University of Pennsylvania
ScholarlyCommons

Operations, Information and Decisions Papers

Wharton Faculty Research

2009

Consumer Returns Policies and Supply Chain Performance

Xuanming Su
University of Pennsylvania

Follow this and additional works at: http://repository.upenn.edu/oid_papers

Part of the Business Administration, Management, and Operations Commons, Marketing Commons, and the Operations and Supply Chain Management Commons

Consumer Returns Policies and Supply Chain Performance

Abstract

This paper develops a model of consumer returns policies. In our model, consumers face *valuation uncertainty* and realize their valuations only after purchase. There is also *aggregate demand uncertainty*, captured using the conventional newsvendor model. In this environment, consumers decide whether to purchase and then whether to return the product, whereas the seller sets the price, quantity, and refund amount.

Using our model, we study the impact of full returns policies (e.g., using 100% money-back guarantees) and partial returns policies (e.g., when restocking fees are charged) on supply chain performance. Next, we demonstrate that consumer returns policies may distort incentives under common supply contracts (such as manufacturer buy-backs), and we propose strategies to coordinate the supply chain in the presence of consumer returns. Finally, we explore several extensions and demonstrate the robustness of our findings.

Disciplines

Business Administration, Management, and Operations | Marketing | Operations and Supply Chain Management

Comments

At the time of publication, author Xuanming Su was affiliated with the University of California. Currently (July 2016), he is a faculty member in the Operation, Information and Decisions Department of the Wharton School at the University of Pennsylvania.

Consumer Returns Policies and Supply Chain Performance

Xuanming Su

Haas School of Business, University of California, Berkeley, CA 94720

Abstract

This paper develops a model of consumer returns policies. In our model, consumers face *valuation uncertainty* and realize their valuations only after purchase. There is also *aggregate demand uncertainty*, captured using the conventional newsvendor model. In this environment, consumers decide whether to purchase and then whether to return the product, while the seller sets the price, quantity, and refund amount.

Using our model, we study the impact of full returns policies (e.g., using 100% money back guarantees) and partial returns policies (e.g., when restocking fees are charged) on supply chain performance. Next, we demonstrate that consumer returns policies may distort incentives under common supply contracts (such as manufacturer buy-backs), and we propose strategies to coordinate the supply chain in the presence of consumer returns. Finally, we explore several extensions and demonstrate the robustness of our findings.

August 2008

1 Introduction

Consumer returns policies are ubiquitous today. The value of products that U.S. consumers return to retailers exceeds $100 billion each year (see Stock, Speh, Shear, 2002). Many consumers have grown accustomed to being able to bring unwanted merchandise back to the store for just about any reason. In the consumer electronics industry, for example, most product returns were due to reasons such as the consumer not liking the product as much as anticipated, not understanding how to use it, or even regretting an impulsive purchase; in fact, only about 5% of consumer returns were truly defective (see Lawton, 2008). With increasing product proliferation, consumers may not know whether a particular item will fit their needs or match their tastes. Returns policies help to stimulate demand by protecting consumers against the possibility of product misfit. However, they may leave firms with excess inventory when items are returned. The proportion of purchases that are returned can range from 11-20% for consumer electronics up to 35% for high fashion apparel (see Guide et al., 2006); these figures are likely to be even higher for Internet and catalogue sales. From an operational standpoint, how should consumer returns policies be designed and what are some of their implications on managing the supply chain? This paper aims to address these issues.

While "no-questions-asked" 100% money-back guarantees are commonly seen, many firms have chosen to tighten these liberal returns policies (see Earnest and Uribarri, 2007, for example). When a consumer returns a product, there may be an explicit restocking fee, or there may be nonrefundable shipping and handling charges. In these cases, the cost of supply-demand mismatches is shared between consumers and firms: the former may incur losses due to product misfit, while the latter may be left with excess stock when sold items end up being returned. In this paper, to address both types of supply-demand mismatches, we develop an integrated model that simultaneously address consumers' purchase decisions and firms' operational decisions. Based on firms' operational policies, we characterize consumers' purchase as well as keep-or-return decisions. Similarly, based on consumers' willingness-to-pay (which are endogenously determined), we study firms' pricing and inventory decisions.

There are two main contributions in this paper. First, we demonstrate that consumer returns policies play an important role in resource allocation. Since consumers are aware of their valuations for the product only after purchase, establishing returns policies makes it possible to realize an ex-post efficient allocation of the product. When returns are not permitted, consumers who find the product unsuitable are required to keep it, even though there are better alternative uses for the product. With returns policies, consumers may return products that can then be channeled to

better uses. However, firms should not go overboard: we demonstrate that full returns (i.e., where consumers obtain a full refund of the selling price) are excessively generous and thus fail to optimize supply chain performance.

The second contribution is to study the supply chain implications of consumer returns policies. In particular, since buy-back contracts are a popular solution to the well-known double marginalization problem, we analyze the interaction between consumer returns and supplier buy-backs in a supply chain setting. We find that buy-back arrangements may induce retailers to implement excessively generous returns policies. In other words, buy-back contracts generate some incentive-distorting effects that have not been previously highlighted. To rectify this problem, we propose several solutions. One is to refine supplier buy-backs so that different credits are applied for unsold items and returned items; however, this requires the supplier to be able to tell them apart and imposes additional monitoring costs. Another approach is to implement direct-to-manufacturer consumer returns policies. Instead of turning to the retailer, consumers who wish to return the product now interact directly with the manufacturer. This may involve additional transaction costs since the retailer is usually in closer proximity to consumers. Yet another alternative is to use sales rebate contracts, under which the manufacturer rewards the retailer for sales. However, this approach requires the manufacturer to monitor sales. We find that rebates should generally be offered for total sales, regardless of whether those units are subsequently returned.

The remainder of this paper is organized as follows. Section 2 provides a literature review. Section 3 introduces the basic model. In Section 4, we analyze the equilibrium outcome under both partial and full returns policies. In Section 5, we study the implications of consumer returns policies on supply chain contracting. We explore some extensions in Section 6 and show that our main results remain unchanged. In Section 7, we provide concluding remarks and offer some directions for future work. All proofs are presented in the Appendix.

2   Literature Review

This paper is most closely related to the emerging area of research on consumer behavior in operations management. By explicitly studying consumer decision processes, this stream of research focuses on modeling demand patterns that depend on firms' operational decisions. For example, Aviv and Pazgal (2008), Elmaghraby, Gulcu and Keskinocak (2008), and Su (2007a) study dynamic pricing problems in which consumers make purchase decisions in anticipation future price changes. Another group of papers, such as Liu and van Ryzin (2008), Su and Zhang (2007a), Cachon and

Swinney (2007), Lai, Debo and Sycara (2007), and Jerath, Netessine, and Veeraraghavan (2007), analyze firms' stocking decisions when consumers anticipate and strategically wait for markdowns, which occur when there is excess inventory. Next, Dana and Petruzzi (2001) and Su and Zhang (2007b) examine scenarios in which the possibility of stock-outs may discourage customer patronage. In all these cases, consumer demand is endogenously determined and depends on the firm's pricing and inventory decisions. The above papers capture consumers' concern over price risk (e.g., reluctance to buy now as prices may become more attractive in the future) and availability risk (e.g., purchase intentions fueled by the threat of subsequent unavailability). In the current work, we focus on fit risk. That is, consumers are uncertain over their valuations, so there may be product misfit if valuations turn out to be low. In this case, consumer returns policies can be used to encourage purchases.

A key modeling ingredient in this paper is valuation uncertainty. When consumers face valuation uncertainty, firms may wish to offer advance purchase discounts to compensate them for bearing risk (see Dana, 1998, and Xie and Shugan, 2001). Further, Gallego and Sahin (2006) show that selling call options on capacity can improve revenues over advance selling. Another related paper is Alexandrov and Lariviere (2007), who study the value of offering reservations to consumers who face valuation uncertainty. Along similar lines, we study the role of consumer returns policies as an "insurance" mechanism when consumers are uncertain over their valuations. There are other papers that incorporate more detailed elements of valuation uncertainty: Yu, Kapuscinski, and Ahn (2005) consider correlated consumer valuations; Koenigsberg, Muller and Vicassim (2008) and Bhargava and Chen (2007) consider multiple consumer types; and Su (2007b) explores the relationship between valuation uncertainty and various psychological/behavioral decision biases. However, in the papers described above, consumer valuation uncertainty is studied in an environment with fixed capacity. In our current work, we also consider stocking decisions, so inventory, though limited, is variable. Other papers that study consumer valuation uncertainty together with variable capacity/inventory include Debo and van Ryzin (2007) and Swinney (2008). The former paper focuses on consumer learning about the underlying unknown quality (which affects all consumers in the same way), whereas the latter paper, like ours, admits ex-post heterogeneity after consumers learn their valuations.

