UC Berkeley

UC Berkeley Previously Published Works

Title

Incentives for retailer forecasting: Rebates vs. Returns

Permalink

<https://escholarship.org/uc/item/2cq9p36b>

Journal

Management Science, 55(10)

ISSN

0025-1909

Authors

Taylor, TA

Xiao, W

Publication Date

2009-10-01

DOI

10.1287/mnsc.1090.1045

Peer reviewed

Incentives for Retailer Forecasting: Rebates vs. Returns

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 studies a manufacturer that sells to a newsvendor retailer who can improve the quality of her demand information by exerting costly forecasting effort. In such a setting, contracts play two roles: providing incentives to influence the retailer's forecasting decision and eliciting information obtained by forecasting to inform production decisions. We focus on two forms of contracts that are widely used in such settings and are mirror images of one another: a rebates contract, which compensates the retailer for the units she sells to end consumers, and a returns contract, which compensates the retailer for the units that are unsold. We characterize the optimal rebates contracts and returns contracts. Under rebates, the retailer, manufacturer, and total system may benefit from the retailer having inferior forecasting technology; this never occurs under returns. Although one might conjecture that returns would be inferior because its provision of "insurance" would discourage the retailer from forecasting, we show that returns are superior.

Key words: forecasting; supply chain contracting; rebates; returns; endogenous adverse selection

History: Received October 12, 2006; accepted March 22, 2009, by Ananth Iyer, operations and supply chain management. Published online in *Articles in Advance* July 20, 2009.

1. Introduction

Because the costs of supply-demand mismatch are so high, firms devote substantial resources to improving their understanding of uncertain demand through forecasting efforts (Laucka 2005). For example, fashion retailers frequently devote substantial resources to conduct costly preseason merchandise tests. In such a test, a retail chain introduces a new product for a short period of time in a subset of the chain's stores prior to the main selling season and uses the observed sales data as a key input into its forecasting process (Fisher and Rajaram 2000). Firms also obtain demand-relevant data by conducting customer surveys and purchase packages of relevant sales or consumer data from third parties. Firms then devote time and resources to analyzing data and consulting internal experts (e.g., sales force) and external experts to develop and refine their forecasts. Because acquiring and processing such data is costly, whether a retailer will find it attractive to exert such costly forecasting effort depends critically on the contractual terms offered by the manufacturer of the product.

As an example of a manufacturer motivated by this concern, IBM PC Co. offered a menu of returns contracts to its resellers with the express intention of providing "a carrot to help the channel better forecast

demand and to more effectively manage inventory." Resellers could choose a contract with a generous returns policy and standard purchasing terms or one with a stingy returns policy but more attractive terms in other dimensions (Zarley 1994).¹

At the outset, it is unclear whether returns contracts are really the best way to encourage forecasting or if there is a better contractual form. More fundamentally, it is unclear whether a manufacturer should design contracts to encourage forecasting in the first place. Although the more precise demand information obtained by retailer forecasting might be beneficial to the overall system, it is not clear that the manufacturer benefits by the retailer's obtaining better demand information because this places the manufacturer at an informational disadvantage relative to the retailer.

We focus on two forms of contracts that are widely used in settings where demand uncertainty is pronounced and forecasting efforts are important in reducing such uncertainty: Under a returns contract, the manufacturer compensates the retailer by buying

¹ Similar menus of returns contracts have been used in the book publishing and pharmaceutical industries (Padmanabhan and Png 1995).

back each unsold unit the retailer has at the end of the selling season. Under a rebates contract, the manufacturer pays the retailer a bonus for each unit the retailer sells.² We focus on these two forms, in part, because they are mirror images of one another: A rebates contract pays the retailer for selling units, and the returns contract pays the retailer for *not* selling units. In addition, although researchers have studied each form of contract, relatively little has been done to compare their effectiveness. The research question we seek to address is the following: In a setting where the retailer can improve her knowledge of demand by exerting forecasting effort, should the manufacturer offer contracts that compensate the retailer for selling units or for *not* selling units?

Intuitively, returns contracts would seem to discourage forecasting because their provision of “insurance” would seem to make precise knowledge of the demand distribution less valuable. In contrast, rebates, instead of providing insurance, essentially provide the retailer with a “lottery”: The retailer does very well if demand turns out to be high, because the retailer receives substantial revenue from both retail-price-paying customers and the bonus-paying manufacturer, but does very poorly if demand turns out to be low. Consequently, precise knowledge of the demand distribution would appear to be more valuable under rebates. This suggests that to the extent that the manufacturer wants to induce the retailer to forecast, he should offer a rebates contract. Our results are exactly to the contrary of this intuition: we show that the manufacturer should offer returns instead of rebates, even when it is optimal to induce forecasting.

In this paper, we study a manufacturer that sells to a newsvendor retailer who can improve the quality of her demand information by exerting costly forecasting effort. We characterize the optimal menus of rebates contracts and returns contracts and compare the manufacturer’s expected profit under each. Under the optimal menu of rebates, the manufacturer cedes profit (information rents) to the retailer and distorts the production quantity downward to ameliorate this loss, which is consistent with the typical adverse selection result. More surprisingly, we show that under the optimal menu of rebates contracts, the retailer, manufacturer, and total system may benefit from the retailer having inferior forecasting technology; in addition, the retailer may overinvest in forecasting. The results differ significantly when the manufacturer instead employs a menu of returns

contracts. In contrast to the rebates case, the optimal menu of returns contracts induces the efficient level of forecasting, and there is no distortion in the production quantity. Furthermore, returns contracts are optimal among all contracts: under the optimal menu of returns contracts, the manufacturer captures the integrated system profit.

Our newsvendor model is appropriate when demand is uncertain, the product’s life cycle is short, and the retail market is sufficiently competitive such that retailers are essentially price takers. The model is appropriate not only for the motivating example of personal computers but also more broadly for products in the electronics, fashion apparel, and toy industries, in which retail competition is intense.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature. Section 3 describes the model. Sections 4, 5, and 6 contain the analysis for the integrated solution, rebates contracts, and returns contracts, respectively. Section 7 provides evidence for the robustness of our main managerial message, that returns are superior to rebates. Section 8 provides concluding remarks. All proofs are provided in the e-companion.³

2. Literature

There is a substantial literature that describes the roles of returns contracts in improving supply chain performance. First, by allowing the retailer to return unsold units to the manufacturer, returns mitigate the retail risk of overstock, thereby boosting the retailer’s order quantity to the manufacturer’s potential benefit. Returns contracts have been shown to be effective instruments for the manufacturer in newsvendor settings (Pasternack 1985), in settings with endogenous retail pricing (Emmons and Gilbert 1998), and in multiperiod settings (Donohue 2000). Second, returns can intensify retail competition to the benefit of the manufacturer (Padmanabhan and Png 1997, 2004). Third, over time, knowledge of the quantity returned in each period allows the manufacturer to better predict end-user demand in future periods (Sarvary and Padmanabhan 2001). However, returns contracts also have limitations; they may not be cost effective when physically returning products is costly or when supply chain members have different salvage values for the unsold products (Tsay 2001). A qualitative discussion of the pros and cons of using returns is provided in Padmanabhan and Png (1995).

Our work contributes to the supply chain literature on returns by identifying a new strategic role that returns play in promoting retailer forecasting efforts

² This kind of rebate is distinct from a consumer rebate, which is paid to the end user. It is also distinct from a reduction in the manufacturer’s unit price (the so-called off-invoice trade discount) in that the “reduction in price” due to the rebate is only realized if the item is sold to an end user.

³ An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

and inducing truthful information sharing. In particular, we demonstrate that a menu of returns contracts is more effective than a menu of rebates contracts in differentiating among privately informed retailers. Arya and Mittendorf (2004) consider a manufacturer-retailer model in which the end market consists of a fixed number of customers with a common reservation price. Prior to placing an order, the retailer privately and costlessly observes partial information about the customers' reservation price. After ordering, the retailer observes the actual reservation price and sets the retail price. Clearly, under a full-returns contract, the retailer either sells the entire stock to the customers at their reservation price or returns it to the manufacturer at the buy-back price, depending on which price is larger. Interestingly, while the Arya and Mittendorf (2004) model and the issues they consider are quite different from ours, they also find that a menu of returns contracts limits the retailer's information rents.

Rebates, another powerful incentive instrument, are also studied in the supply chain contracting literature. Rebate contracts have been shown to be an effective instrument for the manufacturer in settings with endogenous retail pricing (Aydin and Porteus 2008) and in settings with endogenous retailer sales effort (Taylor 2002 and Krishnan et al. 2004).

Our paper differs from the above mentioned work in two respects. First, we study both the returns contract and the rebates contract in a unified model framework, which allows us to compare the effectiveness of the contracts. Second, the retailer in our model can privately acquire demand information via costly forecasting, thereby causing information asymmetry between the two supply chain members. This model feature places our paper in the category often called *endogenous adverse selection*.

