UC Berkeley

UC Berkeley Previously Published Works

Title

Does a manufacturer benefit from selling to a better-forecasting retailer?

Permalink

<https://escholarship.org/uc/item/8fk882bb>

Journal

Management Science, 56(9)

ISSN

0025-1909

Authors

Taylor, TA

Xiao, W

Publication Date

2010-09-01

DOI

10.1287/mnsc.1100.1204

Peer reviewed

Does a Manufacturer Benefit from Selling to a Better-Forecasting Retailer?

Terry A. Taylor

Haas School of Business, University of California, Berkeley, Berkeley, California 94720, taylor@haas.berkeley.edu

Wenqiang Xiao

Stern School of Business, New York University, New York, New York 10012, wxiao@stern.nyu.edu

This paper considers a manufacturer selling to a newsvendor retailer that possesses superior demand-forecast information. We show that the manufacturer's expected profit is convex in the retailer's forecasting accuracy: The manufacturer benefits from selling to a better-forecasting retailer if and only if the retailer is already a good forecaster. If the retailer has poor forecasting capabilities, then the manufacturer is *hurt* as the retailer's forecasting capability improves. More generally, the manufacturer tends to be hurt (benefit) by improved retailer forecasting capabilities if the product economics are lucrative (poor). Finally, the optimal procurement contract is a quantity discount contract.

Key words: supply chain contracting; asymmetric information; forecasting

History: Received November 7, 2008; accepted April 14, 2010, by Aleda Roth, operations and supply chain management. Published online in *Articles in Advance* July 21, 2010.

1. Introduction

Although a retailer's skill at accurately forecasting market demand most obviously and directly impacts the retailer, retailer forecasting accuracy impacts the entire supply chain including the manufacturer that supplies the retailer. The accuracy of a retailer's forecast impacts how much she orders and her in-stock performance, which in turn impact manufacturer profitability (Cederlund et al. 2007).¹

Because they are closer to the end customer, retailers often have better information about market demand than the manufacturer. However, the degree of superiority in forecasting demand varies (1) across retailers and (2) over time. First, some retailers are known to be better forecasters than others. For example, in the retail consumer electronics industry, Circuit City has been plagued by weak forecasting capabilities and trails behind best-in-class retailer Best Buy (Feldman and Cramer 2004, Widlitz 2005). How does a retailer's effectiveness in forecasting influence her

attractiveness to a manufacturer that is selecting a retail partner? When faced with a pool of prospective retailers, *ceteris paribus*, should a manufacturer select a retailer that has strong, weak, or intermediate forecasting capabilities?

Not only do forecasting capabilities vary across retailers, they also vary over time at a single retailer. For example, a manufacturer's retail partner may invest in forecasting capabilities (e.g., by purchasing relevant software) with the intention of improving its forecasting accuracy. Alternately, the retailer may disinvest in forecasting (e.g., by laying off or redeploying forecasting staff), understanding that this action will degrade its ability to accurately forecast demand. What impact should the manufacturer anticipate that such changes by its retail partner will have on the manufacturer's own performance? Should a manufacturer relish and encourage either improved or worsened retailer forecasting accuracy?

It may seem natural that a manufacturer would benefit by selling to a better-forecasting retailer in that by doing so, production decisions can be made with more accurate information, reducing the cost of supply/demand mismatch. However, selling to a better-informed retailer puts the manufacturer at a strategic disadvantage relative to the retailer. The retailer may be able to use her informational advantage to extract a larger portion of the system profit from the manufacturer. The impact of improved retailer forecasting accuracy on manufacturer profit

¹ As evidence that manufacturers care about retailer forecasting accuracy, a number of manufacturers have embarked on initiatives that have improving the accuracy of their retailers' forecasts as a central objective. For example, Fraser (2003) reports on a survey of 120 companies on Collaborative Planning, Forecasting and Replenishment initiatives and puts "improvements in trading partner forecasting accuracy" (p. 75) at the top of the list of benefits anticipated by survey respondents. The benefits listed subsequently flow from this improved accuracy: reduced out-of-stocks, improved service levels and increased sales.

depends on the trade-off between these two factors, and it is this trade-off that we explore in this paper.

This paper considers a manufacturer selling to a newsvendor retailer that possesses superior demand-forecast information. We show that the manufacturer's expected profit is convex in the retailer's forecasting accuracy: The manufacturer benefits from selling to a better-forecasting retailer if and only if the retailer is already a good forecaster. If the retailer has poor forecasting capabilities, then the manufacturer is *hurt* as the retailer's forecasting capability improves. More generally, the manufacturer tends to be hurt (benefit) by improved retailer forecasting capabilities if the product economics are lucrative (poor). Further, the optimal procurement contract is a quantity discount contract.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 describes the model. Sections 4 and 5 contain the analysis for the integrated system and the decentralized system, respectively. Section 6 provides numerical results. Section 7 provides concluding remarks. The proofs are in the appendix.

2. Literature Review

There is substantial literature that studies supply chain settings in which firms have distinct demand information. This literature can be classified into two streams. One stream considers the impact of the truthful sharing of private demand information (e.g., Cachon and Fisher 2000, Lee et al. 2000, Aviv 2001). Demand-information sharing sometimes is not difficult to achieve, e.g., when demand information only consists of historical sales data that are readily verifiable from the information system. In this case, the interesting questions include how to use the shared information to improve supply chain performance and what factors are crucial in affecting the magnitude of the improved performance. However, when demand information also involves the firms' subjective assessment or private knowledge that is not verifiable by a third party, the credibility of truthful information sharing is in doubt because a firm may have incentive to misrepresent its information. (For example, in our setting, the privately informed retailer has incentive to persuade the manufacturer that market demand is weak (regardless of the actual market condition) so as to convince the manufacturer to offer more generous terms (e.g., a lower purchase price).) Indeed, in practice, the scope for opportunistic behavior and lack of trust have proven to be substantial obstacles to demand-forecasting collaboration efforts (Flidner 2003).

A second stream of literature focuses on how self-interested firms behave and interact in the face of private information, and in particular on how contracts

mediate these interactions (Cachon and Lariviere 2001, Arya and Mittendorf 2004, Özer and Wei 2006, Burnetas et al. 2007, Mishra et al. 2007, Ren et al. 2010). A typical theme in this stream of work is to explore how contracts should be designed and then to evaluate the performance of optimal contracts and/or simple and commonly used contracts. See Cachon (2003) and Chen (2003) for reviews. Our work fits within this stream. However, the focus of our work is distinct: we concentrate on the impact of the retailer's forecasting accuracy on the firms' performance.

The paper most closely related to our work is Miyaoka and Hausman (2008). Similar to our work, they study a supply chain where the upstream firm (supplier) sells to a downstream newsvendor (manufacturer) who has private demand-forecast information. Unlike our work, the upstream firm must make a capacity decision. Even so, we share the same objective, which is to evaluate the impact of the downstream firm's forecasting accuracy on the firms' performance. However, they restrict attention to the single wholesale price contract, whereas we study the issue under both the wholesale price contract and the optimal procurement contract, with the emphasis on the latter. Furthermore, our results complement theirs in that they obtain analytical results for the two extreme cases of forecasting accuracy (the downstream firm is either perfectly informed or completely uninformed) and provide numerical results for intermediate cases, whereas we obtain analytical results for the full spectrum, albeit with a simpler model. In particular, we provide a more complete characterization as to how the upstream firm's performance changes in the downstream firm's forecasting accuracy (e.g., the convexity property). The impact of a privately informed firm's forecasting accuracy on supply chain performance has also been discussed, mainly through numerics, in Özer and Wei (2006) and Taylor (2006). In a setting where a manufacturer and retailer share the same demand-forecast information, Iyer et al. (2007) observe that the manufacturer may be better off when forecasting accuracy is poor because this mitigates double marginalization.

Our work is also related to the economics and accounting literature that studies the optimal level of information asymmetry. Lewis and Sappington (1991) and Rajan and Saouma (2006) consider a principal-agent model where the agent privately exerts effort that influences the output. The agent has private information about his cost of effort. The authors examine the impact of the accuracy of the agent's private information on the principal's utility. They establish an "all-or-nothing" result: the principal prefers to deal with either a completely uninformed agent or a perfectly informed agent. Even though our supply chain setting with asymmetric demand information

is quite different, it is interesting that we obtain a roughly parallel result, namely, that a manufacturer facing a pool of retailers prefers to deal with either the best or the worst forecaster.

3. Model

A manufacturer (he) produces a product at unit cost c and sells to a newsvendor retailer (she), who then sells at a fixed retail price p to a market with random demand D in a single selling season. The market demand is normally distributed, i.e., $D \sim N(\mu_0, \sigma_0)$. The salvage value of unsold inventory is normalized to zero.

