior studies report differences in traditional service level across suppliers within an industry. For example, BCG (2015) documents variation in service across CPG suppliers. If traditional service level metrics differ across suppliers within an industry, then, by definition, consistency and recovery will also differ across suppliers within the industry.

Table 1: Supplier performance summary

Supplier	Number of SKUs	Number of weeks	Aggregate quantity of demand	Average SKU-level demand per cycle	SD of SKU-level demand per cycle	Aggregate type 1 service level (%)
CPG	123	128	6,698,446	842	1,605	85.7
Apparel	264	140	2,044,074	55	78	99.2

Table 2: Consistency and recovery estimates for the CPG and apparel suppliers

Supplier	m_{00T}	m_{01T}	m_{10T}	m_{11T}	Consistency $(1-\hat{u}_{1T})$	Recovery $(\hat{u}_{0T})$
CPG	661	263	269	5,390	0.952	0.285
Apparel	190	106	103	36,297	0.997	0.358

Note: The subscript T denotes that the value incorporates all observations for each supplier (i.e., 128 for the CPG supplier, and 140 for the apparel supplier).

MODEL

We model a retailer ordering perfect substitutes (Zinn and Liu 2001; McKinnon et al. 2007) procured from two suppliers with imperfect product availability to satisfy a random demand. The binary Markov model governs each supplier's product availability. The retailer does not know the true consistency and recovery of the suppliers; instead, the retailer develops beliefs about the probability of delivery in the current period (t) using each supplier's performance history over all prior periods (periods $1, \ldots, t-1$). We derive relationships between the maximum-likelihood estimates of consistency and recovery for a supplier and the retailer's order quantity for that supplier. Table 3 summarizes the notation from the model.

The retailer is a newsvendor that sells identical goods procured from two suppliers, A and B, to satisfy a random demand. At the beginning of each period, the retailer places cost-minimizing orders across both suppliers. The retailer then receives each supplier's shipment and fills a random demand, X, with distribution $F(x)$ and density $f(x)$ over the support $[0, \infty)$. Demand not served within the period in which it arrives is lost, and the retailer does not carry excess inventory at the end of a period to the next period. We assume that the variance of X is finite. The retailer faces per-unit costs of overage, c_o, and underage, c_u. The critical fractile, or the retailer's optimal type 1 service level to its own customers, is $\kappa = \frac{c_u}{c_o+c_u}$.

In period t, the retailer orders $r_i(s_i)$ from supplier $i \in \{A, B\}$, where the order is a function of supplier i's state at time $t-1$, denoted s_i. We suppress the functional dependence of r_i on s_i for brevity. The random state of supplier i at time t is $Y_i \in \{0, 1\}$. Therefore, the retailer receives $Y_A r_A + Y_B r_B$ units of product in period t, and the retailer sells $\min\{Y_A r_A + Y_B r_B, X\}$ units in period t.

To model consistency and recovery, we use the binary Markov chain model. The probability that supplier i transitions out of state s is u_{si}. The transition probabilities are unknown to the retailer, and the retailer assumes that the probabilities are independent across all s and i. The retailer must therefore develop beliefs regarding the value of u_{si}. We characterize the retailer's beliefs regarding each supplier's capabilities using the beta distribution (Tomlin 2009). The retailer updates its beliefs using a beta-Bernoulli updating process. At the outset of period 1, the retailer's beliefs regarding u_{si} take a beta distribution with

Table 3: Summary of notation

Symbol	Definition
u_{si}	The probability that supplier i transitions out of state s
m_{svti}	The number of observed transitions from state s to state v for supplier i through time t
X	The retailer's random demand
$F(x)$ and $f(x)$	The distribution and density functions of the retailer's demand
c_o, c_u, and κ	The retailer's overage cost, underage cost, and critical fractile
r_i	The retailer's order for supplier i
Y_i	The random state of supplier i in the current period, t
θ_{ti}	Supplier i's type 1 service level as observed through period t
$P_{s_i i}$	The probability that supplier i delivers in period t assuming the supplier was in state s_i during period $t-1$
$\pi_{s_i i}(\cdot)$	A beta density function representing the retailer's beliefs regarding supplier i's probability of delivery in the current period given prior state, s_i
$\alpha_{s_i i}$ and $\beta_{s_i i}$	The parameters of the retailer's belief distributions
$N(\cdot)$	The retailer's expected cost function for known current supplier states
$C(\cdot)$	The retailer's expected cost function given the retailer's beliefs and unknown current supplier states

parameters $\hat{\alpha}_{si} > 0$ and $\hat{\beta}_{si} > 0$. After observing transition counts of $m_{s,s',t-1,i}$ through period $t-1$ for states s and s' and for supplier i, the retailer's marginal posterior distribution for the probability that the supplier transitions out of state 0, u_{0i}, follows a beta distribution with parameters $a_{0i} = \hat{a}_{0i} + m_{0,1,t-1,i}$ and $b_{0i} = \hat{b}_{0i} + m_{0,0,t-1,i}$. Similarly, the retailer's beliefs about u_{1i} follow a beta distribution with $a_{1i} = \hat{a}_{1i} + m_{1,0,t-1,i}$ and $b_{1i} = \hat{b}_{1i} + m_{1,1,t-1,i}$ (Martin 1967).

We denote the retailer's beliefs about a supplier's delivery probability in the current period, t, via the random variable $P_{s_i i}$. The delivery probability depends on s_i, the supplier's state in period $t-1$. The retailer's beliefs about $P_{s_i i}$ follow a beta distribution with parameters

$$\alpha_{s_i i} = \begin{cases} a_{0i} & s_i = 0 \\ b_{1i} & s_i = 1 \end{cases} \quad \text{and} \quad \beta_{s_i i} = \begin{cases} b_{0i} & s_i = 0 \\ a_{1i} & s_i = 1 \end{cases}$$

The parameters $\alpha_{s_i i}$ and $\beta_{s_i i}$ characterize the retailer's beliefs about the probability of delivery, which is u_{0i} after a stockout and $1 - u_{1i}$ (rather than u_{1i}) after an in-stock. Hence, the dependence of $\alpha_{s_i i}$ and $\beta_{s_i i}$ on $a_{s_i i}$ and $b_{s_i i}$ reverses across the two cases.