There is a stream of literature on consumer returns policies. Davis, Gerstner, and Hagerty (1995) formulate a model in which consumer ex post valuation is a Bernoulli random variable (i.e. the product may either match or not match the buyer's taste). Che (1996) allows consumer

valuations to follow a general distribution and also incorporates risk aversion. Next, there are several papers that focus on how to prevent inappropriate returns from consumers with no intention of keeping their purchases (e.g., buying a video camera before a wedding, or buying a dress for a dance party). Hess, Chu, and Gerstner (1996) develop a model in the context of mail-order catalogs with non-refundable shipping and handling charges. Chu, Gerstner, and Hess (1998) study the distinction between "no-questions-asked" and "verifiable problems only" and find that the former is the optimal solution to handle consumer opportunism. In another paper, Davis, Hagerty, and Gerstner (1998) analyze the optimal level of "hassle" on returns policies. Yalabik, Petruzzi and Chhajed (2005) also study the optimal amount of logistics investment (which reduces the costs of returning) as well as the optimal marketing investment (which decreases the probability of returns). Our current work differs from these previous papers in two main ways. First, while most of the existing literature focus on full refunds, the bulk of our analysis is on partial refunds. In particular, we show that some nonrefundable charges are optimal even when there is no room for consumer opportunism. Second, we consider aggregate demand uncertainty, which has generally been overlooked in previous research on consumer returns. Our model incorporates both individual uncertainty (valuations) as well as aggregate uncertainty (market size).

Several other papers highlight different functions of consumer returns policies. When there are substantial transaction costs in consumer returns, Moorthy and Srinivasan (1995) show that generous returns policies help to signal high quality. In another paper, Heiman, McWilliams, Zhao, and Zilberman (2001) explore the analogy between money-back guarantees and financial put options. In an early paper, Heal (1977) shows that guarantees can be used to implement efficient risk-sharing between the seller and consumers. The spirit of our paper is similar, but we do not require risk aversion; in fact, we show that consumer returns policies are optimal even when all parties are risk-neutral.

There is a body of work on stochastic inventory models with product returns. This literature typically assumes that demand and returns follow Poisson processes, and studies inventory control in such settings. One common approach is to consider situations where items are returned directly to reusable inventory. Cohen, Nahmias and Pierskella (1980) study a model where a fixed fraction of demand is returned. Kelle and Silver (1989) extend this study to cover fixed ordering costs. Fleischmann, Kuik and Dekker (2002) analyze the infinite horizon system and derive the optimal control policy. While the above papers focus on a single location, Decroix, Song and Zipkin (2005) analyze the steady state behavior of a series system, and Decroix and Zipkin (2005) study an

assembly system. An alternative approach of incorporating product returns is to consider a separate remanufacturing facility that processes returned items before they are resold; see Simpson (1978). Inderfurth (1997) extends this analysis to the case of positive but identical lead-times for ordering and remanufacturing. Decroix (2006) studies this type of remanufacturing and inventory control in a series system. For a more comprehensive coverage of the literature on reverse logistics and remanufacturing, readers are referred to the reviews by Fleischmann et. al (1997) and Dekker et. al (2004). Notice that while this literature assumes an exogenous returns process, our approach in this paper is different as we consider consumer returns that depend on the firm's decision variables (such as prices and refunds).

Finally, our work is also related to the literature on buy-back contracts within supply chains. Although our focus is on consumer returns, we will also jointly consider cases where the retailer may return unsold units to the manufacturer (via a buy-back contract). Pasternack (1985) show that channel coordination is possible when the supplier offers to buy back all unsold units at partial credit. However, Emmons and Gilbert (1998) incorporate retailer pricing decisions and show that channel coordination using buy-back contracts may no longer be feasible, unless the retailer can commit to the selling price prior to the selling season; otherwise the retailer has an incentive to set prices too high (see Marvel and Peck, 1995, and Kandel, 1996). Donohue (2000) analyzes a model with forecast updating and shows how to coordinate the supply chain using buy-back contracts. Webster and Weng (2000) study a "risk-free" returns policy, in which both the retailer's and the manufacturer's profits are increased, compared to the situation in which no returns are allowed. Padmanabhan and Png (1997) analyze the interaction between manufacturer returns policies and retail competition. Tsay (2001) demonstrates that markdown allowances may be more appropriate when the retailer is able to secure a higher salvage value. Finally, we stress that all the papers reviewed above consider either supply chain returns policies or consumer returns policies. In contrast, our work integrates both types of returns within a unified framework.

3 Model

Our starting point is the classic newsvendor model. There is a seller who faces random market demand X with distribution F. We interpret market demand as a mass of infinitesimal consumers, each with individual valuations for the product. However, consumers face uncertainty in their own valuations. We assume that customer valuations V are identically and independently drawn from the distribution G. These realizations are not known until the customer purchases the product.

6

We denote $\overline{F} \equiv 1 - F$ and $\overline{G} \equiv 1 - G$. Note that this model captures both aggregate uncertainty (market size) as well as individual uncertainty (independent consumer valuations). Depending on the seller's returns policy, consumers may return the product for a refund if their valuations turn out to be low. The seller's per unit production cost is c and leftover units have salvage value s. We assume that $\mu > c > s$, where we use μ to denote consumers' expected valuation EV.

The seller's decisions include the price of the product p, the stocking quantity q, as well as the refund r to be paid if consumers choose to return the product. In our model, the seller reserves the right to retain a portion of the price paid originally. That is, when consumers return the product, the refund r may be less than the price paid p. When $r < p$, we say that the seller offers partial refunds (or partial returns). In practice, partial refunds have alternative reincarnations: they may appear as restocking fees, or may be disguised as non-refundable shipping charges. In contrast, when $r = p$, we say that the seller offers full refunds (or full returns). The seller's goal is to maximize expected profits.

Consumers, on the other hand, make two decisions sequentially. Initially, they decide whether to purchase the product. If so, they then decide whether to keep it after privately observing their own valuation. Consumers seek to maximize individual expected surplus. Their surplus may be one of the following three cases. If they purchase and keep the product, surplus is ex post valuation minus price paid $v - p$; if they purchase but return the product, surplus is the refund received net of price paid $r - p$; finally, if they do not purchase, or if the product is sold out, they receive zero surplus.

The chronology of events is as follows. First, the seller determines the price p, stocking quantity q, and refund r. Second, market demand X is realized and $\min(X, q)$ units are sold; if demand X exceeds stocking quantity q, consumers who do not obtain the product leave the market. Third, consumers who have bought the product observe their individual valuations and decide whether to keep it or to return it for the refund. Finally, the seller salvages all leftover units at the salvage value s. These leftover units include units that were unsold (which occurs only if demand $X < q$) as well as units that were sold but returned.

The key distinguishing feature of our model is that we consider partial returns. While existing work mainly restricts attention to full returns, we demonstrate below that partial returns generate a number of new results. To make our results more transparent, we let the other modeling elements be as simple as possible. In particular, we focus on a single-period, homogeneous-consumer setting, similar to most of the consumer returns literature. For example, Davis, Gerstner, Hagerty

(1995) assumes that there is a single selling opportunity (single period), and Che (1996) assumes that all consumers face the same uncertainty over their valuations (homogeneous-consumers). In this paper, we adopt Che (1996) as a model benchmark for framing the results and insights of our analysis.

4 Consumer Returns Policies

4.1 Full Returns

We begin our analysis with the case of full refund policies, in which the seller offers to completely refund consumers when the product does not fit their individual needs or tastes. This is essentially a 100% money-back-guarantee offered to ensure consumer satisfaction. In our model, this imposes the restriction that the refund must equal selling price, i.e. $r = p$. Note that earlier work by Che (1996) also restricts attention to full refunds.

With money back guarantees, consumers face no downside risk because they can always return the product and claim a full refund. Regardless of the selling price p, all consumers are willing to buy the product. The seller, on the other hand, needs to set the selling price p and the quantity q. His expected profit is

$$
\begin{aligned}
\Pi_{\text{full refund}}(p, q) &= \underbrace{p\overline{G}(p)E\min(X, q)}_{\text{sold}} + \underbrace{sG(p)E\min(X, q)}_{\text{returned}} \\
&\quad + \underbrace{s(q - E\min(X, q))}_{\text{not sold}} - cq \\
&= (p - s)\overline{G}(p)E\min(X, q) - (c - s)q.
\end{aligned}
$$

(1)

(2)

Here, with full refund $r = p$, it is clear that consumers with valuations $v \geq p$ will opt to keep the product whereas those with valuations $v < p$ will return it. In equation (1), the first term corresponds to revenues from the products that are sold and kept by the consumer, the second term results from salvaging units that are bought but returned by the consumer, the third term comes from unsold units that are salvaged, and the last term is the seller's procurement or production cost. Observe that under this formulation, the seller's pricing and quantity decisions are separable.

One of the key questions in the consumer returns literature is whether it is profitable to allow returns. Would the seller do better if all sales are final and products may not be returned? How should the seller choose between full money-back-guarantees and no-returns policies? We can use our model to compare operational performance under full returns and no returns. Observe that

when returns are not accepted, the highest price consumers are willing to pay is μ, their expected valuation for the product. In this case, the seller faces a standard newsvendor problem with $p = \mu$ and profit function

$$\Pi_{\text{no returns}}(q) = (\mu - s)E\min(X, q) - (c - s)q, \tag{3}$$

so the optimal critical fractile stocking quantity satisfies $\overline{F}(q^*) = \frac{c-s}{\mu-s}$. The following proposition characterizes the optimal full-refund policy (which maximizes (2)) and compares its performance with the optimal no-returns policy (which maximizes (3)).

Proposition 1 *(i) Suppose that the seller offers a full refund policy. Then, the seller's optimal price p^* and quantity q^* are characterized by*

$$p^* = \arg\max_p (p - s)\overline{G}(p), \tag{4}$$

$$\overline{F}(q^*) = \frac{c - s}{(p^* - s)\overline{G}(p^*)}. \tag{5}$$

(ii) When cost c is sufficiently low, or when salvage value s is sufficiently low, the seller prefers no-refund to full-refund.

