

Supply Chain Coordination for False Failure Returns

Mark Ferguson
The College of Management
Georgia Institute of Technology
800 West Peachtree Street
Atlanta, GA 30332-0520
Tel: (404) 894-4330
Fax: (404) 894-6030
Email: Mark.Ferguson@mgt.gatech.edu

V. Daniel R. Guide, Jr.
Dept. of Supply Chain & Information Systems
Smeal College of Business
The Pennsylvania State University
509 Business Administration Bldg
University Park, PA 16802
Tel: (814) 865-6103
Fax: (814) 863-7067
Email: dguide@psu.edu

Gilvan C. Souza
Robert H. Smith School of Business
University of Maryland
College Park, MD 20742
Tel: (301) 405-0628
Fax: (301) 405-8655
Email: gsouza@umd.edu

February 22, 2005

Supply chain coordination for false failure returns

ABSTRACT

False failure returns are products that are returned by consumers to retailers with no functional or cosmetic defect. The cost of a false failure return includes the processing actions of testing, refurbishing if necessary, repackaging, the loss in value during the time the product spends in the reverse supply chain (a time that can exceed several months for many firms), and the loss in revenue because the product is sold at a discounted price. This cost is significant, and is incurred primarily by the manufacturer. Reducing false failure returns, however, requires effort primarily by the retailer, for example informing consumers about the exact product that best fits their needs. We address the problem of reducing false failure returns via supply chain coordination methods. Specifically, we propose a target rebate contract that pays the retailer a specific dollar amount per each unit of false failure returns below a target. This target rebate provides an incentive to the retailer to increase her effort, thus decreasing the number of false failures and (potentially) increasing net sales. We show that this contract is Pareto-improving in the majority of cases. Our results also indicate that the profit improvement to both parties, and the supply chain, is substantial.

1. Introduction

Product returns represent a growing financial concern for firms in the United States and the rest of the world, with recent estimates reaching \$100 billion annually for the United States alone (Stock et al. 2002). Product returns are a result of two phenomena: consumer returns of products to the retailer during a 30, 60 or even 90 day return period, and product overstocks returned to the manufacturer by the retailer. *Consumer* product returns can occur at any time during the product life cycle, and are increasingly important to manufacturers. Hewlett–Packard (HP) recently discovered that the total costs of consumer product returns for North America exceed 2% of total outbound revenue (Reiss 2003).

Product overstocks are the subject of a large body of research, but are not directly related to the problem of consumer product returns since overstocks encompass units that were never sold to the final customer and are only returned at the end of the product lifecycle. Consumer product returns are driven by the ‘consumer is king’ attitude prevalent in the United States and supported by liberal product returns policies at most major retailers. Consumer product returns to the retailer are far less common in the European Union and the rest of the world, but many countries mandate some form of return period for Internet and catalog sales. The problems and costs of consumer product returns are projected to grow and many firms are just beginning to form teams to develop strategies and tactics to reduce the overall costs (Reiss 2003).

Consumer return rates range from 5-9 percent of sales for most retailers and up to 35 percent for fashion apparel (Toktay 2003). A percentage of these returns occur due to true product failure, however, a large percentage of returns have no verifiable functional defect. We refer to this class of consumer product returns as *false failure returns*; returns that have no functional or cosmetic defect. Managers cite a number of reasons why false failure returns occur, including:

installation difficulties, product performance incompatibility with consumer preferences, and remorse (Kumar, Guide and Van Wassenhove 2002). For HP's inkjet printer group, false failure returns can account for up to 80% of their inkjet printer returns (Davey 2001). Since HP's total consumer product returns average slightly higher than 6% of sales, false failure returns average approximately 5% of sales. As of 1999, HP's inkjet printer division handled over 50,000 consumer returns per month in North America (Davey 2001). The problem also persists outside of the high-tech industry. At Bosch Power Tools North America, false failure returns account for 2% of sales (Valenta 2002). In the United Kingdom, manufacturers are seeing an increasing number of consumer returns to retailers disguised as 'defective' products (Helbig 2002). Because of the significant financial impact, manufacturers are interested in reducing false failure returns through improved relations and contracts with retailers.

Not all manufacturers require that retailers return any products returned by consumers. HP and Bosch are highly brand name conscious and have a policy that a product returned for any reason must be returned to their product returns centers. A manufacturer may allow a retailer to sell a returned item if there is nothing wrong with it. However, the determination of whether or not a product is defective is not always clear, and this practice allows gaming by consumers. Best Buy Co. has publicly announced that it will seek to discourage 'devil' customers, customers who frequently return products and then buy them back at open box discounts, from shopping at the store (McWilliams 2004). Other retailers are also getting aggressive about identifying customers with excessive returns by tracking all customer returns through a centralized database and refusing returns after a dollar threshold is reached (Cha 2004). These aggressive practices are controversial with consumers and a sign that retailers are becoming concerned with the costs associated with product returns.

When a manufacturer receives a false failure return, the product is routed through the firm's reverse supply chain involving several testing steps and repackaging before the product can become available for sale at a secondary market, typically at a price discount. The cost of a false failure return includes the processing actions of testing, refurbishing if necessary¹, repackaging, the loss in value during the time the product spends in the reverse supply chain (a time that can exceed several months for many firms), and the loss in revenue because the product is sold at a discounted price. Thus, the cost of false failure returns is significant. The per-unit cost of a false failure return for computer manufacturers, including HP, is around 25% of the product's price.

Table 1 shows information provided by HP about why consumers return products and what actions can be taken to eliminate the cause. The manufacturer can make design changes, but only in the long run (i.e., the next product generation). For the retailer, however, there are a number of actions that may be undertaken in the near term. A retailer may spend extra time with customers and listen to their needs before recommending a particular product. By doing so, customers have a higher probability of purchasing a product that matches their needs the first time. Retailers may also train their sales force so that the proper operating procedures of a product are clearly explained to customers upon purchase. After-sales support by the retailer may also reduce the number of false failure returns from customers who have trouble configuring the new product so that it performs as expected. We focus on the retailer in our research since the manufacturer is already motivated to improve the product design with each generation, but recognizes that design changes will only solve a portion of the problem. Providing an incentive for the retailer to reduce returns provides a fast solution to a sizeable percentage of the reasons products are returned.

¹ While false failure returns still perform functionally, they require cleaning and sometimes updating (e.g., software driver updates) to bring them up to the specifications of the new products being sold when they will reenter the marketplace.

Table 1- Reasons products are returned

Rationale for Return	Percentage	Solutions
True failure	20	Design
Install/Basic Use	27.5	Design, After Sales Support
Performance		Design, Retail Pre-Sales Qualification
Print Speed	25	
Print Quality		
Packaging	2.25	Design, Retail Pre-Sales Qualification
Sales Technique	12.75	Retail Pre-Sales Qualification
Consumer Behavior	12.5	Retail Policy

HP, along with most of their direct competitors, offers a customer returns policy to the retailer where the retailer receives full credit at the unit wholesale price. While the retailer incurs minimal processing and loss of goodwill costs, the manufacturer absorbs the majority of the false failure return costs. In turn, the manufacturer receives the majority of the benefits from reducing false failure returns. The short-term cost to reduce the number of false failure returns, however, may be incurred primarily by the retailer as explained above. Since the retailer incurs all effort-related costs but, not as many benefits from reducing the number of false failure returns, there is a need for contracts that coordinate the supply chain such that interests are aligned.

To address this coordination problem, we propose a *target rebate* contract that pays the retailer a specific dollar amount per each unit of false failure returns below a target T . This target rebate provides an incentive to the retailer to increase her effort of informing consumers about the exact product that best fits their needs, thus decreasing the number of false failure returns. We show that this contract is Pareto-improving in the majority of cases. Additionally, we show that increased retailer's effort produces a magnitude of expected profit improvement per expected return of 31% for the supply chain, where profit is defined as revenues minus the costs associated with false failure returns. Profit improvements for the manufacturer and the retailer depend on the uncertainty in the distribution of false failure returns process.

This paper is organized as follows. In §2, we review the relevant literature. We define our model in §3 and present and discuss a numerical study in §4. In §5, we give examples where

traditional forward contracts may lead to gaming by the retailer that increases the number of false failures. We also offer a forward contract design that eliminates the misalignment in incentives for both the double marginalization problem and the effort to reduce false failure returns. We conclude in §6.

2. Literature review

Our research draws on two separate research streams; closed-loop supply chains and supply chain coordination contracts. In this section we provide an overview of recent work in each area and examine the implications for our research.

Closed-loop supply chains

There is a growing body of literature on closed-loop supply chains, where both the forward and reverse flows of materials are considered. For example, a recent feature issue of *Interfaces* (33(6) 2003) focuses on the practice of closed-loop supply chains in a variety of industry settings. Guide and Van Wassenhove (2003) identify the common processes required by a closed-loop supply chain: product acquisition, reverse logistics, inspection, testing and disposition, remanufacturing, and selling and distribution.

