Incentives for Retailer Forecasting: Rebates versus Returns

Terry A. Taylor and Wenqiang Xiao

Graduate School of Business, Columbia University and Tuck School of Business, Dartmouth College

Stern School of Business, New York University

Abstract

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. We show that the manufacturer is better off offering returns, and, in fact, returns contracts are optimal among all contracts.

Key words: supply-chain contracting; forecasting; rebates; returns; endogenous asymmetric information

October 2006

1 Introduction

Efficiently matching supply with demand is a fundamental concern in supply chain management. By improving the quality of demand information, forecasting efforts play a crucial role in alleviating the problem of supply-demand mismatch. Although information about consumer demand obtained via forecasting is beneficial to the entire supply chain, forecasting is often costly for some supply chain member. For example, fashion retailers often organize a pre-season sales event to obtain information about consumers' tastes for new apparel designs. The sales generated in the event, and more generally the retailer's qualitative assessment of the interests and preferences exhibited by fashion-forward customers, often provide valuable information about what demand the retailer should anticipate in the regular selling season. However, organizing a pre-season sales event is costly for a retailer. 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.

Although the more precise demand information obtained by retailer forecasting is 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. On the other hand, if the retailer has better information about demand, a properly structured contract could allow the production quantity and payments to reflect this more precise information, to the benefit of both firms.

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 rebates contract, the manufacturer pays the retailer a bonus for each unit the retailer sells. Under a returns contract, the manufacturer compensates the retailer by buying back each unsold unit the retailer has at the end of the selling season. 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: 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 its 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. Further, returns contracts are optimal among all contracts: under the optimal menu of returns contracts, the manufacturer captures the integrated system profit.

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 concluding remarks.

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 instrument for the manufacturer in newsvendor settings (Pasternack 1985), in settings with endogenous retail pricing (Emmons and Gilbert 1998, Bernstein and Federgruen 2005, Granot and Yin 2005, Song et al. 2006), and in multi-period settings (Donohue 2000 and Taylor 2001). 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 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, where 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 their 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 2005) and in settings with endogenous retailer sales effort (Taylor 2002 and Krishnan et al. 2004). However, when the retailer exerts sales effort after the demand uncertainty is resolved, Chen and Xiao (2006) show that rebates may actually hurt the manufacturer because the retailer can manipulate the realization of sales.

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. Alternately, by initiating contractual terms, the informed party may credibly convey ("signal") his information to the uninformed. 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 (2005) examine salesforce compensation where the salesperson can acquire costly demand signals prior to the firm's production decision. They show how the cost of information acquisition impacts the relative performance of forecast-based and linear-menu compensation schemes.

The paper closest to our work is Lariviere (2002). In his model, 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 non-forecasting retailer. In contrast, 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, while ours is to compare returns with rebates.

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 $S \in \{H, L\}$. 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 $\overline{F}_S(x) \equiv 1 - F_S(x)$ for $S \in \{N, H, L\}$; 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 $T^r(q, r) : R^+ \times R^+ \to R^+$ 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 $T^b(q, b) : R^+ \times [0, p] \to R^+$ 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 both cases, 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 rebate or buy-back price.

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.

One would typically expect that the transfer price would be increasing in the generosity of the rebate or buy-back price and in the order quantity. As an example of such a menu of contracts, the distributor McKesson offers its retailers the option of a high purchase price coupled with a generous buy-back price or a low purchase price coupled with a stingy buy-back price.

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. Further, 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: 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 non-forecasting 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.

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

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

and the optimal production quantity is $q_N^o = \overline{F}_N^{-1}(c/p)$. 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

$$\Pi_F(q_H, q_L) = \lambda\left[pE\min(q_H, D_H) - cq_H\right] + (1-\lambda)\left[pE\min(q_L, D_L) - cq_L\right].$$

After observing the signal $S \in \{H, L\}$, in setting its production quantity the system again faces a newsvendor problem, and the optimal production quantity is $q_S^o = \overline{F}_S^{-1}(c/p)$. 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 $\Pi_F(q_H^o, q_L^o) \geq \Pi_N(q_N^o)$. Accordingly, it is optimal to forecast if and only if the forecasting cost is sufficiently low. Proposition 1 summarizes the optimal forecasting and production quantity decisions for the integrated system. Let $k^o \equiv \Pi_F(q_H^o, q_L^o) - \Pi_N(q_N^o)$.

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

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

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

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 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

$$\max_{q_N, r_N, t_N} \{t_N - cq_N - r_N E\min(q_N, D_N)\} \tag{OBJ}$$

$$\text{s.t. } R^r(N, N) \geq \lambda \max(R^r(H, N), 0) + (1 - \lambda) \max(R^r(L, N), 0) - k \tag{IC}$$

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

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 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 non-forecasting retailer is willing to accept the contract.