Retailer's orders

We characterize the retailer's optimal order quantities in period t as a function of the suppliers' performances over prior periods. Where supplier i's state during the current period is Y_i, the retailer's cost with respect to random demand is

$$\begin{aligned} N(r_A, r_B, Y_A, Y_B) = c_o \mathrm{E}[Y_A r_A + Y_B r_B - X]^+ \\ + c_u \mathrm{E}[X - Y_A r_A - Y_B r_B]^+ \end{aligned}$$

As the Hessian matrix of $N(r_A, r_B, Y_A, Y_B)$ is positive definite, the retailer's cost is convex in both r_A and r_B for all realizations of supplier states.

Given the retailer's beliefs, the retailer's expected cost function is as follows, where $\bar{p} = 1 - p$:

$$\begin{aligned} C(r_A, r_B) = \int_0^1 \int_0^1 &\big[p_A p_B N(r_A, r_B, 1, 1) + \bar{p}_A p_B N(r_A, r_B, 0, 1) \\ &+ p_A \bar{p}_B N\ (r_A, r_B, 1, 0) \\ &+ \bar{p}_A \bar{p}_B c_u \mathrm{E}[X]\big] \pi_{s_A A}(p_A) \pi_{s_B B}(p_B) \mathrm{d}p_A \mathrm{d}p_B \end{aligned}$$

We denote as $\pi_{s_i i}(p)$ the density function for the beta distribution with parameters $\alpha_{s_i i}$ and $\beta_{s_i i}$. The retailer's expected delivery probability for supplier i is $\mathrm{E}[P_{s_i i}] = \frac{\alpha_{s_i i}}{\alpha_{s_i i} + \beta_{s_i i}}$, which is the simple moving average of the delivery probability conditional on the supplier's prior state. The retailer's expected cost function is convex in r_A and r_B. The first derivative of the expected cost function is

$$\begin{aligned} \frac{\partial}{\partial r_i} C(r_A, r_B) = \mathrm{E}[P_{s_A A}] \mathrm{E}[P_{s_B B}](c_o + c_u) F(r_A + r_B) \\ + \mathrm{E}[P_{s_i i}](1 - E[P_{s_j j}])(c_o + c_u) F(r_i) - \mathrm{E}[P_{s_i i}] c_u \end{aligned}$$

Since $(\partial/\partial r_i)\ C(0, 0) = -c_u \mathrm{E}[p_i] < 0$, the retailer's expected cost function attains a minimum with $r_A > 0$ and $r_B > 0$.

Let $r^*_{s_A, s_B, A}$ and $r^*_{s_A, s_B, B}$ be the retailer's optimal order quantities. For brevity, we suppress the dependence of the order quantity on the prior state and use r^*_A and r^*_B. Solving the first-order conditions $(\partial/\partial r_i)\ C(r_A, r_B) = 0$ for both suppliers yields the following conditions for the optimal orders:

$$\begin{aligned} &\mathrm{E}[P_{s_i i}] F(r^*_A + r^*_B) + (1 - \mathrm{E}[P_{s_i i}]) F(r^*_j) = \kappa \\ &\forall i \in \{A, B\}, j \in \{A, B\}, i \neq j \end{aligned} \tag{1}$$

We note that Equation (1) holds given any distribution for the retailer's beliefs regarding the probability of delivery wherein the beliefs about one supplier are independent of the beliefs about the other. However, it is worth further noting that this cannot be viewed as strictly generalizable to settings where the cost of placing orders itself is substantially high relative to other costs. While this is not the case in our empirical examples, ordering costs may be relevant in other contexts.

Having characterized the retailer's optimal orders as a function of prior supplier performance, we examine how the retailer's orders change with the suppliers' realized performances. We identify how the retailer's optimal order quantity for one supplier changes with the retailer's beliefs regarding that supplier while holding the performance of the other supplier constant.

The belief distribution parameters $\alpha_{s_i i}$ and $\beta_{s_i i}$ act as accumulators of good and bad service for supplier i. Specifically, each transition into state 1 from state s_i by supplier i increments $\alpha_{s_i i}$, whereas each transition into state 0 from state s_i increments $\beta_{s_i i}$. Therefore, $\alpha_{s_i i}$ tracks good service whereas $\beta_{s_i i}$ reflects lapses in service.

> *Proposition 1: The retailer's order quantity in the current period for supplier* i *varies with the retailer's belief distribution regarding supplier* i*'s consistency and recovery. The retailer's order quantity increases as* $\alpha_{s_i i}$ *increases and decreases as* $\beta_{s_i i}$ *increases.*

Proposition 1 characterizes how the retailer's orders change with the suppliers' historical performance. We provide proofs of all propositions in the Appendix. The following proposition describes the relationship between the retailer's orders and the maximum-likelihood estimates of consistency and recovery for the suppliers.

> *Proposition 2: The retailer's order quantity in the current period for supplier* i *increases as the maximum-likelihood estimate of consistency increases when supplier* i *had product available in the prior period. The retailer's order quantity in the current period for supplier* i *increases as the maximum-likelihood estimate of recovery increases when supplier* i *did not have product available in the prior period.*

As Proposition 2 shows, the retailer's ordering decisions vary with the maximum-likelihood estimates of consistency and recovery. However, the retailer's beliefs about the suppliers' service levels are uncertain. One way to describe the uncertainty is the coefficient of variation of the retailer's belief distribution for supplier i. The following proposition relates the retailer's uncertainty about supplier to supplier i's performance.

Proposition 3: The coefficient of variation of the retailer's belief distribution about supplier i*'s delivery probability decreases as* $\alpha_{s_i i}$ *increases and increases as* $\beta_{s_i i}$ *increases.*

Together, Propositions 2 and 3 show that, as the retailer's belief parameters ($\alpha_{s_i i}$ and $\beta_{s_i i}$) change, the retailer's order quantity for supplier i moves in the opposite direction of the coefficient of variation of the supplier's belief distribution for supplier i. In other words, the retailer places larger orders for a supplier as the coefficient of variation of the retailer's beliefs about the supplier's delivery probability decreases. This result parallels the findings of Gerchak and Parlar (1990) and Anupindi and Akella (1993) regarding the relationship between a retailer's orders and the suppliers' actual yield distributions.