General overviews of reverse logistics and remanufacturing are presented by Fleischmann (2001), and Guide (2000). We also refer the reader to the book, edited by Guide and Van Wassenhove (2003), from the First Workshop on Business Aspects of Closed-Loop Supply Chains. Historically, the operational aspects of remanufacturing have received the most attention, with numerous publications dealing with production planning and control (e.g., Ferrer and Ketzenberg 2004; Souza et al. 2002), and inventory control (e.g., Toktay et al. 2000; van der Laan et al. 1999). Ironically, product recovery is often viewed as a narrowly focused, technical operational problem without visibility at the corporate level. At Hewlett-Packard Company,

customer returns were treated as a low-level divisional problem until a thorough analysis showed that the total cost of product returns was equivalent to 2 percent of total outbound sales (Davey 2001). Unfortunately, academic research often tends to reinforce this limited view with its narrow focus on local optimization of operational issues.

More recent research efforts, however, have started to analyze the strategic interaction among the closed-loop supply chain players. Heese et al. (2003) consider the potential competitive advantage for a company actively engaging in product take-back via a game-theoretic model. Majumder and Groenevelt (2001) recommend incentives to the original equipment manufacturer to increase the fraction of remanufacturable products available, or to decrease the costs of remanufacturing. Savaskan et al. (2004) develop a game-theoretic model that addresses the issues of channel choice and coordination of the channel. Debo et al. (2001) investigate the joint pricing and production technology problem of a manufacturer that offers both new and remanufactured products. Ferguson and Toktay (2004) explore how remanufacturing practices can be used strategically to deter competition.

Most of the research focus on product returns to date is on end-of-life product returns or overstocks. Our focus, however is on consumer product returns, an area with little academic research; one exception is the recent work by Souza et al. (2004). They argue that returns are often time sensitive and firms frequently lose much of the value remaining in their returned products by not making quick disposition decisions. Their research focuses on the appropriate reverse supply chain design, responsive or efficient, based on the rate of value decay. Our research is complimentary. Preventing false failure returns increases the revenue from lost sales while reducing the unnecessary expenses of product return, inspection, disposition and distribution for reuse.

Supply chain contracting and returns policies

The pioneer work of Pasternack (1985) shows that if a manufacturer sells products to a retailer under demand uncertainty with a constant wholesale price per unit (higher than the manufacturer's cost but lower than the retail price) the retailer usually stocks less than the supply chain's optimal quantity. From a marketing perspective, Wood (2001) examines how returns policies affect consumer purchase probability and return rates. Wood shows that more lenient policies tend to increase product returns, but that the increase in sales is sufficient to create a positive net sales effect. Other research focuses on the problem of how to set a returns policy between a manufacturer and a retailer and the use of incentives to control the returns flow (Padmanabhan and Png 1997 1995, Pasternack 1985, Davis et al. 1995, Tsay 2001). Choi et al. (2004) study the effect of an e-marketplace on a returns policy in which internet auctions are used to recover value from the stream of product returns.

A considerable amount of research has been devoted to this problem with a focus on appropriate contracts—a transfer payment from the manufacturer to the retailer, which provides incentives for retailers to stock a level of inventory that is optimal for the entire supply chain. Cachon (2003) provides a thorough review on these type contracts. Contracts based on rebates are of particular importance to our research. Based on agency theory, the provision of incentives for managerial effort is necessary when effort is not directly observable. Taylor (2002) offers an application of channel rebates to induce forward supply chain coordination when demand is dependent on the sales effort of the retailer. We propose a similar type contract to induce the supply chain optimal amount of effort to reduce the number of false failure returns; we choose a rebate-type contract because the number of false failure returns per period is dependent on retailer effort *and* random; as a result retailer effort is unobservable by the manufacturer.

There is also a well-established body of research on the use of sales quotas to motivate sales performance. In this research, as in ours, direct effort on the part of the salesperson is not directly

observable (Davis and Farley 1971, Tapiero and Farley 1975, Darmon 1979). Much of the research is concerned with determining the residual effect of past selling effort and estimating the sales-effort-effectiveness functions. There are controlled experiments which show that (1) higher quotas lead to higher levels of effort on the part of sales people (Winer 1973), and (2) a quota set too high leads to a decrease in effort on the part of sales people (Chowdhury 1993). Gaba and Kalra (1999) also shows that with a high quota, sales people opted for more high risk prospects. Other research by Raju and Srinivasan (1996) shows that a basic quota plan performed with little loss in optimality even when there is strong heterogeneity across salespersons and territories. A principle-agent modeling approach is used by both Mantrala et al. (1994) and Mantrala et al. (1997) to determine each salesperson's preferences for income and leisure to design an incentive compatible plan. Both studies show that quotas can be used to effectively motivate salespeople even when direct effort is unobservable. These studies lend support to our contention that our model will motivate retailers to reduce returns, even if their effort is not directly observable.

In the next section, we introduce a mathematical model of the false failure returns coordinating process, and suggest a target rebate contract to coordinate the supply chain.

3. Model

Consider a supply chain comprised of one manufacturer and one retailer. We present a single-period model for ease of exposition, although the model can be easily extended to multiple periods in a stationary setting because there is no dynamic linkage (e.g., beginning inventory) between periods. That is, if all parameters are stationary over an infinite horizon, and the firm maximizes average profit, then the optimal level of effort—our decision variable—is the same across periods. A period, in our context, can be thought of as a quarter.

When a consumer returns a false failure, there are three possible outcomes:

1. The consumer exchanges the returned product with a product from the same manufacturer, which could be of the same or a different model. If the consumer chooses a different model, then we assume, for ease of presentation, that it has an identical price p , wholesale price w , and manufacturing cost c as the returned product's model.
2. The consumer exchanges the returned product with a product from a competitor, with price p_c and wholesale price w_c . The consumer is paid $p - p_c$ if $p_c \leq p$ or pays $p_c - p$ otherwise.
3. The consumer returns the product, is paid p , and walks away with no product.

In each of the three cases, the retailer receives full credit w for the false failure return from the manufacturer, as discussed in the introduction. False failure returns cost the manufacturer and retailer m and r per unit, respectively, including any goodwill cost. We now compute the profit, to the manufacturer and retailer, of reducing *one* false failure return for each of the possible outcomes above. This is shown in Table 2 below, which we explain as follows. In the first outcome, both the retailer and the manufacturer experience a net sale (total sales minus total returns) of one product plus the respective return processing cost. If the false failure return had been avoided, both the manufacturer and retailer would also have experienced a net sale of one product. Consequently, the only profit to each party of avoiding a first outcome's false failure return is its respective return processing cost.

In the second outcome, each party incurs its respective return processing cost, the manufacturer experiences a zero net sale and the retailer experiences a net sale of one competitive product. Consequently, avoiding a second outcome's false failure return implies a profit to the manufacturer of $m + (w - c)$, the sum of its return processing cost plus the margin of the product that the consumer keeps. Avoiding a second outcome's false failure return implies a profit to the retailer of $r + (p - w) - (p_c - w_c)$, the sum of its return processing cost plus the net

margin of the product that the consumer keeps minus the margin of the competitive product, which the retailer would have sold in the second outcome. A similar reasoning holds for the third outcome.

Table 2: Profit of avoiding one false failure return

Outcome	Probability	Profit	
		Manufacturer	Retailer
1: exchange with product from same manufacturer	q_1	m	r
2: exchange with a competitive product	q_2	$m + (w - c)$	$r + (p - w) - (p_c - w_c)$
3: consumer walks away	q_3	$m + (w - c)$	$r + p - w$

From Table 2, we compute the *expected profit* of avoiding *one* false failure return to be $m + (q_2 + q_3)(w - c)$ and $r + (q_2 + q_3)(p - w) - q_2(p_c - w_c)$ for the manufacture and retailer, respectively. Assume for simplicity that the competitive product has the same margin to the retailer, that is, $(p - w) = (p_c - w_c)$. Additionally, define $\delta_m \equiv q_2 + q_3$, and $\delta_r \equiv q_3$, where $0 < \delta_r < \delta_m < 1$; δ_m and δ_r can be interpreted as the sales impact to the manufacture and retailer, respectively, of avoiding one false failure return. The expected profit of avoiding one false failure return is then $m + \delta_m(w - c)$ and $r + \delta_r(p - w)$ for the manufacture and retailer, respectively.

Let ρ , $\rho \geq 1$, denote the retailer's effort to reduce false failure returns. Effort may be thought of as the time the retailer spends working with the customer to choose the right product and demonstrating how to use it. We assume that the retailer's cost of exerting effort is convex in her amount of effort ρ . A simple illustrative case is $a\rho^2/2$, where $a > 0$ can be interpreted as the marginal cost of effort associated with the minimum effort level $\rho = 1$.