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

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

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

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 ($k \geq \Gamma(q_N^o)$), 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 (non-forecasting) integrated system by offering a contract that sells the integrated-system optimal quantity q_N^o 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 ($k \in (\Gamma(q_L^o), \Gamma(q_N^o))$) 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 ($k \leq \Gamma(q_L^o)$), 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 $q_N^* \le q_N^o$.

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-a-vis the principal (manufacturer). In contrast, in our setting when the forecasting cost is low ($k < \Gamma(q_L^o)$), the agent (retailer) earns strictly positive expected profits even though she lacks an informational advantage vis-a-vis the principal (manufacturer): $R^r(N, N) = [\Gamma(q_L^o) - k]/(1 - \lambda) > 0$. Under the optimal contract, the agent 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. 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

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

It is straightforward to verify that 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 principal, 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

$$\max_{(q_H,r_H,t_H),\,(q_L,r_L,t_L)} \{\lambda\,[t_H - cq_H - r_H E\min(q_H, D_H)] + (1-\lambda)\,[t_L - cq_L - r_L E\min(q_L, D_L)]\} \tag{OBJ}$$

$$\text{s.t. } R^r(H,H) \geq R^r(H,L) \tag{IC1}$$

$$R^r(L,L) \geq R^r(L,H) \tag{IC2}$$

$$R^r(H,H) \geq 0 \tag{IR1}$$

$$R^r(L,L) \geq 0 \tag{IR2}$$

$$\lambda R^r(H,H) + (1-\lambda)R^r(L,L) - k \geq R^r(N,H) \tag{IC3}$$

$$\lambda R^r(H,H) + (1-\lambda)R^r(L,L) - k \geq R^r(N,L) \tag{IC4}$$

$$\lambda R^r(H,H) + (1-\lambda)R^r(L,L) - k \geq 0. \tag{IR3}$$

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.

The first step in characterizing the solution to (OBJ)-(IR3) is to simplify the optimization problem by showing that several of the constraints are redundant. First, because $R^r(N,L) \geq R^r(L,L)$, (IC4) and (IR2) imply (IR3). Because $R^r(H,L) \geq R^r(L,L)$, (IC1) and (IR2) imply (IR1). Third, because $R^r(N,H) = \lambda R^r(H,H) + (1-\lambda)R^r(L,H)$, (IC3) implies (IC2). Similarly, (IC4) implies (IC1).

Removing the redundant constraints, (OBJ)-(IR3) simplifies to

$$\max_{(q_H,r_H,t_H),\,(q_L,r_L,t_L)} \{\lambda\,[t_H - cq_H - r_H E\min(q_H, D_H)] + (1-\lambda)\,[t_L - cq_L - r_L E\min(q_L, D_L)]\}$$

$$\text{s.t. } R^r(L,L) \geq 0 \tag{IR2}$$

$$(1-\lambda)R^r(L,L) - k \geq (1-\lambda)R^r(L,H) \tag{IC3}$$

$$\lambda R^r(H,H) - k \geq \lambda R^r(H,L), \tag{IC4}$$

where we use the fact $R^r(N,C) = \lambda R^r(H,C) + (1-\lambda)R^r(L,C)$ (for $C = H, L$) to rewrite (IC3) and (IC4).

To further simplify the manufacturer's problem, we note that (IR2) and (IC4) are binding at the optimal solution. If (IR2) were not binding, then increasing both t_L and t_H by a small amount would increase the objective value without violating the constraints. Similarly, if (IC4) were not binding, increasing t_H by a small amount would increase the objective value without violating the constraints. The two binding constraints imply that an optimal solution has

$$\begin{aligned} t_L &= (p + r_L)E\min(q_L, D_L) \\ t_H &= (p + r_H)E\min(q_H, D_H) - (p + r_L)\Delta(q_L) - k/\lambda, \end{aligned}$$

where $\Delta(q) = E[\min(q, D_H) - \min(q, D_L)]$ is the expected increase in units sold when the retailer has q units and observes the favorable rather than unfavorable demand signal. This simplifies the problem to

$$\max_{(q_H,r_H),\,(q_L,r_L)} \{\lambda\,[pE\min(q_H, D_H) - cq_H - (p + r_L)\Delta(q_L) - k/\lambda] + (1-\lambda)\,[pE\min(q_L, D_L) - cq_L]\}$$

$$\text{s.t. } (p + r_H)\Delta(q_H) - (p + r_L)\Delta(q_L) \geq k/\lambda(1-\lambda). \tag{IC3}$$

Because as r_L decreases, the objective value increases and the constraint is relaxed, it is optimal to set $r_L^* = 0$. Because r_H does not appear in the objective function, we can set r_H to be a large value such that (IC3) is always satisfied for any (q_H, q_L). This leads to an unconstrained optimization problem with only two variables:

$$\max_{q_H,q_L} \{\lambda[pE\min(q_H, D_H) - cq_H] + pE[\min(q_L, D_L) - \lambda\min(q_L, D_H)] - (1-\lambda)cq_L - k\}. \tag{2}$$