Prior to the selling season, the retailer observes a demand forecast $S = D + \varepsilon$, where $\varepsilon \sim N(0, \sigma_1)$ is independent of D . Note that S is an unbiased estimator of D , with the estimation error being normally distributed. It follows from the conjugate property of normal distribution that the posterior demand distribution under the forecast S is also normal (see Winkler 1981), i.e.,

$$D|_S \sim N(a^2\mu_0 + (1-a^2)S, a\sigma_0), \quad (1)$$

where

$$a \equiv \frac{\sigma_1}{\sqrt{\sigma_0^2 + \sigma_1^2}}$$

denotes the fraction of the original demand uncertainty, as measured by the standard deviation, that remains after the forecast is observed. We refer to a as the retailer's *forecasting accuracy parameter*. The lower the value of a , the more accurate the retailer's forecast. In the limiting case where $a = 0$, the forecast perfectly reveals the exact demand. In the opposite limiting case where $a = 1$, the forecast contains no valuable information about demand and the posterior distribution is identical to the prior. For expositional simplicity, we exclude these two extreme cases and restrict attention to $a \in (0, 1)$. In considering the impact of changes in forecasting accuracy, we assume that the distribution of the underlying demand D (i.e., the parameters μ_0 and σ_0) is fixed and only the level of noise in the retailer's forecast (as captured in σ_1) varies.

The retailer privately observes her own forecast. However, the distributions and all other parameters are common knowledge of both the manufacturer and the retailer.² Thus, the manufacturer

knows that the retailer's forecast is normally distributed, i.e., $S \sim N(\mu_0, \sqrt{\sigma_0^2 + \sigma_1^2})$, or equivalently, $S \sim N(\mu_0, \sigma_0/\sqrt{1-a^2})$. It is convenient to rewrite $S = \mu_0 + (\sigma_0/\sqrt{1-a^2})\Theta$, where $\Theta \sim N(0, 1)$. Because there is a one-to-one mapping between S and Θ , we also refer to Θ as the retailer's forecast. Given the retailer's forecast $\Theta = \theta$, the posterior distribution in (1) is

$$D|_{\Theta=\theta} \sim N(\mu_0 + \sqrt{1-a^2}\sigma_0\theta, a\sigma_0). \quad (2)$$

Note that the retailer's forecasting accuracy parameter a impacts both the mean and standard deviation of posterior demand given a forecast: The more accurate the forecast (smaller a), the larger the weight of forecast ($\sqrt{1-a^2}\sigma_0$) in determining the posterior mean, and the smaller the posterior standard deviation ($a\sigma_0$). (Although, for compactness, we represent the retailer's forecast as a scalar $\Theta = \theta$, the retailer's forecast of demand should be thought of as a distribution (2) rather than as a point forecast.)

We assume both the manufacturer and the retailer are risk neutral, maximizing their own expected profits. Because the retailer possesses superior information about the demand, the uninformed manufacturer faces a typical adverse selection problem in contracting with the informed retailer. In such a situation, a procurement contract can be represented by a transfer payment schedule $T(q, \theta)$, which specifies the payment the retailer makes to the manufacturer when the retailer orders q units and reports observing forecast θ .

The sequence of events is as follows: First, the manufacturer specifies a transfer payment schedule $T(q, \theta)$ (without knowing the retailer's forecast), and the retailer (privately) observes the forecast Θ . Second, the retailer orders q units and reports a forecast $\hat{\theta}$; the manufacturer fulfills the retailer's order and receives the payment $T(q, \hat{\theta})$ from the retailer. Third, the market demand D is realized, and the retailer receives sales revenue $p \min(D, q)$. (Although, for generality, we describe the transfer payment schedule as depending on the retailer's reported forecast, we subsequently will show that the optimal transfer payment schedule can be described as depending only on the retailer's order quantity. Thus, the sequence of events simplifies to the manufacturer offering a transfer payment schedule $T(q)$ and the retailer responding by choosing an order quantity q , which more closely resembles managerial practice.)

the retailer has carried in the past). To the extent that this awareness and understanding is reasonably good, the assumption that the manufacturer knows or can infer the retailer's forecasting accuracy is reasonable, at least as an approximation. However, in some cases the manufacturer may lack an understanding of the retailer's forecasting capabilities (e.g., if the manufacturer is unfamiliar with the retailer and relevant public information is scarce); to address this case, a more complex model capturing this additional dimension of private information is required.

²Typically, a manufacturer has some awareness of its (current or prospective) retailer's demand-forecasting investments in computer systems, software, and staff. Further, the manufacturer may have an understanding of the retailer's historic forecasting performance (either directly, or reflected in the retailer's history of stocking too much or too little) and the retailer's familiarity with the manufacturer's product (based, for example, on what products

Our primary goal is to examine the impact of the retailer's forecasting accuracy on the manufacturer's performance under the optimal procurement contract, and we pursue this goal in §§5 and 6. Our main managerial insight is that improvement in the retailer's forecasting accuracy hurts (benefits) the manufacturer when the retailer is a weak (strong) forecaster. At the conclusion of §5, we provide evidence that this main managerial message is robust to the assumption of normally distributed demand, to the assumption of the nonlinear contract form, and to the assumption that the retail price is exogenous.

4. Integrated System

As a benchmark, we first briefly examine the impact of the retailer's forecasting accuracy on the performance of the integrated system, where there is a single decision maker. After observing the demand forecast $\Theta = \theta$, the system faces the demand $D |_{\Theta=\theta}$ with distribution (2). Letting $D_\theta \equiv D |_{\Theta=\theta}$, we can write

$$D_\theta = \mu_0 + \sqrt{1 - a^2} \sigma_0 \theta + a \sigma_0 X, \quad (3)$$

where $X \sim N(0, 1)$. Let $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal density and distribution function, respectively, and $\bar{\Phi}(\cdot) \equiv 1 - \Phi(\cdot)$. The system decides how much to produce by solving the following newsvendor problem:

$$\max_q E_{D_\theta} \{p \min(D_\theta, q) - cq\}, \quad (4)$$

which can be rewritten as (see, e.g., Porteus 2002)

$$\max_z \{(p - c)(\mu_0 + \sqrt{1 - a^2} \sigma_0 \theta) - a \sigma_0 \Gamma(z)\}, \quad (5)$$

where $z \equiv (q - \mu_0 - \sqrt{1 - a^2} \sigma_0 \theta) / (a \sigma_0)$ and $\Gamma(z) \equiv p[\phi(z) - z\bar{\Phi}(z)] + cz$. The first term in (5) is the profit on the mean demand, and the second term is the expected cost of supply/demand mismatch with order quantity $q = \mu_0 + \sqrt{1 - a^2} \sigma_0 \theta + a \sigma_0 z$:

$$a \sigma_0 \Gamma(z) = (p - c) E_{D_\theta} [D_\theta - q]^+ + c E_{D_\theta} [q - D_\theta]^+,$$

where $x^+ \equiv \max(x, 0)$. Because $\Gamma(z)$ is strictly convex and is minimized at $z^l = \bar{\Phi}^{-1}(c/p)$, given the forecast $\Theta = \theta$, the system's optimal production quantity is

$$q^l(\theta) = \mu_0 + \sqrt{1 - a^2} \sigma_0 \theta + a \sigma_0 z^l, \quad (6)$$

and the system's expected profit is

$$\Pi^I(\theta) = (p - c)(\mu_0 + \sqrt{1 - a^2} \sigma_0 \theta) - a \sigma_0 \Gamma(z^l). \quad (7)$$

Because Θ is standard normal, the system's expected profit is

$$\Pi^I = E_\Theta \Pi^I(\Theta) = (p - c)\mu_0 - a \sigma_0 \Gamma(z^l).$$

Intuitively, the integrated system benefits from improved forecasting accuracy (i.e., smaller a): With better forecast information, the system makes a better-informed production quantity decision, which reduces the cost of supply/demand mismatch. In our case this is manifest by the fact that Π^I decreases in a , a consequence of the well-known result that under normal demand, a newsvendor's expected profit decreases in the standard deviation of demand. Further, $(d/da)\Pi^I = -\sigma_0 \Gamma(z^l) = -\sigma_0 p \phi(z^l)$, and $p \phi(z^l)$ is increasing in p and increasing in c on $c \in (0, p/2)$ and decreasing in c on $c \in (p/2, p)$ (Qi and Zhu 2010). The implication is that improvement in forecasting accuracy is of greater value to the integrated system when the price p is high, the cost c is moderate (close to $p/2$), and the underlying demand is volatile (σ_0 is large). Because the impact of forecasting accuracy on an integrated system is well established, our contribution is in the analysis of the decentralized system, to which we turn in the next three sections.³

5. Decentralized System

Although the integrated system always benefits from improved forecasting accuracy, it is not clear whether, in the decentralized system (under the optimal procurement contract), the manufacturer will always benefit from selling to a better-forecasting retailer. In this section, first, we characterize the optimal procurement contract (Propositions 1 and 2). Second, we characterize the impact of the retailer's forecasting accuracy on the manufacturer's profit under this optimal contract (Proposition 3).

In our model setting with asymmetric information, a general procurement contract can be represented by a transfer payment $T(q, \theta)$, which is a function of the retailer's order quantity q and reported forecast θ . It follows from the revelation principle that in our model setting, finding the optimal transfer payment function is equivalent to finding the optimal menu of quantity-payment pairs $\{q(\theta), t(\theta)\}$ that induces retailer truth-telling. The retailer selects from the menu by "reporting" a forecast $\hat{\theta}$, which corresponds to selecting the contract that stipulates $q(\hat{\theta})$ units as the purchase quantity and $t(\hat{\theta})$ as the transfer payment. The menu induces truth-telling if it is in the

³ In the decentralized system, inefficiency arises because, although the manufacturer offers the contract, it is the retailer's right to choose the order quantity. If this decision right of the retailer can be transferred to the manufacturer, then the manufacturer is essentially transformed into a centralized decision maker; the manufacturer achieves the integrated system profit by asking the retailer for her forecast and then dictating that the retailer order the integrated system optimal quantity; and the impact of forecasting accuracy on the manufacturer's expected profit is identical to its impact on the integrated system's.

retailer's interest to report the forecast she actually observed.