In contrast, in the conventional adverse selection setting, it is assumed that one of two parties is *costlessly* endowed with private information. The uninformed party can design contracts to elicit ("screen") information from the informed party. Alternatively, by initiating contractual terms, the informed party may credibly convey ("signal") his information to the uninformed party. Cachon (2003) and Chen (2003) provide excellent reviews of screening and signaling models in the supply chain literature. In an endogenous adverse selection framework, Chen and Xiao (2009) examine salesforce compensation where the salesperson can acquire costly demand signals prior to the firm's production decision.

In a paper close to our work, Lariviere (2002), a manufacturer faces a retailer who is privately informed that her forecasting cost is "high" or "low"; by assumption, the low-cost retailer forecasts and the high-cost retailer does not. The manufacturer offers two

contracts, one intended for the forecasting retailer and one intended for the nonforecasting retailer. In contrast to Lariviere (2002), we assume that the forecasting cost is common knowledge but explicitly make the retailer's forecasting decision endogenous. We focus on the menu of contracts designed to differentiate retailers based on the information they obtain by forecasting. This allows the manufacturer to tailor his production quantity to the observed demand signal. In addition, the main objective of Lariviere (2002) is to compare the performance of the "full-return, partial-credit" contract with the "partial-return, full-credit" contract, whereas ours is to compare returns with rebates.

A companion paper to our own that considers retailers' endogenous forecasting decisions is Shin and Tunca (2009). However, the focus of their paper is very different. They focus on multiple retailers that engage in Cournot competition and a single supply contract, whereas we allow for a menu of contracts and focus on the comparison between returns and rebates in a single-retailer setting.

3. Model

A manufacturer (he) sells to a retailer (she), who in turn sells to a market in which demand is uncertain. The manufacturer must produce, prior to the selling season, at unit cost c . The retail price is fixed at p , and the salvage value of unsold inventory is assumed to be negligible. The end-market demand is a nonnegative random variable D_N with distribution $F_N(\cdot)$ known by both firms. Before production, the retailer can exert forecasting effort, in which case she incurs cost k and observes a signal of the demand, denoted by S . The signal is high $S = H$ with probability λ and is low $S = L$ with probability $1 - \lambda$. The demand conditioned on the signal S is a nonnegative random variable D_S with distribution $F_S(\cdot)$ for $S \in \{H, L\}$. Therefore, the unconditioned demand distribution is

$$F_N(x) = \lambda F_H(x) + (1 - \lambda) F_L(x).$$

The demand conditioned on a high-value signal is stochastically larger than the demand conditioned on a low-value signal, i.e., $F_H(x) \leq F_L(x)$ for all $x \geq 0$. Although both the conditional distributions $F_S(\cdot)$ and the forecasting cost k are common knowledge, only the retailer observes whether she forecasts and, if so, the realized value of the signal. For notational convenience, we define $\bar{F}_S(x) \equiv 1 - F_S(x)$ for $S \in \{H, L, N\}$; for expositional ease, we also assume that the distributions $F_S(\cdot)$ are strictly increasing, but all results hold when this is relaxed.

Under a menu of rebates contracts, the transfer price the retailer pays, $T'(q, r): R^+ \times R^+ \rightarrow R^+$, is a function of her order quantity q and the per-unit rebate r , which

the retailer receives for each unit she sells to end consumers. Under a menu of returns contracts, the transfer price the retailer pays, $T^b(q, b): R^+ \times [0, p] \rightarrow R^+$, is a function of her order quantity q and the per-unit buy-back price b , which the retailer receives for each unit of unsold inventory at the end of the selling season. In each case, the manufacturer offers the menu of contracts, which simply consists of the transfer price function, and the retailer chooses the order quantity and the generosity of the sales-dependent payment (rebate or buy-back price). One would typically expect that the transfer price would be increasing in the generosity of the sales-dependent payment and in the order quantity.

The sequence of events is as follows:

1. The manufacturer offers a menu of rebates or returns contracts.
2. The retailer decides whether to forecast. If the retailer does so, then she observes the realized value of the demand signal (privately).
3. The retailer chooses a contract (order quantity and per-unit rebate or buy-back price) or rejects the contract offer. If the retailer accepts a contract, the manufacturer produces to meet the retailer's order, and the retailer pays the transfer price.
4. Demand is realized, and payments are made according to the chosen contract.

In the integrated system, the revelation of the demand signal reduces the total demand uncertainty and thus leads to a more accurate production decision, reducing the total cost of supply-demand mismatch. In the decentralized system, however, it may not be in the retailer's interest to exert forecasting effort to observe the demand signal. Furthermore, it may not be in the manufacturer's interest to have the retailer possess superior information about demand. Therefore, from the manufacturer's perspective, the questions are as follows: Should the manufacturer induce the retailer to forecast? If so, how should the manufacturer design the contracts to induce forecasting and to take advantage of the retailer's private demand forecast information? Which type of menu of contracts is more effective to this end?

A menu of contracts provides some incentive for the retailer to exert forecasting effort to observe the demand signal because the demand signal is useful in making a better contract choice. It can be demonstrated that because there are only two possible signals, we can without loss of generality restrict attention to menus of two contracts: under rebates, $T^r(q, r) = t_L$ for $q \leq q_L$ and $r \leq r_L$; $T^r(q, r) = t_H$ for $q_L < q \leq q_H$ and $r \leq r_H$; and $T^r(q, r) = \infty$ otherwise. This menu can be written more compactly as $\{(q_H, r_H, t_H), (q_L, r_L, t_L)\}$. Under returns, we can restrict attention to the analogously defined menu $\{(q_H, b_H, t_H), (q_L, b_L, t_L)\}$. The contract with subscript

$S \in \{H, L\}$ is intended for the retailer that has forecasted and observed signal S . The manufacturer may choose to offer a single contract, which is appropriate when the manufacturer intends that the contract be selected by a nonforecasting retailer.

Our analysis is done from the manufacturer's perspective: we characterize the contracts that maximize the manufacturer's expected profit while ensuring that the retailer's expected profits both before and after forecasting are no less than her reservation profit, which we normalize to zero without loss of generality.

The model described here has a number of restrictive assumptions: the forecasting decision is binary, there are only two possible demand signals, forecasting effort only yields demand information for a single product in a single period, the retail price is fixed, and demand is exogenous. In §7 we relax these assumptions and provide evidence that our main managerial message, that returns are superior to rebates, is robust to these assumptions.

4. Integrated System

As a benchmark, we first characterize the solution (forecasting decision and production quantity) that maximizes the expected profit of the integrated system (combined manufacturer and retailer). If the system chooses not to forecast, then it faces a newsvendor problem with demand D_N . The system's expected profit under production quantity q_N is

$$\begin{aligned}\Pi_N(q_N) &= pE \min(q_N, D_N) - cq_N \\ &= p \int_0^{q_N} \bar{F}_N(x) dx - cq_N,\end{aligned}$$

and the optimal production quantity is $q_N^I = \bar{F}_N^{-1}(c/p)$. If the system chooses to forecast, it incurs cost k , observes the demand signal $S \in \{H, L\}$, and then chooses its production quantity q_S . Conditioning on the realized value of the signal S , the system's expected profit excluding the cost of forecasting is

$$\begin{aligned}\Pi_F(q_H, q_L) &= \lambda[pE \min(q_H, D_H) - cq_H] \\ &\quad + (1 - \lambda)[pE \min(q_L, D_L) - cq_L].\end{aligned}$$

Note that if the integrated system did not adjust its production quantity in response to the demand signal, there would be no gain from forecasting: $\Pi_F(q, q) = \Pi_N(q)$. However, after observing the signal $S \in \{H, L\}$, the optimal production quantity is $q_S^I = \bar{F}_S^{-1}(c/p)$, the newsvendor solution under demand distribution F_S . Intuitively, having more precise information about the demand distribution allows the integrated system to choose a production quantity that more accurately trades off the cost of having too much versus having too little, so

$$\Pi_F(q_H^I, q_L^I) \geq \Pi_F(q_N^I, q_N^I) = \Pi_N(q_N^I).$$

Let

$$k^l \equiv \Pi_F(q_H^l, q_L^l) - \Pi_N(q_N^l)$$

denote the increase in expected profit resulting from having the better demand information provided by forecasting. Naturally, it is optimal to forecast if and only if the gain from doing so exceeds the cost. Proposition 1 summarizes the optimal forecasting and production quantity decisions for the integrated system.

PROPOSITION 1. *If $k < k^l$, then the system should forecast and produce q_H^l (q_L^l) when the demand signal is H (L). Otherwise, the system should not forecast and should produce q_N^l .*

5. Rebates Contracts

Now we turn to the decentralized system, where the firms make decisions to maximize their own profits. The purpose of this section is to characterize the optimal menu of rebates contracts that maximizes the manufacturer's expected profit. The manufacturer has the option to offer contracts that induce the retailer either to forecast or not forecast. For the case in which the manufacturer chooses not to induce retail forecasting, we can, without loss of generality, restrict analysis to a menu with a single contract (q_N, r_N, t_N) , under which the retailer pays t_N for q_N units and per-unit rebate r_N . For the case in which the manufacturer chooses to induce the retailer to forecast, the manufacturer offers a menu of distinct contracts to distinguish between retailers that have observed distinct demand signals. We study these two scenarios sequentially to understand which yields greater profit for the manufacturer.