The solution to (2), together with the analysis above, establishes an optimal menu of contracts, which we state formally in the next proposition.[1]

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

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

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

$$\lambda R^r(H,H) + (1-\lambda)R^r(L,L) - k = \lambda R^r(H,L) \tag{3}$$

$$= R^r(N,L) \tag{4}$$

$$= \lambda p \Delta(q_L^*), \tag{5}$$

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 the optimistic-forecast-observing retailer and penalize the pessimistic-forecast-observing retailer. The former re-

[1] 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.

ceives expected profit of $p\Delta(q_L^*) + (1-\lambda)k/\lambda$, while the latter incurs an expected loss of k. To reduce the retailer's information rents, the production quantity for the low-type retailer is distorted downward $q_L^* \le q_L^o$, while there is no distortion for the high-type retailer $q_H^* = q_H^o$, 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 (3). Equation (4) provides an alternate 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 $R^r(L, L) \ge 0$) will yield greater expected profit to the non-forecasting 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 (4)). 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 (5)); the manufacturer bears the cost entirely. The manufacturer's expected profit is

$$\mathcal{M}_F^r = \lambda \left[p \int_0^{q_H^*} \overline{F}_H(x)dx - cq_H^* \right] + (1-\lambda) \left[p \int_0^{q_L^*} \frac{\overline{F}_L(x) - \lambda \overline{F}_H(x)}{1-\lambda} dx - cq_L^* \right] - k. \tag{6}$$

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 (6)) 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 M_N^r increases, and because at $k = 0$, $\mathcal{M}_F^r \ge \mathcal{M}_N^r$, there exists a threshold $k^r \ge 0$ such that

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

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. 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's top panel 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. Figure 1's middle panel 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 ($k \in (0.26, 1.06)$ 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 up.

Third, Figure 1's bottom panel 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 provide 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