The retailer observing a forecast $\Theta = \theta$ is referred to as the type- θ retailer. This retailer faces demand D_θ . If the type- θ retailer chooses the quantity-payment pair $(q(\hat{\theta}), t(\hat{\theta}))$, then her expected profit is

$$\mathcal{R}(\theta, \hat{\theta}) = pE_{D_\theta} \min(D_\theta, q(\hat{\theta})) - t(\hat{\theta}).$$

Let $\mathcal{R}(\theta) \equiv \mathcal{R}(\theta, \theta)$. The optimal menu is the solution to

$$\max_{\{q(\cdot), t(\cdot)\}} E_\Theta[t(\Theta) - cq(\Theta)] \quad (\text{OBJ})$$

$$\text{s.t. } \theta = \arg \max_{\hat{\theta}} \mathcal{R}(\theta, \hat{\theta}), \quad (\text{IC})$$

$$\mathcal{R}(\theta) \geq 0, \quad \text{for every } \theta. \quad (\text{IR})$$

The incentive compatibility (IC) constraint ensures that it is in the best interest of the type- θ retailer to select the quantity-payment pair $(q(\theta), t(\theta))$. The individual rationality (IR) constraint ensures that the retailer accepts the contract because her expected profit by choosing the intended contract is no less than her reservation profit, which without loss of generality, is normalized to zero. We characterize the solution in the following proposition.

PROPOSITION 1. *The optimal menu $\{q^*(\theta), t^*(\theta)\}$ is*

$$q^*(\theta) = \mu_0 + \sqrt{1 - a^2} \sigma_0 \theta + a \sigma_0 z^*(\theta) \quad (8)$$

$$t^*(\theta) = p \left(\mu_0 + \sqrt{1 - a^2} \sigma_0 \theta - a \sigma_0 [\phi(z^*(\theta)) - z^*(\theta) \bar{\Phi}(z^*(\theta))] - \sqrt{1 - a^2} \sigma_0 \int_{-\infty}^{\theta} \Phi(z^*(x)) dx \right),$$

where $z^*(\theta)$ is the unique solution to

$$a[p\bar{\Phi}(z^*(\theta)) - c] - p\sqrt{1 - a^2} \phi(z^*(\theta)) \bar{\Phi}(z^*(\theta)) / \phi(\theta) = 0. \quad (9)$$

Under the optimal menu, the type- θ retailer's expected profit is

$$\mathcal{R}(\theta) = p\sqrt{1 - a^2} \sigma_0 \int_{-\infty}^{\theta} \Phi(z^*(x)) dx, \quad (10)$$

and the manufacturer's expected profit is

$$M = E_\Theta[(p - c)\mu_0 - a\sigma_0 \Gamma(z^*(\Theta))] - p\sqrt{1 - a^2} \sigma_0 \Phi(z^*(\Theta)) \bar{\Phi}(\Theta) / \phi(\Theta). \quad (11)$$

In what follows, we first interpret the optimal menu (Proposition 2) and then turn to the impact of forecasting accuracy on the manufacturer's profit (Proposition 3). To gain a better understanding of the optimal menu, we first note that $q^*(\theta)$ strictly

increases in θ (this follows from Lemma 2 in the appendix). This is intuitive because it simply says an order quantity intended for an optimistic-forecast-observing retailer is greater than that for a pessimistic-forecast-observing retailer. The monotone property of $q^*(\theta)$ implies the existence of its inverse function, denoted by $\theta^*(q)$, i.e., $\theta^*(q^*(\theta)) = \theta$. Consequently, the optimal menu is equivalent to the payment schedule $T^*(q) \equiv t^*(\theta^*(q))$, which, by simply specifying the transfer payment for any given quantity, is a conceptually simpler way to implement the optimal menu of quantity-payment pairs $\{q^*(\theta), t^*(\theta)\}$.

Under payment schedule $T^*(q)$, $(d/dq)T^*(q)$ can be interpreted as the *marginal wholesale price*, because it is the price the retailer pays for the last unit. A payment schedule in which the marginal wholesale price is decreasing in the quantity purchased is a *quantity-discount scheme*, whereas a schedule in which the marginal wholesale price is increasing in the quantity is a *quantity-premium scheme*. In stochastic-demand settings that are distinct from our own in that, inter alia, common information is assumed, Tomlin (2003) and Cachon (2003) show that both quantity-discount schemes and quantity-premium schemes can be effective tools in encouraging efficient quantity decisions to the benefit of individual firms. In principle, in our setting with asymmetric information about demand, it is an open question as to whether the optimal payment scheme exhibits quantity discounts, quantity premia, or a combination of the two.

PROPOSITION 2. (a) *The optimal payment schedule $T^*(q)$ is a quantity-discount scheme:*

$$(d^2/dq^2)T^*(q) < 0.$$

(b) *The marginal wholesale price in the optimal payment scheme, $(d/dq)T^*(q)$, is strictly decreasing in the retailer's forecasting accuracy parameter a .*

Quantity discounts are commonly observed in practice, and distinct explanations have been offered for their use. Quantity-discount schemes have been shown to be effective tools in encouraging larger quantity decisions, to the benefit of firms, in settings with stochastic demand (Tomlin 2003, Cachon 2003) and in settings with deterministic demand but fixed order costs (Weng 1995, Corbett and de Groote 2000, Chen et al. 2001). A buyer that does not internalize a supplier's fixed order processing cost will order frequently in small batches, and so quantity discounts are a natural mechanism to encourage the buyer to order in a fashion that reflects the supplier's economies of scale. Burnetas et al. (2007) show that quantity discounts can be effective in a

setting with asymmetric demand information. Proposition 2(a) provides a stronger result: Quantity discounts emerge endogenously as an optimal response to private demand-forecast information; see Zhang et al. (2010) for a similar result in a considerably different setting.

To see the intuition as to why quantity discounts are optimal, consider the manufacturer's objectives in offering a contract: differentiating among retailers that have observed different signals, encouraging each to purchase roughly the systemwide-efficient quantity (so as to maximize system profit), and extracting a large portion of the surplus from the retailer. The intuition for the optimality of quantity discounts is easiest to see when the retailer, after observing her forecast, still faces considerable uncertainty about demand. In this case, her purchase quantity is sensitive to the marginal wholesale price. Virtually all retailers (all but those observing the most pessimistic forecasts), will anticipate being able to sell the first few units they acquire, so the marginal value of these first units will be approximately the retail price. The marginal value of additional units is decreasing, but the extent of this decrease depends on the retailer's privately observed forecast. Charging a high marginal wholesale price for the first units (nearly the retail price) and charging progressively smaller marginal wholesale prices for larger quantities accomplishes two objectives: First, it makes the (low-quantity) contracts intended for retailers that observed unfavorable forecasts unattractive to retailers that have observed favorable forecasts, which limits the profit the favorable-forecast-observing retailer can extract. More generally, making the marginal wholesale price move in tandem with the marginal value of units to the retailer limits the surplus the retailer can capture. Second, it minimizes the quantity distortion (distortion in quantity away from the systemwide-efficient quantity) for the retailers that have observed favorable forecasts, which is important because potential system profits are the largest (and hence the impacts of quantity distortions most significant) under favorable forecasts.

Proposition 2(b) establishes that as the retailer's forecast accuracy worsens (a increases), the marginal wholesale price in the optimal procurement contract decreases. To see the intuition, consider how the retailer's price sensitivity is impacted by her forecasting accuracy. If the retailer has a very precise sense of what demand will be, her purchase quantity will be insensitive to the marginal wholesale price (so long as the marginal wholesale price is less than the retail price, the retailer will purchase close to the level of demand she anticipates); consequently, it is optimal for the manufacturer to charge a high marginal wholesale price. Conversely, if the retailer has a poor sense

of demand, her purchase quantity will be sensitive to the marginal wholesale price (because she is unsure whether she will sell the units she purchases); consequently, it is optimal to charge a low marginal wholesale price because the positive impact on the purchase quantity more than compensates for the lower per-unit revenue.

The implication of Proposition 2(b) is that as the retailer's forecast accuracy deteriorates, the optimal payment schedule $T^*(q)$ "flattens." Figure 1 depicts the optimal procurement contract as a function of the retailer's forecasting accuracy. When the retailer is a weak forecaster ($a = 0.95$), the optimal contract exhibits substantial quantity discounts, for the reasons described above. In contrast, when the retailer is a strong forecaster ($a = 0.05$), the optimal contract exhibits little in the way of quantity discounts. Expanding on the intuition described above, when a is very small, the retailer that has observed forecast θ knows that demand will be very close to the posterior mean $\mu_0 + \sqrt{1 - a^2}\sigma_0\theta$, and so will purchase almost precisely this quantity so long as the marginal wholesale price is less than the retail price. Accordingly, a contract which specifies a (constant marginal) wholesale price that is slightly less than the retail price, differentiates among retailers that have observed different forecasts, encourages them to purchase nearly-systemwide-efficient quantities, and allows the manufacturer to appropriate nearly all of the system profit. Thus, a contract where the transfer payment is nearly linear in the quantity purchased is optimal.