In the limit, a given type 1 service level may result from different values of consistency and recovery. This relationship may also occur for the finite performance history of the two suppliers over periods 1, ... , $t-1$. The next proposition follows directly from Proposition 2.

Proposition 4: Consider two suppliers with identical historical type 1 service levels. If the suppliers differ in either consistency or recovery, then the retailer's demand for the two suppliers in the current period will differ. If the consistency of supplier i $\in$ {A, B} *exceeds that of the other supplier, supplier* j $\in$ {A, B}, j $\neq$ i, *then the retailer will place a larger order with supplier* i *than with supplier* j *during the period after a fulfillment by both suppliers. Moreover, if the recovery of supplier* j *exceeds that of supplier* i, *then the retailer will place a larger order with supplier* j *during the period after a stockout by both suppliers.*

We conclude that consistency and recovery can influence retailer ordering behavior and cost in several ways. It is therefore plausible that traditional service level metrics provide incomplete information about supply chain performance. To assess the magnitude of the effect of consistency and recovery on retailer orders and cost, we analyze the outcomes of a variety of numerical experiments.

NUMERICAL EXPERIMENTS

We construct numerical experiments to assess the impact of consistency and recovery in practice using cost, demand, and service level parameters. While the propositions of the prior section determine the effect of consistency and recovery on the retailer's orders, the intent of the numerical analysis is to assess the potential magnitude of the effect in practice. Specifically, we construct scenarios where competing suppliers are identical except for differences in consistency and recovery (i.e., the suppliers sell the same product at the same cost). For each combination of parameters, we examine how differences in consistency and recovery across suppliers affect demand from a retailer as well as the retailer's cost.

We model demand using a normal distribution truncated at 0 with mean, μ, set to 1,000 units per period. We model demand uncertainty using a range for the coefficient of variation of demand, with standard deviations, σ, of 50, 100, 500, 1,000, 2,000. This range of standard deviations includes the coefficients of variation of retailer demand we observed at both the CPG and the apparel suppliers (the coefficient of variation of SKU-level demand for the CPG supplier is 1.9 and, for the apparel supplier, 0.7), accounting for the truncation of the demand distribution. We consider several values for the critical fractile, κ, which we vary across the set {0.1, 0.25, 0.5, 0.75, 0.9} (the underage cost exceeds the overage cost when the critical fractile is >0.5). This range of costs includes values for CPG and apparel retailers that sell the products of the suppliers from which we gathered data.

In our experiment, suppliers A and B are identical in terms of traditional inventory service level but differ on consistency and recovery. We model supplier service level by assuming that the retailer has observed supplier A over T_A transitions and supplier B over T_B transitions. To compare different lengths of the relationships between the retailer and its suppliers, we fix T_B at 1,000 and solve the model with $T_A \in \{50, 100, 500, 1{,}000, 2{,}000\}$. We assume that supplier i is in the same state during periods 1 and T_i. This assumption simplifies the analysis by ensuring that $m_{10T_i i} = m_{01T_i i}$ in all scenarios (in general, $m_{10T_i i}$ and $m_{01T_i i}$ can differ by at most 1). By definition, this equality holds as T_i goes to infinity, and relaxing the assumption does not change the results. The retailer's observation of type 1 service level for supplier i is $\theta_{T_i i} = \frac{\hat{u}_{0T_i i}}{\hat{u}_{0T_i i}+\hat{u}_{1T_i i}} = \frac{m_{01T_i i}+m_{11T_i i}}{m_{00T_i i}+2m_{01T_i i}+m_{11T_i i}}$. We examine scenarios wherein type 1 service level equals 0.4, 0.6, 0.8, 0.9, or 0.95 for both suppliers.

We vary consistency and recovery as a function of $m_{01T_i i}$ using the following relationships. The number of transitions from state 0 to state 0 is $m_{00T_i i} = (1-\theta_{T_i i})T_i - m_{01T_i i}$. The number of transitions from state 1 to state 1 is $m_{11T_i i} = \theta_{T_i i}T_i - m_{01T_i i}$. Then, consistency is $\frac{m_{11T_i i}}{m_{10T_i i}+m_{11T_i i}} = 1 - \frac{m_{01T_i i}}{\theta_{T_i i}T_i}$. Further, recovery is $\frac{m_{01T_i i}}{m_{00T_i i}+m_{01T_i i}} = \frac{m_{01T_i i}}{(1-\theta_{T_i i})T_i}$. For a given scenario comprising the cost and demand parameters as well as a type 1 service level, we fix consistency and recovery for supplier B while varying consistency and recovery for supplier A as a function of m_{01T_A}. As the transition counts must be positive, for a given service level and length of history, we have $m_{01T_i i} < \min\{\theta_{T_i i}T_i, (1-\theta_{T_i i})T_i\}$. We vary $m_{01T_i i}$ from 5 to the upper bound, avoiding lower numbers of transitions between states that are not likely for retail supply chains (e.g., a single transition from state 1 to state 0 and back). We fix supplier B's performance such that supplier B's recovery is equal to one-third of its consistency, a relationship that approximates our observations of the CPG and apparel suppliers (see Table 2). Hence, we consider a supplier that varies its consistency and recovery (supplier A) in competition with a supplier with service similar to that of the CPG or apparel supplier (supplier B).

The values of demand, costs, length of supplier history, and service level combine for a total of 875 scenarios. Across these scenarios, the ranges of consistency and recovery for supplier A yield 139,330 numerical experiments. For each experiment, we calculate the retailer's optimal order for both suppliers using the first-order conditions (Equation 1) for all four combinations of prior supplier state (state 0 or 1 for both suppliers). We implemented the numerical experiments using MATLAB version 8.6.0 (MathWorks, Natick, MA).

The numerical experiments reveal that varying the values of consistency and recovery can lead to different market shares for the suppliers, even when the traditional service levels for the two suppliers are identical. Moreover, the experiments show that the retailer may substantially reduce its cost by determining orders based on consistency and recovery instead of type 1 service level. To understand how the retailer's orders for supplier *A* and *B* compare to each other, we examine symmetric scenarios across the two suppliers. Two examples are as follows. First, we compare the retailer's orders for supplier *A* to those for *B* when both suppliers delivered in the prior period. Second, we compare the retailer's orders for supplier *A* when, in the previous period, supplier *A* filled in full while supplier *B* stocked out to the retailer's orders for supplier *B* when supplier *A* stocked out in the previous period while supplier *B* filled in full.