This proposition reveals that while commonly observed in practice, full returns policies can be suboptimal. For instance, when the salvage value is low, the seller should not offer returns policies because the returned items have limited value. Similarly, when costs are low, the seller tends to have high inventory and thus high sales; if returns were allowed, the sheer volume of returns can quickly become unprofitable. These observations serve to caution retailers against blindly offering full returns policies. Previous research has also identified some of these conditions under which full refund policies should be offered. In particular, Che (1996) shows that money back guarantees are desirable when consumers are highly risk averse or when retail costs are high. In another study, Davis, Gerstner, Hagerty (1995) find that the seller prefers to offer money back guarantees when transaction costs (of returning the product) are low or when salvage values are high. However, these earlier results were obtained under simpler models (e.g., Che (1996) assumes deterministic market size, while Davis, Gerstner, Hagerty (1995) assumes Bernoulli distributions for consumer valuations). Therefore, our analysis shows that these insights are robust and continue to hold under more general settings.

Full refunds have an interesting informational property. When consumers face valuation uncertainty, allowing full returns is similar to providing information. To see this, consider two seemingly distinct scenarios: one in which the seller implements full returns policies, and another

9

in which consumers have complete information about their valuations before purchase. In both cases, only individuals with valuations exceeding the price will eventually consume the product. Therefore, by allowing full returns, the seller is essentially endowing consumers with information about their valuations. In this sense, our analysis applies more generally beyond returns policies. For example, the comparison between no-refunds and full-refunds is similar to a comparison between selling early and selling late (when consumers are fully aware of their valuations); see, for example, Courty (2003).

4.2 Partial Returns

In this section, we extend the analysis to the case of partial returns. Here, we allow the refund amount r to be strictly less than the selling price p, and we may interpret the difference as a restocking fee or a nonrefundable charge that the seller imposes. Mathematically, this generalization introduces a free decision variable. It turns out that this new dimension changes the prior analysis dramatically. Specifically, earlier research (as well as our analysis in the previous section) has focused on understanding when returns policies (full refunds) should be offered, and when returns should not be accepted. However, it turns out that once we consider partial refunds, they are always optimal and thus dominate the no-returns scenario. Using our model, we shall highlight the operational benefits of partial returns.

We first consider the consumer's decisions under this allocation mechanism. Recall that the consumer's decision tree involves two sequential decisions: (i) buy or not, and then, if the consumer buys, (ii) keep or return (after observing his valuation). Clearly, at stage (ii), the consumer will choose to keep the product if his valuation is at least as high as the refund ($V \geq r$) and to collect the refund otherwise ($V < r$). Therefore, in stage (i), should the consumer choose to buy, his expected utility will be $E \max(V, r)$. In other words, the consumer's reservation price (i.e., the highest price he is willing to pay for the product) is $E \max(V, r)$. There are two cases to consider: when price p exceeds $E \max(V, r)$, no consumer will want to make a purchase, but when the price is less than or equal to $E \max(V, r)$, every consumer will want to buy. From the seller's perspective, demand is thus

$$D = \begin{cases} X, & \text{if } E \max(V, r) \geq p, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

To induce consumers to buy, the highest price that the seller may charge is $p = E \max(V, r)$. Notice that this price p must exceed the refund r in order for the seller to extract the surplus of consumers who keep the product (with valuations $V \geq r$). Therefore, the monopolist seller will offer partial

10

refunds. Put in another way, when the seller offers full refunds with $r = p$, he is leaving some surplus for the consumers.

There are two important implications of our consumer decision model. First, since consumers consider the returns policy when making purchases, a higher refund r increases willingness-to-pay and allows the seller to charge a higher price $p = E \max(V, r)$. Anderson, Hansen, Simester (2008) provides empirical evidence; for example, they estimate that the option to return women's footwear is worth more than \$15 per purchase. Second, since the returns policy directly determines consumers' keep-or-return decisions, a higher refund r also induces higher return rates $G(r)$; this result was observed in the experimental study by Wood (2001). Together, these two effects sum up the costs and benefits of offering a generous returns policy.

Next, we study the seller's problem. There are three decisions: price p, quantity q, and refund r. The seller's profit function is

$$
\begin{aligned}
\Pi(p, q, r) &= \underbrace{p\overline{G}(r)E\min(X, q)}_{\text{sold}} + \underbrace{(p - r + s)G(r)E\min(X, q)}_{\text{returned}} \\
&\quad + \underbrace{s(q - E\min(X, q))}_{\text{not sold}} - cq \qquad (7) \\
&= [(p - s)\overline{G}(r) + (p - r)G(r)]E\min(X, q) - (c - s)q. \qquad (8)
\end{aligned}
$$

In equation (7), we see that each unit that is sold and kept by the consumer yields revenue p, each returned unit yields $p - r$ from the consumer and s from salvaging it, and each unsold unit yields only the salvage value s. The last term is the seller's procurement or production cost. Observe that under this formulation, the seller's pricing and quantity decisions are separable. The next proposition characterizes the seller's optimal decisions.

Proposition 2 *The seller's optimal price p^*, quantity q^*, and refund r^* are given by*

$$
\begin{aligned}
p^* &= E\max(V, s), \qquad &(9) \\
\overline{F}(q^*) &= \frac{c - s}{p^* - s} = \frac{c - s}{E\max(V, s) - s}, \qquad &(10) \\
r^* &= s. \qquad &(11)
\end{aligned}
$$

There is an intuitive interpretation of this result. First, note that the monopoly seller will extract maximum consumer surplus by charging $p^*(r) = E\max(V, r)$. Next, the optimal quantity follows from a simple newsvendor calculation. What is most interesting is the optimal refund amount: the seller chooses $r^* = s$. Recall that the value of r determines whether consumers decide to keep or return the product. When the refund $r^* = s$, consumers with valuation above the salvage

11

value keep the product while those with lower valuations return it to the seller to be salvaged at s. In our model, the product has two alternative uses (i.e., retained by consumers for valuation V or salvaged at s). Therefore, setting the refund $r^* = s$ implements an efficient ex post allocation of the good. Note that the salvage value s is an important parameter because when $s = 0$, our model reduces to the standard newsvendor model without consumer returns ($r^* = 0$). We also stress that the keep-or-return decisions are made voluntarily by individual consumers. This analysis suggests that returns policies play an important role as an allocation mechanism by creating the correct incentives that help realize ex post efficiency.

The main insight from our analysis is that consumer returns policies provide a useful operational instrument. In particular, the refund amount (equivalently, the restocking fee) can be an important strategic decision. Yet this decision variable has largely been overlooked since many firms focus on full returns, essentially by default. From Proposition 2, we see that the optimal policy involves partial refunds ($r^* < p$), and that the optimal refund equals the salvage value ($r^* = s$). However, we caution that the latter mathematical result should not be taken too literally because it holds only within the boundaries of our model. In contrast, the finding that partial refunds are optimal is robust. Several major retailers have begun to tighten their liberal returns policies, imposing restocking fees and stringent eligibility guidelines for product returns. Intuitively, full refunds tend to offer "too much protection" to consumers, transferring the potential downside of excess inventory and product misfit entirely to the seller. Our analysis instead recommends partial refunds.

Apart from increasing expected profits, partial refunds also induce the seller to stock larger quantities. Despite the apparent stringency of partial returns policies, they lead to higher service levels. This can generate benefits beyond the profit dimension. The next result follows easily from Proposition 2.

Corollary 1 *The seller's optimal profit and optimal quantity is higher under partial refunds than under full refunds.*

This result also has implications on upstream members of the supply chain. For concreteness, consider an exogenous wholesale price contract between a manufacturer and a retailer. In this case, observe that the manufacturer's profits are higher under partial returns compared to full returns. This is because manufacturer profits are directly proportional to the quantity that the retailer orders. Since partial returns policies induce the retailer to place higher orders (as Corollary 1 shows), the manufacturer is better off. This observation suggests that with regards to the choice

between full and partial returns, the incentives of the manufacturer and the retailer are aligned. Both parties are similarly motivated to introduce restocking fees that penalize consumer returns.

We are also interested in the welfare implications of consumer returns policies. In a conventional setting where individuals know their valuations beforehand, the price mechanism allocates the product to consumers with sufficiently high valuations. In our model, with valuation uncertainty, returns policies also serve to allocate the product to consumers with high valuations. However, there is a fundamental difference. With valuation uncertainty, consumer heterogeneity and thus allocation occurs only ex post, whereas in conventional settings, consumer heterogeneity is already present right from the start. It turns out that ex ante and ex post heterogeneity lead to starkly different welfare implications. We define social welfare as the sum of consumer surplus and the seller's profit, which is given by

$$
\begin{aligned}
SW(p,q,r) &= \underbrace{\int_r^\infty vg(v)dv \cdot E\min(X,q)}_{\text{sold}} + \underbrace{sG(r)E\min(X,q)}_{\text{returned}} \\
&\quad + \underbrace{s(q - E\min(X,q))}_{\text{not sold}} - cq \\
&= \int_r^\infty (v-s)g(v)dv \cdot E\min(X,q) - (c-s)q.
\end{aligned}
$$

(12)

(13)

With this definition, the next result follows from showing that the welfare-maximizing policy coincides with the profit-maximizing policy from Proposition 2.

Corollary 2 *The seller's optimal price p^*, quantity q^*, and refund r^* maximize social welfare.*

This result shows that social efficiency can be realized when there is ex post heterogeneity. It is well-known that under ex ante heterogeneity, monopoly pricing generates welfare losses. This disparity suggests that the inefficiency of monopoly stems from ex ante rather than ex post heterogeneity.