Given an effort ρ by the retailer, the number of false failure returns is $X(\rho) \geq 0$, a random variable with cumulative distribution function (cdf) and probability density function (pdf) $F(x|\rho)$ and $f(x|\rho)$ respectively. We assume that F is differentiable and strictly increasing in

x. Further, we assume that $E\{X(\rho)\} = \frac{\beta}{\rho}$, where β is the expected number of false failure returns when the retailer exerts the minimum effort level $\rho = 1$. This assumption models diminishing returns to scale in the retailer's effort, similarly to more commonly studied sales effort models. Thus, the *expected* value of the reduction in false failure return from the minimum effort level $\rho = 1$, for any effort level ρ is $\beta - \frac{\beta}{\rho} = \beta(1 - \frac{1}{\rho})$.

We consider here only the costs and revenues impacted by the retailer's effort to reduce false failure returns. As a result, the coordinated supply chain's expected profit $\Pi(\rho)$ as a function of the retailer's effort is

$$\Pi(\rho) = (m + \delta_m(w - c) + r + \delta_r(p - w))\beta \left(1 - \frac{1}{\rho}\right) - \frac{1}{2}a\rho^2. \quad (1)$$

The first term in (1) is the total expected profit as a result of the reduced number of false failure return for an effort level ρ ; the second term is the retailer's cost of effort. The expected profit (1) is concave in ρ , and is maximized at

$$\rho^c = \left[\frac{m + \delta_m(w - c) + r + \delta_r(p - w)}{a} \beta \right]^{1/3}, \quad (2)$$

where ρ^c denotes the optimal retailer effort in a coordinated supply chain. We assume the interesting case where $(m + r)\beta > a$, such that the total potential cost saving of reducing false failure returns is greater than the retailer's marginal cost a of exerting the minimum effort level. As a result, $\rho^c > 1$ (the coordinated supply chain exerts more than the minimum effort).

In a non-coordinated supply chain, the retailer incurs the cost for reducing false failure returns but is only rewarded with part of the benefit. Thus, it seems unlikely that the optimal level of effort will be observed in the absence of some form of coordination contract. Specifically, the retailer's expected profit $\pi_R(\rho)$ is

$$\pi_R(\rho) = -\frac{a\rho^2}{2} + [r + \delta_r(p - w)]\beta \left(1 - \frac{1}{\rho}\right). \quad (3)$$

In the absence of an incentive by the manufacturer, the retailer exerts the optimal effort level

$$\rho^D = \max \left\{ \left([r + \delta_r (p - w)] \beta / a \right)^{1/3}, 1 \right\}. \quad (4)$$

Clearly, $\rho^C \geq \rho^D$. The manufacturer's expected profit in this decentralized case is

$$\pi_M(\rho^D) = [m + \delta_m (w - c)] \beta \left(1 - \frac{1}{\rho^D} \right). \quad (5)$$

Since the manufacturer enjoys most of the benefits from reducing the number of false failure returns, the manufacturer must share some of these benefits with the retailer in order to achieve a coordinated solution. The manufacturer cannot mandate a given level of effort since effort is not observable due to the randomness in the false failure returns process.

Given that effort is not directly contractible, a second option is for the manufacturer to either not pay the retailer a full credit for returns or charge the retailer a penalty for each false failure return. Not providing a full credit for a false failure return is problematic because false failures are included in the defective product returns and the retailer should not be held responsible for the manufacturer's poor quality. Since the product must go through the reverse logistics process, a significant amount of time may expire before the manufacturer knows how many units from a given return are false failures. Even if we ignore the timing issues, there are other complications with charging the retailer a penalty for each false failure return. Charging a penalty to the retailer for each false failure penalizes the retailer for some effects that are not under the retailer's control. Examples include when the packaging of the product misleads the customer or when the customer incurs buyer's remorse despite the fact that the retailer did everything possible to ensure the customer purchased the correct product. For these reasons, it is unlikely that the retailers would agree to pay a portion of the cost for each false failure.

As a third option, the manufacturer may contract on the number of false failure returns through the use of a target rebate, and thus coordinate the supply chain. In this contract, the retailer receives a reward $\$u$ for every false failure return below a predetermined target level T . A low number of false failure returns is a signal of good effort on the retailer's part to help the

customer purchase the correct product and its operation. A target rebate creates an incentive for the retailer to exert effort in this respect. The contract is based on false failure returns, and not total consumer returns (which includes true failures) because as argued before, the retailer should not be penalized by a manufacturer's poor quality. Although the breakdown of returns (true vs. false failures) is provided by the manufacturer after the testing stage, we believe that it is not in the manufacturer's best interest to over-report the number of false failures because the retailer always has the option of verifying the information, albeit at an additional cost. HP has experimented with different forms of target rebate contracts to reduce the level of false failure returns (Reiss 2003).

Under a target rebate contract, the retailer's expected profit is

$$\pi_R(\rho | T, u) = uE_X \left\{ [T - X(\rho)]^+ \right\} - \frac{a\rho^2}{2} + [r + \delta_r(p - w)]\beta \left(1 - \frac{1}{\rho} \right), \quad (6)$$

and the manufacturer's expected profit is

$$\pi_M(\rho | T, u) = (m + \delta_m(w - c))\beta \left(1 - \frac{1}{\rho} \right) - uE_X \left\{ [T - X(\rho)]^+ \right\}. \quad (7)$$

We find the values of u and T that coordinate this supply chain by equating to (2) the value of ρ that maximizes (6), which can be obtained via first order conditions. First order conditions are sufficient for optimality provided that (6) is concave in ρ ; a sufficient condition is given below.

Proposition 1: If $\partial^2 F(x | \rho) / \partial \rho^2 \leq 0$, then the retailer's expected profit (6) under a target rebate contract is concave in ρ .

Proof: We first show that the first term of (6) is concave in ρ .

$E_X \left\{ [T - X(\rho)]^+ \right\} = \int_0^T (T - x)f(x | \rho)dx = \int_0^T F(x | \rho)dx$, where the last equality follows from integration by parts. If $\partial^2 F(x | \rho) / \partial \rho^2 \leq 0$, then this first term of (6) is concave in ρ , since concavity is maintained in integration. It is easy to check that the last two terms of (6) are

concave in ρ and the sum of a set of concave functions is also concave, therefore (6) is concave in ρ . \square

To calculate u and T via first order conditions, it is necessary to know the distribution of $X(\rho)$. Based on discussions with our contacts at HP, we demonstrate our model using the uniform and normal distribution cases below. The uniform distribution, which has a constant coefficient of variation, allows us to have closed-form solutions while the normal distribution case covers cases with varying degrees of uncertainty (coefficient of variation) in the false failure returns process.

Uniform distribution for number of false failure returns

Consider that $X(\rho) \sim \text{Uniform}\left(0, \frac{2\beta}{\rho}\right)$ (such that $E\{X(\rho)\} = \frac{\beta}{\rho}$). In this case, $F(x|\rho) = \frac{\rho}{2\beta}x$, $\partial^2 F(x|\rho)/\partial\rho^2 = 0$; thus (6) is concave in ρ according to Proposition 1. Then, $E_X\{[T - X(\rho)]^+\} = u \int_0^T (T-x) \cdot \frac{\rho}{2\beta} dx = T^2 u \rho / 4\beta$. Substituting this expression in (6) yields

$$\pi_R(\rho|T, u) = \frac{T^2 u \rho}{4\beta} - \frac{a\rho^2}{2} + [r + \delta_r(p-w)]\beta \left(1 - \frac{1}{\rho}\right). \quad (8)$$

The value ρ^* that maximizes (8) must satisfy the first-order conditions:

$$T^2 u (\rho^*)^2 - 4a\beta (\rho^*)^3 + 4[r + \delta_r(p-w)]\beta^2 = 0. \quad (9)$$

Substituting ρ^* with (2), we find an expression that relates u and T , so that the supply chain is coordinated by the contract:

$$T = \frac{2\beta^{2/3} a^{1/3} (m + \delta_m(w-c))^{1/2}}{u^{1/2} (m + \delta_m(w-c) + r + \delta_r(p-w))^{1/3}}. \quad (10)$$

For a fixed value of u , the corresponding value of T is obtained from (10). The value of u (and T) has to satisfy $T < 2\beta/\rho$, that is, the upper limit in the integral $E_X\{[T - X(\rho)]^+\}$ cannot exceed the upper limit of the uniform distribution. Substituting T from (10) and ρ from (2), yields $u > m + \delta_m(w-c)$. This is interesting since it shows that the value of u that coordinates this

contract is *larger* than the cost of a return to the manufacturer. Still, the manufacturer can be better off with a target rebate contract as Proposition 2 shows below.

Proposition 2: Suppose $X(\rho)$ is uniformly distributed. A (u, T) contract that coordinates the supply chain always makes the retailer better off, and makes the manufacturer better off if and only if $\rho^C \geq 2\rho^D$.

Proof: See Appendix.

Proposition 2 indicates that the manufacturer is better off with a target rebate only if the optimal centralized effort is large enough relative to the retailer's own optimal decentralized effort level. If the retailer's profit of avoiding one false failure return $r + \delta_r(p - w)$ is large enough relative to the marginal effort cost per return a/β , such that $\rho^D > 1$ according to (4), then $\rho^C \geq 2\rho^D$ is equivalent to $m + \delta_m(w - c) \geq 7[r + \delta_r(p - w)]$. This indicates that the target rebate contract makes the manufacturer better off only if its own profit margin of reducing a false failure return is substantially larger than the retailer's.

