INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

MARKETING SCIENCE
Vol. 29, No. 6, November–December 2010, pp. 1071–1085
ISSN 0732-2399 | EISSN 1526-548X | 10 | 2906 | 1071

DOI 10.1287/mksc.1100.0578
© 2010 INFORMS

Optimal Reverse Channel Structure for Consumer Product Returns

Jeffrey D. Shulman
Department of Marketing, Michael G. Foster School of Business, University of Washington, Seattle, Washington 98109, jshulman@u.washington.edu

Anne T. Coughlan
Department of Marketing, Kellogg School of Management, Northwestern University, Evanston, Illinois 60208, a-coughlan@kellogg.northwestern.edu

R. Canan Savaskan
Information Technology and Operations Management Department, Cox School of Business, Southern Methodist University, Dallas, Texas 75275, cebert@smu.cox.edu

Consumers often return a product to a retailer because they learn after purchase that the product does not match as well with preferences as had been expected. This is a costly issue for retailers and manufacturers—in fact, it is estimated that the U.S. electronics industry alone spent $13.8 billion dollars in 2007 to restock returned products [Lawton, C. 2008. The war on returns. *Wall Street Journal* (May 8) D1]. The bulk of these returns were nondefective items that simply were not what the consumer wanted. To eliminate returns and to recoup the cost of handling returns, many retailers are adopting the practice of charging restocking fees to consumers as a penalty for making returns. This paper employs an analytical model of a bilateral monopoly to examine the impact of reverse channel structure on the equilibrium return policy and profit. More specifically, we examine how the return penalty is affected by whether returns are salvaged by the manufacturer or by the retailer. Interestingly, we find that the return penalty may be more severe when returns are salvaged by a channel member who derives greater value from a returned unit. Also, the manufacturer may earn greater profit by accepting returns even if the retailer has a more efficient outlet for salvaging units.

Key words: channels of distribution; pricing; reverse logistics
History: Received: December 9, 2008; accepted: March 11, 2010; processed by Greg Shaffer. Published online in *Articles in Advance* June 24, 2010.

1. Introduction and Literature Review

This paper models channel management issues of consumer product returns. Product returns occur for a number of reasons. Unsold products are returned at the end of a selling season from retailers to manufacturers as a result of overstocking of inventory.[1] Consumers return defective units under warranty to resellers for replacement or repair.[2] Some returns might be the result of "opportunistic behavior," where a consumer buys a product, uses it temporarily, and then returns it for a refund.[3] However, even without inventory overstocks, defective products, or opportunistic behavior, products are also often returned after purchase because consumers learn that the product does not match preferences as well as had been expected. The management of this category of product returns is the focus of this research.

In 2007 alone, the U.S. electronics industry spent $13.8 billion to repackage and resell returned products (Lawton 2008). Of those returned products, 95% were nondefective items that were not what the consumer was expecting. Returns occur at rates of at least 6% for electronics retailers (Strauss 2007) and as high as 35% for catalog retailers (Rogers and Tibben-Lembke 1998). It is clear that product returns from consumers are costing companies a substantial amount of money. What is not as clear is who should pay for the cost and who should take responsibility for the returned units.

Many retailers have adopted restocking fees by which consumers pay a fee to return nondefective items.[4] For example, Best Buy charges a 15%

[1] For example, see the review article by Cachon (2003).

[2] For example, see Chao et al. (2009).

[3] For example, see Chu et al. (1998).

[4] In practice, if the product is returned because a wrong order is delivered or the unit is damaged or is defective, consumers are not charged a restocking by the seller. Such returns are not the focus of this research.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

restocking fee on returns of opened electronics items and 25% for appliances, and many wallpaper retailers charge restocking fees as high as 30%.[5] Although penalizing returns via a restocking fee may benefit the firm by extracting revenue and reducing return volumes, such a policy may also harm the firm by discouraging consumers from trying the product in the first place. In many instances, retailers explain that they charge return penalties to consumers because they themselves receive only partial refunds from the manufacturer.[6] However, although there are manufacturers that offer full or partial refunds for returned nondefective items, not all returns are sent back to the manufacturer. It is sometimes possible for the retailer to put a returned unit back on the shelf and sell it as new.[7] Other times retailers get a fraction of the new-product selling price by selling open-box items or by liquidating returned inventory (Strauss 2007). One of the research questions addressed in this paper is how reverse channel structure (that is, who salvages returned units) affects the return policy offered to consumers.

Danze, a manufacturer of kitchen and bath accessories such as faucets and showerheads, represents an interesting case. Danze uses a third-party liquidator to extract value from returned units. The retailer could use the same third-party liquidator, but instead the returned units are shipped from the retailer back to Danze, creating an extra cost.[8] Although salvaging by the retailer would save the shipping cost from the retailer to the manufacturer, Danze takes on the task of salvaging returned units. PetSafe is another manufacturer that uses the third-party liquidator Channel Velocity even though the extra costs associated with getting consumer returns from the retailer to the manufacturer could be saved by having the retailer use Channel Velocity directly. Our research helps explain why manufacturers such as Danze and PetSafe take returns from the retailer even though it could be more cost-efficient for the retailer to salvage the returned units. We summarize our research objectives with the following research questions:

1. How is the equilibrium return policy offered to consumers affected by the reverse channel structure in processing product returns?
2. Why might product returns be processed by the manufacturer even if doing so decreases the net value of a returned unit?

Our analysis provides an explanation for observed differences in the handling of product returns and leads to insights about how manufacturers and retailers should set prices in forward and reverse channels. Previous research has focused on managing *overstock returns* from the retailer to the manufacturer that arise because of inaccurate stocking decisions with unknown demand. Treating price as exogenous, Pasternack (1985) finds that to induce the order quantity of a vertically integrated channel, the manufacturer should offer partial refunds on all returns or full refunds on a fraction of returns. Kandel (1996) examines how the manufacturer's decision to take back returns depends on its relative advantage in disposing of the unsold inventory, optimal risk allocation, channel members' influence over demand, and information asymmetry. Padmanabhan and Png (1997, 2004) show that a manufacturer's decision to take back returns depends on the degree of competition at the retail level, the level of uncertainty surrounding demand, and the marginal cost of production. Allowing the demand distribution to be price dependent, Emmons and Gilbert (1998) demonstrate that the manufacturer can increase both channel members' profits by offering a positive refund for overstock returns. Iyer and Villas-Boas (2003) prove that the manufacturer is less likely to offer a return policy for overstock returns when the retailer has greater bargaining power. Arya and Mittendorf (2004) find that when the retailer has an information advantage, a manufacturer return policy can be used to elicit a retailer's private information. He et al. (2008) show that the manufacturer can use a return policy to convey information about product demand. Cachon (2003) provides an extensive review of the literature examining inventory decisions and return contracts between the retailer and the manufacturer.

Our research focuses on returns of a different nature: product returns of experience goods (Nelson 1970) from consumers to the retailer that are discovered to be a poor match with consumer preferences. In contrast to returns generated by overstocks, we model

[5] As of June 2010, American Blinds (http://www1.americanblinds.com/control/infopage?page=return.html&master=resourcecenter) and Steves Blinds and Wallpaper (http://www.stevesblindsand-wallpaper.com/sb_ourpolicies.aspx) both listed 30% restocking fees, and Seabrook Wallpaper (http://www.seabrookwallpaper.com/customer-service/) and The Internet Wallpaper Store (http://www.wallpaperstore.com/Returns.htm) both listed 25% restocking fees.

[6] As of June 2010, RehabMart (http://www.rehabmart.com/returns.asp) states in its return policy: "Access-Able, Fabrication Enterprises, Invacare Corporation, Sunrise Medical Corporation, Healthcraft Corporation and others listed below charge RehabMart an exceptional re-stocking fee on all items. RehabMart could not stay in business and serve its customers if it were to incur exceptional losses with re-stocking fees. This is why we must pass these restocking fees along to our customers."

[7] For example, a representative from Saks Fifth Avenue was quoted as saying, "If merchandise looks like it's in saleable condition, and has not been worn, we do put it back on the sales floor" (*Shame on You* 2003).

[8] Danze uses Channel Velocity to salvage returned units. From its website, http://www.channelvelocity.com, this company offers its services to retailers as well.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

a retailer that sets a price and a return penalty (i.e., a restocking fee and/or shipping charges) charged to consumers which affects both sales quantity and the number of units returned. In contrast to an overstock returns setting where retail prices are found to be decreasing with the manufacturer's refund (e.g., Padmanabhan and Png 1997, Emmons and Gilbert 1998, He et al. 2008), we find the equilibrium retail price in dealing with consumer returns increases with the manufacturer's refund.