5.1. No Forecasting

The retailer's expected profit under demand distribution F_S and rebate contract (q_C, r_C, t_C) is

$$R^r(S, C) = (p + r_C)E \min(q_C, D_S) - t_C.$$

We use superscripts to denote the contract form, where r corresponds to the rebate contract. If the retailer is faced with a single contract offer (q_N, r_N, t_N) , does not forecast, and accepts the contract, her expected profit is $R^r(N, N)$. The retailer will accept contract (q_N, r_N, t_N) if and only if her expected profit under the contract exceeds her reservation profit, which we have normalized to zero. Forecasting allows the retailer to more accurately assess her expected profit under a given contract, and so make a better informed decision about whether to accept the contract. Consequently, the retailer's expected profit under forecasting is

$$\lambda \max(R^r(H, N), 0) + (1 - \lambda) \max(R^r(L, N), 0) - k.$$

If the contract were sufficiently generous that the retailer would accept it regardless of the signal she observed, then there would be no gain from forecasting: $\lambda R^r(H, N) + (1 - \lambda) R^r(L, N) = R^r(N, N)$. Forecasting is only valuable when the low signal conveys to the retailer that market conditions are sufficiently poor that the retailer should not carry the manufacturer's product (i.e., should reject the contract): $R^r(L, N) < 0$.

An optimal rebate contract that does not induce forecasting is the solution to

$$\max_{q_N, r_N, t_N} \{t_N - c q_N - r_N E \min(q_N, D_N)\} \quad (\text{OBJ})$$

$$\text{s.t. } R^r(N, N) \geq \lambda \max(R^r(H, N), 0)$$

$$+ (1 - \lambda) \max(R^r(L, N), 0) - k \quad (\text{IC})$$

$$R^r(N, N) \geq 0. \quad (\text{IR})$$

The manufacturer chooses the contract parameters (q_N, r_N, t_N) to maximize his expected profit, subject to two constraints. The left-hand side of the constraints is the retailer's expected profit under contract (q_N, r_N, t_N) when she does not forecast. The incentive compatibility constraint (IC) ensures that it is in the interest of the retailer not to forecast. The individual rationality constraint (IR) ensures that the contract satisfies the retailer's participation constraint, so that the nonforecasting retailer is willing to accept the contract.

Proposition 2 characterizes the solution to the contract design problem (OBJ)–(IR). It is useful to define $\Gamma(q) \equiv (1 - \lambda)p \int_0^q [\bar{F}_N(x) - \bar{F}_L(x)] dx$. Recall that q_N^l and q_L^l are the optimal production quantities for the integrated system with demand D_N and D_L , respectively, and $q_N^l \geq q_L^l$. For any $q \in (q_L^l, q_N^l)$, $\bar{F}_N(q) > c/p$, and $\bar{F}_L(q) < c/p$, which implies that $\Gamma'(q) = (1 - \lambda) \cdot p[\bar{F}_N(q) - \bar{F}_L(q)] > 0$. Therefore, the inverse function $\Gamma^{-1}(\cdot)$ is well defined over the interval $[\Gamma(q_L^l), \Gamma(q_N^l)]$.

PROPOSITION 2. *An optimal rebate contract that does not induce forecasting has quantity, rebate, and transfer payment:*

$$(q_N^*, r_N^*, t_N^*) = \begin{cases} (q_L^l, 0, pE \min(q_L^l, D_N) - [\Gamma(q_L^l) - k]/(1 - \lambda)) & \text{if } k \leq \Gamma(q_L^l), \\ (\Gamma^{-1}(k), 0, pE \min(\Gamma^{-1}(k), D_N)) & \text{if } k \in (\Gamma(q_L^l), \Gamma(q_N^l)), \\ (q_N^l, 0, pE \min(q_N^l, D_N)) & \text{if } k \geq \Gamma(q_N^l). \end{cases}$$

A positive rebate provides an incentive for the retailer to forecast because by doing so the retailer is able to more accurately estimate the expected revenue she would receive from the rebate, which puts her in a better position to decide whether or not to accept the contract. Hence, in offering a contract that discourages forecasting, it is optimal to offer no rebate $r_N^* = 0$.

When the retailer's forecasting cost is high ($k \geq \Gamma(q_N^l)$), in offering a contract the manufacturer does not need to be concerned with adjusting the contract to discourage forecasting (constraint (IC) is irrelevant). The manufacturer captures the profit of the (nonforecasting) integrated system by offering a contract that sells the integrated-system optimal quantity q_N^l for a transfer payment equal to the expected revenue generated by this quantity. However, such a contract is not sustainable as the forecasting cost decreases ($k \in (\Gamma(q_L^l), \Gamma(q_N^l))$) because large contractual quantities make forecasting attractive. Thus, to discourage the retailer from forecasting, the manufacturer must lower the production quantity ($q_N^* = \Gamma^{-1}(k)$ decreases as k decreases); the manufacturer continues to charge the expected revenue generated by the quantity. However, distorting the quantity downward from the optimal quantity is costly to the manufacturer in that he cannot charge a large transfer payment. Consequently, as the forecasting cost drops further ($k \leq \Gamma(q_L^l)$), the manufacturer leaves the contractual quantity unchanged but lowers the transfer payment to discourage the retailer from forecasting. Overall, discouraging forecasting requires distorting the production quantity downward $q_N^* \leq q_N^l$.

Typically, in adverse selection models, the agent (retailer) captures strictly positive expected profit (so-called information rents) only if she has an informational advantage vis-à-vis the principal (manufacturer). In contrast, in our setting when the forecasting cost is low ($k < \Gamma(q_L^l)$), the agent (retailer) earns strictly positive expected profits even though she lacks an informational advantage vis-à-vis the principal (manufacturer): $R^r(N, N) = [\Gamma(q_L^l) - k]/(1 - \lambda) > 0$. Under the optimal contract, the retailer does not forecast, and the manufacturer anticipates that in equilibrium the retailer will not have superior information. Consequently, it is not possessing information that drives rents, but rather the threat of acquiring information.⁴ The intuition is as follows: Acquiring demand information puts the retailer in a better position to evaluate whether the contract is attractive enough to accept, and this information is particularly valuable when the contract is stingy. Consequently, the manufacturer must offer a generous contract to discourage forecasting, thereby ceding profit to the nonforecasting retailer.

A further contrast emerges when one considers the impact of costs on the production quantity. Often, optimal production quantities decrease as costs increase,

and this is true in our setting in that the optimal contracted quantity q_N^* is decreasing in the production cost c . However, we find the opposite occurs with respect to the forecasting cost: The optimal contracted quantity q_N^* is *increasing* in the forecasting cost k . Larger forecasting costs allow the manufacturer to offer larger quantities while still discouraging forecasting.

Under the optimal rebate contract that does not induce forecasting, the manufacturer's expected profit is

$$\mathcal{M}_N^r = \begin{cases} pE \min(q_L^l, D_N) - cq_L^l - [\Gamma(q_L^l) - k]/(1 - \lambda) & \text{if } k \leq \Gamma(q_L^l), \\ pE \min(\Gamma^{-1}(k), D_N) - c\Gamma^{-1}(k) & \text{if } k \in (\Gamma(q_L^l), \Gamma(q_N^l)), \\ pE \min(q_N^l, D_N) - cq_N^l & \text{if } k \geq \Gamma(q_N^l). \end{cases} \quad (1)$$

The subscript N denotes no forecasting, and as before the superscript r denotes the rebates contract form. It is straightforward to verify that $\mathcal{M}_N^r$ is increasing in the forecasting cost k . Intuitively, when the retailer's forecasting cost is higher, it is easier for the manufacturer to induce the retailer not to forecast (i.e., there exists a broader set of contract parameters satisfying the (IC) constraint).