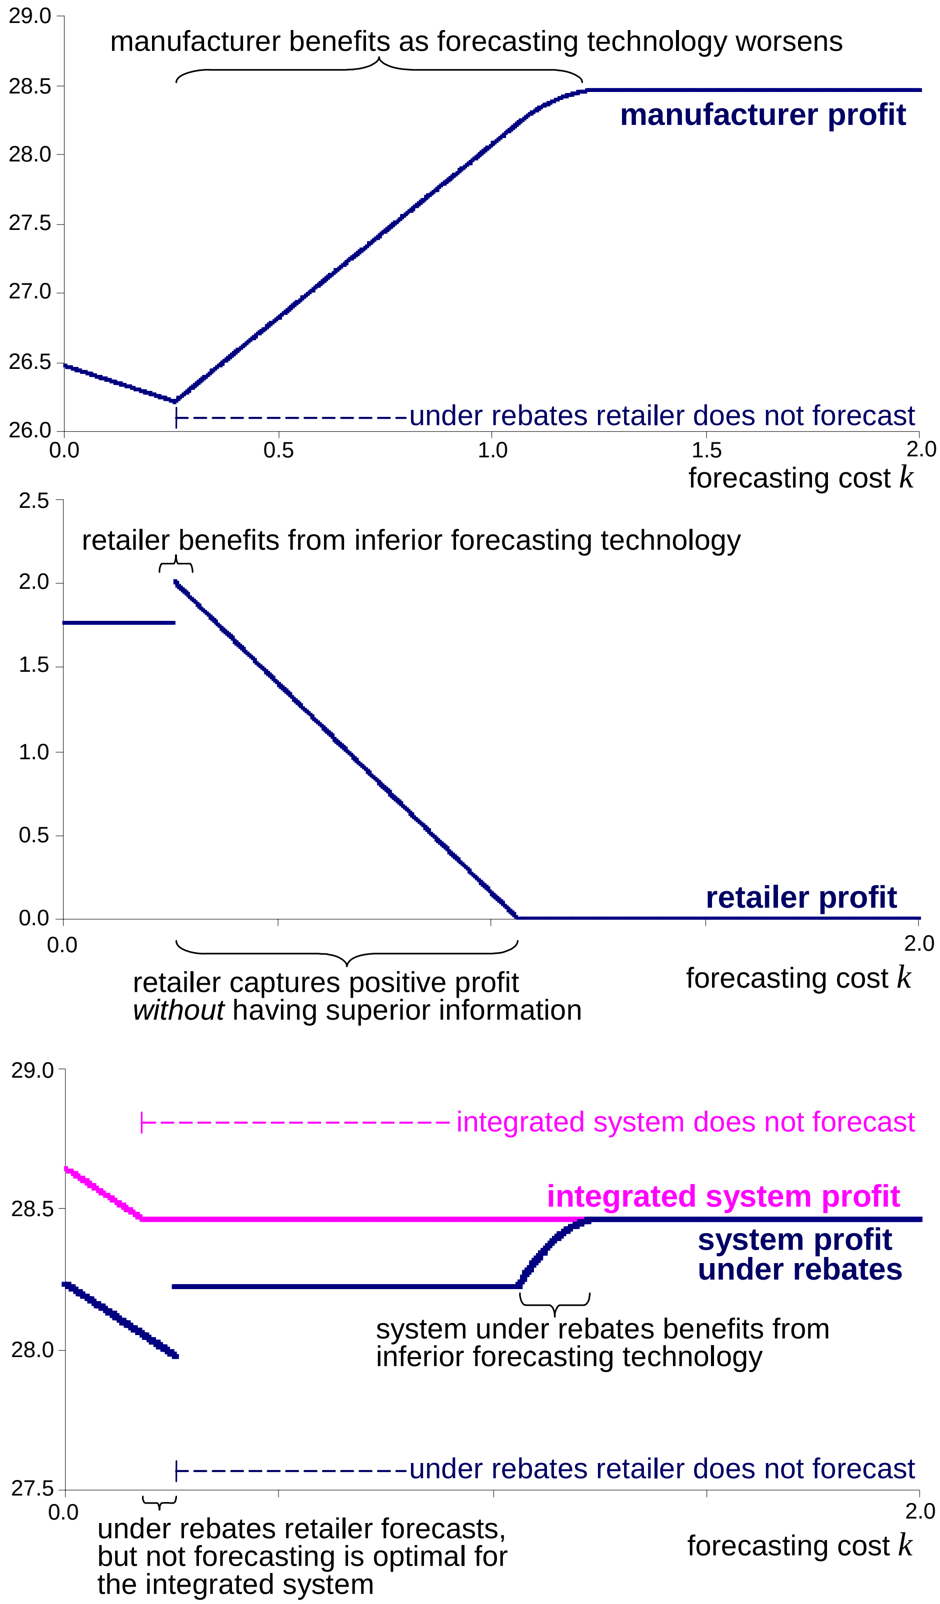

Figure 1: Expected Profits under Optimal Rebate Contracts. Parameters are $p = 20$, $c = 10$, $\lambda = 0.6$, D_H ~exponential(0.10) and D_L ~exponential(0.12).

expected profit under the optimal rebate contracts, i.e., $\mathcal{M}^r = \max(\mathcal{M}^r_F, \mathcal{M}^r_N)$, and let $\mathcal{R}^r$ denote the retailer's expected profit under these contracts.

Proposition 4. *If $ED_H > ED_L$, then there exists $\overline{c}$, $\underline{k}$ and $\overline{k}$ with $\overline{c} > 0$ and $\underline{k} < \overline{k}$ such that if $c < \overline{c}$, 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 $k \in [\underline{k}, \overline{k}]$. If $\lambda \overline{F}_H(x) \geq \overline{F}_L(x)$ for $x > 0$, then there exists $\underline{c} > 0$ and $\widehat{k} > k^r$ such that if $c < \underline{c}$, then the retailer's expected profit is higher when her forecasting cost is higher:*

$$\mathcal{R}^r|_{k \in [0, k^r)} < \mathcal{R}^r|_{k \in (k^r, \widehat{k})}. \tag{7}$$

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 phenomena 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 pre-season sales event). 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 (5) and (6)). If the retailer's forecasting cost is sufficiently high that 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 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's bottom panel 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 (rebate) 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.

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 $q_N = q_N^o$, the manufacturer's expected profit is identical to expected profit of the non-forecasting integrated system $\Pi_N(q_N^o)$. 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 $(q_N, b_N, t_N) = (q_N^o, p, pq_N^o)$.

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

$$\lambda \left[t_H - cq_H - b_H E\max(q_H - D_H, 0)\right] + (1 - \lambda) \left[t_L - cq_L - b_L E\max(q_L - D_L, 0)\right].$$

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

Using the same arguments that are employed to simplify the manufacturer's problem under rebates contracts, the manufacturer's problem under returns contracts simplifies to

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

$$\text{s.t. } R^b(L,L) = 0 \tag{IR2}$$

$$(1-\lambda)R^b(L,L) - k \geq (1-\lambda)R^b(L,H) \tag{IC3}$$

$$\lambda R^b(H,H) - k = \lambda R^b(H,L). \tag{IC4}$$

The two binding constraints (IR2) and (IC4) imply that an optimal solution has

$$\begin{aligned} t_L &= pE\min(q_L, D_L) + b_L E\max(q_L - D_L, 0) \\ t_H &= pE\min(q_H, D_H) + b_H E\max(q_H - D_H, 0) - (p - b_L)\Delta(q_L) - k/\lambda, \end{aligned}$$

which further simplifies the problem to

$$\max_{(q_H,b_H),(q_L,b_L)} \{\lambda\,[pE\min(q_H, D_H) - cq_H - (p-b_L)\Delta(q_L) - k/\lambda] + (1-\lambda)\,[pE\min(q_L, D_L) - cq_L]\}$$

$$\text{s.t. } (p - b_H)\Delta(q_H) - (p - b_L)\Delta(q_L) \geq k/\lambda(1-\lambda). \tag{IC3}$$