Previous research on return policies for consumer returns abstracts away from channel relationships (Davis et al. 1995, 1998; Che 1996; Chu et al. 1998; Matthews and Persico 2005; Shulman et al. 2009; Ofek et al. 2010; Anderson et al. 2009). Our research examines settings where the retailer and the manufacturer are independent channel members. As such, we take into account the fact that the manufacturer's return policy affects the retailer's pricing strategy, which in turn drives demand as well as the consumer's product return behavior.

We examine three reverse channel structures. As a benchmark, we consider a vertically integrated system (VI) in which both the forward (selling) and the reverse (salvaging) channel decisions are made by a single agent, i.e., the manufacturer. We also model a decentralized channel in which there are two possible reverse channel structures. In the first structure, the retailer sells the product as well as processes and salvages returned units (CR). In the second structure, the manufacturer accepts returns from the retailer and salvages the returned units (CRM). For each of these reverse channel structures, we examine two contracting structures: when the manufacturer is able to charge a fixed fee or quantity discount schedule, and when the manufacturer is limited to constant per-unit wholesale and refund rates. Our model examines the interaction between contract structure and reverse channel structure, as well as how each of these affects the retail price, return penalty, sales quantity, and exchanges.

Interestingly, we find that the return penalty may be higher when the retailer salvages the goods than if the manufacturer does—even if the retailer's value for the returned units is higher. Moreover, the manufacturer may earn greater profits from taking back returns itself, even if the retailer has an advantage in salvaging the units.[9] These results are driven by the fact that the retailer's return policy has an impact on demand as well as on the number of units returned. The retailer has an incentive to use return penalties to reduce costs and increase revenue. Such penalties hurt the manufacturer by reducing quantity demanded without enhancing wholesale revenue per unit. The manufacturer's ability to align incentives through the joint use of a wholesale price and refund may outweigh any efficiency loss resulting from handling returned units itself. In fact, when returns pass from the retailer to the manufacturer, the manufacturer is able to induce the retailer to charge the same return penalty to consumers as would be chosen if the manufacturer sold directly to consumers.

The rest of this paper is organized as follows. In §2, we describe the model. In §3, we present the model analysis and results concerning the combined optimal reverse channel structure, equilibrium pricing policies, and refund policies throughout the channel. Finally, in §4, we provide a discussion of the model's intuition and directions for future research.

2. The Model

2.1. Manufacturer

The manufacturer produces two horizontally differentiated products (denoted by subscript j; $j = 0, 1$) located at 0 and 1 on a Hotelling unit line.[10] The manufacturer acts as a Stackelberg leader, producing each product at the same per-unit marginal cost c and charging a per-unit wholesale price for each product w_j to the retailer. If the manufacturer accepts returns from the retailer arising from consumer returns, the manufacturer chooses a refund r_j for each product j to pay to the retailer per unit returned. Each product has the same per-unit net salvage value s for each unit the manufacturer takes back as a return. A positive value of s reflects the manufacturer's ability to resell a returned product through secondary channels at a price higher than the reverse logistics cost of getting back the return and remarketing it. A negative value of s represents a situation where the reverse logistics cost to the manufacturer of remarketing the returned product exceeds any resale value for the good, or the

[9] Kandel (1996) finds that the manufacturer may accept *overstock* returns even if the retailer has a salvaging advantage. The driver in that result is that by taking back returns, the manufacturer is able to realize all channel profits (leaving the retailer with zero profit), essentially becoming a retailer by selling on consignment (because the manufacturer's refund is greater than or equal to its wholesale price). The result is limited by the assumption that the equilibrium order quantity is the same whether or not the manufacturer accepts returns and the assumption that retail price is an exogenous variable. In our model, the manufacturer's return policy has an impact on demand and retail price, leaves the retailer with positive profit, and does not serve as a mechanism equivalent to selling directly to consumers.

[10] Modeling two products, rather than a single product, allows for the possibility that consumers exchange an initial purchase for a product more suited to preferences. As shown in the electronic companion, available as part of the online version that can be found at http://mktsci.pubs.informs.org/, the main results of the paper carry over to a single-product setting.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

returned product is not resalable and the manufacturer disposes of it at some positive cost. A more responsive and operationally efficient reverse logistics process ensures a higher value for s. If product j is exchanged at the retailer level, the replacement good must be produced by the manufacturer at the marginal production cost c and is then sold to the retailer at the wholesale price $w_{\neq j}$. We assume that c is greater than or equal to s.[11] This assumption prevents the unrealistic situation from occurring in which the manufacturer has an incentive to encourage returns to profitably salvage returned units and make an additional sale on the exchange.

2.2. Retailer

For each product j, the retailer chooses a retail price, p_j, and a return penalty, f_j, including any consumer payments for shipping, to charge consumers for a return. The return penalty f_j represents the total financial loss experienced by the consumer for buying and returning a product and thus includes any shipping payments (in purchase or return) that are not reimbursed by the retailer. We examine a scenario in which the retailer may salvage returned units at a net value s_r (symmetric for each product) and another scenario in which the retailer may send returned units to the manufacturer for a net refund equal to r_j. The retailer sends returned units upstream to the manufacturer if and only if $r_j \geq s_r$.

We assume that salvage value is net of all costs associated with getting the product from the consumer to the retailer (and potentially on to the manufacturer), such as shipping paid to third parties, because these costs affect the value of a return. This can be mathematically transformed into a model in which the consumer pays for shipping separately (meaning f_j captures only the actual restocking fee penalty paid to the retailer). Alternatively, we could define h as the exogenous fee paid to third-party shippers, $\underline{s}_r$ as the value to the retailer of having the unit in their possession, and $\underline{f}_j$ as the restocking fee paid by the consumer directly to the retailer. By the definitions in our model, $f_j = \underline{f}_j + h$ and $s_r = \underline{s}_r - h$. Consequently, choosing $\underline{f}_j$ given $\underline{s}_r$ and h is equivalent to our model specification of choosing f_j given s_r.[12] Although this parametric transformation reduces the parameter space and simplifies the exposition of our analysis, it has no impact on our results. Furthermore, this simplification captures the fact that retailers ultimately choose how much it costs a consumer to return a product, and in practice retailers such as http://www.zappos.com choose to pay for the shipping of returned units.

2.3. Consumers

Each risk-neutral consumer is assumed to keep at most one unit of one of the products.[13] In owning a product located at x_j on the Hotelling line, consumer i experiences consumption value equal to $u_i - d|x_j - \theta_i|$, where $d > 0$ measures the disutility per unit deviation from consumer i's ideal taste parameter θ_i, and u_i measures the value consumer i obtains from owning a unit in the product category that perfectly matches his preferences.[14] Consumers are thus differentiated on two dimensions, with u_i and θ_i independently distributed. The taste parameter is distributed as $\theta_i \sim U[0, 1]$, and the parameter u_i takes values between 0 and $\bar{u}$ with equal density (normalized to 1) at each u_i. Modeling a range of u_i values captures the market-expanding effect of relaxing return penalties. In other words, we will show that as the return penalty decreases, the number of consumers who initially purchase the product increases (i.e., the expected utility of buying a product is nonnegative for a greater number of consumers). We assume that consumers are aware of both the price and return penalty before making a purchase decision.[15]

Consumers are equally and completely uninformed about the value of $|x_j - \theta_i|$ for each product before purchase. However, they share a common and known distribution of product fit before the initial purchase: $|x_j - \theta_i| \sim U[0, 1]$. Therefore, consumers who decide to make a purchase in the product category will randomly decide between the two products if the prices and return penalties are symmetric. Otherwise, consumers will purchase the product that offers the greatest expected utility. We recognize that consumers, in reality, may have some prior knowledge about which

[11] As shown in Shulman et al. (2009), this assumption implies that the manufacturer can do no better than to treat the returned unit as new and use it to satisfy demand for new units.

[12] Equivalently, the same logic follows if the retailer pays shipping charges h to move the returns to the manufacturer. Defining the net refund $r = \underline{r} - h$ with the net salvage value $s = \underline{s} - h$ makes choosing the actual refund $\underline{r}$ given $\underline{s}$ and h equivalent to our model specification of choosing r given s.

[13] Risk neutrality by consumers is a common assumption in the literature (e.g., Davis et al. 1998, Matthews and Persico 2005, Shulman et al. 2009). Risk-averse consumers could be modeled in a mean-variance linear utility function with a risk aversion coefficient. This would serve to amplify the existing negative effect of restocking fees on initial sales. Although it would affect the specific return policy, it would not qualitatively alter our results. In the interest of parsimony, risk aversion is excluded from the model.

[14] Clearly, the consumer may be male or female; for ease of exposition, we characterize the consumer as "he."