However, as Figure 1 demonstrates, the optimal contract $T^*(q)$ may exhibit significant nonlinearity. If a nonlinear contract is undesirable, the manufacturer will achieve the same profit by instead offering a menu of *linear* contracts (or, equivalently, a menu of two-part tariffs):

$$T^*(q, \theta) \equiv T^*(q^*(\theta)) + (d/dq)T^*(q^*(\theta)) \cdot [q - q^*(\theta)]. \quad (12)$$

The linear contract intended for the type- θ retailer is simply the straight line that is tangent to the concave curve $T^*(q)$ at $q = q^*(\theta)$. Each θ corresponds to a particular linear contract, which is composed of a per-unit price $(d/dq)T^*(q^*(\theta))$ and a fixed payment $T^*(q^*(\theta)) - (d/dq)T^*(q^*(\theta)) \cdot q^*(\theta)$. The retailer can select a contract with a low per-unit price and a high fixed payment (by "reporting" a large forecast $\hat{\theta}$) or a contract with a high per-unit price and a low fixed payment (by reporting a low $\hat{\theta}$).

Before turning to our main focus—the manufacturer's profit under the optimal procurement contract—it is useful to briefly comment on the retailer's profit. As is standard in adverse selection models, the retailer's informational advantage over the manufacturer translates into profit for the

Figure 1 Optimal Procurement Contract for Different Levels of Retailer Forecasting Accuracy

Note. Parameters are $\mu_0 = 1$, $\sigma_0 = 0.5$, $p = 1$, and $c = 0.2$.

retailer (10). The source of this profit is that a type- θ retailer can threaten to select a contract intended for a retailer that has observed a less optimistic forecast, and discouraging the retailer from doing so requires making the contract intended for the type- θ retailer sufficiently attractive to her. The retailer's expected profit is

$$R = E_{\Theta}[\mathcal{R}(\Theta)] = E_{\Theta} \left[p\sqrt{1-a^2}\sigma_0 \int_{-\infty}^{\Theta} \Phi(z^*(x)) dx \right] \\ = E_{\Theta} [p\sqrt{1-a^2}\sigma_0 \Phi(z^*(\Theta)) \bar{\Phi}(\Theta)/\phi(\Theta)]. \quad (13)$$

This quantity is referred to as the retailer's expected information rent.

We now turn to the second and main topic of this section: how the retailer's forecasting accuracy impacts the manufacturer's expected profit. The manufacturer's expected profit (see Equation (11)) is equal to the expected profit from satisfying the mean demand, $(p-c)\mu_0$, minus the expected cost of supply/demand mismatch, $E_{\Theta}[a\sigma_0\Gamma(z^*(\Theta))]$, and minus the expected information rent captured by the retailer, (13). As the retailer's forecasting accuracy improves, the manufacturer can tailor its contract so that the production quantity reflects this more precise demand information, reducing the expected cost of supply/demand mismatch. On the other hand, as the retailer's informational advantage over the manufacturer increases, it is natural that the information rent captured by the retailer would increase. Whether or not the manufacturer benefits from improved retailer forecasting accuracy depends on the trade-off between the cost of supply/demand mismatch and the retailer's information rent. Because each of these quantities depend on $z^*(\theta)$, to build understanding of the impact of forecasting accuracy parameter a on the manufacturer's profit, we first examine its impact

on $z^*(\theta)$. From (8), $z^*(\theta)$ is the number of standard deviations of safety stock purchased by the θ -type retailer in the optimal contract; we refer to $z^*(\theta)$ as the *safety stock factor*.

It is easy to check that $z^*(+\infty) = z^I$; further, $z^*(\theta)$ is strictly increasing in θ (see Lemma 2 in the appendix). The implication is that only the highest-type retailer orders the system-optimal safety stock and the other types always order less. This is because the manufacturer distorts the quantities downward to limit the information rents earned by the retailer. The result of no distortion for the highest type and downward distortion for other types is a typical result in adverse selection.

LEMMA 1. For every θ , the safety stock factor $z^*(\theta)$ strictly increases in the retailer's forecasting accuracy parameter a .

In other words, as the retailer's forecasting accuracy improves (a decreases), the manufacturer's optimal contract lowers the safety stock factor for every type retailer. The intuition is as follows. As the retailer's informational advantage grows, she is able to more accurately assess the value of various quantities of units to her; consequently, when faced with a fixed menu of contracts, the retailer is able to make a better-informed contract choice, which increases her information rent. To recapture a portion of the retailer's profit, it is optimal (see Proposition 2(b)) for the manufacturer to increase the marginal wholesale price $(d/dq)T^*(q)$ across the full range of quantities q . In response, the retailer selects a contract with a smaller safety stock factor. Pushing down the safety stock factor reduces the retailer's information rent (see (13)). As the retailer's forecasting accuracy improves (a decreases), the manufacturer worries less about distorting downward the safety stock factor

$z^*(\theta)$ (from the systemwide optimum) because doing so has less impact on the retailer's order quantity $q^*(\theta)$ (see (8)) and, as a consequence, on the systemwide efficiency. Therefore, the more accurate the retailer's forecast, the more willing the manufacturer is to distort the safety stock factor to ameliorate the information rent paid to the retailer.

We now turn to the paper's main result.

PROPOSITION 3. (a) *The manufacturer's expected profit under the optimal procurement contract, M , is strictly convex in the retailer's forecasting accuracy parameter a ; there exists a threshold $\bar{a} \in (0, 1)$ such that M is strictly decreasing in a for $a \in (0, \bar{a})$ and strictly increasing for $a \in (\bar{a}, 1)$.*

(b) *The threshold $\bar{a}$ is solely determined by the ratio c/p (i.e., $\bar{a}$ is independent of all other parameters).*

Figure 2 provides a critical supplement to Proposition 3(b): $\bar{a}$ is strictly increasing in c/p .

Whether the manufacturer benefits from improved retailer forecasting accuracy depends on the retailer's current forecasting accuracy parameter a . There are two regimes, $a \in (0, \bar{a})$ and $a \in (\bar{a}, 1)$. The manufacturer benefits from the improved retailer forecasting accuracy if $a \in (0, \bar{a})$, i.e., the retailer's current forecasting accuracy is already very good; whereas the opposite is true if $a \in (\bar{a}, 1)$, i.e., the retailer's current forecasting accuracy is very poor.

First, the manufacturer benefits from improved retailer forecasting accuracy if the retailer is already very good at forecasting, i.e., $a \in (0, \bar{a})$. The intuition is as follows. When the retailer has strong forecasting capabilities, the optimal contract distorts the safety stock factor downward (Lemma 1), which causes $E_\theta[\sigma_0 \Gamma(z^*(\Theta))]$, the normalized cost of supply/demand mismatch, to be relatively large. Therefore,

further improvement in the retailer's forecasting accuracy (reducing a) significantly reduces the total cost of supply/demand mismatch, $E_\theta[a\sigma_0 \Gamma(z^*(\Theta))]$. Because the optimal contract is stingy (characterized by high marginal wholesale prices and low safety stock factors), improved retailer forecasting accuracy has a relatively minor impact on the retailer's information rent. Therefore, the manufacturer benefits from improved retailer forecasting accuracy because the positive impact on reduced supply/demand mismatch cost outweighs any potential negative impact from increased information rents.

Second, the manufacturer is hurt by improved retailer forecasting accuracy if the retailer is very poor at forecasting, i.e., $a \in (\bar{a}, 1)$. The intuition mirrors that of the strong-forecaster case. When the retailer has weak forecasting capabilities, the optimal contract is generous (characterized by low marginal wholesale prices and high safety stock factors). Under a generous contract, increasing the retailer's informational advantage over the manufacturer translates into substantially larger retailer information rent. In contrast, because the safety stock factors are close to the systemwide optimal level, the normalized supply/demand mismatch cost is small, and consequently the savings on the cost of supply/demand mismatch are minor. Consequently, when retailer forecasting accuracy improves, the losses from larger information rent dominate, and the manufacturer is hurt.

The size of the region in which the manufacturer is hurt by selling to a better-forecasting retailer depends on the value of the threshold $\bar{a}$. Because this ratio is restricted to a limited range ($c/p \in (0, 1)$), $\bar{a}$ can be completely characterized for all problem parameters in a simple figure, Figure 2. Figure 2 shows that $\bar{a}$ is strictly increasing in c/p .

So, when should a manufacturer be especially concerned that he will be hurt by selling to a better-forecasting retailer? The region in which this outcome occurs is larger ($\bar{a}$ is smaller) when the retail price p is high and the production cost c is low. The intuition is that when the production cost is a small fraction of the retail price, it is optimal to offer a generous contract so as to encourage the retailer to purchase a large quantity (large safety stock factor). As described immediately above, when the contract is generous, the losses from larger information rent dominate the savings from smaller cost supply/demand mismatch, and so the manufacturer is hurt by improved retailer forecasting.