Normal distribution for number of false failure returns

Now, consider $X(\rho) \sim N(\beta/\rho, \sigma(\rho))$. Denote by $\Phi(\cdot)$ and $\phi(\cdot)$ the cdf and pdf of the standard normal distribution, respectively. Then, $F(x|\rho) = \Phi\left(\frac{x-\beta/\rho}{\sigma(\rho)}\right)$. Proposition 1's condition, $\partial^2 F(x|\rho)/\partial\rho^2 \leq 0$, may not be necessarily satisfied for a general function $\sigma(\rho)$. It is satisfied, however, for two important cases. First, if $\sigma(\rho) \equiv \sigma$ for all ρ , then $\partial^2 F(x|\rho)/\partial\rho^2 = \phi'\left(\frac{x-\beta/\rho}{\sigma}\right) \cdot \frac{\beta}{\sigma\rho^2} - \phi\left(\frac{x-\beta/\rho}{\sigma}\right) \cdot \frac{2\beta}{\sigma\rho^3} \leq 0$, since $\phi(z) > 0$ and $\phi'(z) \leq 0$ for all z . This scenario represents the case where the standard deviation of the number of false failure returns is *independent* of the effort level. A more interesting case, which we focus on this paper, is when $\sigma(\rho) = v\frac{\beta}{\rho}$, where v

is a constant denoting the coefficient of variation of $X(\rho)$; this is representative of the case where the standard deviation of $X(\rho)$ is *proportional* to its mean (and thus, inversely proportional to the effort level). In this case, after some algebra, (6) becomes:

$$\pi_R(\rho|T,u) = u \left\{ T - \frac{\beta}{\rho} \left[1 - vL\left(\frac{T\rho/\beta-1}{v}\right) \right] \right\} - \frac{a\rho^2}{2} + [r + \delta_r(p-w)]\beta \left(1 - \frac{1}{\rho} \right), \quad (11)$$

where $L(\cdot)$ is the standard normal loss function. The first-order condition with respect to ρ results in an implicit equation that can be solved for T , given a value of u or vice-versa:

$$\Phi\left(\frac{T\rho^c/\beta-1}{v}\right) - v\phi\left(\frac{T\rho^c/\beta-1}{v}\right) = \frac{(m + \delta_m(w-c))}{u}, \quad (12)$$

where ρ^c is given by (2), and $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal pdf and cdf, respectively.

Given that $\Phi(\cdot)$ and $\phi(\cdot)$ are necessarily less than one, then a necessary condition for (12) to have a solution is that the right-hand side of (12) is less than one, which implies $u > m + \delta_m(w-c)$, just like in the uniform distribution case. Further, if we define $z = (T\rho^c/\beta-1)/v$, and (arbitrarily) set $u = k(m + \delta_m(w-c))$, where $k > 1$ is a constant, we can rewrite (12) as

$$\Phi(z) - v\phi(z) = 1/k. \quad (13)$$

Denoting the solution of (13) by z^* , then $T = \frac{\beta}{\rho}(vz^* + 1)$. Note that z^* only depends on the coefficient of variation v and the constant k , and thus the design of such a contract is relatively straightforward. Unlike the uniform distribution case, however, the retailer is not always better off under a normal distribution for $X(\rho)$; we discuss this in depth in our comprehensive numerical analysis of §4. First though, we illustrate how the target rebate contract works for both the uniform and normal distribution cases with a numerical example below.

Example 1: Consider $m = 50$, $r = 5$, $p = 200$, $w = 150$, $c = 100$, $\beta = 120$, $a = 120$, $\delta_m = 0.7$, $\delta_r = 0.3$, and $v = 1/\sqrt{3}$ (so that the uniform and normal distribution cases have the same coefficient of variation, $1/\sqrt{3}$). The optimal level of effort, given by (2), is $\rho^c = 5.1$. The coordinated supply

chain's expected profit at the optimal level of effort (1) is $\Pi(\rho^C) = \$ 8,568$. To coordinate the supply chain, different combinations of u and T are possible. A necessary condition for u , under both uniform and normal distributions, is $u > m + \delta_m(w - c) = 85$. We consider five combinations of u and T , for both cases; and present the results in Table 3. Note that the distribution of profits does not depend on the contract parameters for the uniform distribution case. This is not an obvious result, and is a consequence of the unique shape of the uniform distribution—the term $E_X\{(T - \frac{\beta X}{\rho\theta})^+\}$, the only term dependent of (u, T) in the retailer's expected profit function (8), is a linear function of ρ when $X(\rho)$ is uniformly distributed. For the normal distribution case, however, a higher value of u (resulting in a lower T) results in increasing profits for the manufacturer and decreasing profits for the retailer.

Table 3: Different combinations of u and T that coordinate the supply chain in example 1

u	Uniform Distr. for $X(\rho)$			Normal Distr. For $X(\rho)$		
	T	π_R	π_M	T	π_R	π_M
90	49.4	2,366	6,202	50.5	2,805	5,763
95	48.1	2,366	6,202	46.1	2,543	6,026
100	46.9	2,366	6,202	43.4	2,401	6,168
105	45.8	2,366	6,202	41.4	2,309	6,260
110	44.7	2,366	6,202	39.8	2,243	6,325

4. Numerical analysis

In this section, we perform a numerical analysis that has several objectives: 1) to analyze the influence of the distribution of the uncertainty in the returns process $X(\rho)$ on our results; 2) to *quantify* the benefits of coordination, and 3) to identify how the benefits of coordination vary with different model parameter values.

The base numbers are grounded on product prices and returns volume for HP deskjet printers, and the other model parameters are varied over a wide range. We set $p = \$200$ (median price across different inkjet models), $m = 0.25p = \$50$, and a 50% gross margin, which implies $c =$

\$100. We use the two distributions (uniform and normal) for the uncertainty in the returns process studied theoretically. We vary the remaining parameters at four levels each according to the full-factorial experimental design shown in Table 4, which we justify as follows.

We consider values for the coefficient of variation v (when $X(\rho)$ is normally distributed) from 0.15 to 0.45 (we do not consider values for v higher than 0.45 because then there is a significant probability of negative values for $X(\rho)$). We use values for the cost of effort per expected return, a/β from \$0.10 to \$5. We use values for the retailer's return processing cost relative to the manufacturer (r/m) from 2.5% to 10%. While the retailer has some minimal labor cost associated with receiving a false failure, giving the customer credit, and separating the false failure for shipment, the manufacturer incurs all shipping costs, the refurbishing cost, and the re-marketing cost for the refurbished product in the secondary market. We use values for the parameter δ_m , which measures the impact of effort on net sales for the manufacturer, from the 0.3 to 0.9. Similarly, for the retailer we consider values for δ_r/δ_m from 0.3 to 0.9 (since $\delta_r < \delta_m$). Finally, we consider values for the wholesale price as a fraction of the product's price, w/p , from 0.65 to 0.95. The experimental design of Table 4 yields 5,120 cells, 1,024 for the uniform distribution and 4,096 for the normal distribution. Across all cells we select the contract parameter $u = k(m + \delta(w - c))$ with $k = 1.1$; we later report results using different values for k .

Table 4: Experimental design

<i>Parameter</i>	<i>Symbol</i>	<i>Levels</i>
Distribution of $X(\rho)$	-	Uniform, Normal
Coefficient of variation for $X(\rho)$	v	0.15, 0.25, 0.35, 0.45
Cost of effort per expected return	a/β	\$ 0.1, 1.0, 2.5, 5.0
Retailer's return processing cost relative to the manufacturer's	r/m	0.025, 0.05, 0.075, 0.10
Influence of effort on sales – manufacturer	δ_m	0.3, 0.5, 0.7, 0.9
Influence of effort on sales – retailer relative to the manufacturer	δ_r/δ_m	0.3, 0.5, 0.7, 0.9
Wholesale/retail price ratio	w/p	0.65, 0.75, 0.85, 0.95

Our main performance measure is the expected profit improvement per expected return (when $\rho = 1$) for each party or the supply chain as a result of adopting the target rebate contract. For example, for the retailer this is defined as $\Delta_R = (\pi_R(\rho^C | T, u) - \pi_R(\rho^D)) / \beta$, or, in relative terms, $\% \Delta_R = 100\% \Delta_R / |\pi_R(\rho^D) / \beta|$; similarly for the manufacturer and the supply chain. We report relative and absolute values of expected profit improvement here because *both* are important in providing a complete picture of the contract benefits, as we argue below. Detailed results are available upon request.