[15] Alternatively, there may be consumers who are unaware of the restocking fee before purchase. As long as the segment that is aware of the restocking fee is sufficiently large, our results (which rely on the market expansion effect of relaxed return policies) will still hold.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

product they will prefer. We make this assumption to simplify an already complex problem in which the important element is that consumers have uncertainty about product fit, which in turn may trigger a return. If, instead, consumers were partially informed about their product fit, there could be fewer returns by inframarginal consumers, but the price and return penalty would still have an impact on the marginal consumer and thus on the total number of initial purchases and product returns. We also allow for the possibility (occurring with probability $(1-\alpha) < 1$) that a consumer, whose initial beliefs lead him to initially buy, discovers after purchase that neither product fits with preferences. In this case, although the a priori expected utility from purchase is positive, the ex post utility to this consumer of owning either good is in fact revealed to be zero. Meanwhile, with probability $\alpha > 0$, the consumer discovers after purchase that his utility for the product category is indeed positive, although it is possible that the initially purchased product is not the best fit to this consumer's preferences.

Table 1 summarizes our model's variables and definitions.

After making a purchase at the retail price, the consumer learns how well it fits his preferences. Among the fraction α of consumers who have positive product-category utility, a return and exchange is triggered if a consumer realizes that the actual utility of keeping the initially purchased product is less than what could be obtained from returning it and buying an alternative product.[16] The fraction $(1-\alpha)$ of consumers who discover after initial purchase that their category consumption utility is zero, meanwhile, will return the initially purchased product without buying an alternative product.[17] Consumers in the model are assumed not to be loss averse. In the §4, we remark upon the limited impact of incorporating prospect theory (see Kahneman and Tversky 1979) into the consumer utility formulation. The sequence of events is depicted in Figure 1.

[16] In the interest of parsimony, we examine scenarios in which a fraction of the α consumers who do not lose their base valuation for the product offering will choose to purchase the other, better-matching product. We show in the electronic companion that the "return and exchange" behavior is the optimal consumer strategy for a wide array of parameter values (i.e., if $d \leq (2+\sqrt{4-\alpha^2})(c-s)/\alpha$, exchanging offers greater utility than returning without subsequent purchase for those consumers with positive category consumption utility).

[17] Davis et al. (1998) consider the possibility that consumers may act opportunistically and return a product after extracting value from its use before the return. It is estimated that returns resulting from this "free renting" account for less than 5% of all returns (Middleton 2007). In the interest of parsimony, we abstract away from this possibility.

Table 1 Parameters and Decision Variables

Symbol	Definition
c	Manufacturer's marginal cost of product
s	Manufacturer's salvage value of a returned unit (net of costs)
s_r	Retailer's salvage value of a returned unit (net of costs)
d	Consumer disutility per unit of deviation from match with preferences
u_i	Consumer i's reservation utility for perfect match
$\bar{u}$	Upper bound on u_i (i.e., the maximum possible reservation utility)
x_j	Location of product j
θ_i	Consumer i's ideal taste parameter
$(1-\alpha)$	Probability that consumer's ex post utility equals zero for each product
p_j	Retail price for product j
f_j	Consumers' return penalty for product j including shipping costs
w_j	Manufacturer's wholesale price for product j
r_j	Retailer's net refund for product j (the difference between the refund paid by the manufacturer and the shipping costs paid by the retailer)
T	Fixed fee paid by the retailer in a two-part tariff

For a given consumer i who initially purchases product j, we can write the probability $\phi_{ej}(f_j, p_j, p_{-j}; d, \alpha)$ that the consumer exchanges this purchase for the product located at x_{-j}, the probability $\phi_{kj}(f_j, p_j, p_{-j}; d, \alpha)$ that the consumer keeps this purchase, the probability ϕ_{rj} that the consumer returns this purchase without exchange, and the expected utility ex ante of purchasing product j (E_j(utility)), as[18]

$$\begin{aligned}
\phi_{kj}(f_j, p_j, p_{-j}; d, \alpha) &= \alpha\left(\frac{1}{2} + \frac{f_j - p_j + p_{-j}}{2d}\right),\\
\phi_{ej}(f_j, p_j, p_{-j}; d, \alpha) &= \alpha\left(\frac{1}{2} - \frac{f_j - p_j + p_{-j}}{2d}\right),\\
\phi_r &= 1-\alpha, \qquad (1)\\
E_{ij}(\text{utility}) &= \alpha\left(u_i - \frac{p_j + p_{-j} + f_j}{2} - \frac{d}{4}\right.\\
&\qquad \left. + \frac{(f_j - p_j + p_{-j})^2}{4d}\right) - (1-\alpha) f_j.
\end{aligned}$$

Notice from Equation (1) that a higher return penalty increases the likelihood that consumers who make a purchase will keep it rather than exchange it. A higher return penalty also decreases the expected utility of initially making a purchase (if the exchange probability is nonnegative), because consumers anticipate that this penalty will either induce them to keep a product that does not match as well with preferences or be charged when a return or exchange is made.[19]

[18] Proof of this equation is in the electronic companion.

[19] $\partial E_j(\text{utility})/\partial f_j = -(2d - \alpha(f_j - p_j + p_{-j} + d))/(2d) < 0$ for d such that exchanges are nonnegative ($d > f_j - p_j + p_{-j}$).

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

Figure 1 **Sequence of Events and Payoffs**

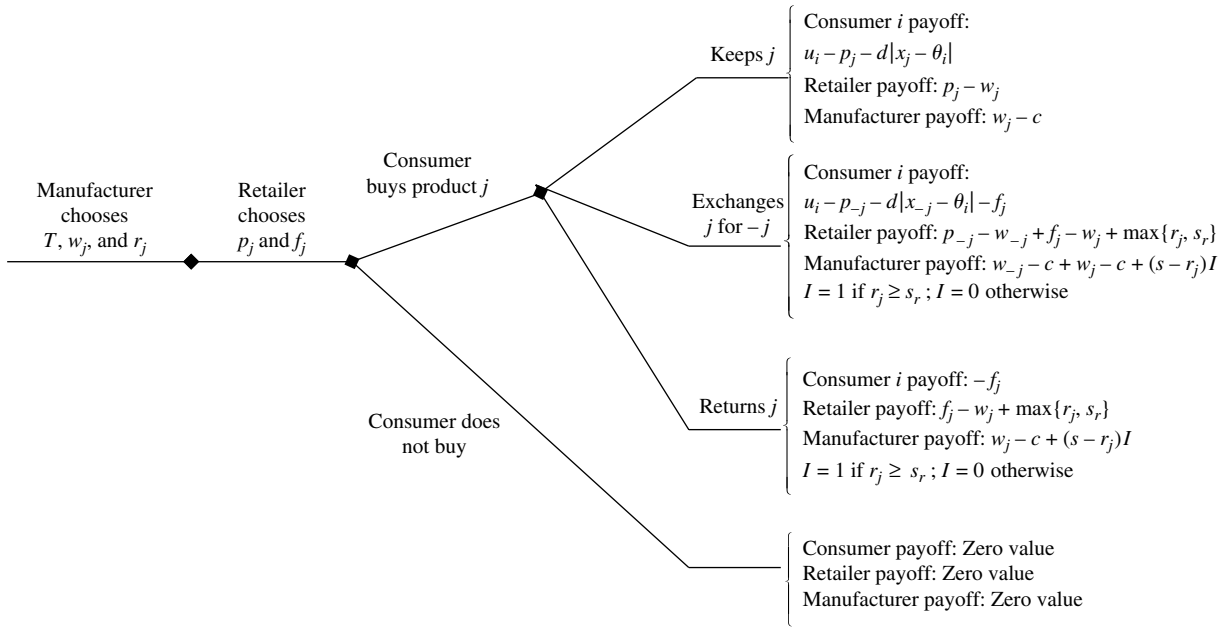

Consumers will purchase the product that offers the greatest utility and will make a purchase initially if and only if their expected utility is greater than or equal to zero, that is, if u_i is sufficiently high. The initial sales of each product j, $q_{bj}(f_j, p_j, f_{-j}, p_{-j}; d, \bar{u}, \alpha)$,[20] can be written as

$$q_{bj}(f_j, p_j, f_{-j}, p_{-j}; d, \bar{u}, \alpha) = \frac{1}{2}\left(\bar{u} - \frac{p_j + p_{-j}}{2} - \frac{d}{4} - \frac{(2-\alpha)f_j}{2\alpha} + \frac{(f_j - p_j + p_{-j})^2}{4d}\right) \quad \text{if } \alpha(f_j + f_{-j})(f_j - f_{-j} - 2p_j + 2p_{-j}) - 2d(f_j - f_{-j})(2-\alpha) = 0,$$

$$q_{bj}(f_j, p_j, f_{-j}, p_{-j}; d, \bar{u}, \alpha) = 0 \quad \text{if } \alpha(f_j + f_{-j})(f_j - f_{-j} - 2p_j + 2p_{-j}) - 2d(f_j - f_{-j})(2-\alpha) < 0, \quad \text{and} \tag{2}$$

$$q_{bj}(f_j, p_j, f_{-j}, p_{-j}; d, \bar{u}, \alpha) = \bar{u} - \frac{p_j + p_{-j}}{2} - \frac{d}{4} - \frac{(2-\alpha)f_j}{2\alpha} + \frac{(f_j - p_j + p_{-j})^2}{4d} \quad \text{if } \alpha(f_j + f_{-j})(f_j - f_{-j} - 2p_j + 2p_{-j}) - 2d(f_j - f_{-j})(2-\alpha) > 0.$$

Equation (2) captures an important phenomenon related to product return policies. Although a lower return penalty (lower f) may result in an increase in returns, there is a market expansion effect of the lenient return policy because more people may be willing to initially try the product, knowing they can return it later.[21] Therefore, a penalty intended to reduce and recoup costs associated with returns may also reduce revenue by forgoing the market expansion effect of a lenient return policy.