Thus, for manufacturers that sell high-margin products (e.g., innovative products (e.g., leading-edge electronics), information goods (e.g., books), or goods with strong brands (e.g., Apple, Nike, Polo), where the production cost is small relative to the retail

Figure 2 Threshold for Manufacturer $\bar{a}$ as Function of the Production Cost to Retail Price Ratio c/p

price), there is a wider range of retailer forecasting abilities for which the manufacturer is hurt by marginal improvements. For manufacturers that sell low-margin products (e.g., mature computer hardware, where the production cost is high relative to the retail price), there is a smaller range of retailer forecasting abilities for which the manufacturer is hurt by marginal improvements.

Proposition 3 also speaks to the setting in which there is a discrete pool of prospective retailers, which requires an understanding that goes beyond the impact of marginal changes in forecasting accuracy. Consider a manufacturer that is selecting a retailer to distribute his product from a pool of N retailers. Retailer i has forecasting accuracy parameter a_i , and $0 < a_1 < a_2 < \dots < a_N < 1$. Because the manufacturer's expected profit is convex in a , it is optimal for the manufacturer to select the strongest ($a = a_1$) or weakest ($a = a_N$) forecaster. It is straightforward to check that as $a \rightarrow 0$, $M \rightarrow (p - c)\mu_0$, which is equal to the integrated system expected profit Π^I with $a = 0$. This is clearly the upper bound on the maximum expected profit that the manufacturer could possibly achieve. Therefore, if the pool includes a very strong forecaster (a_1 is sufficiently close to zero), then the manufacturer should select the strongest forecaster ($a = a_1$). If the retailers are all sufficiently weak forecasters (e.g., $a_1 \geq \bar{a}$), then the manufacturer should select the weakest forecaster ($a = a_N$). All other things being equal (i.e., holding the forecasting accuracy of each retailer fixed), this scenario is more likely to occur when the production cost is small relative to the retail price.

The implication is that manufacturers ought to avoid blindly seeking out retailers with strong forecasting capabilities. If the production cost is high or the pool of retailers contains a strong forecaster, then this naively appealing approach will serve the manufacturer well. However, if the production cost is low and the pool of retailers is not as strong, the manufacturer may benefit by selling to a retailer with inferior forecasting capabilities. We are not the first to point out that the manufacturer may benefit by selling to a retailer with inferior forecasting capabilities. Taylor (2006) and Miyaoka and Hausman (2008) provide numerical examples in which the manufacturer's expected profit is decreasing and then increasing in the retailer's forecasting accuracy; we complement this work by establishing the convexity result analytically.

The convexity of the manufacturer's profit function also has implications for the value of demand-forecast information to the manufacturer. In our base setting, the retailer has more information about market demand than the manufacturer. However, in some settings the manufacturer may be able, through its own efforts, to obtain this additional information,

eliminating the retailer's informational advantage (i.e., make the demand forecast Θ common knowledge). This may be the case, for example, when the additional demand information is set of historical sales data (which the manufacturer can piece together by working with its partners, or perhaps directly purchase) or a third-party market demand analysis. Corollary 1 characterizes the value to the manufacturer of acquiring the additional forecast information possessed by the retailer, as a function of the accuracy of that information.

COROLLARY 1. *There exists $\underline{a} < \bar{a}$ such that the manufacturer's gain in expected profit from observing the forecast Θ is strictly increasing in a for $a \in (0, \underline{a})$ and is strictly decreasing for $a \in (\bar{a}, 1)$.*

Acquiring forecast information to eliminate the retailer's informational advantage is the most valuable when the retailer's informational advantage is moderate. Intuitively, one might expect that the value of eliminating the retailer's informational advantage would be increasing in the size of the informational advantage, because a large informational advantage should translate into large informational rent for the retailer. This conjecture is incorrect because, as noted above, when the retailer's forecasting accuracy is very high, the optimal contract allows the manufacturer to capture nearly the integrated system profit. That is, in this extreme case, the manufacturer eliminates the information rents almost completely even though the retailer has superior information. This result suggests that it is not the retailer's superior information that drives the information rents, but rather the extent to which the incentive compatibility (IC) constraint has bite, i.e., the extent to which achieving efficient quantity self-selection requires distorting contracts intended for pessimistic-forecast-observing retailers so that they are unappealing to optimistic-forecast-observing retailers. When the retailer is a very strong forecaster, very little of this distortion is required because a contract with a high marginal wholesale price and small quantity is naturally unappealing to an optimistic-forecasting retailer.

An implication of Corollary 1 is that if forecasting efforts (efforts that reveal forecast Θ) are costly to the manufacturer, it is optimal for the manufacturer to exert these efforts if and only if the retailer's forecasting accuracy is moderate: $a \in (a_l, a_h)$, where a_l and a_h satisfy $0 < a_l \leq \underline{a} \leq \bar{a} \leq a_h < 1$. (The precise value of the thresholds a_l and a_h will depend on the manufacturer's cost of forecasting.)

Stepping back, this paper's main analytical result is that under the optimal procurement contract, the manufacturer's expected profit is convex in the retailer's forecasting accuracy parameter a . This implies that the

manufacturer's profit is decreasing and then increasing in any measure of forecasting accuracy that is a monotone function of a (or σ_1): Improvement in the retailer's forecasting accuracy hurts (benefits) the manufacturer when the retailer is a weak (strong) forecaster.

In the online appendix (provided in the e-companion),⁴ we provide evidence that this main managerial message is robust to the assumption of normally distributed demand, to the assumption of the nonlinear contract form, and to the assumption that the retail price is exogenous. There is one caveat in each of the latter two cases. Regarding the normal assumption, in the online appendix, we provide an alternative demand model in which the random variables associated with demand are not required to be normally distributed, but only to have an increasing failure rate. Our main result, Proposition 3(a), extends to this alternative formulation. Regarding the contractual form, under the optimal wholesale price contract (i.e., the contract in which the payment is linear in the purchase quantity), the convexity result of Proposition 3(a) extends. The caveat for this case is that for reasonable parameters (e.g., coefficient of variation $\sigma_0/\mu_0 < \sqrt{\pi}/2 \simeq 1.25$), the region in which the manufacturer is hurt by improved retailer forecasting does not exist. Regarding the retail price assumption, in the online appendix, we generalize the model in §3 to allow the retail price to be endogenous, assuming that the demand curve is isoelastic. Parallel to Proposition 3(a), there exists a threshold $\tilde{a} \in [0, 1)$ such that the manufacturer's expected profit under the optimal procurement contract is strictly decreasing in a for $a \in (0, \tilde{a})$ and strictly increasing for $a \in (\tilde{a}, 1)$. The caveat for this case is that the region in which the manufacturer benefits from improved retailer forecasting may not exist (i.e., it may be that $\tilde{a} = 0$); however, in the online appendix, we observe numerically that in the vast majority of cases, this region does exist.

Having examined the impact of the retailer's forecasting accuracy on the manufacturer analytically, we next turn to a numerical study that examines the impact of retailer forecasting accuracy on the supply chain in broader terms.

6. Numerical Study

So far, we have focused on the impact of the retailer's forecasting accuracy on the manufacturer's profit. In this section, we expand our study to examine the impact of the retailer's forecasting accuracy on the retailer's profit and the decentralized total system's profit.

At the outset, it is unclear whether the retailer benefits from having improved forecasting accuracy. On one hand, having more accurate demand information allows the retailer to make a better order-quantity decision to alleviate the cost of supply/demand mismatch. On the other hand, a more accurately forecasting retailer faces less demand uncertainty, and consequently her purchase quantity is less sensitive to her acquisition cost; consequently, the profit-seeking manufacturer may respond by offering stingier contractual terms. Whether the retailer benefits from improved forecasting accuracy depends on the trade-off between these two factors. To investigate the impact of the retailer's forecasting accuracy on her expected profit and the expected profit of the decentralized total system, we conducted a numerical study. We fixed $p = 1$ and $\mu_0 = 1$, and varied the other three parameters: $\sigma_0 \in \{0.10, 0.12, 0.14, \dots, 0.30\}$; $c \in \{0.20, 0.25, 0.30, \dots, 0.80\}$; $a \in \{0.01, 0.03, 0.05, \dots, 0.99\}$. Thus, we tested in total 7,150 instances. For each instance, we computed under the optimal procurement contract the retailer's expected profit R and the manufacturer's expected profit M . Below, first, we summarize our numerical findings. Second, we discuss the intuition for the impact of the retailer's forecasting accuracy on the retailer's and decentralized total system's expected profits. Third, we discuss the managerial implications of the numerical findings.