The statistics for *relative* expected profit improvement $\% \Delta$ for the retailer, manufacturer, and supply chain are reported in Table 5, tabulated by the distribution of $X(\rho)$: uniform, normal over all cases, and normal for each value of v . We do not report the maximum values for $\% \Delta$ because they occur in cases where the non-coordinated supply chain has such a small positive expected profit that any improvement has a huge relative impact and relative numbers can be misleading. A similar phenomenon occurs in Table 5—the extremely small minimum values of $\% \Delta_R$ in the normal distribution case (first row) occurs when $\pi_R(\rho^D)$ has a small absolute value. Again, this points out to the need for also reporting *absolute* values of expected profit improvement, which we do below.

Table 5: Statistics for expected profit improvement $\% \Delta$ across experimental design

Statistic	Party	Distribution of $X(\rho)$					
		Uniform	Normal Overall	Normal $\nu = 0.15$	Normal $\nu = 0.25$	Normal $\nu = 0.35$	Normal $\nu = 0.45$
Minimum	Retailer	10%	$-\infty^a$	$-\infty^a$	$-\infty^a$	-2652%	-307%
	Manufacturer	-31	-22%	2%	0%	-11%	-22%
	Supply chain	1%	1%	1%	1%	1%	1%
25 th perc.	Retailer	68%	-43%	-398%	-93%	12%	45%
	Manufacturer	-2%	9%	16%	12%	7%	2%
	Supply chain	13%	13%	13%	13%	13%	13%
Median	Retailer	217%	10%	-81%	-5%	46%	140%
	Manufacturer	8%	26%	37%	29%	21%	14%
	Supply chain	31%	31%	31%	31%	31%	31%
75 th perc.	Retailer	699%	94%	-17%	10%	169%	470%
	Manufacturer	37%	71%	91%	81%	67%	50%
	Supply chain	115%	115%	115%	115%	115%	115%

^aVery large negative number

Table 5 reveals that *relative* expected profit improvements are significant—the median values for $\% \Delta_M$ are 8% and 26%, and median values for $\% \Delta_R$ are 217% and 10%, for the uniform and normal distribution cases, respectively. The median value of supply chain expected profit improvement is 31%, which does not depend on the distribution of $X(\rho)$. We caution that our analysis is restricted only to profit per false failure return as impacted by the retailer’s effort, not *total* profit (forward + reverse supply chain) for each party. As expected, the retailer is better off ($\% \Delta_R > 0$) in *all* uniform distribution cases, while the manufacturer is better off in more than 50% of cases—the 25th percentile and median for $\% \Delta_M$ are -2% and 8% , respectively. For the normal distribution case, the manufacturer is better off in more than 75% of cases, and is always better off if $\nu \leq 0.25$; the retailer is better off in the majority of our test cases. The retailer’s (manufacturer’s) improvement increases (decreases) as the coefficient of variation of $X(\rho)$, ν increases—the median value of $\% \Delta$ is -81% , -5% , 46% , and 140% (37% , 29% , 21% , and 14%) for $\nu = 0.15, 0.25, 0.35$, and 0.45 , respectively. The higher the level of uncertainty ν in the returns process, the higher the *expected* number of returns below the target T , ceteris paribus, as (11) shows. Because the retailer is not penalized when returns are above the target, but is paid

(by the manufacturer) \$u per return below the target, the retailer's (manufacturer's) expected profit increases (decreases) as uncertainty in the returns process increases, *ceteris paribus*.

Next we study how *absolute* expected profit improvement Δ is related to the factors in our experimental design. For each factor, we average Δ across *all* experiments for each factor level. We focus on the normal distribution results and note that the results are similar for the uniform distribution, except that the magnitude of improvement is higher for the retailer in the uniform case. The results are shown in Figure 1 for factors v (left) and a/β (right), Figure 2 for r/m (left) and w/p (right), and Figure 3 for δ_m (left) and δ_r/δ_m (right). We note that the average expected profit improvement for the supply chain is equal to the sum of improvements across the retailer and the manufacturer. We have already commented on the effect of v on $\% \Delta$, now shown graphically for Δ in Figure 1 (left). Figure 1 (right) shows that expected profit improvement for all parties *decrease* as the retailer's normalized effort cost a/β decreases. A lower value of a/β means that the retailer already has an incentive to exert a significant effort level in the decentralized case ρ^D ; thus the contract does not provide an incentive for a significantly higher effort level.

Figure 1: Average expected profit improvement (\$) per expected return as a function of the coefficient of variation of X (left) and the retailer's effort cost per expected return (right)




Figure 2 (left) shows that the retailer's false failure processing cost relative to the manufacturer's r/m impact minimally the performance of the contract; this is because r is small relatively to the other revenue and cost figures. Figure 2 (right) shows that the supply chain's expected profit improvement increases with the wholesale price relative to product price w/p ; this is because a higher w/p implies a higher profit margin for the manufacturer, who has the most to gain from reducing false failures; the opposite is true for the retailer.

Figure 3 (left) shows that the impact of the parameter δ_m on Δ is relatively mild for the supply chain, however, expected profit improvement for the manufacturer (retailer) decreases (increases) as δ_m increases. From §3 and Table 2, $\delta_m = q_2 + q_3 = 1 - q_1$ is essentially the probability that the customer does *not* exchange a product with a product from the same manufacturer upon return; as this probability increases, the benefits of reducing false failures for the manufacturer decreases. Finally, Figure 3 (right) shows that expected profit improvement decreases as δ_r / δ_m increases across all parties and the supply chain. From §3 and Table 2, $\delta_r = q_3$ is the probability that the customer walks away with nothing after returning a false failure. At high levels of this probability, the retailer already has an incentive to decrease false failures in the decentralized case (and thus ρ^D is high) because false failure returns result in zero net sales for the retailer—thus, the benefit of a coordinating contract is lower.

Finally, we analyze the impact of choosing different levels of $u = k(m + \delta(w - c))$ in the performance of the contract. We study values of $k = 1.1, 1.3, 1.5,$ and 1.7 , and repeat, for each value of k , all 4,096 experiments in Table 4 for the normal distribution case; recall from example 1 that different values of u do not impact the performance of the contract and / or expected profit distribution between the two parties with a uniform distribution. As expected from example 1, the mean value of Δ_M increases as k increases: 18.7, 20.0, 20.6, and 21.1 for $k = 1.1, 1.3, 1.5,$ and 1.7 , respectively. Conversely, the mean value of Δ_R decreases as k increases: 1.3, 0.0, -0.6 ,

and -1.1 , respectively. The mean value of $\Delta_{SC} = 20.0$ does not depend on k . This indicates that a lower value of k , and consequently a lower value of u and a corresponding higher value of T , is indicated if *both* parties are to benefit from the contract. In summary, our numerical results indicate that a target rebate contract may result in significant improvements in the reverse supply chain's expected profit.

Figure 2: Average expected profit improvement (\$) per expected return as a function of the retailer's processing cost relative to the manufacturer r/m (left), and wholesale price (right)


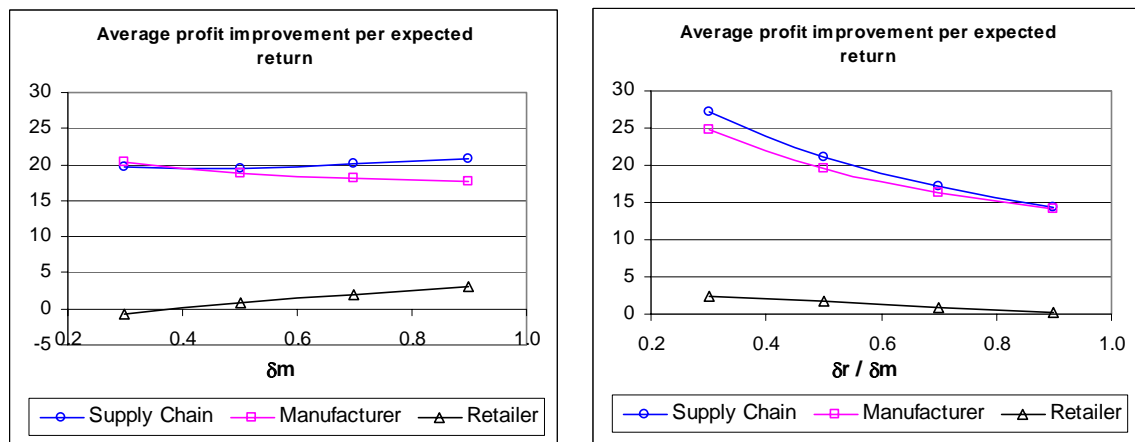


Figure 3: Average expected profit improvement (\$) per expected return as a function of δ_m (left) and δ_r/δ_m (right)



5. Forward contracts

In this section we study the impact of the false failure problem on traditional forward contracts used between a manufacturer and a retailer. The first part explores how some common forward contracts give the retailer an incentive to actually increase rather than decrease the number of false failure returns. The second part assumes the manufacturer is willing and able to negotiate a new forward contract with the retailer. Thus, we give suggestions on how a forward contract can be constructed that solves both the double marginalization problem and the effort incentive problem of reducing false failure returns.