Because as b_L increases or b_H decreases, constraint (IC3) is relaxed and the objective (weakly) increases, an optimal solution has $b_L = p$ and $b_H = 0$. This simplifies the problem to two separate optimization problems: an unconstrained maximization problem in q_L and a constrained maximization in q_H. The solution to each is straightforward, and combining this solution with the analysis above, establishes an optimal menu of contracts, which we state formally in the next proposition.[2]

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

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

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

[2] As a technical point, we define $\Gamma^{-1}(k) = \min\{q : \Gamma(q) = k\}$. If (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, in which case, $\Gamma^{-1}(k)$ is well defined.

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. The optimistic-forecast-observing retailer receives expected profit of $k(1-\lambda)/\lambda$, and the pessimistic-forecast-observing retailer incurs an expected loss of k, so that in expectation the forecasting retailer's profit is zero. 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 $q_L^* = q_L^o$ and upward distortion at the top $q_H^* \geq q_H^o$, where the inequality is strict when the forecasting cost is high $k > \Gamma(q_H^o)$. 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 $k < k^o$, then the optimal menu has contracts with quantities, buy-back prices, and transfer payments*

$$(q_H^*,\quad b_H^*,\quad t_H^*) \quad = \quad (q_H^o,\quad 0,\quad pE\min(q_H^o, D_H) - k/\lambda)$$
$$(q_L^*,\quad b_L^*,\quad t_L^*) \quad = \quad (q_L^o,\quad p,\quad pq_L^o),$$

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

$$(q_N^*,\quad b_N^*,\quad t_N^*) \ = \ (q_N^o,\quad p,\quad pq_N^o),$$

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., $k \geq k^o$). 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., $k \leq \Gamma(q_H^o)$), such a threat is not credible, and so the optimal menu of returns contracts achieves the integrated system expected profit for the manufacturer. Because $\Gamma(q_H^o) \geq k^o$, 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, by reducing the transfer payments the contract can be adapted 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 with 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 a 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 (5) 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 information rents. However, when the forecasting cost is high, the manufacturer must distort the production quantity upward in order 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 when there are two possible signals, the optimal menu of returns contracts is powerful enough to overcome the natural weakness of failing to encourage forecasting.

Indeed, Proposition 6 is a remarkably strong result. It says that among all contracts, returns contract are optimal (i.e., no other form of contract can result in strictly larger profit). However, although our model places minimal restrictions on the information that is conveyed by the demand signal, our assumption that there are only two possible signals is strong. The next section provides evidence that the result that returns contracts are optimal is not driven by this assumption.

6.4 Extension to Arbitrary Number of Demand Signals

As in many models that consider contracts between parties with private information, we have relied on the assumption that the private information is of a binary character, so as to facilitate analytical insights. In this section, we explore how our main conclusion depends on this assumption. Specifically, we explore how the comparison between rebates and returns extends to the general case with $n \geq 2$ possible values of the demand signal. We show analytically that the returns contracts continue to achieve the integrated system expected profit for the manufacturer when the forecasting cost is either small or large; we also conduct an extensive numerical study for the case of $n = 3$ and find that in every instance, the optimal returns menu achieves the integrated system expected profit for the manufacturer. These results, although incomplete, seem to indicate the robustness of our conclusion that returns dominate rebates when the retailer's forecasting decision is endogenous.

It is straightforward to generalize the model presented in §3 to allow for an arbitrary number of demand signals. Suppose the demand signal S has n possible values $S \in \{1, 2, .., n\}$. The probability of observing signal i is λ_i and $\sum_{i=1}^{n} \lambda_i = 1$. A larger signal corresponds to a stochastically larger demand distribution: $\overline{F}_i(x) \leq \overline{F}_j(x)$ for $i \leq j$. Proposition 1, which characterizes the optimal forecasting and production decisions for the integrated system, extends when the definitions are extended in the natural way. For example, the optimal production quantity in the integrated system when the realized signal is i is $q_i^o = \overline{F}_i^{-1}(c/p)$.