Supplier market shares

We examine how the retailer's orders for the two suppliers vary as a function of supplier *A*'s consistency and recovery for all combinations of prior supplier states. Figure 2 plots supplier *A*'s order as a percentage of the retailer's total order (i.e., the sum of the orders for both suppliers) for two levels of demand uncertainty (coefficient of variation, ρ, equals 0.5 or 1) and cost (the critical fractile, κ, equals 0.25 or 0.75). The service level the retailer observed from both suppliers over the 1,000 period history is 0.8. Supplier *B*'s consistency is 0.92, and supplier *B*'s recovery is 0.31. Supplier *A*'s consistency and recovery vary across their respective potential ranges, [0.75, 0.98] and [0.02, 0.99], respectively. The length of the suppliers' histories is equal, that is, $T_A = T_B = 1{,}000$. Figure 2 demonstrates that the retailer's orders for the two suppliers fluctuate substantially across different values of consistency and recovery, even when holding type 1 service level constant. Therefore, suppliers with similar service levels under common measures may experience substantial differences in market share when consistency and recovery impact retailer decisions.

Figure 2 reveals relationships that hold for each combination of demand uncertainty and cost. As supplier *A*'s consistency increases, the retailer shifts its orders toward supplier *A* when both suppliers were in-stock during the prior period (upper-left panel of Figure 2). The change in market share under this condition is substantial: Supplier *A* moves from receiving an order close to zero to receiving nearly the full demand from the retailer. Nonetheless, as supplier *A*'s consistency increases, supplier *A*'s recovery must decrease due to the constant type 1 service level. Therefore, supplier *A*'s order from the retailer after both suppliers stock out decreases with supplier *A*'s consistency (lower-right panel).

When, during the prior period, one supplier was in stock while the other stocked out, we observe two distinct patterns. In the case where the supplier was stocked out while its competitor was in stock (lower-left panel), we see that supplier *A*'s order from the retailer relative to supplier *B*'s first decreases and then increases with supplier *A*'s consistency. This result is due to the combination of consistency and recovery—as supplier *A*'s consistency increases, its recovery decreases. Therefore, supplier *A*'s order from the retailer when supplier *A* was stocked out while supplier *B* was in-stock decreases as supplier *A*'s recovery decreases. Similarly, supplier *B*'s order from the retailer when supplier *B* was stocked out while supplier *A* was in stock decreases as supplier *A*'s consistency increases. In the case where the supplier was in-stock while its competitor stocked out

Figure 2: How consistency and recovery affect the retailer's orders at different levels of demand uncertainty and costs (type 1 service level is 0.8).

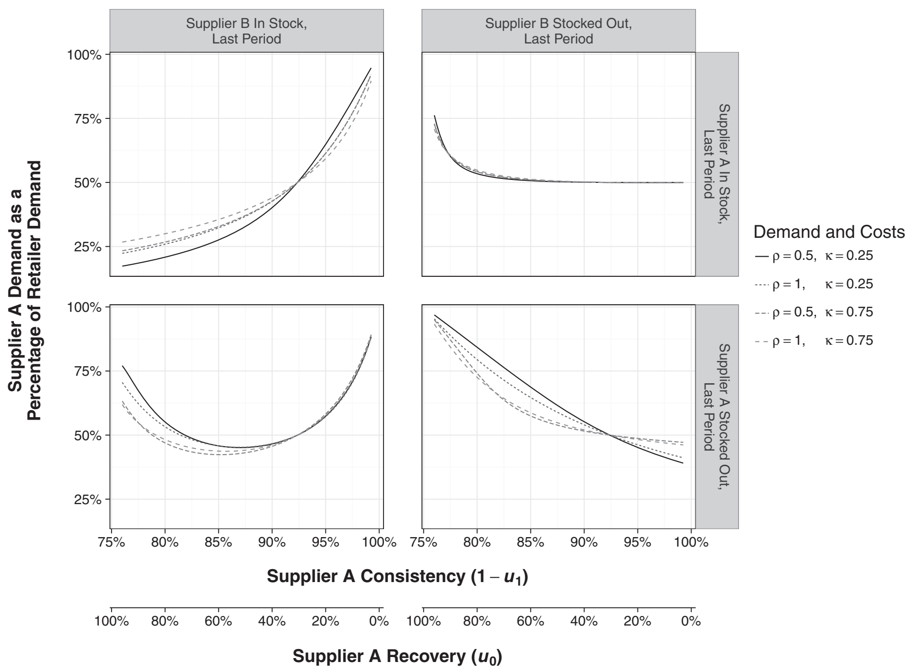

(upper-right panel), we observe that supplier A's order from the retailer decreases relative to supplier B's order as consistency increases. As supplier A's recovery decreases, supplier B's order from the retailer when supplier B was in-stock and supplier A was stocked out increases.

The relationships observed above hold for each combination of experimental parameters. Figure 3 illustrates this result for different type 1 service levels. The demand depicted by this figure has a coefficient of variation, ρ, of 0.5 and a critical fractile, κ, of 0.25. As service level increases, the ranges of potential values of consistency recovery decrease; however, the general trade-offs we observe continue to hold.

As the retailer's orders change with consistency and recovery for all combinations of supplier states, we must assess how the retailer's long-run orders—and, hence, the supplier's market shares—change with consistency and recovery. We calculate the retailer's expected orders for a given supplier performance history without conditioning on the prior state. The probability that supplier i is in state 1 during the prior period is $\theta_{T_i i}$, and the probability that supplier i is in state 0 during the prior period is $(1-\theta_{T_i i})$.

The expected order for supplier i combines the probability of each arrangement of prior supplier states with the order that supplier i receives from the retailer in each case: $\theta_{T_i i}[\theta_{T_j j} r^*_{11i} + (1-\theta_{T_j j}) r^*_{10i}] + (1-\theta_{T_i i})[\theta_{T_j j} r^*_{01i} + (1-\theta_{T_j j}) r^*_{00i}]$. Figure 4 plots the expected retailer order for supplier A as a percentage of the expected retailer total retailer order. The figure incorporates two levels of demand uncertainty ($\rho = 0.5$ and $\rho = 1$) as well as two levels of cost ($\kappa = 0.25$ and $\kappa = 0.75$). The figure depicts a constant service level of 0.8.