5.1. Incentives to game traditional sales contracts

Up to this point, we have assumed that all false failure returns are unintentional on the retailer's part and are due primarily to the retailer's lack of effort in helping customers. In this section, we explore whether there are incentives for the retailer to *intentionally* encourage false failure returns (a practice we term "gaming"). For example, HP offers a buy-back sales contract with its retailers where any unsold product may be returned to HP for a *partial* refund at the end of the product's last selling season (the time when the demand rate for a particular product decreases to a point where the retailer no longer wishes to replenish it; this happens when an older model is being phased out to make room for its replacement version.) If the retailer is offered a *full* refund for a false-failure return (the current practice), then the retailer has an incentive to push a slow-selling product (during its last selling season) to consumers, fully knowing that the product does not fully meet the consumer's needs and thus has a high probability of being returned. Thus, a relevant question is: Does our proposed target rebate contract encourage or discourage the retailer to game her sales-based contract?

We examine several of the most popular sales-based contracts to determine where incentives for gaming exist at the end of the selling season, and show how our proposed target

rebate contract for false failure returns helps mitigate these gaming incentives. Let g_m represent the loss of goodwill to the manufacturer and g_r represent the loss of goodwill to the retailer if the customer purchases a product that does not meet their needs. Note that this is a different penalty cost than the more traditional loss of goodwill due to unsatisfied demand. Let v represent the per unit salvage value of any unsold product for the retailer at the end of the product's last selling season.

Consider the simplest contract, the wholesale price-only contract, where the manufacturer charges the supplier a fixed wholesale price for each unit purchased. In general, this contract does not coordinate the supply chain, but it is easy to monitor and still common in practice. In the absence of a target rebate contract, the incentive to game occurs when the retailer is left with unsold inventory at the end of the last selling season and the retailer's salvage value is less than the net benefit from a false failure return (the wholesale price minus the loss of goodwill to the retailer), i.e. $v \leq w - g_r$. Our target rebate contract changes this condition to

$$v \leq w - g_r - u \text{ if } T < \frac{\beta}{\rho^c}, \quad (14)$$

which reduces the incentive for gaming.

Table 6 summarizes the results for the wholesale price-only contract along with three commonly used forward coordination contracts. The first coordination contract is a buy-back or revenue sharing contract (Cachon, 2003 shows the two are equivalent under most simple settings). In a buy-back contract, the manufacturer sells the product to the retailer at the per unit wholesale price w , but pays the retailer b , $b \leq w$, for each unsold unit at the end of the last selling season. The second contract is a quantity flexibility contract where the retailer receives a full refund from the manufacturer on any leftover inventory at the end of the selling season that is below a predetermined target level λq , where q is the retailer's order quantity over the product's last selling season and $\lambda \in [0,1]$ is a contract parameter. To define this contract, let $Y \in [0,\Omega]$ be

a random variable representing demand over the last selling season, with pdf $f_Y(y)$. The third contract is a quantity discount contract where the wholesale price is a decreasing function of the order quantity. Since this contract induces the retailer to order more by lowering the wholesale price instead of repurchasing unsold units, the condition for gaming is the same as the wholesale price-only contract except that the single price w is replaced by the average wholesale price paid, \bar{w} , in (14) where $\bar{w} < w$.

Table 6: The effect of a Target Rebate contract on the incentive to game forward contracts

Contract Type	Without Target Rebate	With Target Rebate	Reduces Gaming
Wholesale price-only	$v \leq w - g_r$	$v \leq w - g_r - u$ if $T < \frac{\beta}{\rho^c}$	Yes
Buy-back and revenue sharing	$b + v \leq w - g_r$	$b + v \leq w - g_r - u$ if $T < \frac{\beta}{\rho^c}$	Yes
Quantity flexibility	$v \int_{\lambda_q}^{\Omega} f_Y(y) dy \leq w - g_r$	$v \int_{\lambda_q}^{\Omega} f_Y(y) dy \leq w - g_r - u$ if $T < \frac{\beta}{\rho^c}$	Yes
Quantity discount	$v \leq \bar{w} - g_r$	$v \leq \bar{w} - g_r - u$ if $T < \frac{\beta}{\rho^c}$	Yes

One last contract type that is sometimes used in practice is the sales rebate contract. This contract is similar to the target rebate contract discussed in this paper, except that it offers a reward for each unit *sold* above a predetermined target threshold. The exact contractual definition of *sale* results in very different incentives for gaming. If a sale is defined as total sales, including units eventually returned, a clear incentive exists (above even the wholesale price only contract) when the retailer is above her target threshold and receives an additional bonus for each product sold. If a sale is defined as net-sales, where the number of false failure returns is subtracted out at the end, then the use of a sales rebate contract should not add any additional incentive for gaming. Unfortunately, sales rebate contracts do not coordinate the supply chain in isolation and are thus often used in combination with other coordination contracts when sales are dependent on the sales effort of the retailer. Thus, the incentive for gaming may still exist with

these contracts but can be reduced by specifying net-sales in the conditions of the sales rebate contract.

Summarizing, our analysis reveals that traditional sales-based contracts designed to coordinate the forward supply chain may induce gaming at the end of a product's last selling season. We argue, however, that our target rebate contract, designed to coordinate false failure returns in the reverse chain, reduces these gaming incentives when they exist. Further, we believe that the low volumes typically found during the last selling season combined with the side benefits of our target rebate contract implies that there is no need to incur the cost of altering existing contracts to eliminate the gaming incentives. If a manufacturer and a retailer are establishing a new relationship however, it is better to address the misalignment of incentives in the original forward contract. Thus, in the next section we offer a modification to a traditional forward contract that eliminates this misalignment and achieves the coordinated solution on the reduction of false failure returns.

5.2. False Failure Coordinated Forward Contract

In our experience talking with manufacturers who suffered from high levels of false failure returns, there is a strong reluctance to change the forward contracts already in place with the retailers. There are several reasons for this reluctance. The retailer often standardizes forward contracts so any manufacturer that refuses to abide by the standard contract faces the danger of having its product line dropped by the retailer. Also, renegotiating a forward contract is an expensive and time consuming process and the major retailers often do not fully trust the manufacturers so they may refuse a new forward contract even if it is Pareto improving for all parties. Our proposed target rebate contract avoids these problems because it results only in additional revenue for the retailer who has complete discretion on whether or not to alter her

behavior. It also does not require any modification to the current forward contract, making it attractive for the established supply chain relationships.

Despite the difficulties described above, a single forward contract, that also solves the incentive problem on the amount of effort to spend on reducing false failures, is the ideal and there may be situations where negotiating such a contract is feasible. Thus, we now explore how such a contract should work. Before presenting our forward contract, we define some additional notation. Let D be random demand for a single selling period and let $G(y) = P(D \leq y)$. The retailer must order a quantity q before knowing the true demand. For simplicity, assume all unmet demand is lost, any goodwill cost for a lost sale is zero, and the salvage value for unsold inventory is v . Let $S(q)$ be the retailer's expected sales where

$$S(q) = E_D \{ \min(q, D) \} = q - \int_0^q G(y) dy.$$

Finally, let $I(q)$ be the expected left over inventory, $I(q) = q - S(q)$. In the absence of false failure returns, the coordinated supply chain's expected profit is

$$\bar{\Pi}(q) = pS(q) + vI(q) - cq. \quad (15)$$

To include the effect of false failures on the supply chain's expected profit, we first must denote our expected number of false failures to be a function of q . To do so, we set

$$E\{X(\rho, q)\} = \frac{\beta(q)}{\rho}. \quad (16)$$

We further assume that the number of false failures is a fixed percentage of sales (which depends on the effort level); thus $X(\rho, q) = \gamma(\rho) \min(q, D)$ where $\min(q, D)$ is the sales quantity and $\gamma(\rho)$ is a parameter value that depends on the effort level. Taking expectations gives

$$E\{X(\rho, q)\} = \gamma(\rho)S(q). \quad (17)$$

Equating (16) to (17) gives

$$\frac{\beta(q)}{\rho} = \gamma(\rho)S(q),$$

where $\gamma(\rho) = \gamma/\rho$ and γ represents the constant percentage of false failures (relative to actual sales). Thus, $X(\rho, q) = \gamma S(q)/\rho$. Incorporating the cost of false failures in the coordinated supply chain's expected profit gives

$$\bar{\Pi}(q, \rho) = \left(p - \frac{\eta\gamma}{\rho} \right) S(q) + vI(q) - cq - \frac{a\rho^2}{2}, \quad (18)$$

$$\text{where } \eta = (m + \delta_m(w - c) + r + \delta_r(p - w)).$$

The supply chain's expected profit contains the normal cost of falling above or below demand plus the additional cost of the false failure returns and the effort spent to reduce them. Assuming that the expected profit (18) is well-behaved and jointly concave in q and ρ , the optimal order quantity and effort is found by solving

$$1 - F(q) = \frac{(c - v)}{\left(p - \frac{\eta\gamma}{\rho} - v \right)} \quad (19)$$

for the optimal order quantity as a function of effort and

$$\rho = \sqrt[3]{\frac{\eta\gamma S(q)}{a}} \quad (20)$$

for the optimal effort level as a function of the order quantity. Equations (19) and (20) can then be solved simultaneously resulting in the optimal decisions q^C and ρ^C . Although closed form solutions do not exist for the general distribution cases, (19) and (20) can be solved numerically for given distributions of demand and the number of false failures.