Now consider the use of returns contracts. By the argument in §6.1, if it is optimal for the centralized system not to forecast, i.e., $k \geq k^o$, 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 (i.e., $k \leq \lambda_n p \min\left(\Delta(q_n^o, n-1), \Delta(q_n^o, N)\right)$, where $\Delta(q,i) \equiv E[\min(q, D_n) - \min(q, D_i)]$), the same outcome is achieved by offering a no-returns contract intended for the retailer that has observed the most favorable signal and $n-1$ full-return contracts (with different quantities) intended for the retailers observing the other signals: an optimal menu of contracts $\{(q_i^*, b_i^*, t_i^*)\}_{i=1,..,n}$ has $q_n^* = q_n^o$, $b_n^* = 0$, $t_n^* = pE\min(q_n^o, D_n) - k/\lambda_n$ and $q_i^* = q_i^o$, $b_i^* = p$, $t_i^* = pq_i^o$ for $i = 1, 2, ..., n-1$.

To see the rationale, first note that the menu prescribes the integrated-system production quantity for each retailer type. It remains to show that under this menu, it is in the retailer's interest to forecast and choose the contract intended for the revealed demand signal (i.e., the menu satisfies the incentive compatibility constraints), the retailer's expected profit before and after forecasting is no lower than her reservation level (i.e., the menu satisfies the individual rationality constraints), and the retailer's expected profit before forecasting is equal to her reservation profit (i.e., the manufacturer captures the entire system profit). Under the menu, the retailer that has observed the most favorable signal finds the no-returns contract attractive because it yields a positive expected profit of k/λ_n while all the remaining full-returns contracts yield zero expected profit (where we exclude the sunk forecasting cost). Because $k \leq \lambda_n p \Delta(q_n^o, n-1)$, a forecasting retailer that has observed any other signal has no interest in the no-returns contract because it yields her negative expected profit (and the contract intended for her yields zero expected profit). Thus, each forecasting retailer selects the intended contract. Because the forecasting retailer incurs cost k and then receives a positive payoff of k/λ_n with probability λ_n and receives a zero payoff otherwise, the forecasting retailer's expected profit is zero. The retailer can do no better than this by choosing a contract without forecasting: the no-returns contract yields negative profit to the non-forecasting retailer (because $k \leq p\lambda_n \Delta(q_n^o, N)$) and all the remaining full-returns contracts yield zero expected profit. To summarize, we have the following proposition.

Proposition 7. *Suppose there are $n \geq 2$ possible demand signals. If $k \geq k^o$ or $k \leq \lambda_n p \min(\Delta(n-1, q_n^o), \Delta(N, q_n^o))$, then under the optimal menu of returns contracts, the manufacturer's expected profit is the integrated system expected profit.*

The above proposition demonstrates that the dominance result shown in the case of $n = 2$ signals continues to hold in the general case, when the forecasting cost is large or small. However, the proposition is silent when the forecasting cost is of an intermediate value. To further test the robustness of the dominance result, we conducted an extensive numerical study for the case of $n = 3$. The parameters are as follows: $p = 10$; $c = \{2, 5, 8\}$; $\lambda_3 = \{0.1, 0.3, 0.5, 0.8\}$; $\lambda_2 = \{0.1(1-\lambda_3), 0.5(1-\lambda_3), 0.8(1-\lambda_3)\}$; D_i is a Normal(μ_i, σ^2) random variable, truncated such that its probability mass is distributed over $x \geq 0$, where $\mu_1 = 2(1-\theta)$, $\mu_2 = 2$, $\mu_3 = 2(1+\theta)$, $\theta = \{0.2, 0.5, 0.8\}$, and $\sigma = 0.5$; and $k = \{0, 0.05k^o, 0.1k^o, ..., 0.95k^o\}$. In each of the 2160 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 ($b_1^* = p$).

7 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 result is that in a setting with endogenous forecasting, the manufacturer should offer returns contracts rather than rebates contracts. Although rebates are very effective at encouraging forecasting, they are much less effective at cheaply differentiating among forecasting retailers that have observed different signals. In contrast, returns contracts are highly effective: a combination of analytical and numerical results suggests that the optimal menu of returns contracts achieves the integrated system expected profit for the manufacturer, and so is optimal among all possible contractual forms.

To the extent that returns contracts are optimal among all contractual forms, it is worth commenting on the value of studying rebate contracts. First, it is natural to compare the effectiveness of these two contractual forms because they are mirror images of one other, rewarding the retailer either for selling or not selling units. Second, rebates contracts are widely used in practice in settings where retailers can improve their understanding of demand by exerting costly forecasting effort. It is possible that for historical or other reasons firms employ rebates contracts despite the existence of other contract forms that would yield greater profit. Alternately, there may be other factors outside our model (e.g., sales effort), that explain their use.