5.2. Forecasting

Suppose the manufacturer chooses to induce retailer forecasting. From the revelation principle, we can, without loss of generality, restrict attention to a menu of two contracts $\{(q_H, r_H, t_H), (q_L, r_L, t_L)\}$, where the contract with subscript $S \in \{H, L\}$ is intended for the retailer that has forecasted and observed signal S . An optimal menu of rebate contracts that induces forecasting is the solution to

$$\begin{aligned} \max_{(q_H, r_H, t_H), (q_L, r_L, t_L)} & \left\{ \lambda[t_H - cq_H - r_H E \min(q_H, D_H)] + (1 - \lambda) \right. \\ & \left. \cdot [t_L - cq_L - r_L E \min(q_L, D_L)] \right\} \quad (\text{OBJ}) \\ \text{s.t. } & R^r(H, H) \geq R^r(H, L) \quad (\text{IC1}) \\ & R^r(L, L) \geq R^r(L, H) \quad (\text{IC2}) \\ & R^r(H, H) \geq 0 \quad (\text{IR1}) \\ & R^r(L, L) \geq 0 \quad (\text{IR2}) \\ & \lambda R^r(H, H) + (1 - \lambda) R^r(L, L) - k \\ & \geq R^r(N, H) \quad (\text{IC3}) \\ & \lambda R^r(H, H) + (1 - \lambda) R^r(L, L) - k \\ & \geq R^r(N, L) \quad (\text{IC4}) \\ & \lambda R^r(H, H) + (1 - \lambda) R^r(L, L) - k \\ & \geq 0. \quad (\text{IR3}) \end{aligned}$$

⁴ Although this phenomenon (as well as the other phenomena described in this section) is shown only in the setting where the manufacturer chooses not to induce forecasting, these phenomena persist when we make the manufacturer's decision of whether to induce forecasting endogenous.

The first four constraints ensure that after forecasting and observing the demand signal, the retailer selects the intended contract. Constraints (IC1) and (IC2) ensure that it is more attractive to select this contract rather than the contract intended for the retailer that observed the other demand signal. Constraints (IR1) and (IR2) ensure that the intended contract satisfies the retailer's interim participation constraint so that after forecasting and observing a signal, the retailer will accept the contract; the forecasting cost is excluded because it is sunk when the retailer is making her decision to accept or reject the contract. The last three constraints ensure that the retailer forecasts. Constraints (IC3) and (IC4) ensure that the retailer is better off forecasting and choosing the intended contract than not forecasting and selecting either contract. Constraint (IR3) ensures that the contracts satisfy the retailer's ex ante participation constraint so that the retailer is better off forecasting and accepting the intended contract rather than not forecasting and rejecting the contracts.

If (OBJ)–(IR3) does not have a solution, then it is optimal to *not* induce forecasting; Proposition 3 addresses the case where (OBJ)–(IR3) does have a solution. Let

$$\Delta(q) \equiv E[\min(q, D_H) - \min(q, D_L)]$$

denote the expected increase in units sold when the retailer has q units and observes the favorable rather than unfavorable demand signal.

PROPOSITION 3. *An optimal menu of rebate contracts that induces forecasting has quantities, rebates, and transfer payments:*

$$\begin{aligned} q_L^* &= \arg \max_{q_L \geq 0} \left\{ p \int_0^{q_L} [\bar{F}_L(x) - \lambda \bar{F}_H(x)] dx - (1 - \lambda) c q_L \right\}, \\ r_L^* &= 0, \\ t_L^* &= p E \min(q_L^*, D_L), \\ q_H^* &= q_H^l, \\ r_H^* &= k / \lambda (1 - \lambda) \Delta(q_H^l), \\ t_H^* &= (p + r_H^*) E \min(q_H^*, D_H) - p \Delta(q_L^*) - k / \lambda. \end{aligned}$$

To reward forecasting and distinguish between retailers that have observed different signals, the manufacturer offers extreme contracts: one contract with a large rebate, quantity, and transfer payment and one contract with no rebate and a small quantity and payment. The large-rebate contract will only be attractive to a retailer that has an optimistic demand forecast, whereas the no-rebate contract with its small quantity and transfer payment appeals to the pessimistic retailer.

Typically, in adverse selection models, the party with superior information captures positive expected profit (information rents). This result continues to

hold in our setting, despite the fact that it is costly for the retailer to obtain her information advantage. Under the optimal menu of contracts, the retailer's expected profit is

$$\lambda R^r(H, H) + (1 - \lambda) R^r(L, L) - k = \lambda R^r(H, L) \quad (2)$$

$$= R^r(N, L) \quad (3)$$

$$= \lambda p \Delta(q_L^*), \quad (4)$$

which is strictly positive when $F_L(0) = 0$ and $F_H(x) < F_L(x)$ for some sufficiently small x . Although the retailer receives positive expected profit, the contract is designed to reward only the optimistic-forecast-observing retailer: Excluding the cost of forecasting, the pessimistic-forecast-observing retailer receives zero profit, and the optimistic-forecast-observing retailer receives expected profit of $p \Delta(q_L^*) + k / \lambda$. To reduce the retailer's information rents, the production quantity for the low-type retailer is distorted downward $q_L^* \leq q_L^l$, whereas there is no distortion for the high-type retailer $q_H^* = q_H^l$, which follows the typical adverse selection results. In the typical adverse selection setting, the source of the informed party's expected profit is that when her private information is favorable (which occurs with probability λ in our model) the informed party can select the contract intended for the unfavorably informed party, and this explanation is consistent with (2). Equation (3) provides an alternative explanation for the retailer's expected profit in our setting with endogenous asymmetric information. Any contract that appeals to the pessimistic-forecast-observing retailer (i.e., satisfies $R^r(L, L) \geq 0$) will yield greater expected profit to the nonforecasting retailer (who possesses a more optimistic forecast). Indeed, the retailer's expected profit under the optimal menu is the profit she would receive by not forecasting and then selecting the contract intended for the pessimistic-forecast-observing retailer (Equation (3)). In this sense, surprisingly, one can interpret the source of the retailer's profit as being the retailer's threat of *not* acquiring information.

Notably, the retailer's expected profit is independent of her forecasting cost (see (4)); the manufacturer bears the cost entirely. The manufacturer's expected profit is

$$\begin{aligned} \mathcal{M}_F^r &= \lambda \left[p \int_0^{q_H^*} \bar{F}_H(x) dx - c q_H^* \right] + (1 - \lambda) \\ &\quad \cdot \left[p \int_0^{q_L^*} \frac{\bar{F}_L(x) - \lambda \bar{F}_H(x)}{1 - \lambda} dx - c q_L^* \right] - k. \end{aligned} \quad (5)$$

5.3. Forecasting or No Forecasting

We now turn to whether the manufacturer should offer contracts that encourage or discourage retailer forecasting. The manufacturer's expected profit under

the optimal menu of contracts that induces forecasting is $\mathcal{M}_F^r$ (see (5)), and under the optimal menu of contracts that does not induce forecasting is $\mathcal{M}_N^r$ (see (1)). The manufacturer should offer contracts that induce forecasting if and only if $\mathcal{M}_F^r > \mathcal{M}_N^r$. Because as k increases, $\mathcal{M}_F^r$ strictly decreases and $\mathcal{M}_N^r$ increases, and because at $k = 0$, $\mathcal{M}_F^r \geq \mathcal{M}_N^r$, there exists a threshold $k^r \geq 0$ such that

$$\mathcal{M}_F^r > \mathcal{M}_N^r \quad \text{if and only if } k < k^r.$$

Thus, when forecasting is cheap, the manufacturer should offer a menu with a generous-rebate contract and a no-rebate contract (see Proposition 3) to induce retail forecasting. When forecasting is costly, the manufacturer should offer a single no-rebate contract (see Proposition 2) that discourages the retailer from forecasting.

We next turn to the impact of the retailer's forecasting cost on the expected profit of the firms. The insights on this point are most easily illustrated by a numerical example; Figure 1 depicts the expected profits of the manufacturer, retailer, and total system as a function of the retailer's forecasting cost. Several observations are noteworthy. First, should the manufacturer prefer a retailer with a lower cost of forecasting or a retailer with a higher cost of forecasting? A retailer with a lower forecasting cost can be easily motivated to forecast, collecting valuable demand information that leads to a better match between supply and demand. This suggests that the manufacturer should benefit from a reduction in the retailer's forecasting cost. Indeed, under the optimal rebates contract, a decline in the forecasting cost results in a strict increase in the manufacturer's expected profit, provided that the current forecasting cost is low ($k < k^r$). However, if the forecasting cost is moderately high ($k > k^r$), a reduction in the forecasting cost decreases the manufacturer's expected profit. Figure 1(a) depicts this result; in the figure, $k^r = 0.26$. The intuition as to why the manufacturer is hurt when the retailer's forecasting technology improves is that when the forecasting cost is moderately high, it is optimal for the manufacturer to discourage the retailer from forecasting, and it is easier to do so when forecasting is costly for the retailer.

Second, one might conjecture that, even if the manufacturer could benefit by the retailer's having an inferior forecast technology, the retailer would never benefit (because having inferior forecasting technology means that it is more difficult for the retailer to obtain information that would be helpful in making a contract choice decision). Figure 1(b) shows that this conjecture is false. The retailer's expected profit jumps up when the forecasting cost crosses the threshold k^r . Interestingly, when the forecasting cost is moderate

Figure 1 Expected Profits Under Optimal Rebate Contracts

Note. Parameters are $p = 20$, $c = 10$, $\lambda = 0.6$, $D_H \sim \text{exponential}(0.10)$, and $D_L \sim \text{exponential}(0.12)$.