From Figure 4, we observe that—for all levels of demand uncertainty and cost—supplier A's expected order from the retailer does not exceed supplier B's expected order until supplier A's consistency exceeds supplier B's consistency (fixed at 0.92 in this example). Therefore, consistency may, when type 1 service level is relatively high, be a more important driver of market share than recovery. This result has implications for the design of retail supply chains: For example, managers may wish to focus on reliability of supply rather than on the ability to quickly cover shortages when trade-offs between the two exist. On the other hand, when consistency is low (toward the left side of Figure 4), improving recovery may be more valuable to the supplier.

Different supplier history lengths affect our observations regarding changes in consistency and recovery. If the retailer has had a shorter relationship with supplier A than supplier B, and thus has less data about supplier A than supplier B, how does the difference affect the retailer's orders? Proposition 3 shows that the retailer's order for a supplier increases as the retailer's uncertainty about that supplier's performance decreases. For the scenarios we study, we find that the effect of the length of the retailer–supplier relationship is small relative to the impact of changes in consistency and recovery. As an example, we consider all scenarios with a supplier type 1 service level of 0.8, a coefficient of variation of demand equal to 0.5, and a critical fractile of 0.75. We find that supplier A's steady-state average market share across all possible values of consistency and recovery improves by 1.8% when supplier A's history increases from 100 to 2,000 periods. Hence, suppliers with a longer record relative to their competitors can earn a greater market share, although the effect may be small relative to the effect of changes in consistency and recovery on market share.

Retailer cost and profitability

To determine the benefit to retailers of using consistency and recovery, we consider two cases for each numerical scenario. In the first case, the retailer places optimal orders based on consistency and recovery. In the second case, the retailer places optimal orders according to type 1 service level. By comparing the retailer's cost across the two scenarios, we can determine the benefit to the retailer of using consistency and recovery instead of type 1 service level.

A retailer that relies on type 1 service level alone does not respond to state-dependent supplier information. Hence, we use the conditions from Equation (1) with the expected probability of delivery equal to the type 1 service level to calculate the orders from a retailer that does not employ consistency and recovery. We calculate the steady-state retailer cost for both cases for the retailer's orders using the same method as for the steady-state supplier market shares.

Figure 3: How consistency and recovery affect the retailer's orders at different type 1 service levels ($\rho = 0.5$, $\kappa = 0.25$).

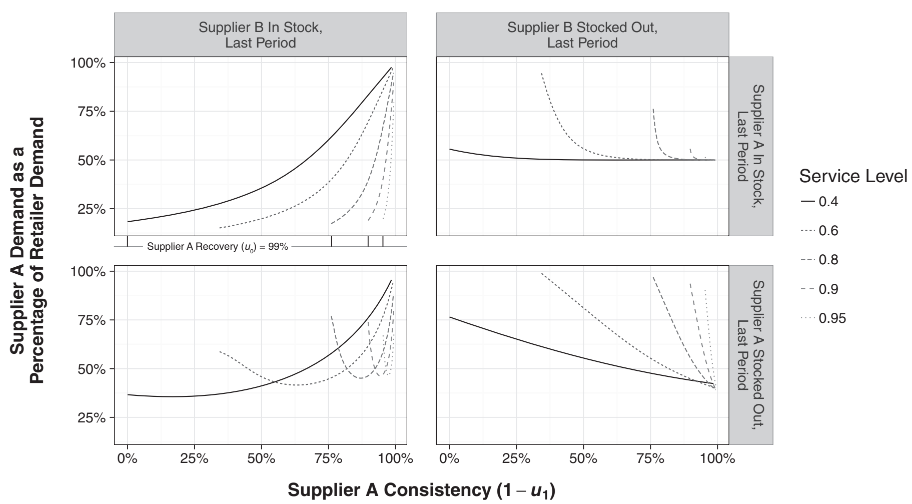

Figure 4: How consistency and recovery affect the retailer's expected orders at different levels of demand uncertainty and costs (type 1 service level is 0.8).

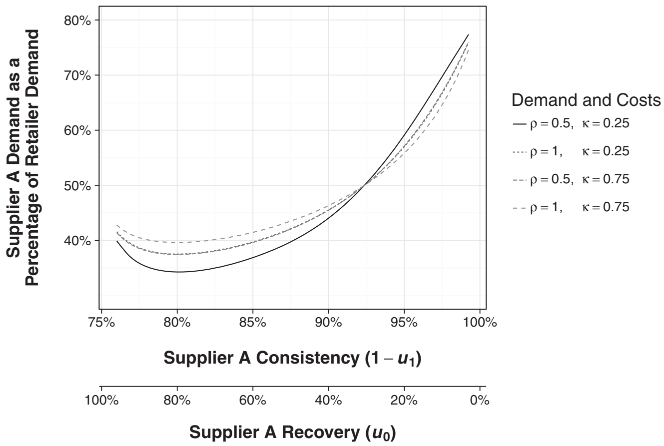

Across all numerical experiments, the retailer reduces its cost by 6.5%, on average, by ordering based on consistency and recovery instead of type 1 service level. Figure 5 depicts the cost savings for a coefficient of variation of demand of 0.5 or 1, a critical fractile of 0.25 or 0.75, and a type 1 service level of 0.8. We observe that the cost savings is highest when consistency and recovery depart substantially from a generalized view of type 1 service level, that is, toward the extreme values of recovery and consistency (which are again emblematic of our empirical observations of CPG and apparel suppliers). The cost savings decreases as the retailer's demand uncertainty increases. Further, the cost savings increases as the retailer's critical fractile increases.

The retailer's average cost savings tends to be larger for higher supplier type 1 service levels. For example, the retailer's average cost savings for a supplier type 1 service level of 0.6 is 7.3%. In contrast, the retailer's average cost savings for a type 1 service level of 0.8 is 10.5%. These results reveal opportunities for retailers to reduce the cost of supply uncertainty, and, hence, to improve profitability, using consistency and recovery to determine orders.