For the non-coordinated supply chain, the retailer's expected total profit is

$$\bar{\pi}_R(q, \rho) = \left[p - (r + \delta_r(p - w)) \frac{\gamma}{\rho} \right] S(q) + vI(q) - wq - \frac{a\rho^2}{2}. \quad (21)$$

The manufacturer's expected total profit for a non-coordinated chain is

$$\bar{\pi}_M(q, \rho) = (w - c)q - (m + \delta_m(w - c)) \frac{\gamma}{\rho} S(q). \quad (22)$$

The second part of (22) is the additional cost incurred by the manufacturer due to false failure returns. A contract that induces the retailer to expend a ρ^C level of effort must assign the total cost of false failures to the retailer. Thus, the retailer must be charged a cost of $m + \delta_m(w - c)$ for each false failure returned to the manufacturer. To induce the retailer to accept the new contract, the retailer must be compensated for assuming this extra cost. This can be done through a lump sum payment or by choosing the exchange price w low enough so that the retailer's expected profit is higher than it would be without the contract.

While assigning all the false failure related cost to the retailer solves one problem, it does not solve the double-marginalization problem where the retailer's optimal order quantity is less than q^C . Cachon (2003) discusses many contract types that solve this piece of the problem. We choose a buy-back contract to demonstrate our forward coordination contract since it is one of the more commonly used in practice. Adding the partial rebate for unsold units (b , $b \leq w$) and the total false failure cost to the retailer's expected profit gives

$$\bar{\pi}_R(q, \rho) = \left(p - \frac{\eta\gamma}{\rho} \right) S(q) + (v + b)I(q) - w_b q - \frac{a\rho^2}{2}. \quad (23)$$

The manufacturer's expected profit is

$$\bar{\pi}_M(q, \rho) = (w_b - c)q - bI(q). \quad (24)$$

Pasternack (1985) shows that a range of buy-back parameter values $\{w_b, b\}$ solve the double-marginalization problem and that every possible profit allocation is feasible. Therefore, quantity and effort coordination can be achieved by choosing the set of values that result in the retailer's choice of q^C and ρ^C as well as a Pareto improvement in the expected profits for both the retailer and the manufacturer over their expected profits in the absence of the contract.

6. Conclusion

In this paper we propose a target rebate contract to coordinate false failure returns, which are consumer returns returned to the retailer that have no functional or cosmetic defect. This contract stipulates a payment of $\$u$ for each false failure return below a target T . We propose a simple, but insightful model, which reveals the substantial benefits that can be derived by enticing retailers to exert a higher level of effort to decrease the number of false failure returns. This higher level of effort has two benefits for the supply chain. First, it reduces the overall processing cost for false failure returns, which is significant considering that a false failure return may cost 25% of the product price. Second, a higher level of effort may result in a higher level of net sales (sales that are not returned).

We show that our contract has a number of attractive properties, including Pareto–improvement in a majority of cases. We find that the magnitude of profit (associated with false failure returns) improvement per expected return as a result of the increased retailer’s effort is very significant, with a median value of 10%, 26% and 31% for the retailer, manufacturer, and supply chain respectively. For the retailer (manufacturer), the profit improvement per expected return increases (decreases) with the uncertainty in the returns process. This is because as uncertainty increases, *ceteris paribus*, the retailer’s (manufacturer’s) profit increases (decreases) due to an increase in the expected number of false failure returns below the contract target T .

Our research is a first step in understanding how a target rebate contract provides incentives to the retailer to improve her effort in the sales process targeted at reducing the number of false failure returns. We use a stylized model that summarizes the retailer’s effort by a single variable, a common approach in the current body of contract research. We also discuss the implications of using quantity-based coordination contracts on the false failure return problem, and offer guidance on which ones provide the least incentive to game. One possibility for future research is to generalize our results to other settings. Another is to monitor the implementation of such a

contract in practice, and compare the improvements against the theory. In any event, our research indicates that significant savings may be obtained by reducing the number of false failure returns.

Appendix: Proof of Proposition 2

Proposition 2: Suppose $X(\rho)$ is uniformly distributed. A (u, T) contract that coordinates the supply chain always makes the retailer better off, and makes the manufacturer better off if $\rho^C \geq 2\rho^D$.

Proof: First, consider the manufacturer. The manufacturer is better off with the (u, T) contract if $\pi_M(\rho^C | T, u) \geq \pi_M(\rho^D)$. The manufacturer's expected profit (7) can be rewritten as

$$\pi_M(\rho^C | T, u) = (m + \delta_m(w - c))\beta \left(\frac{\rho^C - 1}{\rho^C} \right) - \frac{T^2 u \rho^C}{4\beta}. \quad (25)$$

We can rewrite (9) as $T^2 u (\rho^C)^2 = 4a\beta(\rho^C)^3 - 4[r + \delta_r(p - w)]\beta^2$, and from (2),

$$(\rho^C)^3 = (m + \delta_m(w - c) + r + \delta_r(p - w))\beta / a. \text{ Thus,}$$

$$T^2 u (\rho^C)^2 = 4(m + \delta_m(w - c) + r + \delta_r(p - w))\beta^2 - 4[r + \delta_r(p - w)]\beta^2 = 4\beta^2(m + \delta_m(w - c)).$$

Substituting this in (25) yields:

$$\pi_M(\rho^C | T, u) = (m + \delta_m(w - c))\beta \left(\frac{\rho^C - 1}{\rho^C} \right) - \frac{\beta(m + \delta_m(w - c))}{\rho^C} = (m + \delta_m(w - c))\beta \left(\frac{\rho^C - 2}{\rho^C} \right)$$

Again, we need $\pi_M(\rho^C | T, u) \geq \pi_M(\rho^D)$, where $\pi_M(\rho^D)$ is given by (5). Simple algebra shows that this is true if $\rho^C \geq 2\rho^D$.

Now, we consider the retailer. Define $\Delta_R = \pi_R(\rho^C | T, u) - \pi_R(\rho^D)$. For the retailer to be better off, we need $\Delta_R \geq 0$. From (9), $T^2 u / 4\beta = a\rho^C - [r + \delta_r(p - w)]\beta / (\rho^C)^2$, which can be substituted into (8), yielding

$$\begin{aligned}\pi_R(\rho | T, u) &= \frac{a(\rho^C)^2}{2} - \frac{2[r + \delta_r(p-w)]\beta}{\rho^C} + [r + \delta_r(p-w)]\beta \\ &= \frac{a(\rho^C)^3 - 4[r + \delta_r(p-w)]\beta}{2\rho^C} + [r + \delta_r(p-w)]\beta.\end{aligned}\quad (26)$$

Now, from (2), $(\rho^C)^3 = (m + \delta_m(w-c) + r + \delta_r(p-w))\beta/a$, which can be substituted in (26), yielding

$$\pi_R(\rho | T, u) = \frac{\beta(m + \delta_m(w-c) - 3[r + \delta_r(p-w)])}{2\rho^C} + [r + \delta_r(p-w)]\beta. \quad (27)$$

In the decentralized case, note that (3) can be rewritten as

$$\pi_R(\rho^D) = \frac{-a(\rho^D)^3 - 2[r + \delta_r(p-w)]\beta}{2\rho^D} + [r + \delta_r(p-w)]\beta. \quad (28)$$

But $\rho^D = \max\left\{\left([r + \delta_r(p-w)]\beta/a\right)^{1/3}, 1\right\}$. First, consider the case where $\rho^D > 1$; in this case $\rho^D = \left([r + \delta_r(p-w)]\beta/a\right)^{1/3}$, which implies $a(\rho^D)^3 = [r + \delta_r(p-w)]\beta$; substituting this in (28) yields

$$\pi_R(\rho^D) = -\frac{3[r + \delta_r(p-w)]\beta}{2\rho^D} + [r + \delta_r(p-w)]\beta. \quad (29)$$

Taking the difference (27) – (29) yields

$$\begin{aligned}\Delta_R &= \pi_R(\rho^C | T, u) - \pi_R(\rho^D) = \frac{\beta[m + \delta_m(w-c) - 3[r + \delta_r(p-w)]]}{2\rho^C} + \frac{3[r + \delta_r(p-w)]\beta}{2\rho^D} \stackrel{\rho^C \geq \rho^D}{\geq} \\ &\quad \frac{\beta[m + \delta_m(w-c) - 3[r + \delta_r(p-w)]]}{2\rho^C} + \frac{3[r + \delta_r(p-w)]\beta}{2\rho^C} = \frac{\beta[m + \delta_m(w-c)]}{2\rho^C} \geq 0.\end{aligned}\quad (30)$$