($k \in (0.26, 1.06)$ in Figure 1), the retailer captures positive expected profit *without* having superior information. Here it is the threat of acquiring information that leads to positive expected profit for the retailer, and this threat is the most formidable when the forecast cost is not too high, which explains why the retailer's expected profit jumps.

Third, Figure 1(c) shows that when the forecasting cost is moderately high ($k \in (1.06, 1.25)$ in Figure 1), the total system benefits as the forecasting cost increases. This stands in sharp contrast to the result for the integrated system, where expected profit is always decreasing in the forecasting cost.

The next proposition provides sufficient conditions under which the manufacturer, retailer, and total system benefit when the retailer's forecasting

cost increases. Let $\mathcal{M}^r$ denote the manufacturer's expected profit under the optimal rebate contracts, i.e., $\mathcal{M}^r = \max(\mathcal{M}_F^r, \mathcal{M}_N^r)$, and let $\mathcal{R}^r$ denote the retailer's expected profit under these contracts.

PROPOSITION 4. *If $ED_{H_-} > ED_L$, then there exists $\bar{c}$, $\underline{k}$, and $\bar{k}$ with $\bar{c} > 0$ and $\underline{k} < \bar{k}$ such that if $c < \bar{c}$, under the optimal rebate contracts, the expected profit of the manufacturer $\mathcal{M}^r$ and the total system $\mathcal{M}^r + \mathcal{R}^r$ are strictly increasing in the retailer's forecasting cost $k \in [\underline{k}, \bar{k}]$. If $\lambda \bar{F}_H(x) \geq \bar{F}_L(x)$ for $x > 0$, then there exists $\underline{c} > 0$ and $\hat{k} > k^r$ such that if $c < \underline{c}$, then the retailer's expected profit is higher when her forecasting cost is higher:*

$$\mathcal{R}^r|_{k \in [0, k^r]} < \mathcal{R}^r|_{k \in (k^r, \hat{k})}. \quad (6)$$

The conditions in the proposition on the demand distributions ensure that the forecast signal conveys some information. When the production cost c is small, there is little value to having production decisions informed by more precise demand information. Consequently, for a large range of forecasting costs k , it is optimal to induce no forecasting. As Figure 1 illustrates, when this region is large, the manufacturer and retailer can be hurt by the retailer having a lower forecasting cost.

We now summarize the managerial implications of our results regarding the impact of the forecasting cost under rebate contracts. First, manufacturers ought not blindly seek out retailers with low forecasting costs. Being able to forecast cheaply puts the retailer in the position where she will walk away from contracts that are not generous, when her forecasting efforts reveal that market conditions are unfavorable. Thus, low forecasting costs can force the manufacturer to offer generous contracts. Proposition 4 indicates that this phenomenon occurs when the production cost is low, so manufacturers operating in such environments should be particularly careful about selecting a retailer with a low forecasting cost. A manufacturer may be able to reduce the retailer's forecasting cost by providing direct assistance to the retailer (e.g., by providing materials or other support for a preseason merchandise test). Our results indicate that the marginal benefit to the manufacturer of doing so is positive when the retailer is already good at forecasting ($k < k^r$); otherwise, it is negative.

A second implication is that retailers ought not blindly devote resources to reduce their forecasting costs (improve their forecasting technology). If the retailer's forecasting cost is sufficiently low that it is optimal for the manufacturer to induce forecasting, then the retailer receives no benefit from reducing its forecasting cost; the manufacturer receives the entire benefit (see (4) and (5)). If the retailer's forecasting cost is sufficiently high that it is optimal for the manufacturer not to induce forecasting, then the retailer

may be hurt by reducing its cost of forecasting. Proposition 4 indicates that this concern is particularly acute when the production cost is low and the retailer learns a lot about demand by forecasting (the signals correspond to quite different demand distributions). (For a related result in a setting with no-rebate contracts and exogenous asymmetric information, see Taylor 2006.) We should caution that the results here rely on the assumption that the retailer's forecasting cost is common knowledge.⁵

We conclude this subsection by noting the impact of decentralization on forecast investment. Decentralized systems are often characterized by underinvestment relative to the integrated system optimal investments, which in our setting would suggest that decentralization leads to underinvestment in forecasting. Figure 1(c) shows that the opposite may occur. When the forecasting cost is moderately low ($k \in (0.18, 0.26)$ in Figure 1), decentralization leads to *overinvestment* in forecasting: It is optimal for the manufacturer to offer contracts that induce the retailer to forecast, even though not forecasting is optimal for the integrated system. The intuition is that discouraging forecasting is costly to the manufacturer because it compels the manufacturer to lower the contractual quantity and transfer payment. Discouraging forecasting is especially costly when the forecasting cost is not too high.

6. Returns Contracts

In the previous section, we established that when the manufacturer has the option to compensate the retailer for selling units, he should do so only when he wants to encourage the retailer to forecast, in which case he offers a menu of two contracts, only one of which compensates the retailer for selling units. In this section, we turn to addressing the question of whether (and if so, how) the manufacturer should compensate the retailer for not selling units (paying the retailer for each unit of unsold inventory). More importantly, we address whether the manufacturer should offer contracts that compensate the retailer for selling units (rebates) or not selling units (returns). Following the same sequence as the previous section, we begin with the case where the manufacturer induces the retailer not to forecast.

⁵ When the retailer's forecasting cost is private information, then the manufacturer, in designing contracts, needs to differentiate retailers not only based on the information that retailers obtain through forecasting (as in our setting) but also based on the retailers' forecasting cost. See Lariviere (2002) for a treatment that focuses on differentiating retailers based on their private forecasting cost information. It is possible that our results, which consider a single dimension of asymmetric information (demand information obtained by the retailer), may change when this second dimension of asymmetric information (retailer's forecasting cost) is incorporated.

6.1. No Forecasting

Under a returns contract, the manufacturer buys back the retailer's unsold inventory at the end of the selling season, paying b for each unit. The retailer's expected profit when she faces demand distribution F_S under returns contract (q_C, b_C, t_C) is

$$R^b(S, C) = pE \min(q_C, D_S) + b_C E \max(q_C - D_S, 0) - t_C.$$

We refer to a returns contract (q_N, b_N, t_N) with buy-back price $b_N = p$ and transfer payment $t_N = pq_N$ as a full-returns contract because the retailer receives a full refund of her per-unit purchase price for any unsold quantity. Under any full-returns contract, the retailer receives zero profit under any demand distribution. Therefore, under a full-returns contract, the retailer has no incentive to forecast. If the manufacturer offers a full-returns contract with quantity $q_N = q_N^I$, the manufacturer's expected profit is identical to the expected profit of the nonforecasting integrated system $\Pi_N(q_N^I)$. Because this is the largest profit that the manufacturer could capture under a contract that does not induce forecasting, if the manufacturer chooses not to induce forecasting, it is optimal for him to offer the full-returns contract $(q_N, b_N, t_N) = (q_N^I, p, pq_N^I)$.

6.2. Forecasting

Suppose the manufacturer offers a menu of two returns contracts, denoted by $\{(q_H, b_H, t_H), (q_L, b_L, t_L)\}$, to induce the retailer to forecast. To derive the optimal contract parameters that maximize the manufacturer's profit, we follow an approach similar to that in §5.2.

When the retailer forecasts and chooses the intended contract, the manufacturer's expected profit is

$$\begin{aligned} & \lambda[t_H - cq_H - b_H E \max(q_H - D_H, 0)] \\ & + (1 - \lambda)[t_L - cq_L - b_L E \max(q_L - D_L, 0)]. \end{aligned}$$

The manufacturer's contract design problem has the same set of constraints (IC1)–(IR3), where $R^b(S, C)$ replaces $R^r(S, C)$.

If the manufacturer's returns contract design problem (OBJ)–(IR3) does not have a solution, then it is optimal to not induce forecasting. Proposition 5 addresses the case in which (OBJ)–(IR3) does have a solution. As a technical point, we define $\Gamma^{-1}(k) \equiv \min\{q: \Gamma(q) = k\}$.

PROPOSITION 5. *An optimal menu of returns contracts that induces forecasting has quantities, buy-back prices, and transfer payments:*

$$\begin{aligned} (q_H^*, b_H^*, t_H^*) &= (\max(q_H^I, \Gamma^{-1}(k)), 0, \\ & pE \min(\max(q_H^I, \Gamma^{-1}(k)), D_H) - k/\lambda), \\ (q_L^*, b_L^*, t_L^*) &= (q_L^I, p, pq_L^I). \end{aligned}$$

As in the case with rebates contracts (see Proposition 3), to reward forecasting and distinguish between retailers that have observed different signals, the manufacturer offers extreme contracts: a full-returns contract and a no-returns contract. The no-returns contract is attractive to a retailer that has an optimistic demand forecast, whereas the full-returns contract appeals to the pessimistic retailer.