Figure 5: Comparing the retailer's cost when determining orders through consistency and recovery instead of type 1 service level (type 1 service level is 0.8).

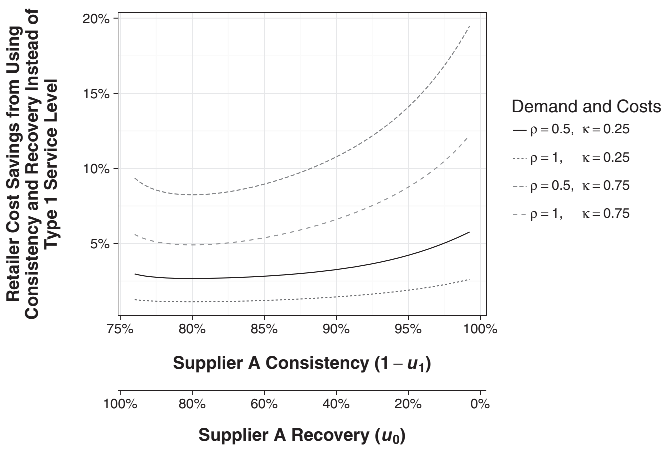

CONCLUSION

Prior research finds that supplier reliability comprises several aspects of performance, including a supplier's consistency in providing product as well as a supplier's ability to recover from a stockout. Nonetheless, common inventory service level metrics do not capture these aspects of service. We propose a stylized model of supplier performance that incorporates consistency and recovery. In a retail supply chain context, we illustrate how to estimate consistency and recovery using data from suppliers in two distinct industries. We observe that suppliers to retailers exhibit marked differences across their ability to consistently deliver a product and their ability to recover from a stockout. The differences in the values of consistency and recovery within both suppliers support prior arguments in the literature that consistency and recovery represent distinct facets of supplier performance that supply chain parties should monitor and manage.

Based on these results, we model a retailer ordering from suppliers that have differing values of consistency and recovery. In keeping with practice, the retailer does not know the true service level of its suppliers but must instead develop beliefs using each supplier's performance. Therefore, the retailer faces uncertainty about both demand and the suppliers' service levels. The model shows that suppliers that appear similar based on common inventory service level metrics may receive very different orders from retailers due to contrasts in consistency and recovery. This suggests that suppliers that perform similarly according to traditional metrics could earn different market shares among their retailer customers based on their consistency and recovery capabilities.

We assess the magnitude of the differences in market shares across suppliers with similar type 1 service levels using numerical experiments. In these experiments, we fix the performance of one supplier to represent the relationship between consistency and recovery we observed in practice, and we vary the other supplier's consistency and recovery. We conduct the numerical experiments using ranges of demand uncertainty, cost, and service level that reflect a variety of supply chain contexts. These experiments suggest that varying both consistency and recovery has material impacts on retailer orders.

We find that the retailer can decrease its inventory system cost by placing orders based on consistency and recovery instead of type 1 service level. The average cost reduction we observe is 6.5%. The retailer's cost reduction (1) increases as the supplier's type 1 service level increases, (2) decreases as the retailer's demand uncertainty increases, and (3) increases as the retailer's optimal service level to its own customers increases.

Our study is not without limitations. First, we employ a stylized model of consistency and recovery wherein a supplier's performance in the current period depends only upon its performance in the prior period. In practice, supplier performance may exhibit more complex dependencies that future research could explore. Moreover, researchers could examine whether alternate measures of consistency and recovery outperform those presented herein and test the extent to which these measures augment supplier scorecard measures commonly used in practice.

Research could also extend our model to incorporate features such as multiple products, ordering costs, and inventory carrying. Future models could examine how dimensions of supplier performance affect a retailer's decision to single- or multisource. Further, research could explore the impact that an industry's competitive intensity may have on the consistency and recovery framework. It is plausible that the impact of consistency and recovery differs depending on a multitude of product category or industry characteristics (e.g., brand strength and availability of substitutes). Finally, future research could use laboratory experiments to explore in detail how consistency and recovery affect individuals' ordering decisions.

Managerial implications

Our research reveals the potential value of managing inventory using the metrics of consistency and recovery, which capture distinct dimensions of supplier performance. As many common service level metrics combine these dimensions, retailers and suppliers should consider evaluating supply chain performance in terms of consistency and recovery. Suppliers that use traditional inventory service level metrics to benchmark against competitors' performance may miss aspects of supplier service level that materially affect market shares. Moreover, retailers may increase profitability using consistency and recovery in addition to type 1 service level to determine their orders for suppliers.

Through numerical experiments informed by data from suppliers in two industries, we find that adjusting supplier consistency and recovery may cause significant changes in retailer orders. This result holds even when the suppliers maintain the same type 1 service levels. The presence of such distinctions has both tactical and strategic implications and may suggest different approaches to bolstering demand for those firms with a broader view of service. Specifically, depending on the capabilities of a supplier, there appear to be opportunities to invest further in one aspect of performance—even at the expense of another—provided that the supplier can maintain its overall service level. In other words, there exist opportunities to pursue meaningful service differentiation, even within a given type 1 inventory service level. Managers should consider these opportunities not only in making stocking and supply chain design decisions but also when considering other factors that interact with service level, for example, assortment choices.

Finally, retailers can reduce the cost of managing inventory by placing orders based on consistency and recovery in addition to type 1 service level. Unlike type 1 service level and related measures, consistency and recovery allow retailers to systematically incorporate a supplier's current product availability into the immediate ordering decision. The retailer can thus reduce the cost of supply uncertainty and improve profitability by better allocating its orders across suppliers.

ACKNOWLEDGMENTS

This article benefited from discussions with Ananth Raman as well as with numerous industry partners. We thank the editors and four anonymous reviewers for their constructive comments and guidance.