To conclude this section, we point out a unique feature of partial returns. Under partial returns, both the seller and consumers may incur losses at the end of the day. On one hand, the seller faces a more complex inventory management problem and may be left with excess inventory when consumers return unwanted products. On the other, consumers with valuations $V \in (r^*, p^*)$ will have to keep the product at a loss. Therefore, our analysis suggests that social efficiency and profit maximization require consumers and the seller to jointly shoulder the burden of potential demand-supply mismatches.

5 Supply Chain Contracts

Now, we turn to the supply chain setting. There are two firms, a manufacturer and a retailer. The manufacturer supplies units to the retailer, who sells to end-consumers and makes all related operational decisions. Under an exogenous wholesale price contract (i.e., the manufacturer sells to the retailer at per unit wholesale price w), we adopt a similar model as the previous section except that the seller's production cost c is now re-interpreted as the retailer's procurement cost w. In this section, we study the impact of consumer returns on the performance of several supply chain contractual mechanisms. Specifically, we are interested in achieving supply chain coordination; that is, we hope to design mechanisms under which the retailer's decentralized decisions will attain maximum profit for the supply chain.

5.1 Buy-back Contracts

One common alternative to handle consumer returns is to send them back to the manufacturer. Here, we consider the situation where there is a buy-back contract between the manufacturer and retailer, on top of the returns policy that the retailer offers to consumers. It is useful to consider both types of returns within a single model framework.

In our setting, the manufacturer sells to the retailer at wholesale price w, but agrees to buy back left over units at the buy-back price b. These contract parameters are exogenously given. The buy-back clause serves as an explicit mechanism through which the retailer may return leftover units to the manufacturer. It is well-known that appropriately chosen parameters w, b can coordinate the supply chain in the standard setting without consumer returns. In other words, buy-back contracts align the retailer's incentives in the interests of the supply chain and generate maximum total profits. Further, under the buy-back contract, these profits can be arbitrarily allocated between supply chain members. Unfortunately, this is no longer the case when consumers face uncertainty in their valuations before trying out the product.

To see this, observe that with manufacturer buy-backs, the retailer's profit function is

$$\Pi_R(p, q, r) = [(p - b)\overline{G}(r) + (p - r)G(r)]E\min(X, q) - (w - b)q. \tag{14}$$

In contrast, the total supply chain profit function, as a function of the retailer's decisions, is

$$\Pi_T(p, q, r) = [(p - s)\overline{G}(r) + (p - r)G(r)]E\min(X, q) - (c - s)q. \tag{15}$$

Following the same logic as in Proposition 2, it can be shown that the supply chain optimal refund amount (which maximizes Π_T) is $r^* = s$ as it ensures that items are allocated efficiently

ex post. However, the retailer's optimal refund amount (which maximizes Π_R) is $r^* = b$ and it is determined by the buy-back price. This is not surprising because while the salvage value is s from the supply chain's perspective, it is b from the perspective of the retailer who sends excess units back to the manufacturer. Therefore, when the manufacturer offers attractive buy-back terms, the retailer is likely to offer excessively generous returns policies; this in turn implies that the refund and price both become inefficiently high, leading to inefficient ex post allocations of the product to a selective group of high-valuation consumers. This is precisely what happens when supply chains attempt to design buy-backs to achieve coordination, which requires $b > s$ (see Cachon, 2003). In conventional models without consumer valuation uncertainty, buy-backs mitigate the retailer's risk of having excessive inventory, induce higher order quantities, and combats the double marginalization problem. However, in the present situation, buy-back contracts can distort the retailer's price and refund decisions. In fact, our analysis suggests that unless $b = s$ (which is unlikely), buy-back contracts will fail to maximize supply chain efficiency.

In our model, we stress that supply chain coordination is attained only when two conditions are satisfied. First, we require the decentralized retailer to stock the supply-chain-optimal quantity. That is, we need to eradicate the familiar double marginalization problem. But this alone is not sufficient. The second condition is that the decentralized retailer must implement the supply-chain-optimal returns policy. Only then can we achieve an ex-post efficient allocation (from the system wide perspective). For example, under the class of buy-backs considered above, it is possible to choose contract parameters to achieve quantity alignment, but the decentralized retailer will offer excessive refunds, i.e. $r^* = b > s$. As a result, the total supply chain profit (15) is not maximized.

Our result is consistent with the general observation in Cachon (2003) that, while contracts with two parameters (such as the simple buy-back contract with wholesale price w and buy-back price b) are sufficient to coordinate the supply chain along a single dimension (quantity q), additional free parameters are required when more actions are taken into account. For example, there are supply chain models that consider, on top of stocking quantities, pricing or sales effort. Similarly, in our case, we consider consumer returns policies. Our general line of thought suggests that in order to achieve coordination in these cases, we need to refine supply contracts and include additional free contract parameters.

5.2 Differentiated Buy-backs

One possible approach to refine buy-back contracts to achieve coordination is to distinguish between *unsold* units and *returned* units. We propose using so-called "differentiated" buy-back contracts. This contract has three parameters $w, b,$ and l. Under this contract, the manufacturer sells to the retailer at the wholesale price w but agrees to buy back all remaining units. The retailer receives a credit of b for each new (unsold) unit, but for returned items, the credit is reduced by l per unit. In other words, the manufacturer buys back unsold units at price b and buys back returned units at price $b - l$. Under this contract, it can be shown that the retailer's optimal refund amount is $r^* = b - l$ (details omitted as they are similar to Proposition 2), so the "penalty" parameter l can be chosen such that $b - l = s$ to induce supply chain optimal outcomes. In other words, the buy-back credit should be s for returned items, and the buy-back price for new items b can then be independently chosen to coordinate the supply chain. This is made precise in the next proposition. Let $\overline{\phi} \equiv 1 - \phi$.

Proposition 3 *The differentiated buy-back contract with the following parameters w, b, l induces the retailer to set price, quantity, and refund to maximize supply chain profit and also allocates a fraction $\phi \in [0, 1]$ of the total profits to the retailer.*

$$w = \phi c + \overline{\phi} E \max(V, s) \tag{16}$$

$$b = \phi s + \overline{\phi} E \max(V, s) \tag{17}$$

$$b - l = s \tag{18}$$

Notice that the expressions for w and b above resemble the coordinating parameters of the standard buy-back contract with no valuation uncertainty (see, e.g, Cachon, 2003). This suggests that as long as returned items are distinguished and we set $b - l = s$, the active incentive mechanisms of the buy-back contract remains largely unchanged.

We stress that the differentiated buy-back contract is effective to the extent that the manufacturer is able to differentiate between unsold and returned units. If the retailer is able to claim the full buy-back rate b for returned units, then the arrangement above will not work. This may the case in fashion retailing, since a returned item remains very much similar to unsold items. However, for packaged goods, returned (opened) items can be distinguished from unsold items, and differentiated buy-back contracts may be used.

5.3 Direct-to-Manufacturer Returns

Alternatively, the standard buy-back contract of Section 5.1 can be combined with having consumers return the product directly to the manufacturer (rather than to the retailer). In a related but different context of remanufacturing used products into new ones, Savaskan, Bhattacharya and Van Wassenhove (2004) consider an arrangement in which the manufacturer (rather than the retailer) is responsible for collecting used products. Here we adopt a similar model. Suppose that consumers buy the product from the retailer at price p, and they may return the product directly to the manufacturer for a refund of r (from the manufacturer). We assume that the refund amount r is chosen by the manufacturer, while the price p and quantity q are chosen by the retailer. As in Section 5.1, there is a simple buy-back contract with wholesale price w and buy-back price b. Given these parameters, we seek the Nash equilibrium in the game between the retailer and the manufacturer. Their profit functions are given respectively by:

$$\Pi_R(p,q,r) = (p-b)E\min(X,q) - (w-b)q, \tag{19}$$

$$\Pi_M(p,q,r) = \underbrace{w\overline{G}(r)E\min(X,q)}_{\text{sold}} + \underbrace{(w-r+s)G(r)E\min(X,q)}_{\text{returned}}$$
$$+ \underbrace{(w-b+s)(q-E\min(X,q))}_{\text{not sold}} - cq \tag{20}$$
$$= [(b-s)\overline{G}(r) + (b-r)G(r)]E\min(X,q) - [(c-s)-(w-b)]q. \tag{21}$$

In (19), the retailer faces a standard newsvendor profit function since he no longer handles returns from consumers. Next, (20) consists of the following cash flows for the manufacturer: each sold unit yields revenue w; each returned unit generates wholesale revenue w, salvage revenue s and refund r to the consumer; each unsold unit accounts for net revenue $w-b$ from the retailer and salvage revenue s. These terms sum up to the manufacturer profit function in (21). The total supply chain profit function (i.e. sum of retailer and manufacturer profits) is the same expression as (15) given above.

The next proposition demonstrates that in equilibrium, supply chain coordination can be achieved using direct-to-manufacturer returns, combined with a buy-back contract.