Our numerical observations for the impact of the forecasting accuracy parameter a on the retailer's and decentralized total system's expected profit in several ways parallel those of our main analytical result, Proposition 3, which addresses the impact on the manufacturer's expected profit. In particular, we observed that for every fixed value of σ_0 and c , there exists a threshold $a^r \in (0, 1)$ such that R is strictly increasing in a for $a \in (0, a^r)$ and strictly decreasing for $a \in (a^r, 1)$: as the retailer's forecasting accuracy improves, her expected profit increases and then decreases. For every fixed value of σ_0 and c , there exists a threshold $a^t \in (0, 1)$ such that the total system's expected profit under the optimal procurement contract $M + R$ is strictly decreasing in a for $a \in (0, a^t)$ and strictly increasing for $a \in (a^t, 1)$: as forecasting accuracy improves, system profit first decreases and then increases. Under the assumption that R ($M + R$, respectively) is unimodal, which is the case in every instance in the numerical study, one can establish analytically a result parallel to Proposition 3(b): the threshold a^r (a^t , respectively) is solely determined by the ratio of the production cost to the retail price c/p . Figure 3 supplements Figure 2 by depicting the thresholds a^r and a^t , in addition to $\tilde{a}$. There are two regimes: if $c/p < 0.4$, then $\tilde{a} < a^r$; otherwise, $\tilde{a} > a^r$. Figure 4 illustrates the impact of the forecasting accuracy parameter on the retailer's, manufacturer's, and decentralized total system's expected profits.

⁴ An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Figure 3 Thresholds for Manufacturer $\bar{a}$, Retailer a^r , and Decentralized Total System a^t as Functions of the Production Cost to Retail Price Ratio c/p

The retailer's expected profit increases and then decreases as her forecasting accuracy improves. As noted above, an improvement in the retailer's forecasting accuracy has two effects: First, the more precise forecast information allows the retailer to make a better quantity choice. Second, the manufacturer responds to the retailer's improved forecasting accuracy by offering a more stingy contract (Proposition 2(b)). The negative impact of the second effect outweighs (is outweighed by) the positive impact of

the first effect when the retailer is a strong (weak) forecaster. To build intuition, it is helpful to consider the extreme cases. As noted previously, when the retailer is a very strong forecaster ($a \approx 0$), the optimal contract achieves nearly the entire integrated system profit for the manufacturer, leaving very little profit for the retailer. When the retailer is a very weak forecaster ($a \approx 1$), the integrated system optimal quantity does not depend on the retailer's privately observed forecast (from (6), $q^l(\theta) \approx \mu_0 + \sigma_0 z^l$ for all θ); consequently, the manufacturer can extract nearly the entire integrated system profit by offering a contract in which the price of purchasing this quantity is the expected revenue it generates, and other quantities are priced sufficiently high as to be made unattractive to the retailer. Consequently, as a strong forecaster's forecast accuracy improves (small a decreases) or a weak forecaster's accuracy worsens (large a increases), the retailer's expected profit decreases toward zero.

As the retailer's forecasting accuracy improves, system profit first decreases and then increases. Thus, the intuitive result that the total system benefits from improved forecasting accuracy (which holds for the integrated system, as discussed in §4) is reversed in the decentralized system when the retailer is a poor forecaster. The impact of forecasting accuracy on the total system profit follows the pattern of its impact on the manufacturer, not the retailer; this follows because

Figure 4 Under the Optimal Procurement Contract, Expected Profits of the Retailer R , Manufacturer M , and the Total Decentralized System $M + R$ as Functions of the Retailer's Forecasting Accuracy Parameter

Note. Parameters are $p = 1$, $c = 0.2$, $\mu = 1$, and $\sigma_0 = 0.2$.

the manufacturer, as the Stackelberg leader, has a bigger impact on system profit than the retailer.⁵

We now turn to the managerial implications of our numerical findings. A retailer can improve her forecasting accuracy by exerting effort to acquire and process demand-relevant data.⁶ The numerical results suggest that the retailer should be wary of improving her forecast accuracy. The retailer should be particularly concerned when she is already a strong forecaster; under such circumstances, even when the cost of improving her forecast accuracy is ignored, improved accuracy results in a loss to the retailer. (Taylor 2006 establishes similar results in a much simpler model.) To see this graphically, observe that in Figure 4 the retailer's expected profit decreases as her forecasting accuracy improves (a decreases) when the retailer is already a strong forecaster ($a < a^r$). Figure 4 also illustrates how the manufacturer's and retailer's interests align or diverge with respect to improved retailer forecasting accuracy. When the retailer is already a very strong forecaster (Region 1, which corresponds to $a < \bar{a}$), improved retailer forecasting accuracy has divergent impacts on the firms' profits: the retailer is hurt, but the manufacturer benefits. When the retailer is a weak forecaster (Region 3, which corresponds to $a > a^r$), improved retailer forecasting has the opposite effect on the firms: the retailer benefits, but the manufacturer is hurt. When the retailer is a moderately skilled forecaster (Region 2, which corresponds to $a \in (\bar{a}, a^r)$), the interests of the firms are aligned: both firms are hurt by improved retailer forecasting accuracy. Figure 4 depicts the case where the production cost is small ($c/p < 0.4$, so that $\bar{a} < a^r$). When the production cost is large ($c/p > 0.4$, so that $\bar{a} > a^r$), the results are identical with one exception. When the retailer is a moderately skilled forecaster (Region 2, which corresponds to $a \in (a^r, \bar{a})$ in the high production cost regime), again the interests of the firms are aligned, but in this case both firms benefit by improved retailer forecasting.

The numerical results suggest the following managerial implications: When the retailer is a poor forecaster, she has incentive to improve her forecasting accuracy, but the manufacturer would be hurt by such improvements and so might be inclined to frustrate the retailer's forecast improvement efforts. When the retailer is a strong forecaster, the manufacturer would

like the retailer to invest in improving her forecasting accuracy, but the retailer concerned only with her own profit lacks the incentive to do so. Only when the production cost is large and the retailer is a moderately skilled forecaster, are the interests of the two firms' favorably aligned: both manufacturer and retailer would like the retailer to improve her forecasting accuracy (provided that the cost of doing so is not too high). Perhaps our most surprising numerical observation is that improved retailer forecasting accuracy can simultaneously hurt both firms. Our numerical results suggest that this occurs if and only if the production cost is small ($c/p < 0.4$, so that $\bar{a} < a^r$) and the retailer is a moderately skilled forecaster ($a \in (\bar{a}, a^r)$).

Before concluding, it is worth pointing out three caveats. First, our results rely on the assumption that the manufacturer knows or can infer the retailer's forecasting accuracy. If the retailer is able to improve her accuracy without the manufacturer's knowledge, the results may differ. Second, in some contexts, there may be a practical limit for how accurate a retailer can become in forecasting. In our setting this would correspond to imposing a lower limit on the retailer's forecasting parameter, restricting a to $a \in (\hat{a}, 1)$, where $\hat{a} > 0$. If the level of inherent demand uncertainty that cannot be resolved through retailer forecasting efforts is large (i.e., $\hat{a}$ is large), then our results may qualitatively change. For example, if $\hat{a} > \max(a^r, \bar{a}, a^l)$, then the effect of improved retailer forecasting accuracy would always be to benefit the retailer and hurt the manufacturer and the total decentralized system. Third, our results rely on the assumption that the only contractual mechanism between the firms is the quantity-based procurement contract. When the retailer is a sufficiently strong forecaster ($a < a^l$), the positive impact of improved retailer forecasting accuracy on the manufacturer outweighs the negative impact on the retailer: the total system's expected profit $M + R$ increases. Consequently, if the firms can make transfer payments that are contingent on the retailer's forecasting accuracy (e.g., the manufacturer provides direct support to help the retailer improve her forecasting capabilities), then the firms may be able to overcome the incentive misalignment that prevents the firms from capturing this gain when the firms rely only on quantity-based contracts.

7. Discussion

Our main finding is that a manufacturer's expected profit is convex in the forecasting accuracy of its retail partner. This convexity result lends insight into two managerial questions. First, when faced with a pool of prospective retailers, *ceteris paribus*, should a manufacturer select a retailer that has strong, weak, or

⁵ Further, we observed that for every fixed value of σ_0 and c , the difference between the integrated system profit and the decentralized total system profit $\Pi^I - (M + R)$ is first increasing and then decreasing as the retailer's forecasting accuracy parameter a increases in $a \in (0, 1)$. Thus, the gain in system profit from centralization $\Pi^I - (M + R)$ is largest when the retailer's forecasting accuracy is moderate.

⁶ See Taylor and Xiao (2009) and Shin and Tunca (2010) for analyses that formally model a retailer's forecasting effort decision.

intermediate forecasting capabilities? The convexity result implies: if none of the retailers is sufficiently strong, the manufacturer should choose the weakest forecaster; otherwise, the manufacturer should choose the strongest forecaster.

Second, does a manufacturer benefit when his retail partner improves her forecasting capabilities? Our convexity result implies that a manufacturer benefits by improved forecasting at its retail partner if and only if the retailer is already a good forecaster. To the extent that two retailers are quite distinct in their forecasting capabilities, our model predicts that a marginal improvement in forecasting capabilities at these two retailers would have an opposite effect on their manufacturer-partners. Improved forecasting by a strong-forecaster retailer makes a “good” situation better (for the manufacturer), whereas the same improvement by a weak-forecaster retailer makes a “bad” situation worse (for the manufacturer).

We establish that the optimal procurement contract exhibits quantity discounts: quantity discounts emerge endogenously as an optimal response to the retailer’s private demand-forecast information. Under the optimal contract, the manufacturer tends to be hurt by improved retailer forecasting when the product economics are lucrative. Conversely, the manufacturer tends to benefit by improved retailer forecasting when the product economics are poor. We conclude that a manufacturer should be most concerned about improvements in retailer forecasting accuracy when the retailer is a poor forecaster and the product economics are lucrative.