APPENDIX

Proof of Proposition 1

We prove the proposition for the relationship between r_i^* and $\alpha_{s_i i}$. The proof of the relationship between r_i^* and $\beta_{s_i i}$ is similar. The first-order conditions from Equation (1) are

$$\mathrm{E}[P_{s_j j}]F(r_A^* + r_B^*) + (1 - \mathrm{E}[P_{s_j j}])F(r_i^*) = \kappa$$

and

$$\mathrm{E}[P_{s_i i}]F(r_A^* + r_B^*) + (1 - \mathrm{E}[P_{s_i i}])F(r_j^*) = \kappa$$

where $\mathrm{E}[P_{s_i i}] = \frac{\alpha_{s_i i}}{\alpha_{s_i i} + \beta_{s_i i}}$. Solving the first of the two conditions for $F(r_A^* + r_B^*)$ and substituting the result into the second condition yields

$$F(r_i^*) = \frac{\alpha_{s_j j}\beta_{s_i i}}{\alpha_{s_i i}\beta_{s_j j}}\left[\left(\frac{\alpha_{s_i i}\beta_{s_j j}}{\alpha_{s_j j}\beta_{s_i i}} - 1\right)\kappa + F(r_j^*)\right]$$

or

$$\beta_{s_i i}F(r_j^*) + \alpha_{s_i i}\left[\kappa + \frac{\beta_{s_j j}}{\alpha_{s_j j}}[\kappa - F(r_i^*)]\right] - (\alpha_{s_i i} + \beta_{s_i i})\kappa = 0$$

The latter equation is of the form $\psi_\alpha(r_i^*, \alpha_{s_i i}) = 0$. Hence, we obtain the derivative of r_i^* with respect to $\alpha_{s_i i}$ via implicit differentiation. The partial derivatives with respect to $\alpha_{s_i i}$ and r_i^* are

$$\frac{\partial}{\partial \alpha_{s_i i}}\psi_\alpha(r_i^*, \alpha_{s_i i}) = \frac{\beta_{s_j j}}{\alpha_{s_j j}}[\kappa - F(r_i^*)]$$

and

$$\frac{\partial}{\partial r_i^*}\psi_\alpha(r_i^*, \alpha_{s_i i}) = -\frac{\alpha_{s_i i}\beta_{s_j j}}{\alpha_{s_j j}}f(r_i^*)$$

Therefore,

$$\frac{\mathrm{d}r_i^*}{\mathrm{d}\alpha_{s_i i}} = -\frac{\frac{\partial\psi_\alpha(r_i^*, \alpha_{s_i i})}{\partial \alpha_{s_i i}}}{\frac{\partial\psi_\alpha(r_i^*, \alpha_{s_i i})}{\partial r_i^*}} = \frac{\kappa - F(r_i^*)}{\alpha_{s_i i}f(r_i^*)}$$

Substituting for κ using the first-order condition gives

$$\frac{\mathrm{d}r_i^*}{\mathrm{d}\alpha_{s_i i}} = \frac{1}{\alpha_{s_i i}f(r_i^*)}\mathrm{E}[P_{s_j j}][F(r_i^* + r_j^*) - F(r_i^*)] > 0$$

Proof of Proposition 2

The retailer's estimate of recovery for supplier i, $\hat{u}_{0i}$, increases as the count of transitions from state 0 to 1 increases ($m_{0,1,t-1,i}$) and decreases in the count of transitions from state 0 to 0 ($m_{0,0,t-1,i}$). Further, α_{0i} increases in $m_{0,1,t-1,i}$ whereas β_{0i} decreases with $m_{0,0,t-1,i}$. Therefore, α_{0i} increases with $\hat{u}_{0i}$ while β_{0i} decreases with $\hat{u}_{0i}$. By Proposition 1, r_i^*, increases as $\hat{u}_{0i}$ increases. A similar proof yields the relationship between r_i^* and consistency.

Proof of Proposition 3

The coefficient of variation of the retailer's belief distribution is

$$v_{s_i i} = \frac{\sqrt{\mathrm{Var}[P_{s_i i}]}}{\mathrm{E}[P_{s_i i}]} = \sqrt{\frac{\beta_{s_i i}}{\alpha_{s_i i}(\alpha_{s_i i} + \beta_{s_i i} + 1)}}$$

Since

$$\frac{\partial v_{s_i i}}{\partial \alpha_{s_i i}} = -\frac{\sqrt{\beta_{s_i i}}(2\alpha_{s_i i} + \beta s_i i + 1)}{2[\alpha_{s_i i}(\alpha_{s_i i} + \beta_{s_i i} + 1)]^{3/2}} < 0$$

$v_{s_i i}$ decreases with $\alpha_{s_i i}$. Further, since

$$\frac{\partial v_{s_i i}}{\partial \beta_{s_i i}} = \frac{\alpha_{s_i i} + 1}{2\sqrt{\alpha_{s_i i}\beta_{s_i i}(\alpha_{s_i i} + \beta_{s_i i} + 1)^{3/2}}} > 0$$

$v_{s_i i}$ increases with $\beta_{s_i i}$.

REFERENCES

Anderson, T.W., and Goodman, L.A. 1957. "Statistical Inference About Markov Chains." *The Annals of Mathematical Statistics* 28(1):89–110.

Anupindi, R., and Akella, R. 1993. "Diversification Under Supply Uncertainty." *Management Science* 39(8):944–63.

Bakshi, N., and Kleindorfer, P. 2007. "Co-Opetition and Investment for Resilience in Global Supply Chains." Working Paper, The Wharton School.

BCG 2015. "Time to Shift Gears: Top Trends in the CPG Supply Chain." https://www.bcgperspectives.com/content/articles/consumer-products-supply-chain-management-time-shift-gears-top-trends/

Burke, G.J., Carrillo, J.E., and Vakharia, A.J. 2009. "Sourcing Decisions With Stochastic Supplier Reliability and Stochastic Demand." *Production and Operations Management* 18(4):475–84.

Chen, M., Xia, Y., and Wang, X. 2010. "Managing Supply Uncertainties Through Bayesian Information Update." *IEEE Transactions on Automation Science and Engineering* 7(1):24–36.

Christopher, M. 2005. *Logistics and Supply Chain Management*. 3rd ed. Harlow, UK: Prentice Hall.

Craig, N., Dehoratius, N., and Raman, A. 2016. "The Impact of Supplier Inventory Service Level on Retailer Demand." *Manufacturing & Service Operations Management* 18(4):461–74.

Craighead, C.W., Blackhurst, J., Rungtusanathamn, M.J., and Handfield, R.B. 2007. "The Severity of Supply Chain Disruptions: Design Characteristics and Mitigation Capabilities." *Decision Sciences* 38:131–56.