Proposition 4 *Suppose that consumers return the product directly to the manufacturer. The buy-back contract given below induces the retailer and the manufacturer to set price, quantity, and refund to maximize supply chain profit and also allocates a fraction $\phi \in [0,1]$ of the total profits to the retailer.*

$$w = \phi c + \overline{\phi} E\max(V,s) \tag{22}$$

$$b = \phi s + \overline{\phi} E \max(V, s) \qquad (23)$$

However, this arrangement suffers from a similar drawback as before: it is not viable if the manufacturer is unable to distinguish "fresh" items from returned items. To see this, notice that the manufacturer's optimal refund to consumers is $r^* = s$ while the buy-back price is $b > s$. This suggests that the retailer has an incentive to "intercept" the direct-to-manufacturer returns by offering consumers some refund amount $r' \in (s, b)$, which exceeds the manufacturer's refund $r^* = s$, and then claiming full credit $b > r'$ on these items. Unless the manufacturer is able to identify the returned items, such strategic behavior on the part of the retailer may limit the usefulness of direct-to-manufacturer returns.

5.4 Sales Rebates

Fortunately, it is possible to coordinate the supply chain even when returned items and new items are indistinguishable. However, we need to impose another condition: the manufacturer must be able to monitor the retailer's sales. Then, we show that using sales rebates is one possible way to coordinate the supply chain.

Under the sales rebate contract, the manufacturer sells to the retailer at wholesale price w but offers a rebate of u per unit for all units that are sold (whether or not they are eventually returned). However, the retailer takes full responsibility for unsold or returned units; this eliminates the incentive-distorting effects of buy-backs discussed earlier. Under this sales rebate contract, the retailer's profit function is

$$
\begin{aligned}
\Pi_R(p, q, r) &= \underbrace{(p + u)\overline{G}(r)E\min(X, q)}_{\text{sold}} + \underbrace{(p + u - r + s)G(r)E\min(X, q)}_{\text{returned}} \\
&\quad + \underbrace{s(q - E\min(X, q))}_{\text{not sold}} - wq \qquad (24) \\
&= [(p - s + u)\overline{G}(r) + (p - r + u)G(r)]E\min(X, q) - (w - s)q. \qquad (25)
\end{aligned}
$$

We stress that the rebate u is earned for both sold units as well as returned units. The total supply chain profit function, as before, is given by (15).

The next result states that there is a sales rebate contract that coordinates the supply chain.

Proposition 5 *The sales rebate contract with parameters satisfying the condition below induces the retailer to set price, quantity, and refund that maximize supply chain profit.*

$$u = \frac{E\max(V, s) - s}{c - s} \cdot (w - c) \qquad (26)$$

Using the condition given above, the sales rebate can be chosen to attain supply chain optimal profits. Further, the wholesale price w parameter can be adjusted to allocate profits between the manufacturer and the retailer; for example, setting $w = c$ gives the entire profit share to the retailer, while increasing w allocates a larger profit share to the manufacturer. Another possible approach to manipulate the profit allocation is to institute a sales target. That is, the rebate u is paid only for units sold in excess of the sales target; see Taylor (2002). Yet another approach, of course, is to use fixed transfer payments. Under all these approaches, we find that the sales rebate is a powerful contractual mechanism that can attain supply chain coordination without requiring any party to distinguish between new and returned units. Furthermore, since the rebate is paid for all sold units, all the information that is required is point-of-sales data that electronically record the total sales volume.

We close with a final remark about sales rebates. Suppose that the sales rebate u is given only for units that are sold eventually (i.e., sold and retained by the consumer). Using similar logic as before, observe that under this contract, the retailer's optimal refund is $r^* = s - u < s$. Again, the retailer's optimal refund amount is distorted away from the supply chain optimal choice of $r^* = s$, but this time the retailer's chosen refund is too low. This is because the retailer now has an increased incentive to discourage returns as each unit retained by the consumer earns an additional amount u. This example stresses the fact that for supply chain coordination, the sales rebate must be paid for all sold units without consideration of subsequent returns.

5.5 Summary

The main insights from this section is that supply chain coordination in the presence of consumer returns policies creates complications that can be resolved only with increased monitoring by the manufacturer. Although we have examined various types of supply contracts, the key behind these strategies lies in their informational requirements. Specifically, the manufacturer is required to possess "special" monitoring abilities. We considered two such possibilities. The first possibility is to require the manufacturer to distinguish between new and returned units. Given this ability, we have seen that differentiated buy-backs and direct-to-manufacturer returns can help coordinate the supply chain. The second possibility is to require the manufacturer to monitor sales volume. In this case, sales rebates can coordinate the supply chain. We stress that it is these monitoring abilities that are essential, rather than the structure of the contracts analyzed above.

6 Extensions

6.1 Heterogeneous consumers

Our analysis has assumed that consumers are homogeneous ex ante. Although they have different valuations ex post, these valuations follow the same distribution G. Put differently, all consumers have identical probabilistic assessments of their valuations ex ante. However, in reality, different consumers may have different pre-purchase estimates of how much they will eventually value the product. In this extension, we consider multiple "types" of customers with different valuation distributions. For simplicity, we sketch the basic ideas for the case of two consumer types, $\theta = L, H$, with proportions π_θ (so $\pi_L + \pi_H = 1$). Low-types ($\theta = L$) have valuations V_L distributed according to $G_L(\cdot)$ and high-types ($\theta = H$) have valuations V_H distributed according to $G_H(\cdot)$. We assume that $G_H(\cdot)$ stochastically dominates $G_L(\cdot)$. Intuitively, while all consumers are uncertain of their valuations at the outset, high-types can expect to have higher valuations. With this new two-type model, how does our analysis change?

There are two cases to consider. First, the seller may find it profitable to sell only to high-types. Then, this situation reduces to our single-type model as before. Second, the seller may wish to sell to both types. In this case, the seller should charge no more than $p^*(r) = E \max(V_L, r)$, so that low-types (and thus high-types) are willing to buy, and furthermore, the proportion of consumers who eventually return the product is $\pi_L G_L(r) + \pi_H G_H(r)$. Compared to our basic model, the main difference here is that it is optimal for the seller to offer more generous refunds. Specifically, it can be shown that the optimal refund $r^* > s$ is higher than that obtained in our basic single-type model (where we showed that $r^* = s$). The intuition proceeds as follows. Similar to our basic model, increasing the refund r has two opposing effects on profits: it allows the seller to charge a higher price, but at the same time, it induces higher return rates. However, while the price is determined by the low type, return rates are determined by the entire population (and are thus on average lower than that of the low type). This implies that the cost of offering a higher refund (in the form of increased returns) is lower when there are heterogeneous consumers. Hence, the seller can afford to offer more generous refunds.

More generally, the discussion above suggests that an important driver of optimal refund policies is the divergence between *marginal* valuations and *average* valuations. In markets where average valuations (among all those who buy) are substantially larger than the marginal valuation (of the most reluctant buyer), optimal refund amounts tend to be high. Within such environments,

heterogeneity may explain the generous refund policies commonly observed in practice.

Finally, we also find that our results on supply chain contracting continue to hold when there are heterogeneous consumers. In particular, the contracts described in Sections 5.2 to 5.4 can be used to coordinate the supply chain. However, as we elaborate below, the coordinating contracts are more complex when there are multiple consumer types. This is not surprising because we now need to additionally ensure that the retailer chooses to target the "correct" group of consumers. We find that the same contracts discussed earlier can be used to coordinate the supply chain, as long as their parameters are contingent on the chosen price and refund. A similar approach has been adopted in Bernstein and Federgruen (2005), who show that price-contingent buy-backs can coordinate the price-setting newsvendor, but they do not consider consumer returns. The details of our analysis are presented in Appendix B.

6.2   Processing costs

Product returns are often associated with processing costs. For example, there may be inspection, re-packaging and shipping costs. Furthermore, for products with short life cycles, returns can lose up to 30% of their value due to processing delays (see Guide et. al, 2006). There is an extensive body of research that studies the cost structure of reverse product flows. Tibben-Lembke and Rogers (2002) highlight major differences between forward and reverse logistics, and Andel and Aichlmayr (2002) estimate that product returns may cost as much as three to four times more than outbound shipments.

In our basic model, we have assumed that returned units and unsold units are identical and can be salvaged at the same value. Here, we describe an extension that considers the processing costs of returns. Suppose that the seller incurs marginal cost y on each returned unit before salvaging it at s. Then, the residual value of each returned unit is effectively $s - y$. (Note that the salvage value of unsold units remain at s as before.) In this case, following the same logic as above, it is easily shown that the optimal refund amount that maximizes both seller profits and social welfare is $r = s - y$. In other words, setting the refund $r = s - y$ is both supply chain optimal as well as socially optimal. This implies that our earlier results still hold, as long as we consider the "effective residual value" of returned products $s - y$ in place of the salvage value s in our analysis above. The logistical or processing cost (represented by y) is thus an important parameter. As processing costs decrease, the effective salvage value of returned units will increase. Blackburn et al. (2004) present several strategies for controlling costs in reverse supply chains.

Efficient reverse logistic processes, through cost reduction, increase the value of product returns and thus increase the refund that sellers can afford to give. It then follows from our model that the seller can charge a higher price in the first place. As profit margins increase, service levels and overall profits also increase. This discussion suggests that as reverse product flows in the supply chain become less costly, the corresponding forward flows also become more profitable.

6.3   Different salvage capabilities

In our description of the supply chain setting, we assume that the manufacturer and the retailer possess identical salvaging capabilities. However, in reality, they may face different salvage values. For instance, individual members of the supply chain may have access to different salvage markets; the manufacturer may turn to discounters while the retailer may simply mark down unsold items. There may also be transportation costs when left-over items are shipped from one location to another. Therefore, in a more comprehensive model, we would consider different salvage values s_M and s_R for the manufacturer and the retailer. With differing salvage capabilities, supply chain coordination requires not only choosing the optimal policy parameters (i.e. price, quantity, and refund), but also salvaging left-over units by the more efficient party. This point has been addressed by Tsay (2001) and similar methods can be employed here. Essentially, from a supply chain perspective, the salvage value is $s = \max(s_M, s_R)$. As we briefly describe below, our analysis can be adapted and the key insights remain unchanged.