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Acknowledgments

The authors are grateful to the associate editor, referees, and seminar participants at the University of California, Berkeley; University of California, San Diego; University of North Carolina, Chapel Hill; University of Pennsylvania; University of Washington; and Washington University for helpful comments.

Appendix

Lemma 2 is useful in the proofs of Lemma 1 and Propositions 1–3.

LEMMA 2. $z^*(\theta)$, the solution to (9), is unique. $z^*(\theta)$ strictly increases in θ , a , and p , and strictly decreases in c .

PROOF OF LEMMA 2. By dividing the both sides of (9) by $ap\bar{\Phi}(z^*(\theta))$, we have

$$1 - \frac{c}{p\bar{\Phi}(z^*(\theta))} - \frac{\sqrt{1-a^2}}{a} \frac{\phi(z^*(\theta))}{\bar{\Phi}(z^*(\theta))} \frac{\bar{\Phi}(\theta)}{\phi(\theta)} = 0. \quad (14)$$

Let $A(z) \equiv c/p\bar{\Phi}(z) + [(\sqrt{1-a^2})/a](\phi(z))/\bar{\Phi}(z)(\bar{\Phi}(\theta)/\phi(\theta))$. Clearly, $A(z)$ strictly increases in z . Because $A(-\infty) = c/p < 1$ and $A(+\infty) = +\infty > 1$, there exists a unique solution $z^*(\theta)$ that satisfies (14) for each θ . Because $A(z)$ strictly decreases in θ , a and p (strictly increases in c), the solution $z^*(\theta)$ strictly increases in θ , a and p and strictly decreases in c . $\square$

PROOF OF PROPOSITION 1. The proof proceeds as follows. The bulk of the proof is devoted to identifying a solution to the relaxed contract design problem in which the (IC) constraint is replaced by the corresponding first-order necessary condition. We then observe that the solution to this relaxed problem satisfies the constraints of the original problem. Using (3), we can write

$$\begin{aligned} \mathcal{R}(\theta, \hat{\theta}) &= pE_{D_\theta} \min(D_\theta, q(\hat{\theta})) - t(\hat{\theta}) \\ &= p(\mu_0 + \sqrt{1-a^2}\sigma_0\theta) \\ &\quad + pa\sigma_0 E_X \min\left\{X, \frac{q(\hat{\theta}) - \mu_0 - \sqrt{1-a^2}\sigma_0\theta}{a\sigma_0}\right\} - t(\hat{\theta}) \\ &= p(\mu_0 + \sqrt{1-a^2}\sigma_0\theta) \\ &\quad - pa\sigma_0 \left[\phi\left(z(\hat{\theta}) + \frac{\sqrt{1-a^2}}{a}(\hat{\theta} - \theta)\right) - \left(z(\hat{\theta}) + \frac{\sqrt{1-a^2}}{a}(\hat{\theta} - \theta)\right) \right. \\ &\quad \left. \times \bar{\Phi}\left(z(\hat{\theta}) + \frac{\sqrt{1-a^2}}{a}(\hat{\theta} - \theta)\right) \right] - t(\hat{\theta}), \quad (15) \end{aligned}$$

where $z(\theta) \equiv [q(\theta) - \mu_0 - \sqrt{1-a^2}\sigma_0\theta]/(a\sigma_0)$ and the last equality follows from (3). It follows from (IC) and the Envelope Theorem that

$$\begin{aligned} \mathcal{R}'(\theta) &= \frac{\partial \mathcal{R}(\theta, \hat{\theta})}{\partial \theta} \Big|_{\hat{\theta}=\theta} \\ &= p\sqrt{1-a^2}\sigma_0\Phi(z(\theta)) \quad (\text{by (15)}), \end{aligned}$$

which by integration, leads to $\mathcal{R}(\theta) = \mathcal{R}(-\infty) + p\sqrt{1-a^2}\sigma_0 \int_{-\infty}^{\theta} \Phi(z(x)) dx$. Note that $\mathcal{R}(\theta)$ increases in θ . Clearly, at the optimal solution, $\mathcal{R}(-\infty) = 0$. Hence we have

$$\mathcal{R}(\theta) = p\sqrt{1-a^2}\sigma_0 \int_{-\infty}^{\theta} \Phi(z(x)) dx. \quad (16)$$

By definition of $\mathcal{R}(\theta)$,

$$\begin{aligned} \mathcal{R}(\theta) &= \mathcal{R}(\theta, \theta) \\ &= p(\mu_0 + \sqrt{1-a^2}\sigma_0\theta) \\ &\quad - pa\sigma_0 [\phi(z(\theta)) - z(\theta)\bar{\Phi}(z(\theta))] - t(\theta). \quad (17) \end{aligned}$$

From (16) and (17), we can express $t(\theta)$ by using $z(\cdot)$ as follows:

$$\begin{aligned} t(\theta) &= p(\mu_0 + \sqrt{1-a^2}\sigma_0\theta) - pa\sigma_0 [\phi(z(\theta)) - z(\theta)\bar{\Phi}(z(\theta))] \\ &\quad - p\sqrt{1-a^2}\sigma_0 \int_{-\infty}^{\theta} \Phi(z(x)) dx. \quad (18) \end{aligned}$$

Substituting $t(\theta)$ in (OBJ) with the right-hand side of the above equation, (OBJ) can be rewritten as a function of $z(\cdot)$:

$$\begin{aligned} E_\theta[t(\Theta) - c(\mu_0 + \sqrt{1-a^2}\sigma_0\Theta + a\sigma_0z(\Theta))] \\ &= E_\theta[(p-c)(\mu_0 + \sqrt{1-a^2}\sigma_0\Theta) - a\sigma_0\Gamma(z(\Theta)) \\ &\quad - p\sqrt{1-a^2}\sigma_0\Phi(z(\Theta))\bar{\Phi}(\Theta)/\phi(\Theta)], \end{aligned}$$

which by pointwise optimization, maximized at $z^*(\theta)$ for every θ , where $z^*(\theta)$ is uniquely (see Lemma 2) determined from (9). The corresponding $t^*(\theta)$ can then be determined from (18).

Clearly, the solution $(z^*(\theta), t^*(\theta))$ constructed above yields an upper bound on the manufacturer's expected profit, and also satisfies (IR) and the first-order necessary condition of (IC). From Lemma 2, $z^*(\theta)$ increases in θ , which is a sufficient condition to ensure that the solution $(z^*(\theta), t^*(\theta))$ satisfies (IC). Therefore, $(z^*(\theta), t^*(\theta))$ solves the manufacturer's contract design problem (OBJ)–(IR). $\square$

PROOF OF PROPOSITION 2. (a) Note that

$$\frac{dt^*(\theta)}{d\theta} = p\sigma_0\bar{\Phi}(z^*(\theta))\left[\sqrt{1-a^2} + a\frac{dz^*(\theta)}{d\theta}\right] \quad (19)$$

and

$$\frac{dq^*(\theta)}{d\theta} = \sigma_0\left[\sqrt{1-a^2} + a\frac{dz^*(\theta)}{d\theta}\right]. \quad (20)$$

By definition of $T^*(q)$,

$$\begin{aligned} \frac{dT^*(q)}{dq} &= \frac{dt^*(\theta)}{d\theta} \bigg/ \frac{dq^*(\theta)}{d\theta} \\ &= \frac{p\sigma_0\bar{\Phi}(z^*(\theta))[\sqrt{1-a^2} + a(d/d\theta)z^*(\theta)]}{\sigma_0[\sqrt{1-a^2} + a(d/d\theta)z^*(\theta)]} \quad (\text{by (19) and (20)}) \\ &= p\bar{\Phi}(z^*(\theta)), \end{aligned} \quad (21)$$

where θ is such that $q = q^*(\theta)$. We can also write $(d/dq)T^*(q) = p\bar{\Phi}(Z(q))$, where $Z(q)$ is the z value corresponding to q . Clearly, $Z(q)$ is strictly increasing q . Hence, $(d^2/dq^2)T^*(q) = -p\phi(Z(q))Z'(q) < 0$.

(b) It follows from (21) that $(\partial^2/\partial q\partial a)T^*(q) = -p\phi(z^*(\theta)) \cdot (\partial/\partial a)z^*(\theta)$, where θ is such that $q = q^*(\theta)$. Because $(\partial/\partial a)z^*(\theta) > 0$ (from Lemma 2), $(\partial^2/\partial q\partial a)T^*(q) < 0$. $\square$

PROOF OF LEMMA 1. See Lemma 2. $\square$

PROOF OF PROPOSITION 3. (a) First we establish the convexity property. By the Envelope Theorem,

$$\frac{dM}{da} = E_{\Theta}\left[-\sigma_0\Gamma(z^*(\Theta)) + \frac{pa\sigma_0}{\sqrt{1-a^2}}\Phi(z^*(\Theta))\frac{\bar{\Phi}(\Theta)}{\phi(\Theta)}\right]. \quad (22)$$

Note that $\Gamma'(z) = -p\bar{\Phi}(z) + c$. It follows from (9) that $p\bar{\Phi}(z^*(\theta)) - c > 0$ for every θ . Thus

$$\Gamma'(z^*(\theta)) < 0,$$

which together with the fact that $z^*(\theta)$ strictly increases in a (see Lemma 2), implies that $\Gamma(z^*(\theta))$ strictly decreases in a . Note that the second part of $(d/da)M$ is clearly strictly increasing in a . Consequently, $(d/da)M$ strictly increases in a ; therefore, M is strictly convex in a .