Typically, in adverse selection models, the party with superior information captures positive expected profit (information rents), and §5.2 shows that the result continues to hold in our setting with endogenous asymmetric information and rebate contracts. Contrast emerges, however, in our setting with returns contracts. Under the optimal returns menu, the retailer's expected profit is zero. Excluding the cost of forecasting, the pessimistic-forecast-observing retailer receives zero profit, and the optimistic-forecast-observing retailer receives expected profit of k/λ . One can think of this profit as the retailer's ex post information rents. The contract is designed so that the expected value of the ex post rents generated from the information obtained by forecasting is exactly equal to the cost of acquiring that information (the forecasting cost k), so the retailer is left with zero expected profit (zero ex ante information rents). In the remainder of the paper we use *information rents* to mean these ex ante rents, inclusive of the forecasting cost. To see the intuition for why the returns results diverge from those for rebates, recall that there the driver behind the retailer's receiving rents was the retailer's threat of selecting the contract intended for the pessimistic-forecast-observing retailer when the retailer observed the optimistic forecast. The optimal returns menu empties this threat of its power because the full-returns contract is unappealing to the optimistic-forecast-observing retailer (it yields her zero profit).

Further contrast with the typical adverse selection result emerges when the forecasting cost is high. The typical adverse selection result is that the second best contract is characterized by downward distortion in the contract for the low-type agent and “no distortion at the top” (i.e., for the high-type agent). In contrast, in our setting there is no distortion at the bottom $q_L^* = q_L^I$ and upward distortion at the top $q_H^* \geq q_H^I$, where the inequality is strict when the forecasting cost is high $k > \Gamma(q_H^I)$. The intuition for the latter result is that the contract must be designed to discourage the retailer from not forecasting and then selecting the contract intended for the high-type retailer (constraint (IC3)). When the forecasting cost is high, it becomes difficult to encourage the retailer to forecast. However, increasing the quantity in the no-returns contract makes this contract unattractive to any retailer that is not confidently optimistic (where confidence comes from forecasting and observing a favorable signal).

6.3. Forecasting or No Forecasting

The structure of the manufacturer's optimal decision as to whether he should offer returns contracts that encourage or discourage retailer forecasting follows the intuitive structure from the rebates case: It is optimal to offer contracts that induce forecasting if and only if the forecast cost is below a threshold. More surprisingly, this threshold coincides with the threshold for the integrated system, and the optimal menu of returns contracts allows the manufacturer to capture the entire integrated system profit.

PROPOSITION 6. *An optimal menu of returns contracts is given by the following. If $k < k^l$, then the optimal menu has contracts with quantities, buy-back prices, and transfer payments*

$$(q_H^*, b_H^*, t_H^*) = (q_H^l, 0, pE \min(q_H^l, D_H) - k/\lambda),$$

$$(q_L^*, b_L^*, t_L^*) = (q_L^l, p, pq_L^l),$$

which induces the retailer to forecast. Otherwise, the optimal menu has a single contract

$$(q_N^*, b_N^*, t_N^*) = (q_N^l, p, pq_N^l),$$

which induces the retailer not to forecast. Under the optimal menu, the manufacturer's expected profit is the integrated system expected profit.

It is optimal for the manufacturer to offer a full-returns contract. If forecasting is expensive, this is all the manufacturer offers. If forecasting is cheap, the manufacturer also offers a no-returns contract with a larger quantity so as to encourage the retailer to forecast.

The optimal returns menu achieves the integrated system expected profit for the manufacturer. The intuition is the easiest to see when the forecasting cost is large (i.e., $k \geq k^l$). As noted in §6.1, by offering a single properly designed full-returns contract, the manufacturer can always induce the retailer to not forecast and to purchase the associated integrated system quantity. To see the intuition as to why the manufacturer continues to achieve the integrated system expected profit when the forecasting cost is small, recall that the optimal menu of returns contracts that induces forecasting only requires distortion away from the integrated system production quantities when the retailer can credibly threaten not to forecast. When the forecasting cost is small (i.e., $k \leq \Gamma(q_H^l)$), such a threat is not credible, so the optimal menu of returns contracts achieves the integrated system expected profit for the manufacturer. Because $\Gamma(q_H^l) \geq k^l$, there is no intermediate range of forecast costs where the optimal returns menu fails to achieve the integrated system profit for the manufacturer. (Although, admittedly, the full-returns contracts in Proposition 6 are not common in practice, they are consistent, as an approximation, with the offering of a menu of contracts that includes a generous returns option.) Although the

menu of contracts in Proposition 6 allocates the entire system profit to the manufacturer, the contract can be adapted (by reducing the transfer payments) to arbitrarily allocate the profit between the firms.

Proposition 6 demonstrates that returns contracts are superior to rebates contracts: The manufacturer should offer contracts that compensate the retailer for not selling units rather than for selling units. To understand how robust this result is, we elucidate the strengths and weaknesses of each kind of contract.

Rebates contracts are very effective in encouraging the retailer to forecast but are less effective in differentiating between optimistic-forecast-observing and pessimistic-forecast-observing retailers. The intuition for why rebates are effective in encouraging the retailer to forecast is that rebates contracts provide the retailer a "lottery" with rich payoffs when demand is high and poor payoffs when demand is low, so the offer of a rebate contract makes precise knowledge of the demand distribution valuable to the retailer. However, as discussed in §5.2, rebates are ineffective in cheaply distinguishing between different types of retailers. The optimal menu of rebates contracts cedes information rents to the retailer because a contract designed for a pessimistic-forecast-observing retailer yields positive profit to an optimistic-forecast-observing retailer. These information rents are due to differentiating between retailer types rather than inducing the retailer to forecast (this is evident from the fact that the information rents (4) do not depend on the forecasting cost). Thus, to the extent that the manufacturer uses the rebates contracts to induce forecasting, the departure from the system optimum is caused by screening information, not by inducing information acquisition.

In contrast, returns contracts are very effective at differentiating between retailers but may be less effective in encouraging the retailer to forecast. As noted in the discussion following Proposition 5, by offering a menu with a full-returns contract, the manufacturer is able to differentiate between retailer types without ceding (ex ante) information rents. However, when the forecasting cost is high, the manufacturer must distort the production quantity upward to encourage forecasting. Thus, to the extent that the manufacturer uses the returns contracts to induce forecasting, the source of deviation from the system optimum is not from screening information but from inducing information acquisition. Intuitively, because returns provide insurance, such a contract discourages forecasting by making precise knowledge of the demand distribution less valuable. However, offering a menu of returns contracts that includes a no-returns contract mitigates this problem because the more precise demand information obtained by forecasting is

valuable under such a contract. Proposition 6 indicates that the optimal menu of returns contracts is powerful enough to overcome the natural weakness of failing to encourage forecasting.

7. Robustness

Proposition 6 is a remarkably strong result. It says that among all contracts, returns contracts are optimal (i.e., no other form of contract can result in strictly larger manufacturer profit). However, although our model places minimal restrictions on the information that is conveyed by the demand signal, our model makes other assumptions that are restrictive: there are only two possible demand signals, the retailer's forecasting effort decision is binary, and only a single product is relevant. This section provides evidence that the result that returns contracts are optimal is not driven by these assumptions. Furthermore, it provides evidence that our main managerial message, that returns are superior to rebates, is also preserved when the transfer payment is restricted to be linear in the order quantity (as in the commonly studied contracts that feature a wholesale price) and when the retailer's main decision is a sales effort decision or a retail price decision rather than a forecasting effort decision. Proofs of the assertions in this section are provided in the e-companion.

7.1. Arbitrary Number of Demand Signals

So as to facilitate analytical insights, we have relied on the assumption, common to many models that consider contracts between parties with private information, that the private information is of a binary character. In this section, we extend the model to the case with $n \geq 2$ possible demand signals: $S \in \{1, 2, \dots, n\}$. The probability of observing signal i is λ_i , where $\sum_{i=1}^n \lambda_i = 1$, and a larger signal corresponds to a stochastically larger demand distribution: $F_j(x) \leq F_i(x)$ for $i \leq j$. We begin by establishing that the result that returns contracts are optimal among all contracts continues to hold, provided that the forecasting cost is large or small. In the notation below, the random variable D_i is the demand conditioned on observing signal i ; D_N is the unconditioned demand, k^l is the increase in integrated system expected profit from having the better demand information provided by forecasting, and $q_n^l = \bar{F}_n^{-1}(c/p)$ is the optimal production quantity in the integrated system under the most favorable signal.

PROPOSITION 7. *Suppose there are $n \geq 2$ possible demand signals. If $k \geq k^l$ or $k \leq \lambda_n p \min(E[\min(q_n^l, D_n)] - \min(q_n^l, D_{n-1}]), E[\min(q_n^l, D_n)] - \min(q_n^l, D_N)]$, then under the optimal menu of returns contracts, the manufacturer's expected profit is the integrated system expected profit.*

By the argument in §6.1, if it is optimal for the integrated system not to forecast, i.e., $k \geq k^l$, then offering

a single full-returns contract achieves the integrated system expected profit for the manufacturer. At the other extreme, if the forecasting cost is small, a menu of n contracts, with individual contracts specifying either no-returns or full-returns, achieves the integrated system profit for the manufacturer.