When the manufacturer is more efficient ($s = s_M \geq s_R$), the desired outcome requires the manufacturer to take back and salvage leftover units. This can be implemented using the buy-back arrangements discussed above. Notice that in the efficient buy-back contracts described above, we always have $b \geq s \geq s_R$. This shows that under an efficient buy-back contract, the retailer will always prefer to send remaining units back to the manufacturer rather than to salvage them independently.

When the retailer is more efficient ($s = s_R \geq s_M$), the retailer should be the party who salvages remaining units. In this case, markdown allowances can be used: the retailer receives credit for unsold items that have to be marked down or salvaged. Here, the manufacturer does not regain physical possession of those units and merely compensates the retailer during low demand realizations. In fact, we may view markdown allowances as buy-backs. This is because from the retailer's point of view, a buy-back price b is equivalent to a markdown allowance m as long as $m = b - s_R$. Therefore, with this expression, we may apply the results above to obtain efficient

markdown allowances in situations where the salvaging is undertaken by the retailer. However, we caution that implementing markdown allowances involves substantial informational and monitoring costs. These issues become more difficult to deal with when we consider more complex contracts, such as differentiated markdown allowances for unsold and returned units (similar to differentiated buy-backs).

6.4 Consumer travel cost

So far we have assumed that consumers do not incur any cost when returning products. In reality, consumers may have to make an extra trip to the store, and this takes up time, effort, and even monetary resources. Suppose that consumers who return the product incur cost h. More generally, we may think of h as a "hassle" cost, as suggested by Davis, Hagerty, and Gerstner (1998). This implies that when the potential refund is r, only consumers with valuations less than $r - h$ will return the product, and the ex-ante willingness to pay is $E\max(V, r - h)$.

In this case, should the optimal refund r^* take consumer travel cost into account? It turns out that our previous result $r^* = s$ remains unchanged. It can be shown that the first-best, seller-optimal, and supply-chain optimal refund amounts all equal the salvage value s as before. This is because the role of the optimal refund is to induce consumers to voluntarily generate an efficient ex-post allocation of the product, after observing their own valuations. When there is a travel cost h, the ex-post efficient allocation is for consumers to retain the product if and only if their valuation exceeds $s - h$. This is made possible when the refund equals s as before. Therefore, from the perspective of the seller, whether in a centralized or decentralized scenario, the consumer's travel cost h is irrelevant when setting the optimal refund amount $r^* = s$. Since this key result is unchanged, the strategies recommended above can still be employed to coordinate the supply chain.

6.5 Forgetful consumers

In our analysis, we assume that upon observing their valuations, consumers rationally choose between keeping or returning the product. When the refund amount is r, consumers with valuations below r will return it. Anecdotal evidence (and personal experience) may suggest that consumers who plan on returning a product sometimes end up not doing so. For instance, the consumer may have missed the return deadline, or may simply have forgotten about it. Here, we model consumer bounded rationality in the following way (see Su, 2008, for another application of bounded

rationality in a newsvendor context). Suppose that when the consumer wishes to return the product, there is only a probability ϕ that it will be carried out. How will our earlier results change?

Observe that in this case, consumers' ex-ante willingness to pay is reduced to $p^*(r) = \int_r^\infty vg(v)dv + \int_0^r (\phi r + \overline{\phi} v)g(v)dv$. Even though the potential refund is r, there is a chance that the consumer will fail to claim it, so the expected valuation from the purchase is less than $E\max(V, r)$. Therefore, given refund r, the seller will charge the price $p^*(r)$ given here.

After some calculations, we can then show that the seller-optimal and first-best values of r are once again $r^* = s$. When consumers are boundedly rational in the above sense, this key result remains unchanged and our results above apply similarly. Here, apart from lowering the selling price, the main effect of bounded rationality is to reduce the supply chain profit and social welfare in equilibrium. This is inevitable because consumers can no longer be relied upon to realize the ex-post efficient allocation after observing their valuations.

There is a subtle distinction that should be highlighted. Here, we have assumed that consumers are aware of their own boundedly rational tendencies and adjust their willingness to pay accordingly. If they were not as sophisticated, they may naively believe that they will manage to return the product as planned. In other words, consumers may believe that they will return the product for sure, even though they will manage to do so only with probability ϕ. The result is that consumers are now willing to pay $p^*(r) = E\max(V, r)$ for the product as before. In this case, it can be shown that while the first-best refund remains at $r^{FB} = s$, the seller's optimal refund is increased to $r^* > s$. This is because the seller has an incentive to exploit consumer naivete by offering generous returns policies. Since intended returns may not eventually occur, such policies generate benefits to the seller without the corresponding costs.

7 Conclusion

In this paper, we have examined the situation in which consumers face uncertainty in their valuation for products. This uncertainty is resolved only after purchase. With consumer returns policies, the consumer can then decide whether to keep or return the product. In this environment, the seller makes price and quantity decisions and also designs an appropriate returns policy. Using this model, we put forth two main findings.

- Full refunds are too generous and do not optimize supply chain performance. The refund amount is a useful operational instrument that has largely been overlooked. In general,

the optimal refund is less than the selling price (i.e., partial refunds are optimal). This arrangement helps to realize an efficient ex-post allocation of the product.

- Supplier buy-backs may induce retailers to adopt inefficient consumer returns policies. We propose several alternatives that can coordinate the supply chain in the presence of consumer returns: (i) use a different buy-back rate for new and returned items, (ii) have consumers return products directly to the manufacturer, and (iii) offer sales rebates to the retailer. However, these suggestions require the manufacturer to have additional monitoring capabilities.

This analysis can be extended in several directions. First, many stores offer a wide selection of different products, and consumers who return one may elect to buy another. In fact, when consumers return unwanted products, some stores offer store credit only. Analyzing consumer purchase and exchange decisions would require a multi-product model. Second, many stores permit consumer returns up to a certain time limit. For products with short life cycles, delays can be costly: over time, the product loses value quickly and consumers may also exit the market. To understand how the duration of returns policies should be set, a dynamic model would most likely be necessary. Third, our model can be extended to incorporate resale. In practice, it is not unusual for firms to put returned items back into circulation. The option value of resale casts a new light on how returns policies should be designed. It would also be interesting to study consumer choice over new and returned products, since resale may lead to cannibalization of demand. Fourth, our model can be used to study the reverse supply chain more fully by incorporating processes such as inspection, distribution, remanufacturing and marketing. In all these activities, customer behavior models (for both end-consumers as well as commercial customers) can be used to study incentives of the parties involved. Finally, our model can be extended to incorporate retail competition. Does competitive pressure induce firms to offer 100% money-back-guarantees, or is it still viable to levy some restocking fees? It would be interesting to understand how market structure influences the design of consumer returns policies. We believe that all the above areas present fruitful opportunities for future research.

Appendix A: Proofs

Proof of Proposition 1 (i) Since the price p enters only in the term $(p - s)\overline{G}(p)$, the optimal p^* should maximize this term. Then, given p^*, the optimal q^* follows from the standard critical fractile solution of the conventional newsvendor model.

(ii) Consider the extreme case where $s = 0$. In this case, the seller's profit function is $p\overline{G}(p)E\min(X,q) - cq$ under full-returns but it is $\mu E\min(X,q) - cq$ under no-returns. Observe that $\mu \geq p\overline{G}(p)$ for any p so the seller prefers to offer no returns. By continuity, the seller prefers no-returns when s is sufficiently small. Similarly, at the extreme case where $c = 0$, we must have $s = 0$ since $0 \leq s \leq c$, so the seller prefers no-returns. Again, by continuity, the seller prefers no-returns when c is sufficiently small. ∎

Proof of Proposition 2 The first step in maximizing seller revenue is to find the optimal p^*, r^* that maximizes the expression $[(p-s)\overline{G}(r)+(p-r)G(r)]$, subject to the constraint $p \leq E\max(V,r)$ so that consumers are willing to buy in the first place. Clearly the seller should charge the maximum possible price $p = E\max(V,r)$, so we have an expression in terms of only r:

$$
\begin{align}
(p-s)\overline{G}(r) + (p-r)G(r) &= p - s\overline{G}(r) - rG(r) \tag{27}\\
&= E\max(V,r) - s\overline{G}(r) - rG(r) \tag{28}\\
&= rG(r) + \int_r^\infty vg(v)dv - s\overline{G}(r) - rG(r) \tag{29}\\
&= \int_r^\infty (v-s)g(v)dv. \tag{30}
\end{align}
$$

This is maximized when $r = s$, so we have $r^* = s$ and $p^* = E\max(V,s)$. Finally, we solve the resulting newsvendor problem in q to obtain $\overline{F}(q^*) = \frac{c-s}{p^*-s}$ as desired. ∎

Proof of Corollary 1 Let Π_P^*, Π_F^* denote the seller's optimal profits under partial and full refunds. Let q_P^*, q_F^* denote the seller's optimal quantities under partial and full refunds. Let p_P^*, p_F^* denote the seller's optimal prices under partial and full refunds.