3. Model Analysis and Results

As a benchmark to the later analysis of decentralized channel systems, we first model pricing and return policy decisions in a vertically integrated system that can choose the salvaging technology of the manufacturing level (s) or the retailing level (s_r), whichever is greater. The vertically integrated system's objective function can be written as

$$\begin{aligned}\max_{p_0, f_0, p_1, f_1} \pi^{\mathrm{VI}} = {} & q_{b0}(f_0, f_1, p_0, p_1; d, \bar{u}, \alpha) \\ & \cdot ((p_0 - c) + (-p_0 + f_0 + \max\{s, s_r\}) \\ & \quad \cdot (\phi_{e0}(f_0, p_0, p_1; d, \alpha) + \phi_r) + (p_1 - c) \\ & \quad \cdot \phi_{e0}(f_0, p_1, p_0; d, \alpha)) \\ & + q_{b1}(f_0, f_1, p_0, p_1; d, \bar{u}, \alpha) \\ & \cdot ((p_1 - c) + (-p_1 + f_1 + \max\{s, s_r\}) \\ & \quad \cdot (\phi_{e1}(f_1, p_0, p_1; d, \alpha) + \phi_r) + (p_0 - c) \\ & \quad \cdot \phi_{e1}(f_1, p_1, p_0; d, \alpha)).\end{aligned}$$

[20] Note that initial sales are in units rather than in probability because u_i takes values between 0 and $\bar{u}$ with equal density at each value of u_i, where we do not impose an upper bound on $\bar{u}$.

[21] $\partial q_{bj}(f_j, p_j, f_{-j}, p_{-j}; d, \bar{u}, \alpha)/\partial f_j = -(2d - \alpha(f_j - p_j + p_{-j} + d))/(2\alpha d) < 0$ for d such that exchanges are nonnegative.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

The first expression in the profit function of the vertically integrated system denotes the number of customers who purchase product 0 initially. This is multiplied by the average margin that this purchase generates for the vertically integrated system. The margin expression explicitly takes into account the possibility of returns with and without subsequent product purchase (i.e., exchange). This is summed with the profit from selling product 1 initially.

We compare and contrast the optimal channel decisions of the vertically integrated systems to those in the decentralized systems. We examine the optimal reverse channel structure in a decentralized channel (i.e., whether the manufacturer or the retailer should be responsible for extracting value from returned units) under two scenarios: (i) when the manufacturer charges the retailer a fixed fee and per-unit wholesale price (or equivalently offers a quantity discount schedule) along with a per-unit refund, and (ii) when the manufacturer is limited just to the per-unit wholesale price and refund components in its contract with the retailer. In each structure, both the retailer and the manufacturer are assumed to be risk neutral and can rationally forecast the number of returns. Hence, we assume that the manufacturer and retailer have enough production capacity and inventory, respectively, to supply demand. Therefore, in this case, the manufacturer's production technology can be either make-to-stock or make-to-order without affecting model results. The vertically integrated structure and the two decentralized channel structures are depicted in Figure 2.

3.1. Fixed Fees or Quantity Discount Schedules

First, we examine a two-part tariff wholesale contract, under which the manufacturer charges the retailer a fixed fee (T) and wholesale prices (w_j) to acquire units of each product $j \in \{0, 1\}$ to resell to consumers.

Figure 2 **Reverse Channel Structures**

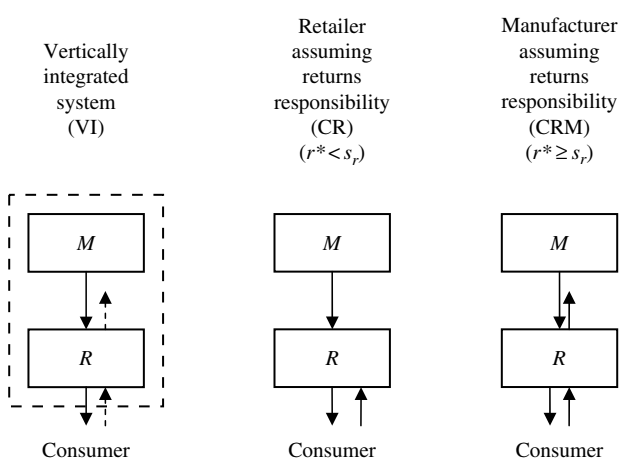

We allow as well for the manufacturer to offer a refund rate for each unit returned of product $j \in \{0, 1\}$ from the retailer to the manufacturer (r_j). With this contracting mechanism, the outcome of the vertically integrated system is achievable. The symmetric equilibrium outcome is described in Table 2 and proven in the electronic companion. It is also shown in the electronic companion that asymmetric choices (i.e., different policies for product 0 than for product 1) will not improve profit.

Because our focus is on developing an understanding of which channel member should be responsible for salvaging returned units, we restrict our attention to parameter values such that exchanges are positive: $d > c - \max\{s, s_r\}$.[22] Interestingly, the optimal return penalty in the vertically integrated system simply passes on all costs associated with the exchange to consumers; i.e., $f^{\mathrm{VI}} = (c - \max\{s, s_r\})$. The optimal reverse channel structure with a manufacturer-to-retailer contract involving a fixed fee, per-unit wholesale price, and per-unit refund rate is described in Proposition 1.

Proposition 1. *Consider a manufacturer who can charge a fixed fee, a per-unit wholesale price to the retailer, and a per-unit refund rate. If and only if $s > s_r$, then the manufacturer earns greater profit from the reverse channel structure in which the manufacturer salvages returns than the reverse channel structure in which the retailer salvages returns. The equilibrium is described in Table* 2.

Proof. See the electronic companion.

Proposition 1 confirms the intuition that when full coordination is achievable, the manufacturer will choose to allocate the responsibility for salvaging returns to the channel member who can extract the greatest salvage value. This contracting form implies that the manufacturer accepts returns if and only if it can extract greater value from returned units than the retailer can. As such, when this contracting mechanism is feasible, we would not observe manufacturers selling returned units to salvagers or liquidators, because allocating this activity to the retailer instead would save the cost of shipping returned units from the retailer to the manufacturer.

In practice, the manufacturer may not always offer fixed fees or quantity discounts.[23] Therefore, we examine the equilibrium in a decentralized channel when wholesale and refund contracts involve only a per-unit price, but no fixed fee. We identify the optimal reverse channel choice of a manufacturer selling

[22] From Equation (1), exchanges are positive if $d > f$; given the equilibrium values in Table 2, this implies immediately that exchanges are positive if $d > c - \max\{s, s_r\}$.