The purpose of this paper is to take a first step towards understanding the implications for contract design of a retailer's ability to improve her demand information by exerting costly forecasting effort. There are several directions in which future research can enrich this understanding. First, we assume that the forecast decision is binary. Relaxing this assumption would allow a more refined study of how much forecasting effort the manufacturer should induce. Second, we consider a fairly general class of contracts which allows the transfer payment to be a general function of the quantity purchased. Focusing on contracts with linear transfer payments would provide insights into simpler, easier to implement contracts. Third, we assume that the demand distribution is exogenous. However, not only can retailers improve their understanding of their demand distribution by forecasting, they also can directly influence this distribution by exerting sales effort. A richer model could provide insights into understanding the interaction between these two decisions and the consequent implications for contract design.

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. 2005. The effects of retailer vs. consumer rebates in a supply chain. Working paper, University of Michigan.

Bernstein, F., A. Federgruen. 2005. Decentralized supply chains and competing retailers under demand uncertainty. *Management Sci.* **51**(1) 18-29.

Cachon, G. 2003. Supply chain coordination with contracts. S. Graves and T. De Kok, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management.* Elsevier, Amsterdam.

Chen, F. 2003. Information sharing and supply chain coordination. S. Graves and T. De Kok, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management.* Elsevier, Amsterdam.

Chen, F., W. Xiao. 2005. Salesforce incentives, information acquisition, and production planning.

Working paper, Columbia University.

Chen, F., W. Xiao. 2006. Is using target rebates good for manufacturers? Working paper, Columbia University.

Donohue, K. L. 2000. Efficient supply contracts for fashion goods with forecast updating and two production models. *Management Sci.* **46**(11) 1397-1411.

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.

Granot, D., S. Yin. 2005. On the effectiveness of returns policies in the price-dependent newsvendor model. *Naval Res. Logist.* **52**(8) 765-779.

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.

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 Letters* **12**(4) 341-350.

Song, Y., S. Ray, S. Li. 2006. Structural properties of buyback contracts for price-setting newsvendors. Working paper, McGill University.

Taylor, T. A. 2001. Channel coordination under price protection, mid-life returns, and end-of-life returns in dynamic markets. *Management Sci.* **47**(9) 1220-1234.

Taylor, T. A. 2002. Coordination under channel rebates with sales effort effect. *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.

Appendix

Proof of Proposition 2. Let $\pi_N \equiv R^r(N,N)$. Clearly, $R^r(H,N) = \pi_N + (p+r_N)\int_0^{q_N}[\overline{F}_H(x) - \overline{F}_N(x)]dx$ and $R^r(L,N) = \pi_N - (p+r_N)\int_0^{q_N}[\overline{F}_N(x) - \overline{F}_L(x)]dx$. Thus, the contract design problem (OBJ)-(IR) can be rewritten as

$$\max_{q_N, r_N, \pi_N} \left\{p\int_0^{q_N}\overline{F}_N(x)dx - cq_N - \pi_N\right\} \qquad \text{(OBJ}'\text{)}$$

$$\text{s.t. } (1-\lambda)\min\left\{0, \pi_N - (p+r_N)\int_0^{q_N}[\overline{F}_N(x) - \overline{F}_L(x)]dx\right\} + k \geq 0 \qquad \text{(IC}'\text{)}$$

$$\pi_N \geq 0. \qquad \text{(IR}'\text{)}$$

The above problem (OBJ′)-(IR′) can be further simplified as follows. First, because decreasing r_N relaxes (IC′) without any impact on (OBJ′) and (IR), an optimal solution has $r_N^* = 0$. Second, because $k \geq 0$, (IC′) can be simplified to $(1-\lambda)\left\{\pi_N - p\int_0^{q_N}[\overline{F}_N(x) - \overline{F}_L(x)]dx\right\} + k \geq 0$. Third, because the objective function is decreasing in π_N, at an optimal solution, either (IC′) and (IR′) binds , i.e., $\pi_N = \max\{0, p\int_0^{q_N}[\overline{F}_N(x) - \overline{F}_L(x)]dx - k/(1-\lambda)\}$. Consequently, (OBJ′)-(IR′) reduces to an unconstrained optimization problem with only one variable:

$$\max_{q_N} \left\{p\int_0^{q_N}\overline{F}_N(x)dx - cq_N - \max\{0, [\Gamma(q_N) - k]/(1-\lambda)\}\right\}. \qquad \text{(OBJ}''\text{)}$$

It is straightforward to verify the solution q_N^* defined in Proposition 2 maximizes the objective in (OBJ″). ■

Proof of Proposition 4. Recall that for $k \in (\Gamma(q_L^o), \Gamma(q_N^o))$, the manufacturer's expected profit under the optimal menu that induces no forecasting $\mathcal{M}_N^r$ is strictly increasing in k, while the retailer's expected profit remains constant at zero. Hence, if $k^r < \Gamma(q_N^o)$, then both the manufacturer's and system's expected profits under the optimal menu increase in k for $k \in [k^r, \Gamma(q_N^o)]$. Therefore, it suffices to show that $k^r < \Gamma(q_N^o)$ holds when c is sufficiently small. Because $\mathcal{M}_N^r$ increases in k for $k \leq \Gamma(q_N^o)$ and equals to $p\int_0^{q_N^o}\overline{F}_N(x)dx - cq_N^o$ for $k \geq \Gamma(q_N^o)$, a sufficient condition to ensure