Dada, M., Petruzzi, N.C., and Schwarz, L.B. 2007. "A Newsvendor's Procurement Problem When Suppliers are Unreliable." *Manufacturing and Service Operations Management* 9(1):9–32.

Dana, J.D., and Petruzzi, N.C. 2001. "Note: The Newsvendor Model With Endogenous Demand." *Management Science* 47 (11):1488–97.

Davis-Sramek, B., Germain, R., and Stank, T.P. 2010. "The Impact of Order Fulfillment Service on Retailer Merchandising Decisions in the Consumer Durables Industry." *Journal of Business Logistics* 31(2):215–30.

Duffy, M. 2004. "How Gillette Cleaned Up Its Supply Chain." *Supply Chain Management Review* 8(3):20–27.

Federgruen, A., and Yang, N. 2009. "Optimal Supply Diversification Under General Supply Risks." *Operations Research* 57(6):1451–68.

Gerchak, Y., and Parlar, M. 1990. "Yield Randomness, Cost Tradeoffs, and Diversification in the EOQ Model." *Naval Research Logistics* 37(3):341–54.

Gurnani, H., Ramachandran, K., Ray, S., and Xia, Y. 2013. "Ordering Behavior Under Supply Risk: An Experimental Investigation." *Manufacturing & Service Operations Management* 16(1):61–75.

Liu, S., So, K.C., and Zhang, F. 2009. "Effect of Supply Reliability in a Retail Setting With Joint Marketing and Inventory Decisions." *Manufacturing & Service Operations Management* 12(1):19–32.

Malmbak, V., and Albaum, M. 2007. "Market and Brand Analysis." *TextilWirthschaft*.

Martin, J.J. 1967. *Bayesian Decision Problems and Markov Chains*. New York, NY: John Wiley & Sons.

McKinnon, A.C., Mendes, D., and Nababteh, M. 2007. "In-Store Logistics: An Analysis of On-Shelf Availability and Stockout Responses for Three Product Groups." *International Journal of Logistics: Research & Applications* 10(3):251–68.

Nahmias, S. 2008. *Production and Operations Analysis*. 6th ed. New York, NY: McGraw-Hill.

Ozekici, S., and Parlar, M. 1999. "Inventory Models With Unreliable Suppliers in a Random Environment." *Annals of Operations Research* 91:123–36.

Parlar, M., Wang, Y., and Gerchak, Y. 1995. "A Periodic Review Inventory Model With Markovian Supply Availability." *International Journal of Production Economics* 42(2):131–36.

Porteus, E.L. 1986. "Optimal Lot Sizing, Process Quality Improvement and Setup Cost Reduction." *Operations Research* 34(1):137–44.

Ross, S.M. 2007. *Introduction to Probability Models*. 9th ed. Burlington, MA: Elsevier Academic Press.

Sheffi, Y. 2007. *The Resilient Enterprise*. Cambridge, MA: MIT Press.

Solomon, M.B. 2012. *Slow Steam Ahead. DC Velocity*, April, 37–39.

Song, J.S., and Zipkin, P.H. 1996. "Inventory Control With Information About Supply Conditions." *Management Science* 42(10):1409–19.

Stank, T.P., Goldsby, T.J., Vickery, S.K., and Savitskie, K. 2003. "Logistics Service Performance: Estimating Its Influence on Market Share." *Journal of Business Logistics* 24(1):27–55.

Su, X., and Zhang, F. 2009. “On the Value of Commitment and Availability Guarantees When Selling to Strategic Consumers.” *Management Science* 55(5):713–26.

Swait, J., and Erdem, T. 2002. “The Effects of Temporal Consistency of Sales Promotions and Availability on Consumer Choice Behavior.” *Journal of Marketing Research* 39(3):304–20.

Tomlin, B. 2009. “Impact of Supply Learning When Suppliers Are Unreliable.” *Manufacturing and Service Operations Management* 11(2):192–209.

Turner, M.D. 2011. “Be Prepared to Bounce Back: Building a Resilient Supply Chain.” https://www.strategyand.pwc.com/media/file/Strategyand_Be-prepared-to-bounce-back.pdf.

Yano, C.A., and Lee, H.L. 1995. “Lot Sizing With Random Yields: A Review.” *Operations Research* 43(2):311–34.

Zinn, W., and Liu, P.C. 2001. “Consumer Response to Retail Stockouts.” *Journal of Business Logistics* 22(1):49–71.

SHORT BIOGRAPHIES

Elliot Bendoly (PhD Indiana University) is Associate Dean for Undergraduate Students and Programs, and Fisher College of Business Distinguished Professor in the Management Sciences, at The Ohio State University. He holds the title Operations Management Distinguished Scholar at the Academy of Management, and was the Caldwell Research Fellow at Emory University. He has published 53 peer-reviewed articles, with 24 in the UT Dallas/FT listings (ISI h-index 23; Google h-index 35). Elliot is the author of “Visual Analytics for Management” (www.ma-vis.com) and “Excel Basics to Black Belt” (www.excel-blackbelt.com), and his LinkedIn discussion forums, “Excel Basics to Blackbelt” and “Operations Management in Practice,” boast 28,000+ and 22,000+ members, respectively.

Nathan Craig (PhD Harvard Business School) is an Assistant Professor at The Ohio State University’s Fisher College of Business. His research interest is the use of novel data science techniques to improve managerial decisions and knowledge, with a focus on the retail industry. Nathan also holds master and bachelor degrees from The Ohio State University. He is a recipient of the 2016 Ralph Gomory Best Industry Studies Paper Award.

Nicole DeHoratius (DBA Harvard Business School) is an Adjunct Faculty Member at the University of Chicago. She also received her MSc from the University of Sussex as a Rotary Ambassadorial Scholar, and her AB, magna cum laude, from Harvard College. She has taught operations management, service operations, and supply chain management to executives in programs and companies around the globe. In 2015 and 2016, she earned the Rotman School of Management Teaching Award for exemplifying the characteristics students most value in their professors. The Manufacturing & Service Operations Management (M&SOM) Society honored her with its Best Paper Award in 2009 and she is a recipient of the 2016 Ralph Gomory Best Industry Studies Paper Award.