[23] In personal interviews, store managers from a large department store and a large discount retailer claimed that none of their suppliers offered quantity discounts or charged a fixed fee. These stores carry clothing, food, electronics, and home furnishing items.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

Table 2 **Equilibrium Price, Return Penalty, Quantities, and Profit in a VI System, Also Achievable Under a (Fixed Fee, Wholesale Price, Refund Rate) Wholesale Contract**

Term	Equilibrium value
Retail price	$p^{\text{VI}} = \frac{\bar{u}}{2} + \frac{(c-\max\{s,s_r\})^2}{8d} + \frac{2(c+\max\{s,s_r\})-d}{8} - \frac{(1-\alpha)(c-\max\{s,s_r\})}{2\alpha}$
Return penalty	$f^{\text{VI}} = c - \max\{s,s_r\}$
Quantity sold initially	$q_b^{\text{VI}} = \frac{\bar{u}}{2} + \frac{(c-\max\{s,s_r\})^2}{8d} - \frac{2(3c-\max\{s,s_r\})+d}{8} - \frac{(1-\alpha)(c-\max\{s,s_r\})}{2\alpha}$
Exchange probability	$\phi_e^{\text{VI}} = \alpha\left(\frac{1}{2} - \frac{(c-\max\{s,s_r\})}{2d}\right)$
Channel profit	$\pi^{\text{VI}} = \frac{(\alpha(c-\max\{s,s_r\})^2 - (4+2\alpha)cd + 2d((2-\alpha)\max\{s,s_r\}+2\alpha\bar{u}) - \alpha d^2)^2}{64d^2}$
Wholesale price	$w^{\text{VI}} = c$
Refund rate	$r^{\text{VI}} = s$
Fixed fee	$T^{\text{VI}} = \frac{(\alpha(c-\max\{s,s_r\})^2 - (4+2\alpha)cd + 2d((2-\alpha)\max\{s,s_r\}+2\alpha\bar{u}) - \alpha d^2)^2}{64d^2}$

through an independent retailer, and we characterize the equilibrium outcome in each reverse channel structure.

3.2. Wholesale Contracts with Per-Unit Wholesale Price and Per-Unit Refund

In a decentralized channel with a wholesale contract involving a per-unit wholesale price charged to the retailer, the manufacturer chooses between two reverse channel structures. In either case, consumers return unwanted products to the retailer; however, from that point on, the manufacturer may allocate responsibility for salvaging returned units to the retailer or choose to salvage the returned units itself, offering the retailer a refund per unit returned upstream. In each case, the manufacturer is the Stackelberg leader choosing a per-unit wholesale price w_j to charge the retailer for each product j and, when accepting returns, a per-unit refund r_j for each unit of product j returned from the retailer to the manufacturer. When the manufacturer does not accept returns from the retailer, the effective refund rate is $r_j = 0$. Ultimately, the manufacturer chooses the optimal reverse channel structure and sets the wholesale price (and the refund rate to the retailer), taking the outcomes in each reverse channel structure into account.

For each unit of product j initially sold, the retailer earns the unit margin $(p_j - w_j)$. Each consumer who buys product j initially faces a probability $\phi_{ej}(f_j, p_j, p_{-j}; d, \alpha)$ that he will exchange the unit and a probability ϕ_r that he will return the unit without exchanging it. In expectation, total exchanges of each product j therefore equal $\phi_{ej}(f_j, p_j, p_{-j}; d, \alpha) \cdot q_{bj}(f_j, p_j, f_{-j}, p_{-j}; d, \bar{u}, \alpha)$, and total returns (without an exchange) are equal to $\phi_r \cdot q_{bj}(f_j, p_j, f_{-j}, p_{-j}; d, \bar{u}, \alpha)$. For each unit of product j returned, the retailer pays out a consumer refund equal to $(p_j - f_j)$ and then returns it to the manufacturer for a refund r_j or salvages the returned unit at a net value s_r (if the manufacturer does not accept returns or if $r_j < s_r$). If the return is part of an exchange, the retailer also sells the unit for which the initial purchase is exchanged at a unit margin $(p_j - w_j)$.

The retailer chooses prices and return penalties to maximize its profit. The retailer's objective function can be written as

$$\begin{aligned}\max_{p_A, f_A, p_1, f_1} \pi_{\text{ret}} = {} & q_{b0}(f_0, f_1, p_0, p_1; d, \bar{u}, \alpha) \\ & \cdot ((p_0 - w_0) + (-p_0 + f_0 + \max\{s_r, r_0\}) \\ & \quad \cdot (\phi_{e0}(f_0, p_0, p_1; d, \alpha) + \phi_r) + (p_1 - w_1) \\ & \quad \cdot \phi_{e0}(f_0, p_1, p_0; d, \alpha)) \\ & + q_{b1}(f_0, f_1, p_0, p_1; d, \bar{u}, \alpha) \\ & \cdot ((p_1 - w_1) + (-p_1 + f_1 + \max\{s, r_1\}) \\ & \quad \cdot (\phi_{e1}(f_1, p_0, p_1; d, \alpha) + \phi_r) + (p_0 - w_0) \\ & \quad \cdot \phi_{e1}(f_1, p_1, p_0; d, \alpha)),\end{aligned}$$

where

$$\phi_{ej}(f_j, p_j, p_{-j}; d, \alpha), \quad \phi_r, \quad \text{and} \quad q_{bj}(f_j, f_{-j}, p_j, p_{-j}; d, \bar{u}, \alpha)$$

are defined as in Equations (1) and (2).

The retailer will only return units of product j to the manufacturer if the refund to the retailer, r_j, is at least as large as the retailer's own salvage value; hence we have the term $\max\{s_r, r_j\}$ in the retailer's optimization problem. We first solve for the unique symmetric solution in which decision variables remain constant across products. In the electronic companion,

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

we show that asymmetric choices do not improve firm profit. The reaction functions defining the retailer's profit-maximizing price and return penalty are functions of w_j and r_j:

$$p_j(w_j, r_j) = \frac{\bar{u}}{2} + \frac{(w_j - \max\{s_r, r_j\})^2}{8d} + \frac{2(w_j + \max\{s_r, r_j\}) - d}{8} - \frac{(1-\alpha)(w_j - \max\{s_r, r_j\})}{2\alpha}, \tag{3}$$

$$f_j(w_j, r_j) = w_j - \max\{s_r, r_j\} \quad \text{for } j \in \{0, 1\} \text{ when } \{w_j = w_{-j},\ r_j = r_{-j}\}.$$

The retailer's best response function illustrates a key difference between consumer returns and overstock returns. Holding wholesale price constant, the retail price is increasing in the refund offered by the manufacturer to the retailer.[24] This is in contrast to the finding that retail prices are decreasing in the manufacturer's refund for overstock returns (e.g., Padmanabhan and Png 1997, Emmons and Gilbert 1998, He et al. 2008). The manufacturer's refund to the retailer has a different impact on retail price in consumer returns than overstock returns because of its impact on the return penalty that the retailer charges to consumers. The retailer facing a lower cost of returns (because of a greater refund from the manufacturer for items returned up channel) in turn charges a lower return penalty to consumers. This raises consumers' expected utility of purchase and thereby their willingness to pay, allowing the retailer to optimally charge a higher price in the presence of a higher refund.

Whether or not the manufacturer accepts responsibility for returned units, the manufacturer earns its margin $(w_j - c)$ on units of product j sold initially as well on units for which consumers exchange their original purchases. If it is optimal to accept responsibility for returns, the manufacturer also pays out a refund r_j and earns the salvage value s for each exchange. The manufacturer's objective function is then

$$\begin{aligned}\max_{w_0, r_0, w_1, r_1, I_z} \quad \pi_{\text{mfgr}} = {} & q_{b0}(f_0(w_j, r_j), p_0(w_j, r_j), \\ & p_1(w_j, r_j); d, \bar{u}, \alpha) \\ & \cdot ((w_0 - c) + (w_1 - c - I_z \cdot (r_0 - s)) \\ & \cdot \phi_{e0}(f_0(w_j, r_j), p_0(w_j, r_j), \\ & p_1(w_j, r_j); d, \alpha) - I_z \cdot \phi_r \cdot (r_0 - s)) \\ & + q_{b1}(f_1(w_j, r_j), p_0(w_j, r_j), \\ & p_1(w_j, r_j); d, \bar{u}, \alpha) \\ & \cdot ((w_1 - c) + (w_0 - c - I_z \cdot (r_0 - s)) \\ & \cdot \phi_{e1}(f_0(w_j, r_j), p_0(w_j, r_j), \\ & p_1(w_j, r_j); d, \alpha) - I_z \cdot \phi_r \cdot (r_1 - s)), \\ \text{s.t.} \quad & I_z(r_j - s_r) \geq 0,\end{aligned}$$

where I_z is the manufacturer's choice of reverse channel structure and is an indicator variable equal to one if the manufacturer accepts and salvages returned units from the retailer (CRM reverse channel structure) and equal to zero otherwise (CR reverse channel structure), and the constraint $I_z(r_j - s_r) \geq 0$ ensures that the manufacturer choosing the CRM reverse channel structure actually encourages the retailer to return the units rather than salvage them. For I_z equal to zero, the refund rate is effectively equal to zero. The manufacturer's choice of reverse channel structure is a choice variable rather than a corner solution of the condition $r < s_r$ because the two optimization problems are not equivalent. The refund rate r is chosen under the assumption that the returned unit is worth s and is salvaged by the manufacturer. When the manufacturer chooses to have the retailer salvage units, each returned unit has a value s_r, which is captured by the retailer (rather than the manufacturer), and the wholesale price is set accordingly.