$k^r < \Gamma(q_N^o)$ is

$$\Gamma(q_N^o) > \overline{\Psi}(c),$$

where

$$\begin{aligned} \overline{\Psi}(c) \quad \equiv \quad & \lambda \left[p \int_0^{q_H^*} \overline{F}_H(x) dx - c q_H^* \right] + (1 - \lambda) \left[p \int_0^{q_L^*} \frac{\overline{F}_L(x) - \lambda \overline{F}_H(x)}{1 - \lambda} dx - c q_L^* \right] \\ & - \left[p \int_0^{q_N^o} \overline{F}_N(x) dx - c q_N^o \right]. \end{aligned}$$

Suppose that $ED_H > ED_L$. Because $\overline{\Psi}(c)$ and $\Gamma(q_N^o)$ are continuous in c, $\lim_{c \to 0} \Gamma(q_N^o) = \lambda(1 - \lambda) p E[D_H - D_L] > 0$, and $\lim_{c \to 0} \overline{\Psi}(c) \leq -\lambda p E[D_H - D_L] < 0$, there exists $\bar{c} > 0$ such that $k^r < \Gamma(q_N^o)$ for $c \leq \bar{c}$. This proves the first part of the proposition.

Now suppose $\lambda \overline{F}_H(x) \geq \overline{F}_L(x)$ for $x > 0$; this implies $q_L^* = 0$. The retailer's expected profit under the optimal menu that induces forecasting is (from (5)) $\lambda p \Delta(q_L^*) = 0$, while her expected profit under no forecasting is $\max(\Gamma(q_L^o) - k, 0)/(1 - \lambda)$. Thus, if $k^r < \Gamma(q_L^o)$, then (7) holds with $\widehat{k} = \Gamma(q_L^o)$. Thus, it is sufficient to show that $k^r < \Gamma(q_L^o)$ holds when c is sufficiently small. Because $\mathcal{M}_N^r = pE \min(q_L^o, D_N) - c q_L^o - [\Gamma(q_L^o) - k]/(1 - \lambda)$ for $k \leq \Gamma(q_L^o)$, if $k^r \leq \Gamma(q_L^o)$, then

$$k^r = \frac{1 - \lambda}{2 - \lambda} \underline{\Psi}(c),$$

where

$$\underline{\Psi}(c) \equiv \lambda \left[p \int_0^{q_H^*} \overline{F}_H(x) dx - c q_H^* \right] - \left[p \int_0^{q_L^o} \overline{F}_N(x) dx - c q_L^o - \Gamma(q_L^o)/(1 - \lambda) \right].$$

Hence, $k^r < \Gamma(q_L^o)$ if

$$\frac{2 - \lambda}{1 - \lambda} \Gamma(q_L^o) > \underline{\Psi}(c).$$

Because $\underline{\Psi}(c)$ and $\Gamma(q_L^o)$ are continuous in c, $\lim_{c \to 0} \Gamma(q_L^o) = \lambda(1 - \lambda) p E[D_H - D_L] > 0$, and $\lim_{c \to 0} \underline{\Psi}(c) = \lambda p E[D_H - D_L]$, there exists $\underline{c} > 0$ such that $k^r < \Gamma(q_L^o)$ for $c \leq \underline{c}$. ■

Proof of Proposition 6. Based on Proposition 5, it suffices to show that $\Gamma(q_H^o) \geq k^o$, which follows from

$$
\begin{aligned}
k^o &= \Pi_F(q_H^o, q_L^o) - \Pi_N(q_N^o) \text{ (by definition of } k^o\text{)} \\
&= \lambda[\Pi_F(q_H^o, q_L^o) - \Pi_N(q_N^o)] + (1-\lambda)[\Pi_F(q_H^o, q_L^o) - \Pi_N(q_N^o)] \\
&\leq \lambda[\Pi_F(q_H^o, q_L^o) - \Pi_N(q_H^o)] + (1-\lambda)[\Pi_F(q_H^o, q_L^o) - \Pi_N(q_L^o)] \\
&\quad \text{(by definition of } q_N^o\text{)} \\
&= \lambda(1-\lambda)p\int_{q_L^o}^{q_H^o}[\overline{F}_H(x) - \overline{F}_L(x)]dx \text{ (by definition of } \Pi_F \text{ and } \Pi_N\text{)} \\
&\leq \lambda(1-\lambda)p\int_0^{q_H^o}[\overline{F}_H(x) - \overline{F}_L(x)]dx \\
&= \Gamma(q_H^o).
\end{aligned}
$$

■