The other case to consider is if $\rho^D = 1$; in this case $\pi_R(\rho^D) = -a/2$. Thus, from (27),

$$\Delta_R = \pi_R(\rho^C | T, u) - \pi_R(\rho^D) = \frac{\beta(m + \delta_m(w-c) - 3[r + \delta_r(p-w)])}{2\rho^C} + [r + \delta_r(p-w)]\beta + \frac{a}{2}. \quad (31)$$

If the first term on the right-hand side of (31) is positive, then $\Delta_R \geq 0$ because the remainder two terms of (31) are positive. If, on the other hand, the first term on the right-hand side of (31) is negative, then we can write

$$\begin{aligned}
\Delta_R &= \frac{\beta(m + \delta_m(w - c) - 3[r + \delta_r(p - w)])}{2\rho^c} + [r + \delta_r(p - w)]\beta + \frac{a}{2} \stackrel{\rho^c \geq 1}{\geq} \\
&= \frac{\beta(m + \delta_m(w - c) - 3[r + \delta_r(p - w)])}{2} + [r + \delta_r(p - w)]\beta + \frac{a}{2} \\
&= \frac{\beta[m + \delta_m(w - c) - 3[r + \delta_r(p - w)]]}{2} + \frac{2[r + \delta_r(p - w)]\beta}{2} + \frac{a}{2} = \\
&= \frac{\beta[m + \delta_m(w - c) + r + \delta_r(p - w)] - 2[r + \delta_r(p - w)]\beta + a}{2} = \\
&= \frac{a}{2} \left\{ \frac{\beta[m + \delta_m(w - c) + r + \delta_r(p - w)]}{a} - 2 \frac{[r + \delta_r(p - w)]\beta}{a} + 1 \right\} \\
&= \frac{a}{2} \left\{ (\rho^c)^3 - 2 \frac{[r + \delta_r(p - w)]\beta}{a} + 1 \right\} \stackrel{(*)}{\geq} \frac{a}{2} \{ (\rho^c)^3 - 2 + 1 \} = \frac{a}{2} \{ (\rho^c)^3 - 1 \} \stackrel{\rho^c \geq 1}{\geq} 0
\end{aligned}$$

where in (*) we have used the fact that if $\rho^D = \max\left\{\left(\frac{[r + \delta_r(p - w)]\beta}{a}\right)^{1/3}, 1\right\}$, and $\rho^D = 1$, then necessarily $\frac{[r + \delta_r(p - w)]\beta}{a} \leq 1$. This completes the proof. \square

References

- Cachon, G. 2003. Supply chain coordination with contracts. Chapter 6 of *Handbooks in operations research & management science*, Vol. 11, A. G. de Kok and S. C. Graves (eds.), Elsevier B. V.
- Cha, A.E. 2004. Some shoppers find fewer happy returns: Databases limit ‘excessive’ exchanges. *The Washington Post*, November 7 2004, A.01.
- Choi, T., D. Li, H. Yan, 2004. Optimal return policy for supply chain with e-marketplace. *International Journal of Production Economics* 88, 205-227.
- Chowdhury, J. 1993. The motivational impact of sales quotas in effort. *Journal of Marketing Research* 30 28-41.

- Darmon, R. 1979. Setting sales quotas with conjoint analysis. *Journal of Marketing Research* **16** 133-140.
- Davey, S. 2001. Personal communication with the authors. World wide product returns manager, Inkjet Products Group, Hewlett-Packard Company.
- Davis, O., J. Farley. 1971. Allocating sales force effort with commissions and quotas. *Management Science* **18** P55-P-63.
- Davis, S., E. Gerstner, M. Hagerty. 1995. Money back guarantees in retailing: Matching products to consumer tastes. *Journal of Retailing* **71** 7-22.
- Debo, L., B. Toktay, L. Van Wassenhove. 2001. Market segmentation and technology selection for remanufacturable products. INSEAD Working Paper, 2001/47/TM/CIMSO 18.
- Ferguson, M., B. Toktay. 2004. The effect of external competition on recovery strategies. Georgia Institute of Technology College of Business working paper
- Ferrer, G., M. Ketzenberg. 2004. Value of information in remanufacturing complex products. *IIE Transactions* **36** 265-277.
- Fleischmann, M. 2001. *Quantitative models for reverse logistics*. Lecture Notes in Economics and Mathematical Systems, Volume 501, Springer-Verlag, Berlin.
- Gaba A., A. Kalra. 1999. Risk behavior in response to quotas and contests. *Marketing Science* **18** 417-434.
- Guide, Jr. V.D.R. 2000. Production planning and control for remanufacturing. *Journal of Operations Management* **18** 467-483.
- Guide, Jr., V.D.R., L. Van Wassenhove, 2001. Managing product returns for remanufacturing. *Production and Operations Management* **10**(2): 142-155.

- Guide, Jr. V.D.R., L. Van Wassenhove. 2003. Closed-loop supply chains. In *Business Aspects of Closed-Loop Supply Chains* (V.D.R. Guide, Jr., L.N. Van Wassenhove, eds.). Carnegie Mellon University Press, Pittsburgh PA USA.
- Helbig, Barbara. 2002. Personal communication with the authors. Solution and services manager, Equipment Management and Remarketing, Hewlett-Packard GmbH.
- Heese, H., K. Cattani, G. Ferrer, W. Gilland, A. Roth. 2004. Competitive advantage through take-back of used products. *European Journal of Operational Research* in-press.
- Kumar, N., V.D.R. Guide, Jr., L. Van Wassenhove. 2002. *Managing product returns at Hewlett Packard*, teaching case 05/2002-4940, INSEAD.
- Majumder, P., H. Groenevelt. 2001. Competition in remanufacturing. *Production and Operations Management* **10** 125-141.
- Mantrala, M., P. Sinha, A. Zoltners. 1994. Structuring a multiproduct sales quota-bonus plan for a heterogeneous salesforce: A practical model-based approach. *Marketing Science* **13** 121-144.
- Mantrala, M., K. Raman, R. Desiraju. 1997. Sales quota plans: Mechanisms for adaptive learning. *Marketing Letters* **8** 393-405.
- McWilliams, G. 2004. Minding the store: Analyzing customers, Best Buy decides not all are welcome; retailer aims to outsmart dogged bargain-hunters, and coddle big spenders. *The Wall Street Journal* November 8, 2004 A.1.
- Padmanabhan, V., I. Png. 1997. Manufacturer's returns policies and retail competition. *Marketing Science* **17** 81-94.
- Padmanabhan, V., I. Png. 1995. Returns policies: Make money by making good. *Sloan Management Review* **37**(1) 65-72.

- Pasternack, B. A. 1985. Optimal pricing and return policies for perishable commodities. *Marketing Science*, **4**(2) 166-176.
- Raju, J., V. Srinivasan. 1996. Quota-based compensation plans for multiterritory heterogeneous salesforces. *Management Science* **42** 1454-1462.
- Reiss, H. 2003. Personal communication with the authors. Vice-President & General Manager, Equipment Manufacturing & Remarketing, Hewlett-Packard Co.
- Savaskan R., S. Bhattacharya, L. Van Wassenhove. 2004. Channel choice and coordination in a remanufacturing environment. *Management Science* **50**, 239-252.
- Souza, G., V. D. R. Guide, Jr., L. Van Wassenhove, and J. Blackburn. 2004. Time value of commercial product returns. Working paper, INSEAD.
- Souza, G., M. Ketzenberg, V.D.R. Guide, Jr. 2002. Capacitated remanufacturing with service level constraints. *Production and Operations Management* **11**, 231-248.
- Stock, J., T. Speh, H. Shear. 2002. Many happy (product) returns. *Harvard Business Review* **80**(7) 16-17.
- Tapiero, C., J. Farley. 1975. Optimal control of sales force effort in time. *Management Science* **21** 976-085.
- Taylor, T. 2002. Supply chain coordination under channel rebates with sales effort effects. *Management Science* **48** 992-1007.
- Toktay, B., L. Wein, Z. Stefanos. 2000. Inventory management of remanufacturable products. *Management Science* **46** 1412-1426.
- Toktay, B. 2003. Forecasting product returns. In: *Business Aspects of Closed-Loop Supply Chains* (V.D.R. Guide, Jr., L.N. Van Wassenhove, eds). Carnegie Mellon University Press, Pittsburgh, PA USA.

Tsay, A. 2001. Managing retail channel overstock: Markdown money and return policies.

Journal of Retailing **77** 457-492.

Valenta, R. 2002. Personal communication with the authors. Director, Product Service, Robert

Bosch Tool Corporation.

van der Laan, E., M. Salomon, R. Dekker, L. Van Wassenhove. 1999. Inventory control in

hybrid systems with remanufacturing. *Management Science* **45** 733-747.

Winer, L. 1973. The effect of product sales quotas on sales force productivity. *Journal of*

Marketing Research **10** 180-183

Wood, S. 2001. Remote purchase environments: The influence of return policy leniency on two-

stage decision processes. *Journal of Marketing Research* **38** 157-169.