First, it is clear that $\Pi_P^* \geq \Pi_F^*$ since full refunds is a special case of partial refunds; the former is a constrained optimal solution of the latter model.

Next, to show that $q_P^* \geq q_F^*$, from (10) and (5), it suffices to show that $E\max(V,s) - s \geq (p_F^* - s)\overline{G}(p_F^*)$. Note that $p_F^* \geq s$. Therefore we have

$$
\begin{align}
[E\max(V,s) - s] - (p_F^* - s)\overline{G}(p_F^*) &= \int_s^\infty (v-s)g(v)dv - \int_{p_F^*}^\infty (p_F^* - s)g(v)dv \\
&= \int_s^{p_F^*} (v-s)g(v)dv + \int_{p_F^*}^\infty (v-p_F^*)g(v)dv \geq 0,
\end{align}
$$

which completes the proof. ∎

Proof of Corollary 2   Observe that the refund $r^{FB} = s$ maximizes $\int_r^\infty (v - s)g(v)dv$. Next, we substitute $r^{FB} = s$ into (13) and maximize the corresponding newsvendor profit function to obtain $\overline{F}(q^{FB}) = \frac{c-s}{\int_s^\infty (v-s)g(v)dv}$. Finally, we note that any $p \leq E\max(V, r^{FB})$ induces consumers to try out the product and thus implements the first-best outcome. We may thus pick the welfare-maximizing price to be $p^{FB} = E\max(V, s)$.

From Proposition 2, observe that $p^{FB} = p^* = E\max(V, s)$, $r^{FB} = r^* = s$, and $\overline{F}(q^{FB}) = \overline{F}(q^*) = \frac{c-s}{\int_s^\infty (v-s)g(v)dv} = \frac{c-s}{E\max(V,s)-s}$. This completes the proof.   ∎

Proof of Proposition 3   From Proposition 2, the supply chain optimal decisions that we would like to implement are $p^{SC} \equiv E\max(V, s)$, $q^{SC} \equiv \overline{F}^{-1}\left(\frac{c-s}{E\max(V,s)-s}\right)$, and $r^{SC} \equiv s$. The retailer's profit function under the differentiated buy-back contract (w, b, l) is

$$\Pi_R(p, q, r) = [(p - b)\overline{G}(r) + (p - r - l)G(r)]E\min(X, q) - (w - b)q. \tag{31}$$

Therefore, the retailer's optimal choices are $p^* = E\max(V, b - l)$, $q^* = \overline{F}^{-1}\left(\frac{w-b}{E\max(V,b-l)-b}\right)$, and $r^* = b - l$. Under the contractual parameters given in the proposition, we have $b - l = s$, so we have $p^* = p^{SC}$ and $r^* = r^{SC}$, i.e. the price and refund are coordinated. Now notice that for these contractual parameters, we have

$$w - b = \phi(c - s), \tag{32}$$

$$E\max(V, b - l) - b = \phi[E\max(V, s) - s], \tag{33}$$

so this yields $q^* = q^{SC}$ and thus the quantity is also coordinated. Next, from (31), the retailer's optimal profit function is

$$[(p^* - b)\overline{G}(r^*) + (p^* - r^* - l)G(r^*)]E\min(X, q^*) - (w - b)q^* \tag{34}$$

$$= (p^* - b)E\min(X, q^*) - (w - b)q^* \tag{35}$$

$$= [E\max(V, b - l) - b]E\min(X, q^*) - (w - b)q^* \tag{36}$$

$$= \phi[E\max(V, s) - s]E\min(X, q^*) - \phi(c - s)q^* \tag{37}$$

$$= \phi\{(p^{SC} - s)E\min(X, q^{SC}) - (c - s)q^{SC}\} \tag{38}$$

$$= \phi\{[(p^{SC} - s)\overline{G}(r^{SC}) + (p^{SC} - r^{SC})G(r^{SC})]E\min(X, q^{SC}) - (c - s)q^{SC}\}, \tag{39}$$

which is a fraction ϕ of the optimal total supply chain profit. It remains to note that for any $\phi \in [0, 1]$, we have $w \geq b$ (which holds since $c \geq s$), $l \geq 0$ (which holds since $E\max(V, s) \geq s$), and $b - l \geq s$. With these conditions, the differentiated buy-back can be implemented. The proof is now complete.   ∎

Proof of Proposition 4 Let p^{SC}, q^{SC}, r^{SC} denote the supply chain optimal decisions. We will show that the retailer choosing p^{SC}, q^{SC} and the manufacturer choosing r^{SC} is a Nash equilibrium. Given $r^{SC} = s$, from the profit function in (19), it is clear that the retailer will choose price $p^* = p^{SC} = E\max(V, s)$. Further, we also see that the retailer's equilibrium quantity satisfies $\overline{F}(q^*) = \frac{w-b}{p^*-b} = \frac{w-b}{E\max(V,s)-b} = \frac{\phi(c-s)}{\phi(E\max(V,s)-s)} = \overline{F}(q^{SC})$. Finally, note that given the retailer's price $p^* = p^{SC}$ and quantity $q^* = q^{SC}$, the manufacturer will choose $r^* = r^{SC} = s$. This is because, if he deviates to $r < s$, no one will buy and he earns zero profit; similarly, if he deviates to $r > s$, his profits will also decrease (because total supply chain profit decreases from the optimum, while the retailer's profit increases since he can now charge a higher price with a higher refund r). Therefore, the supply chain optimal decisions constitute a Nash equilibrium. Under these choices, from (19), the retailer profit is $\Pi_R = (p^{SC} - b)E\min(X, q^{SC}) - (w-b)q^{SC} = \phi\{(p^{SC} - s)E\min(X, q^{SC}) - (c-s)q^{SC}\}$, which is a fraction ϕ of optimal supply chain profits. ∎

Proof of Proposition 5 Let p^{SC}, q^{SC}, r^{SC} denote the supply chain optimal decisions. Under the sales rebate contract, since the rebate u is paid whether or not the unit is returned, the equilibrium price and refund coincide with the supply chain optimal choices, so $p^* = p^{SC} = E\max(V, s)$ and $r^* = r^{SC} = s$. Under the given w, u, we also see that the equilibrium quantity satisfies $\overline{F}(q^*) = \frac{w-s}{p^*+u-r^*} = \frac{w-s}{u+E\max(V,s)-s}$, while the supply chain optimal quantity satisfies $\overline{F}(q^{SC}) = \frac{c-s}{E\max(V,s)-s}$. The two quantities coincide when $\frac{w-s}{u+E\max(V,s)-s} = \frac{c-s}{E\max(V,s)-s}$, which can be reduced to the expression given in the proposition. ∎

Appendix B: Supply Chain Contracts When Consumers Are Heterogeneous

In this appendix, we show that the contracts described in Sections 5.2 to 5.4 can be used to coordinate the supply chain when consumers are heterogeneous in their valuation distributions as modeled in Section 6.1. Here, contract parameters depend on the chosen price and refund.

Let $\hat{\pi}$ denote the fraction of consumers who purchase, so $\hat{\pi} = 1$ if all consumers purchase and $\hat{\pi} = \pi_H$ when only high-types purchase. Further, let $\hat{G}(\cdot)$ denote the distribution function for aggregate valuations. That is, when both types purchase, we have $\hat{G}(\cdot) = \pi_L G_L(\cdot) + \pi_H G_H(\cdot)$, but when only high-types purchase, we have $\hat{G}(\cdot) = G_H(\cdot)$. Let $\phi \in [0, 1]$ denote the retailer's profit share.

Differentiated buy-backs   Similar to Section 5.2, we first show that the following differentiated buy-back contract will coordinate the supply chain:

$$w(p,r) \;=\; \phi c + \overline{\phi}p - \overline{\phi}(r-s)\hat{G}(r), \tag{40}$$

$$b(p,r) \;=\; \phi s + \overline{\phi}p - \overline{\phi}(r-s)\hat{G}(r), \tag{41}$$

$$b(p,r) - l(p,r) \;=\; s. \tag{42}$$

To see why, let us substitute the contract parameters into the retailer's profit function. For convenience, we introduce the notation $z = q/\hat{\pi}$ and express quantity decisions in terms of z rather than q. After some algebraic manipulation we obtain

$$\Pi_R(p,z,r) \;=\; \hat{\pi} \cdot \{[(p-b(p,r))(1-\hat{G}(r)) + (p-r-l(p,r))\hat{G}(r)]E\min(X,z) - (w(p,r)-b(p,r))z\}$$

$$= \; \hat{\pi} \cdot \phi \cdot \{[p - s(1-\hat{G}(r)) - r\hat{G}(r)]E\min(X,z) - (c-s)z\} = \phi \cdot \Pi_T(p,z,r),$$

which is a fixed fraction of total profit Π_T. Therefore, the retailer's incentives are perfectly aligned with the supply chain's incentives.

Direct-to-manufacturer returns   When consumers return products directly to the manufacturer, the buy-back contract below can coordinate the supply chain.

$$w(p,r) \;=\; \phi c + \overline{\phi}p + \phi(r-s)\hat{G}(r), \tag{43}$$

$$b(p,r) \;=\; \phi s + \overline{\phi}p + \phi(r-s)\hat{G}(r). \tag{44}$$

As above, we can substitute these parameters into the retailer and manufacturer profit functions to show that each is a fixed proportion of total profit, for all combinations of decisions. Observe that when the manufacturer handles returns directly, the contract parameters $w(p,r)$ and $b(p,r)$ are increasing in the refund r, whereas the reverse is true, as in (40) and (41), when the retailer is responsible.