It may appear that returns clearly benefit both the manufacturer and retailer if the manufacturer allocates responsibility for salvaging the units to the retailer, because the retailer gets the salvage value s_r from a returned unit and the manufacturer gets to make another sale on exchanged units. However, absent a return penalty, returns create a cost for the retailer. The retailer's profit from selling a unit of product j that is kept by a consumer is $p_j - w_j$. The retailer's profit from selling a good that is exchanged (with a full refund) is $(p_j - w_j) - (p_j - s_r) + (p_{-j} - w_{-j}) = p_{-j} - w_j - w_{-j} + s_r$. Thus, for any $p_{-j} - w_j - w_{-j} + s_r < p_j - w_j$, the retailer would earn less profit from a unit exchanged than a unit kept by the consumer, unless this cost is passed to consumers in the form of a return penalty. When firm choices are symmetric across products, this condition simplifies to $s_r < w_j$. When consumers are charged a return penalty, the manufacturer's initial sales go down (because all consumers' expected utility of making an initial purchase goes down), and thus the manufacturer is not strictly better off if there are more exchanges unless the retailer does not charge a return penalty.

[24] For $r \geq s_r$, $\partial p(w, r \mid r > s_r)/\partial r = (2d - \alpha(w - r + d))/4\alpha d$. Note that at $f(w, r) = w - r$, the probability of exchange from Equation (1) is equal to $\phi_e = \alpha(d - (w - r))/2d$, which is greater than zero iff $d > w - r$. Therefore $\partial p(w, r \mid r > s_r)/\partial r > 0$ for equilibrium values for which the exchange probability is nonnegative.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

Table 3 Equilibrium Under a {Wholesale Price, Refund Rate} Wholesale Contract

Term	Equilibrium value[a]
Retail price	$p^* = \frac{\bar{u}}{2} + \frac{(w^* - \max\{r^*, s_r\})^2}{8d} + \frac{2(w^* + \max\{r^*, s_r\}) - d}{8} - \frac{(1-\alpha)(w^* - \max\{r^*, s_r\})}{2\alpha}$
Return penalty	$f^* = w^* - \max\{r^*, s_r\}$
Quantity sold initially	$q_b^* = \frac{\bar{u}}{2} - \frac{w^* - \max\{r^*, s_r\}}{2\alpha} - \frac{d^2 + 2d(w^* + \max\{r^*, s_r\}) - (w^* - \max\{r^*, s_r\})^2}{8d}$
Exchange probability	$\phi_e^* = \alpha\left(\frac{1}{2} - \frac{(w^* - \max\{r^*, s_r\})}{2d}\right)$
Retailer profit	$\pi_{\text{ret}}^* = \frac{(d^2\alpha - \alpha(w^* - \max\{r^*, s_r\})^2 + 2d((2+\alpha)w^* - (2-\alpha)\max\{r^*, s_r\} - 2\alpha\bar{u}))^2}{64\alpha d^2}$
CRM manufacturer profit ($I_z = 1$)	$\pi_{\text{mfgr}}^{\text{CRM}} = \frac{(\alpha(c-s)^2 + d(2s(2-\alpha) + \alpha(4\bar{u} - d) - 2c(2+\alpha)))^2}{128\alpha d^2}$
CRM wholesale price ($I_z = 1$)	$w^{*,\text{CRM}} = c + \frac{\bar{u}}{2} + \frac{(c-s)^2}{8d} - \frac{2(3c-s)+d}{8} - \frac{(1-\alpha)(c-s)}{2\alpha}$
CRM refund rate ($I_z = 1$)	$r^{*,\text{CRM}} = s + \frac{\bar{u}}{2} + \frac{(c-s)^2}{8d} - \frac{2(3c-s)+d}{8} - \frac{(1-\alpha)(c-s)}{2\alpha}$
CR manufacturer profit ($I_z = 0$)	$\pi_{\text{mfgr}}^{\text{CR}} = \frac{(w^* - c)(\alpha(s_r - w^*) + d(2+\alpha))(\alpha(w^* - \max\{r^*, s_r\})^2 - \alpha d^2)}{16\alpha d^2} + \frac{(w^* - c)(\alpha(s_r - w^*) + (2+\alpha)d)((2-\alpha)s_r + 2\alpha\bar{u} - (2+\alpha)w^*)}{8\alpha d}$
CR wholesale price ($I_z = 0$)[b]	$w^{*,\text{CR}} = \frac{1}{4}\left(c + 3\left(s_r + d + \frac{2d}{\alpha}\right)\right) - \frac{3^{2/3}}{12}\left(\frac{A}{12(B + \sqrt{B^2 - 3\alpha^6 A^3})^{1/3}} + \frac{(B + \sqrt{B^2 - 3\alpha^6 A^3})^{1/3}}{3^{1/3}\alpha^2}\right)$
CR refund rate ($I_z = 0$)	$r^{*,\text{CR}} = 0$

[a]In the equations for retail price, return penalty, quantity sold initially, exchange probability, and retailer profit, $w^* = w^{*,\text{CRM}}$ and $r^* = r^{*,\text{CRM}}$ when the channel structure is CRM; $w^* = w^{*,\text{CR}}$ and $r^* = r^{*,\text{CR}}$ when the channel structure is CR. In the equation for CR manufacturer profit above, $w^* = w^{*,\text{CR}}$ and $r^* = r^{*,\text{CR}}$.

[b]$A \equiv 44d^2 + 4\alpha d(-3c + 11d + 3s_r) + \alpha^2(3(c - s_r)^2 + 19d^2 + 38ds_r - 6cd - 32\bar{u})$ and $B \equiv -9\alpha^3(-2d + a(c - d - s_r))(-12d^2 - 4ad(c + 3d - s_r) + a^2((c-s)^2 - 7d(d + 2s_r) - 2cd + 16\bar{u}))$.

The model is solved recursively; the symmetric equilibrium results are as presented in Table 3.[25]

Proof. See the electronic companion.

Although the reverse channel structure is in itself a choice of the manufacturer, we examine how this choice affects the equilibrium return penalty charged to consumers. The following observation is derived from comparing the equilibrium return penalty f when the manufacturer accepts responsibility for returned units to that when the retailer salvages returns.

Observation 1. Consider a situation in which the manufacturer charges the retailer a per-unit wholesale price but not a fixed fee. If $s > s_r$, then the return penalty f charged to consumers is greater when the retailer salvages returned units than when the manufacturer accepts and salvages returned units from the retailer. There exists a critical salvage value for the manufacturer, $\tilde{s} < s_r$, such that $\tilde{s} < s < s_r$ also implies that the return penalty f charged to consumers is greater when the retailer salvages returned units than when the manufacturer accepts and salvages returned units from the retailer. If $s \leq \tilde{s}$, then the return penalty f charged to consumers is lower when the retailer salvages returned units than when the manufacturer accepts and salvages returned units from the retailer.

Proof. See the electronic companion.

Observation 1 shows that it is possible for the equilibrium return penalty to be *higher* in the reverse channel structure for which the returned units have *greater* value—specifically, it is possible for the retailer to charge a higher return penalty when salvaging returns than if the manufacturer salvages returns, even when the retailer's salvage value exceeds that of the manufacturer. In contrast, previous research (Davis et al. 1995, Matthews and Persico 2005, Shulman et al. 2009) shows the more expected result that sellers with higher salvage values charge lower return penalties. Holding constant the channel member who handles returns, the current model also predicts a negative relationship between salvage value and return penalty. Furthermore, if the manufacturer has a higher salvage value than the retailer ($s > s_r$), this negative relationship between salvage value and return penalty

[25] Note that it is shown in the electronic companion that asymmetric firm choices will not improve profit for either the manufacturer or the retailer.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

also holds in comparison across reverse channel structures. However, the negative relationship between salvage value and return penalty does not always hold when comparing the reverse channel structure if $s < s_r$; a retailer whose salvage value for returned units exceeds that of the manufacturer may actually charge a higher return penalty when it is the channel member responsible for returns processing than when the less efficient manufacturer is responsible for salvaging returned units (i.e., when $\tilde{s} < s < s_r$). To understand the intuition, consider the manufacturer's refund when salvaging returned units. The manufacturer pays more to the retailer for a return than the product's salvage value. Therefore, there is a range of potential salvage values for which the refund offered by the manufacturer to the retailer is greater than the retailer's own salvage value, and the retailer responds by offering a more generous return penalty to consumers. In the following proposition, we describe when the manufacturer would choose to accept and salvage returned units from the retailer and when responsibility for returns will be allocated to the retailer.