Proposition 7 is silent when the forecasting cost is of an intermediate value, so to further test the robustness of the optimality of returns contracts result, we conducted an extensive numerical study for the case of $n = 3$. The parameters are as follows: $p = 10$; $c \in \{2, 5, 8\}$; $\lambda_3 \in \{0.1, 0.3, 0.5, 0.8\}$; $\lambda_2 = \{0.1(1 - \lambda_3), 0.5(1 - \lambda_3), 0.8(1 - \lambda_3)\}$; D_i is a $\text{Normal}(\mu_i, \sigma^2)$ random variable, truncated such that its probability mass is distributed over $x \geq 0$, where $\mu_1 = 2(1 - \theta)$, $\mu_2 = 2$, $\mu_3 = 2(1 + \theta)$, $\theta \in \{0.2, 0.5, 0.8\}$, and $\sigma = 0.5$; and $k \in \{0, 0.05k^l, 0.10k^l, \dots, 0.95k^l\}$. In each of the 2,160 instances, the optimal menu of returns contracts achieves the integrated system expected profit for the manufacturer. In each instance, the optimal menu consists of a no-returns contract for the highest-type retailer ($b_3^* = 0$), a partial-returns contract for the intermediate-type ($b_2^* \in [0, p]$), and a full-returns contract for the lowest-type ($b_1^* = p$). These results, while incomplete, provide support for the suggestion that the result that returns contracts are optimal is robust to the assumption in the base model that there are only $n = 2$ possible demand signals.

7.2. Continuous Forecasting Effort

Our base model is appropriate when the retailer's main forecasting decision is of a binary character, as when the decision is whether to conduct a pre-season merchandise test, whether to conduct a customer survey, or whether to purchase a standardized package of relevant data. However, in other circumstances, the key decision may be more of an incremental nature: how much in the way of time and resources should a firm devote to analyzing data and consulting internal experts (e.g., sales force) and external experts so as to refine its forecast.

To address this second setting, in this section, we extend the model so that the retailer's forecasting effort decision is continuous. The random end-market demand depends on a market condition $M \in \{H, L\}$ as follows: Under market condition M , the demand distribution is $G_M(\cdot)$; as before, demand conditioned on the high-value market condition is stochastically larger than under the low-value condition, i.e., $G_H(x) \leq G_L(x)$. The market condition is high $M = H$ with probability λ and low $M = L$ with probability $1 - \lambda$. The retailer exerts forecasting effort $e \in [0, 1]$ to acquire a private signal S , which is correlated with M in the following way:

$$P(M = H | S = H) = e + (1 - e)\lambda,$$

$$P(M = L | S = L) = e + (1 - e)(1 - \lambda).$$

It is easy to verify that S has the same distribution as M , i.e., $S = H$ with probability λ and $S = L$ with probability $1 - \lambda$. As the retailer exerts greater forecasting effort e , the quality of the signal S in estimating the market condition M improves, in that the probability that the signal corresponds to the market condition increases. (The special case in which the forecasting effort decision is restricted to $e \in \{0, 1\}$ corresponds to our base model: the decision not to forecast corresponds to $e = 0$ in that then the signal contains no information about the market condition; the decision to forecast corresponds to $e = 1$ in that then the signal fully reveals the market condition.) The retailer's forecasting cost, $K(e)$, is increasing and convex in her forecasting effort $e \in [0, 1]$. We retain our assumptions from the base case that all distributions and the cost function are common knowledge, but the retailer's forecasting effort and signal are unobservable to the manufacturer.

Our main theoretical result extends to this setting: Under the optimal menu of returns contracts, the manufacturer's expected profit is the integrated system expected profit. The returns menu is sufficiently powerful not only to induce the retailer to order the optimal system quantities but also to exert precisely the same forecasting effort as the integrated system.

7.3. Multiple Periods and Multiple Products

In our base model, retailer forecasting effort only improves demand information for a single product. Our single-period model is appropriate when retailer forecasting effort is aimed at improving demand information over a short time horizon, namely for a product's selling season. However, retailers may instead invest in longer-term forecasting capabilities (e.g., infrastructure investments in people, processes, and systems that collect, analyze, and interpret demand-relevant data); such investments translate into improved demand information not only for the immediate product but also for subsequent products. To address this setting, we extend the model to the case with $T \geq 2$ periods, where the retailer sells a distinct product in each period. If the retailer invests in forecasting capabilities in period τ , then in period τ and each subsequent period she observes the demand signal prior to making her contract choice decision. Our main theoretical result extends to this setting: By offering a properly designed menu of returns contracts for each product, the manufacturer captures the integrated system expected profit.

Even when the retailer's forecasting effort is aimed at improving near-term demand information, a single-product model may be inadequate. The single-product model is appropriate when the retailer carries a core product that constitutes the bulk of her sales revenue or when the retailer carries multiple products but her forecasting efforts are product-specific

(e.g., conducting a preseason merchandise test for a particular product or conducting a product-specific consumer survey). However, in other settings, the retailer's forecasting investment generates valuable demand information across a number of products, as when a preseason merchandise test or customer survey informs the retailer's understanding of demand across an entire product category. In this case, the forecasting decision has to be made by trading off the fixed forecasting cost with the aggregated benefits from having improved demand information across multiple products. Again, our main theoretical result extends to this setting: By offering a properly designed menu of returns contracts for each product, the manufacturer captures the integrated system expected profit.

7.4. Linear Transfer Payment

In our base model, the transfer payment is allowed to be a general nonlinear function of the retailer's order quantity. Industry practice provides support for this modeling choice: Nonlinear pricing (in the form of quantity discounts) is commonly used in industries where the sales-dependent payments of returns or rebates are also used, such as publishing, pharmaceuticals, and computer hardware (see Padmanabhan and Png 1995, Taylor 2002, Dunehew et al. 2005, Faletta 2006). For example, this kind of nonlinear pricing is used alongside returns in book distribution (Garment 1997) and alongside rebates in computer hardware distribution (Woelbern 2007). In practice, quantity discount contracts often have quantity breakpoints which make ordering at (or slightly above) those breakpoints considerably more attractive than other quantity levels. The fixed quantity contract we propose can be interpreted as a quantity discount scheme with a maximum order quantity (see the step-function form of transfer payment described in §3). Even so, linear transfer payments are also common in practice, as exhibited by the use of contracts that feature a per-unit wholesale price.

We next explore the extent to which our main managerial message—that returns are superior to rebates—continues to hold when the transfer payment is restricted to be a linear function of the retailer's order quantity. Under a wholesale price and rebates contract (w, r) , the retailer decides the order quantity q , pays wq , and receives a per-unit rebate r for every unit sold to end consumers. The wholesale price and returns contract (w, b) is defined similarly.

As in the base case, when the cost of forecasting is sufficiently high $k \geq k^I$ that it is optimal for the integrated system not to forecast, then a single full-returns contract achieves the integrated expected profit for the manufacturer, and so returns dominate rebates.

However, when the cost of forecasting is smaller $k \in (0, k^I)$, returns does not dominate in all cases. For a partial analytical characterization of the optimal menu of each contractual form, see the e-companion. Here we describe the results of a numerical study that compares the performance of the contractual forms. The parameters are as follows: $p = 10$; $c \in \{2, 5, 8\}$; $\lambda \in \{0.1, 0.3, 0.5, 0.8\}$; D_S is a Normal(μ_S, σ^2) random variable for $S \in \{H, L\}$, truncated such that its probability mass is distributed over $x \geq 0$, where $\mu_L = 2(1 - \theta)$, $\mu_H = 2(1 + \theta)$, $\theta \in \{0.2, 0.5, 0.8\}$, and $\sigma = 0.5$; and $k \in \{0, 0.05k^I, 0.10k^I, \dots, 0.95k^I\}$.

Although the universal dominance result that holds under the nonlinear transfer payment does not carry over the case of linear transfer payment, out of the 720 tested instances, rebates are better than returns in only 2.9% of the instances. The parameters of these 21 instances share some common features: high production cost c , large difference between two possible demand distributions (large θ), and moderate value of the forecasting cost k . In the e-companion we discuss the intuition as to why rebates are superior in this (very small) parameter region but inferior elsewhere.

We conclude by stating the two primary conclusions of this analysis. First, the overall thrust of our managerial message, that manufacturers are better off offering returns contracts than rebates contracts, is supported by the numerical study. Second, the observation in §6.3 that returns contracts are optimal among all contracts depends crucially on the provision of nonlinear transfer payments. When returns contracts are coupled with linear transfer payments, the optimal menu of returns contracts does not guarantee the integrated system expected profit for the manufacturer. Indeed, in the interesting case where the forecasting cost is sufficiently small $k \in (0, k^I)$ that forecasting is optimal for the integrated system, the linear transfer price restriction strictly reduces the manufacturer's expected profit. In contrast, in the classical setting without retailer forecasting, returns coupled with a linear transfer payment is sufficient to maximize the total system profit and allocate it arbitrarily (Pasternack 1985).