Sales rebates   Similar to Section 5.4, we also find that sales rebate contracts satisfying the following condition will coordinate the supply chain.

$$u(p,r) = \frac{(p-s) - (r-s)\hat{G}(r)}{c-s} \cdot (w-c) \tag{45}$$

Substituting this condition into the retailer's profit function yields an expression that is a fixed proportion of the total supply chain profit.

Finally, observe that if we set $r = s$ and $p = E\max(V,s)$ (which hold when consumers are homogeneous), each of the contracts above reduce to their corresponding counterparts obtained in Sections 5.2 to 5.4.

References

Alexandrov, A., M.A. Lariviere. 2007. Are reservations recommended? Working paper.

Andel, T., M. Aichlmayr. 2002. Turning returns into cash. *Transportation Distribution.* 43(8): 29-39.

Anderson, E.T., K. Hansen, D. Simester. 2008. The option value of returns: Theory and empirical evidence. *Mktg. Sci.* Forthcoming.

Aviv, Y., A. Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *M&SOM.* 10(3): 339-359.

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

Bhargava, H.K., R.R. Chen. 2007. Sale timing under buyer valuation uncertainty: When to endure information disadvantage? Working paper.

Blackburn, J.D., V.D.R. Guide, G.C. Souza, L.N. Van Wassenhove. 2004. Reverse supply chains for commercial returns. *California Mgmt. Rev.* 46(2): 6-22.

Cachon, G. 2003. Supply chain coordination with contracts. *Handbooks in OR&MS.* Graves, de Kok, eds.

Cachon, G. P., R. Swinney. 2007. Purchasing, pricing, and quick response in the presence of strategic consumers. Working paper.

Che, Y-K. 1996. Customer return policies for experience goods. *Jour. Indus. Econ.* 44: 17-24.

Chu, W., E. Gerstner, J.D. Hess. 1998. Managing dissatisfaction: How to decrease consumer opportunism by partial refunds. *Jour. Service Res.* 1(2): 140-155.

Cohen, M., S. Nahmias, W. Pierskalla. 1980. A dynamic inventory system with recycling. *Naval Res. Logist.* 27: 289-296.

Courty, P. 2003. Ticket pricing under demand uncertainty. *Jour. Law & Econ.* 46(2): 627-652.

Dana, J. 1998. Advance-purchase discounts and price discrimination in competitive markets. *Jour. Pol. Econ.* 106(2):395-422.

Dana, J.D., N.C. Petruzzi. 2001. The newsvendor model with endogenous demand. *Mgmt. Sci.* 47(11): 1488-1497.

Davis, S., E. Gerstner, M. Hagerty. 1995. Money back guarantees in retailing: Matching products to consumer tastes. *Jour. Retailing.* 71(1): 7-22.

Davis, S., M. Hagerty, E. Gerstner. 1998. Return policies and the optimal level of "hassle." *Jour. Econ. & Busi.* 50: 455-460.

Debo, L., G. van Ryzin. 2007. Inventories and consumer search behavior when product quality is uncertain. Working paper.

DeCroix, G., J.S. Song, P. Zipkin. 2005. A series system with returns: Stationary analysis. *Oper. Res.* 53: 350-362.

DeCroix, G., P. Zipkin. 2005. Inventory management for an assembly system with product or component returns. *Mgmt. Sci.* 51(8): 1250-1265.

DeCroix, G. 2006. Optimal inventory policy for a multi-echelon inventory system with remanufacturing. *Oper. Res.* 54(3): 532-543.

Dekker, R., M. Fleischmann, K. Inderfurth, L. van Wassenhoven. 2004. *Reverse Logistics for CLosed-Loop Supply Chains.* Springer, New York.

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

Earnest, L., A.G. Uribarri. 2007. Costco halts liberal electronics return policy. *Los Angeles Times.* February 28. Page C1.

Elmaghraby, W., A. Glc, P. Keskinocak. 2008. Designing optimal preannounced markdowns in the presence of rational customers with multiunit demands. *M&SOM.* 10(1): 126-148.

Emmons, H., S.M. Gilbert. 1998. The role of returns policies in pricing and inventory decisions for catalogue goods. *Mgmt. Sci.* 44(2): 276-283.

Fleischmann, M., J.M. Bloemhof-Ruwaard, R. Dekker, E. van der Laan, J.A.E.E. van Nunen, L.N. Van Wassenhove. 1997. Quantitative models for reverse logistics: A review. *Eur. J. Oper. Res.* 103(1): 1-17.

Fleischmann, M., R. Kuik, R. Dekker. 2002. Controlling inventories with stochastic item returns: A basic model. *Eur. J. Oper. Res.* 138: 6375.

Gallego, G., O. Sahin. 2006. Inter-temporal valuations, product design and revenue management. Working paper.

Guide, V.D.R. Jr., G.C. Souza, L.N. Van Wassenhove, J.D. Blackburn. 2006. Time value of commercial product returns. *Mgmt. Sci.* 52(8): 1200-1214.

Heal, G. 1977. Guarantees and risk-sharing. *Rev. Econ. Stud.* 44(3): 549-560.

Heiman, A., B. McWilliams, J. Zhao, D. Zilberman. 2002. Valuation and management of money-back guarantee options. *Jour. Retailing.* 78: 193-205.

Hess, J.D., W. Chu, E. Gerstner. 1996. Controlling product returns in direct marketing. *Mktg. Letters.* 7(4): 307-317.

Inderfurth, K. 1997. Simple optimal replenishment and disposal policies for a product recovery system with leadtimes. *OR Spektrum.* 19: 111-122.

Jerath, K., S. Netessine, S. K. Veeraraghavan. 2007. Revenue management with strategic customers: Last minute selling and opaque selling. Working paper.

Kelle, P., E. Silver. 1989. Purchasing policy of new containers considering the random returns of previously issued containers. *IIE Trans.* 21: 349354.

Koenigsberg, O., E. Muller, N.J. Vicassim. 2008. EasyJet pricing strategy: Should low-fare airlines offer last-minute deals? *Quantitative Mktg & Econ.* 6(3): 279-297.

Lai, G., L. G. Debo, K. Sycara. 2007. Impact of price matching policy on pricing, inventory investment and profit with strategic consumers. Working paper.

Lawton, C. 2008. The War on Returns. *Wall Street Jour.* May 8. Page D1.

Liu, Q., G. van Ryzin. 2008. Strategic capacity rationing to induce early purchases. *Mgmt. Sci.* 54(6): 1115-1131.

Marvel, H.P., J. Peck. 1995. Demand uncertainty and returns policies. *Internat. Econ. Rev.* 36(3): 691-714.

Moorthy, S., K. Srinivasan. 1995. Signaling quality with a money-back guarantee: The role of transaction costs. *Mktg. Sci.* 14(4): 442-466.

Padmanabhan, V., I. Png. 1997. Manufacturer's returns policies and retail competition. *Mktg. Sci.* 16: 81-94.

Pasternack, B. 1985. Optimal pricing and returns policies for perishable commodities. *Mktg. Sci.* 4:166-176.

Savaskan, R.C., S. Bhattacharya, L.N. Van Wassenhove. 2004. Closedloop supply chain models with product remanufacturing. *Mgmt. Sci.* 50(2): 239-252.

Simpson, V. 1978. Optimum solution structure for a repairable inventory problem. *Oper. Res.* 26: 270281.

Stock, J., T. Speh, H. Shear. 2002. Many happy (product) returns. *Harvard Busi. Rev.* 80(7): 16-17.

Su, X. 2007a. Intertemporal pricing with strategic customer behavior. *Mgmt. Sci.* 53(5): 726-741.

Su, X. 2007b. A model of consumer inertia with applications to dynamic pricing. Working paper.

Su, X. 2008. Bounded rationality in newsvendor models. *M&SOM.* Forthcoming.

Su, X., F. Zhang. 2007a. Strategic consumer behavior, commitment, and supply chain performance. *Mgmt. Sci.* Forthcoming.

Su, X., F. Zhang. 2007b. On the value of inventory information and availability guarantees when selling to strategic consumers. Working paper.

Swinney, R. 2008. Who has the hot product? Selling to strategic consumers when product value is uncertain. Working paper.

Taylor, T.A. 2002. Supply chain coordination under channel rebates with sales effort effects. *Mgmt. Sci.* 48(8): 992-1007.

Tibben-Lembke, R.S., D.S. Rogers. 2002. Differences between forward and reverse logistics in a retail environment. *Supply Chain Mgmt: An Internat. Jour.* 7(5): 271-282.

Tsay, A.A. 2001. Managing retail channel overstock: Markdown money and returns policies. *Jour. Retailing.* 77: 457-492.

Webster, S., Z.K. Weng. 2000. A risk-free perishable items returns policy. *MSOM.* 2(1): 100-106.

Wood, S.L. 2001. Remote purchase environments: The influence of return policy leniency on two-stage decision processes. *Jour. Mktg. Res.* 38(2): 157-169.

Xie, J., S. M. Shugan. 2001. Electronic tickets, smart cards, and online prepayments: When and how to advance sell. *Mktg. Sci.* 20(3) 219-243.

Yabalik, B., N.C. Petruzzi, D. Chhajed. 2005. An integrated product returns model with logistics and marketing coordination. *Eur. J. Oper. Res.* 161: 162-181.

Yu, M., R. Kapuscinski, H.-S. Ahn. 2005. Advance selling to homogeneous customers. Working paper.