Proposition 2. *Consider the situation in which the manufacturer charges the retailer a per-unit wholesale price but not a fixed fee. If $s > s_r$, then the manufacturer accepts and salvages returned units from the retailer. There exists a critical salvage value for the manufacturer, $\bar{s} < s_r$, such that $\bar{s} < s < s_r$ also implies the manufacturer will accept and salvage returned units from the retailer. If $s \leq \bar{s}$, then the manufacturer chooses to allocate responsibility for salvaging returned units to the retailer.*

Proof. See the electronic companion.

As one might expect, Proposition 2 shows that if the manufacturer is more efficient than the retailer in salvaging returned units ($s > s_r$), then the manufacturer would optimally take responsibility for this task.[26] Surprisingly, the manufacturer may find it optimal to salvage returned units even if the retailer can extract greater value from returned units ($s < s_r$), as long as the difference in salvage values is not too great. The manufacturer will choose to have the retailer salvage returned units only if the retailer has a sufficiently significant advantage in salvage value over the manufacturer. The driver of this result is demonstrated by the following proposition comparing the equilibrium outcome with a contract involving a per-unit wholesale price (but no fixed fee) to that of a vertically integrated system.

[26] If the manufacturer were able to resell a returned product from one region into a different market region as a new product, this could be a reason for the manufacturer's salvage value s to be greater than the retailer's salvage value, s_r. We thank the editor for this insight.

Proposition 3. *Consider a situation in which the manufacturer charges the retailer a per-unit wholesale price but not a fixed fee. If the manufacturer's salvage value is greater than the retailer's salvage value ($s \geq s_r$) the manufacturer can replicate the return penalty charged to consumers in a vertically integrated channel by accepting product returns from the retailer. The retail price charged to consumers will be distorted upward and manufacturer profit will be distorted downward from a vertically integrated channel. Otherwise, when $s < s_r$, the return penalty charged to consumers will be greater in a decentralized channel than in a vertically integrated system.*

Proof. See the electronic companion.

Proposition 3 shows that total channel profit with per-unit wholesale pricing to the retailer (but no fixed fee) is less than in a vertically integrated system achievable with the addition of a fixed fee to the retail contract. Although the per-unit wholesale pricing contract results in double marginalization on initial sales ($p^* > p^{\mathrm{VI}}$), the manufacturer may eliminate double marginalization on returns ($f^* = f^{\mathrm{VI}}$) by accepting returned units when the manufacturer's salvage value is the same as in the vertically integrated channel ($s > s_r$). The manufacturer essentially polices itself and the system by choosing w and r so as not to distort the retailer's return penalty choice; specifically, when the manufacturer buys back the returned units, the manufacturer refunds to the retailer the margin earned on new units $w^* - c$ in addition to its own salvage value for the returned unit s. That is, $r^* = (w^* - c) + s$. This refund level induces the retailer to set its return penalty so that the *return rate* is the same as in the vertically integrated channel. The higher price dampens initial sales ($q_b^* < q_b^{\mathrm{VI}}$), and therefore fewer total unit returns occur than in the vertically integrated system ($\phi_e^* q_b^* < \phi_e^{\mathrm{VI}} q_b^{\mathrm{VI}}$).

Therefore, the apparently counterintuitive result of Proposition 2—in which the manufacturer may choose to accept returns even if its salvage value is lower than the retailers'—stems from the negative externality created by the retailer's return policy. A restrictive return penalty charged by the retailer provides revenue to the retailer but adversely affects the total number of units sold. When the retailer salvages returns, the equilibrium values of p, f, and w are distorted relative to the vertically integrated channel. The manufacturer can internalize this externality by accepting responsibility for returns through its joint use of wholesale price and a generous refund rate for returned units. Of course, it could still be more efficient to allocate the salvaging task to the retailer but only if the retailer's salvage value superiority is great enough. If the retailer's salvage advantage is not too great, the gain resulting from the manufacturer's ability to strategically set the wholesale price and refund can more than compensate

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

for the efficiency loss resulting from moving the function to the manufacturer. The manufacturer optimally encourages the retailer to salvage returned units only if the retailer has a significant advantage in its ability to extract value from returned units.[27]

In summary, when fixed fees or quantity discount schedules are not used in the wholesale contract, we have defined conditions under which the manufacturer may better approximate the outcome of a coordinated channel by choosing a reverse channel structure in which it takes over ultimate responsibility for processing consumer returns rather than leaving this responsibility to the retailer. The reason for this is that the gain in coordination of channel decision making can outweigh the efficiency loss that can result from the manufacturer's performance of this important channel function. Our results show that taking back the returns-processing function and combining this channel *structure* decision with an appropriate channel *contracting* decision is a partial substitute for using a fixed fee or equivalent quantity discount schedule.

4. Discussion

Product returns represent an enormous cost, greatly affecting manufacturers and retailers alike. Learning to manage product returns successfully can greatly increase a company's profitability. Our research examines consumer returns of nondefective products that are discovered by consumers not to match with preferences. These returns are different from overstock returns in that the return policy is interrelated with product demand itself. The return penalty charged to consumers has several effects: it recoups costs associated with returns from consumers and it prevents a number of returns from occurring, but it also reduces consumers' willingness to pay for initial purchases because of uncertainty about whether or not the product matches with preferences. In a channel context, the choice of return penalty adds an additional layer of complexity because it may asymmetrically affect the retailer and the manufacturer. The retail price and return penalty charged to consumers are set to maximize retailer profit, but the return penalty shifts the demand curve as well, thereby affecting the manufacturer's quantity sold.

One of our key findings is that even if the retailer is more efficient than the manufacturer at salvaging returned units, it is possible for the retailer to charge a higher return penalty when salvaging returned units than if the salvaging were done by the manufacturer. When the retailer salvages returned units without the help of the manufacturer, the retailer misses out on the generous refund from the manufacturer. Although a greater salvage value generally implies a more generous return policy, this effect may be reversed by the manufacturer's incentive to coordinate the pricing in the reverse channel.

Intuition would suggest that the salvaging of returned units should be done as efficiently as possible, but our analysis shows in contrast that the manufacturer may earn greater profit by taking back product returns even if the retailer is more efficient at salvaging the units. The ensuing loss in efficiency can be more than compensated for by the gain in profitability as a result of manufacturer's internalization of the negative externality that is otherwise created when the retailer handles returns. A retailer that handles returns penalizes consumers for returning products more heavily than the optimal penalty the manufacturer would have chosen; this can reduce the manufacturer's profit, even if the retailer has the advantage in terms of salvage value.

If the manufacturer is able to "sell the firm" to the retailer through the use of a fixed fee or perfectly developed quantity discount schedule, then the manufacturer will be able to extract all profit from the retailer regardless of reverse channel structure. The subsequent fully-coordinated channel will charge a return penalty to consumers exactly equal to the cost to the channel of a product return. However, it may not always be feasible for the manufacturer to extract all profit through the use of a fixed fee. If fixed fees (or the corresponding quantity discount schedule) cannot be implemented, the manufacturer can use a return policy to approximate the results of a vertically integrated channel. When offering per-unit wholesale and refund rates, the manufacturer offers a generous refund above its own value for the returned unit to induce the retailer to set the optimal return penalty. Although the return penalty of the vertically integrated channel can be replicated when the manufacturer accepts returns, double marginalization in pricing persists, reducing not only total quantity sold but also total product returns and channel profit relative to a vertically integrated channel outcome.

Our model suggests when each reverse channel structure is likely to be observed in equilibrium. In the paper's electronic companion, we provide some illustrative examples to show that across a broad spectrum of parameter space, the manufacturer is more likely handle returns itself (i.e., doing so earns greater profit than having the retailer salvage returns) when consumer disutility for a mismatch (d) or retailer salvage

[27] Iyer and Villas-Boas (2003) find that the manufacturer may increase profit by taking back overstock returns because it reduces extreme double marginalization in price. From Propositions 2 and 3, we see that taking back consumer returns may also be profitable for the manufacturer through a related mechanism: the reduced marginalization on product returns (i.e., the reduced return penalty).

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

value (s_r) are lower, or when consumer utility from owning a product in the category ($\bar{u}$), marginal production cost (c), or manufacturer salvage value (s) are higher. Otherwise, the manufacturer will optimally decline returns and allocate that salvage responsibility to the retailer.