The second conclusion has two implications. First, the practical implication is that the retailer's ability to improve her knowledge of demand by forecasting makes nonlinear transfer pricing an attractive lever to a manufacturer offering a returns contract. Second, the conclusion provides theoretical support for our study of nonlinear transfer prices in that the commonly studied wholesale price and returns contract is suboptimal, but the simple relaxation of the linear transfer price restriction yields a contract that is optimal among all contracts.

7.5. Endogenous Demand

We have focused on the retailer's decision to exert effort to acquire better demand information. However, in some settings, the key decision facing the retailer regards influencing demand directly (e.g., through sales effort or the retail price) instead of learning about exogenous demand (forecasting). We consider sales effort and retail price selection sequentially. The scenario in which sales effort is the key decision occurs when activities like advertising and promotions substantially impact demand and when the costs of these activities are much higher than the costs of forecasting. The scenario in which the retail price is the key decision occurs when the retailer has pricing power, demand is sensitive to the retail price, and the demand information available to the retailer is insensitive to her efforts.

To address the sales-effort scenario, we modify the base model as follows: After the manufacturer offers the menu of contracts, but before the retailer costlessly observes the signal of demand, the retailer exerts sales effort, which is captured in our setting by choosing $\lambda \in [0, 1]$, the probability that the actual demand distribution is favorable. The cost of sales effort is $\mathcal{K}(\lambda)$, which is assumed to be increasing and strictly convex.

Our base model demonstrates that when the key decision facing the retailer is a forecasting decision, returns are superior to rebates. One might conjecture that when the key decision facing the retailer is sales effort, the result would be reversed: Because rebates reward the retailer for selling units whereas returns reward the retailer for not selling units, the manufacturer should be better off offering rebates because that instrument better rewards and encourages the retailer's sales effort. Surprisingly, this conjecture is false: The manufacturer's expected profit is greater under the optimal menu of returns contracts than under the optimal menu of rebates contracts. The intuition is twofold. First, to counter the conjecture that returns are ineffective in encouraging sales effort, consider a menu with a no-returns contract coupled with a transfer payment providing a relatively low per-unit acquisition cost and a full-returns contract coupled with a high per-unit acquisition cost. Under this menu, the market is considerably more attractive to the retailer when the demand forecast is favorable because then the retailer can take advantage of the low acquisition cost, no-returns contract. Second, the intuition from the base model, that returns are more effective than rebates at cheaply differentiating between retailers that have observed different demand signals, extends. This, taken together with the effectiveness of the *menu* of returns contracts in encouraging sales effort, is what establishes that

our main managerial message persists in the sales-effort setting: returns are superior to rebates. However, under retailer sales effort, the result that the optimal returns menu achieves the integrated system profit for the manufacturer no longer obtains: returns are more effective in providing appropriate incentives for forecasting effort than sales effort.

When the retailer influences the demand distribution by setting the retail price rather than through sales effort, the results described above continue to hold. To address the retail-pricing scenario, we modify the base model as follows: After costlessly observing the demand signal $S \in \{H, L\}$ and selecting a contract, the retailer sets the retail price p , which impacts the distribution of random demand $D_S(p)$. Under endogenous retail pricing, the two results from the sales effort case continue to hold: First, returns are superior to rebates. Second, the manufacturer's expected profit under the optimal returns menu is strictly lower than the integrated system expected profit.

8. Discussion

This paper compares the effectiveness of rebates and returns contracts when a retailer can improve her demand information by exerting costly forecasting effort. We show that under the optimal menu of rebates contracts, the retailer, manufacturer, and total system may benefit from retailer having inferior forecasting technology; in addition, the retailer may overinvest in forecasting. Under returns, none of these results occur.

Our main managerial message is that in a setting with endogenous forecasting, a manufacturer should offer returns contracts rather than rebates contracts. Although rebates contracts are effective in encouraging forecasting, they are much less effective than returns contracts in cheaply differentiating among forecasting retailers that have observed different signals, and it is this fact that leads to the superiority of returns contracts. To establish the robustness of the conclusion that returns are superior, we provide evidence that it holds for somewhat general models of forecasting effort (continuous effort choice, multiple demand signals), when the transfer price is nonlinear in the order quantity or linear (except in a few rare cases), when there are multiple periods, when there are multiple products, and when the retailer's key decision is a sales effort decision (or a retail price decision) rather than a forecasting effort decision. Indeed, in several of these settings we provide analytical and/or numerical results that support an even stronger conclusion: the optimal menu of returns contracts achieves the integrated system expected profit for the manufacturer and so is optimal among all possible contractual forms.

9. 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, Martin Lariviere, Robert Shumsky, and seminar participants at the Chinese University of Hong Kong, Hong Kong University of Science and Technology, INSEAD, New York University, Northwestern University, Southern Methodist University, Stanford University, University of California-Berkeley, and University of California-Santa Cruz for helpful comments.

References

- Arya, A., B. Mittendorf. 2004. Using returns policies to elicit retailer information. *RAND J. Econom.* 35(3) 617–630.
- Aydin, G., E. L. Porteus. 2008. Manufacturer-to-retailer versus manufacturer-to-consumer rebates in a supply chain. N. Agrawal, S. Smith, eds. *Retail Supply Chain Management*. Springer, New York, 237–270.
- Cachon, G. P. 2003. Supply chain coordination with contracts. S. C. Graves, A. G. de Kok, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management*. Elsevier, Amsterdam, 229–339.
- Chen, F. 2003. Information sharing and supply chain coordination. S. C. Graves, A. G. de Kok, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management*. Elsevier, Amsterdam, 341–421.
- Chen, F., W. Xiao. 2009. Provision of incentives for information acquisition: A comparison between forecast-based contracts and menus of linear contracts. Working paper, Columbia University, New York.
- Donohue, K. L. 2000. Efficient supply contracts for fashion goods with forecast updating and two production models. *Management Sci.* 46(11) 1397–1411.
- Dunehew, A. R., H. Stokes, J. Newberry, D. Evangelista. 2005. Contract price integrity for pharmaceuticals. *Amer. J. Health-System Pharm.* 62 422–429.
- Emmons, H., S. M. Gilbert. 1998. Note: The role of returns policies in pricing and inventory decisions for catalogue goods. *Management Sci.* 44(2) 276–283.
- Faletra, R. 2006. Get ready for 2006 to become the year of the loyalty discussion. *CRN* (January 9) 58.
- Fisher, M., K. Rajaram. 2000. Accurate retail testing of fashion merchandise: Methodology and application. *Marketing Sci.* 19(3) 266–278.
- Garment, L. 1997. Mid-list crisis. *Forbes* (October 20) 118–119.
- Krishnan, H., R. Kapuscinski, D. Butz. 2004. Coordinating contracts for decentralized supply chains with retailer promotional effort. *Management Sci.* 50(1) 48–63.
- Lariviere, M. A. 2002. Inducing forecast revelation through restricted returns. Working paper, Northwestern University, Evanston, IL.
- Laucka, K. 2005. The accuracy trap. *Supply Demand Chain Executive* online supplement (July 5), <http://www.sdcexec.com/online/archives.jsp?siteSection=22&date=2005-July>.
- Padmanabhan, V., I. P. L. Png. 1995. Returns policies: Make money by making good. *Sloan Management Rev.* 37(1) 65–72.
- Padmanabhan, V., I. P. L. Png. 1997. Manufacturer's returns policies and retail competition. *Marketing Sci.* 16(1) 81–94.

- Padmanabhan, V., I. P. L. Png. 2004. Reply to “Do returns policies intensify competition?” *Marketing Sci.* **23**(4) 614–618.
- Pasternack, B. A. 1985. Optimal pricing and return policies for perishable commodities. *Marketing Sci.* **4**(2) 166–176.
- Sarvary, M., V. Padmanabhan. 2001. The informational role of manufacturer returns policies: How they can help in learning the demand. *Marketing Lett.* **12**(4) 341–350.
- Shin, H., T. Tunca. 2009. The effect of competition on demand forecast investments and supply chain coordination. Working paper, Stanford University, Stanford, CA.
- Taylor, T. A. 2002. Coordination under channel rebates with sales effort effects. *Management Sci.* **48**(8) 992–1007.
- Taylor, T. A. 2006. Sale timing in a supply chain: When to sell to the retailer. *Manufacturing Service Oper. Management* **8**(1) 23–42.
- Tsay, A. A. 2001. Managing retail channel overstock: Markdown money and return policies. *J. Retailing* **77**(4) 457–492.
- Woelbern, M. 2007. Private conversation with Ingram Micro executive.
- Zarley, C. 1994. IBM PC co. tries to determine real demand: Want ThinkPads? Assure no returns. *Comput. Reseller News* (October 31) 1.