To show the existence of $\bar{a} \in (0, 1)$, it suffices to show that $(d/da)M|_{a \rightarrow 1^-} > 0$ and $(d/da)M|_{a \rightarrow 0^+} < 0$. It follows from (9) that as $a \rightarrow 1^-$, $z^*(\theta) \rightarrow z^l$ for every θ . Therefore, by (22), we have $(d/da)M|_{a \rightarrow 1^-} = +\infty > 0$. Similarly, as $a \rightarrow 0^+$, $z^*(\theta) \rightarrow -\infty$ for every θ . Therefore, by (22), we have $(d/da)M|_{a \rightarrow 0^+} = -\infty < 0$.

(b) From (22), the definition of $\Gamma(\cdot)$, and the fact that $\bar{a}$ is the unique solution to $(d/da)M = 0$, $\bar{a}$ is the unique solution to

$$E_{\Theta}\left[-\phi(z^*(\Theta)) + z^*(\Theta)\bar{\Phi}(z^*(\Theta)) - \frac{c}{p}z^*(\Theta) + \frac{a}{\sqrt{1-a^2}}\Phi(z^*(\Theta))\frac{\bar{\Phi}(\Theta)}{\phi(\Theta)}\right] = 0. \quad (23)$$

From (9), for every θ , $z^*(\theta)$ depends solely on the ratio c/p and a (i.e., $z^*(\theta)$ is independent of all other parameters). Therefore, in terms of the model primitives, the left-hand side of (23) depends only on c/p and a . Because for any fixed c/p , the solution to (23), $\bar{a}$, is unique, $\bar{a}$ is determined solely by the ratio c/p (i.e., $\bar{a}$ is independent of all other parameters). $\square$

PROOF OF COROLLARY 1. The optimization problem for the manufacturer that observes forecast θ can be written as

$$\begin{aligned} \max_{\{T(q), q\}} \quad & \{T(q) - cq\} \\ \text{s.t.} \quad & q \in \arg \max_{\hat{q}} \{pE_{D_{\theta}} \min(D_{\theta}, \hat{q}) - T(\hat{q})\}, \\ & pE_{D_{\theta}} \min(D_{\theta}, q) - T(q) \geq 0. \end{aligned}$$

The retailer's expected profit when she chooses quantity q is $pE_{D_{\theta}} \min(D_{\theta}, q) - T(q)$; the first constraint ensures that it is in the retailer's best interest to select quantity q ; and the second constraint ensures that the retailer accepts the contract because her expected profit by choosing the contract is no less than her reservation profit. Under the payment schedule $T(q) = pE_{D_{\theta}}[\min(D_{\theta}, q^l(\theta))]$ for $q \leq q^l(\theta)$ and $T(q) = pq$ for $q > q^l(\theta)$, it is optimal for the retailer to order the integrated system quantity $q^l(\theta)$; the retailer cannot increase her expected profit by selecting a distinct order quantity. Under this optimal order quantity $q^l(\theta)$, the manufacturer achieves the integrated system profit $T(q^l(\theta)) - cq^l(\theta) = pE_{D_{\theta}}[\min(D_{\theta}, q^l(\theta))] - cq^l(\theta)$. Therefore, in expectation over the forecast Θ , the manufacturer's expected profit is the integrated system expected profit $E_{\Theta}[T(q^l(\theta)) - cq^l(\theta)] = \Pi^l$.

Consequently, the manufacturer's gain in expected profit from observing the forecast is $\Pi^l - M$. Because $M \leq \Pi^l$, $\lim_{a \rightarrow 0^+} M = \Pi^l$, M is strictly convex in a (from Proposition 3), and $(d/da)\Pi^l = -\sigma_0\Gamma(z^l)$, $\lim_{a \rightarrow 0^+} (d/da)M < -\sigma_0\Gamma(z^l)$. This together with the facts that $(d/da)M|_{a=\bar{a}} = 0$ and M is strictly convex and continuous in a , implies the existence of a unique $\underline{a} \in (0, \bar{a})$ satisfying $(d/da)M|_{a=\underline{a}} = -\sigma_0\Gamma(z^l)$. Therefore, $(d/da)[\Pi^l - M] > 0$ for $a \in (0, \underline{a})$ and $(d/da)[\Pi^l - M] < 0$ for $a \in (\underline{a}, 1)$. $\square$

References

- Arya, A., B. Mittendorf. 2004. Using returns policies to elicit retailer information. *RAND J. Econom.* 35(3) 617–630.
- Aviv, Y. 2001. The effect of collaborative forecasting on supply chain performance. *Management Sci.* 47(10) 1326–1343.
- Burnetas, A., S. M. Gilbert, C. E. Smith. 2007. Quantity discounts in single-period supply contracts with asymmetric demand information. *IEEE Trans.* 39(5) 465–479.
- Cachon, G. P. 2003. Supply chain coordination with contracts. A. G. de Kok, S. C. Graves, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management*, Chap. 6. Elsevier B.V., Amsterdam.

- Cachon, G., M. Fisher. 2000. Supply chain inventory management and the value of shared information. *Management Sci.* **46**(8) 1032–1048.
- Cachon, G. P., M. A. Lariviere. 2001. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Sci.* **47**(5) 629–646.
- Cederlund, J. P., R. Kohli, S. A. Sherer, Y. Yao. 2007. How Motorola put CPFR into action. *Supply Chain Management Rev.* **11**(7) 28–35.
- Chen, F. 2003. Information sharing and supply chain coordination. A. G. de Kok, S. C. Graves, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management*, Chap. 7. Elsevier B.V., Amsterdam.
- Chen, F., A. Federgruen, Y.-S. Zheng. 2001. Coordination mechanisms for a distribution system with one supplier and multiple retailers. *Management Sci.* **47**(5) 693–708.
- Corbett, C. J., X. de Groot. 2000. A supplier's optimal quantity discount policy under asymmetric information. *Management Sci.* **46**(3) 444–450.
- Feldman, J., J. Cramer. 2004. Hardlines retail. Report, April 26, SG Cowen, New York.
- Fliedner, G. 2003. CPFR: An emerging supply chain tool. *Indust. Management Data Systems* **103**(1) 14–21.
- Fraser, J. 2003. CPFR—Status and perspectives. D. Seifert, ed. *Collaborative Planning, Forecasting, and Replenishment: How to Create a Supply Chain Advantage*. AMACON, New York, 70–93.
- Iyer, G., C. Narasimhan, R. Niraj. 2007. Information and inventory in distribution channels. *Marketing Sci.* **53**(10) 1551–1561.
- Lewis, T., D. Sappington. 1991. All-or-nothing information control. *Econom. Lett.* **37**(2) 111–113.
- 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.
- Mishra, B. K., S. Raghunathan, X. Yue. 2007. Information sharing in supply chains: Incentives for information distortion. *IIE Trans.* **39**(9) 863–877.
- Miyaoka, J., W. H. Hausman. 2008. How improved forecasts can degrade decentralized supply chains. *Manufacturing Service Oper. Management* **10**(3) 547–562.
- Özer, Ö., W. Wei. 2006. Strategic commitment for optimal capacity decision under asymmetric forecast information. *Management Sci.* **52**(8) 1238–1257.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.
- Qi, F., K. Zhu. 2010. Endogenous information acquisition in supply chain management. *Eur. J. Oper. Res.* **201**(2) 454–462.
- Rajan, M., R. Saouma. 2006. Optimal information asymmetry. *Accounting Rev.* **81**(3) 677–712.
- Ren, Z. J., M. A. Cohen, T. H. Ho, C. Terwiesch. 2010. Information sharing in a long-term supply chain relationship: The role of customer review strategy. *Oper. Res.* **58**(1) 81–93.
- Shin, H., T. Tunca. 2010. The effect of competition on demand forecast investments and supply chain coordination. *Oper. Res.* Forthcoming.
- Taylor, T. A. 2006. Sale timing in a supply chain: When to sell to the retailer. *Manufacturing Service Oper. Management* **8**(1) 23–42.
- Taylor, T. A., W. Xiao. 2009. Incentives for retailer forecasting: Rebates vs. returns. *Management Sci.* **55**(10) 1654–1669.
- Tomlin, B. 2003. Capacity investments in supply chains: Sharing the gain rather than sharing the pain. *Manufacturing Service Oper. Management* **5**(4) 317–333.
- Weng, Z. K. 1995. Channel coordination and quantity discounts. *Management Sci.* **41**(9) 1509–1521.
- Widlitz, S. 2005. CC: SWAT teams and test catalog may be new catalysts—neutral. Report, July 21, Fulcrum Global Partners, New York.
- Winkler, R. L. 1981. Combining probability distributions from dependent information sources. *Management Sci.* **27**(4) 479–488.
- Zhang, H., M. Nagarajan, G. Sošić. 2010. Dynamic supplier contracts under asymmetric inventory information. *Oper. Res.*, ePub ahead of print, August 17, <http://or.journal.informs.org/cgi/content/abstract/opre.1100.0810v1>.