Several facets of our model merit further discussion. These findings were developed in a model of heterogeneous consumers (in terms of their reservation utility u_i) who are completely uninformed a priori about their preferences between two products. Alternatively, one could model a situation where some consumers are more informed than others or some have a stronger preference for one of the products than for the other. When such asymmetries are introduced into our modeling framework, one would expect to see fewer returns. Furthermore, had we instead assumed that a single product was offered to consumers who are homogeneous and uninformed before purchase (but vary in their valuation discovered after purchase), then the results would be different: namely, the return penalty would be lower (and the manufacturer would earn the greatest profit) in the channel structure that allocates returns to the channel member with the greatest salvage value.[28] Our results, on the other hand, are preserved when consumers are heterogeneous in their prepurchase valuation (u_i) or when consumers can choose between two products. In either of these cases, the retailer's choice of return penalty has an asymmetric impact on the retailer and the manufacturer. We have chosen to incorporate both complexities because this not only more accurately depicts reality, but it also generates a more parsimonious model. In the electronic companion, we show how the qualitative results of Observation 1 and Proposition 2 are preserved in a single-product setting.

Our model also treats gains equivalently to losses in the eyes of the consumer. However, prospect theory suggests that losses (i.e., return penalties) are weighted more heavily than gains (i.e., the amount of consumption utility greater than the price). When prospect theory is captured by a loss aversion parameter t such that any payment of a return penalty, f, results in a loss in utility tf (i.e., instead of experiencing loss $-f$, consumers experience loss of $-tf$ with $t \geq 1$), the qualitative results of Propositions 1–3 are preserved, although the specific price and return penalty are affected.[29]

In reality, there are additional complexities that suggest avenues for future research in this area. For example, consumers may vary in their return probabilities. This variance can be due to variation in their initial level of information or due to the possibility that a return will be made without a subsequent purchase. With either of these modifications, however, the same forces are at play as in our model. A return penalty would still reduce returns, pass the costs of returns on to consumers, and have a negative effect on demand. Without a manufacturer's return policy, the retailer's incentives to reduce costs associated with returns would still ignore the resulting impact on the manufacturer's profit from initial sales. It would logically follow that in this more complex setting, the manufacturer would still offer a generous return policy that has the same effect as in Observation 1 and Proposition 2. The additional complexities may lead to distortion of the return penalty in a CRM reverse channel structure relative to the solution in a vertically integrated system. However, it logically follows that the return penalty in the CRM reverse channel structure would be closer to that in the vertically integrated channel than is the penalty in the CR reverse channel structure.

The model abstracts away from the possibility that there may be loss of goodwill, future sales, or accessorial sales that a retailer incurs by charging a consumer a return penalty. In a companion model available from the authors, we capture a per-unit-returned loss in retailer profit as a function of the return penalty. Allowing for this possibility changes the specific prices and return penalties chosen in equilibrium. It yields the additional result that the return penalty in both the vertically integrated structure and the structure where the manufacturer handles returns can actually be higher when the retailer's future or accessorial profit is diminished by the return penalty (if and only if the disutility of mismatch is sufficiently low). However, our qualitative results still hold.

Another model extension could consider the possibility of diminishing salvage value as a function of the number of units returned, perhaps because of capacity constraints or marketplace forces. Adding this to our model would increase the strictness of the retailer's return policy to control the supply of returned units. In turn, this would decrease the manufacturer's profits because a stricter return policy would lower the consumer's initial willingness to pay for a product and therefore decrease final demand.[30]

There are also other contracting structures that could be examined. For example, a manufacturer could seek to alter the retailer's incentive to set a high return penalty by making the per-unit wholesale price an increasing function of the return penalty

[28] This can easily be shown by extending the model of Matthews and Persico (2005).

[29] It is straightforward to follow the analysis contained in this paper with the augmented demand equation that results from an augmented expected utility function.

[30] We thank the editor for this insight.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.

the retailer charges.[31] Although we could not find instances of the use of such contingent contracts in real-world channels with product returns, this or various other contracting options could be examined in future research. Moreover, one could examine a dynamic model in which the number of returned units impacts the manufacturer–retailer negotiations in the following period. In a dynamic model, the retailer may use the threat of salvaging returned units without the manufacturer (when the manufacturer earns greater profit from accepting returns from the retailer) as a means of extracting more favorable terms from the manufacturer. A careful analysis of long-term payoffs is necessary to determine whether such holdup is an equilibrium strategy.

Future research could also consider the possibility that the retailer itself could run a resale market for returned and/or used goods that compete directly with the market for new goods. This extension applies to situations in which the retailer has the both the legal right and the operational capability to retest and refurbish a returned product so that it is again sellable. The specific equilibrium return policies in such an extension will depend on both the level of differentiation between returned and new units as well as the consumer sensitivity to these differences. Our research implicitly assumes that there are distinct, separable consumer segments (either geographically or in quality preferences) such that any resale of returned units is operated through an alternative channel that does not compete with the primary, new-unit channel. When the retailer's salvaging of returned units via resale instead cannibalizes the manufacturer's new-unit sales, there is an additional incentive for the manufacturer to accept returns from the retailer, even when the retailer has the salvage advantage, relative to our model.

In sum, this paper adds to our understanding of optimal product returns management in a marketing channel. The reverse channel structure has a substantive impact on the price and return penalty ultimately charged to consumers. In addition to cost structure considerations, the strategic interaction between the manufacturer and the retailer plays an important role in determining how consumer returns should be managed. Manufacturers cannot ignore the consumer returns facing a retailer that occur through no fault of the manufacturer. Our model highlights the fact that the reverse channel and the forward channel are closely related and both should be carefully considered in determining the firm's profit-maximizing channel strategy.

[31] We thank the area editor for suggesting this possibility to us.

5. Electronic Companion

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

Acknowledgments

The authors thank Eric T. Bradlow, the associate editor, two anonymous reviewers, J. Miguel Villas-Boas, Fabio Caldieraro, Gary Erickson, Daniel Putler, Beril Toktay, and Charles Weinberg for helpful comments on an earlier version of this paper.

References

Anderson, E. T., K. Hansen, D. Simester. 2009. The option value of returns: Theory and empirical evidence. *Marketing Sci.* **28**(3) 405–423.

Arya, A., B. Mittendorf. 2004. Using return policies to elicit retailer information. *RAND J. Econom.* **35**(3) 617–630.

Cachon, G. P. 2003. Supply chain coordination with contracts. S. C. Graves, A. G. de Kok, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management*. North Holland, Amsterdam, 229–340.

Chao, G. H., S. M. R. Iravani, R. C. Savaskan. 2009. Quality improvement incentives and product recall cost sharing contracts. *Management Sci.* **55**(7) 1122–1138.

Che, Y.-K. 1996. Consumer return policies for experience goods. *J. Indust. Econom.* **44**(1) 17–24.

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

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

Davis, S., M. Hagerty, E. Gerstner. 1998. Return policies and the optimal level of "hassle." *J. Econom. Bus.* **50**(5) 445–460.

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

He, C., J. Marklund, T. Vossen. 2008. Vertical information sharing in a volatile market. *Marketing Sci.* **27**(3) 513–530.

Iyer, G., J. M. Villas-Boas. 2003. A bargaining theory of distribution channels. *J. Marketing Res.* **40**(1) 80–100.

Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* **47**(2) 263–291.

Kandel, E. 1996. The right to return. *J. Law Econom.* **39**(1) 329–356.

Lawton, C. 2008. The war on returns. *Wall Street Journal* (May 8) D1.

Matthews, S. A., N. Persico. 2005. Information acquisition and the excess refund puzzle. Working paper, University of Pennsylvania, Philadelphia.

Middleton, D. 2007. Stores cracking down on returns; Target is one of many retailers tightening policies and implementing restocking fees. *Florida Times-Union* (May 19) A-1.

Nelson, P. 1970. Information and consumer behavior. *J. Political Econom.* **78**(2) 311–329.

Ofek, E., Z. Katona, M. Sarvary. 2010. "Bricks and clicks": The impact of product returns on the strategies of multichannel retailers. *Marketing Sci.*, ePub ahead of print July 29, http://mktsci.journal.informs.org/cgi/content/abstract/mksc.1100.0588v1.

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 retail 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.

Rogers, D. S., R. S. Tibbens-Lembke. 1998. *Going Backwards: Reverse Logistics Trends and Practices*. Reverse Logistics Executive Council, Reno, NV.

Shame On You: Victoria's Secret, Saks and Macy's. 2003. First broadcast November 19, 2003 by WCBS (New York). Produced by Arnold Diaz.

Shulman, J. D., A. T. Coughlan, R. C. Savaskan. 2009. Optimal restocking fees and information provision in an integrated demand-supply model of product returns. *Manufacturing Service Oper. Management* **11**(4) 577–594.

Strauss, M. 2007. Report on business: Canadian; inventory management. *The Globe and Mail* (January 15) B1.

INFORMS holds **copyright** to this article and distributed this copy as a courtesy to the author(s).
Additional information, including rights and permission policies, is available at http://journals.informs.org/.