



Risk sensitivity in distribution channel partnerships: implications for manufacturer return policies

Andy A. Tsay*

Leavey School of Business, Santa Clara University, Santa Clara, California 95053-0382, USA

September 21, 2000; revised March 30, 2001, September 1, 2001

Abstract

The manufacturer return policy is widely regarded as a means for channel partners to share risk. However, existing studies of this popular institutional practice use frameworks that assume risk-neutrality of all parties.

This report analyzes how sensitivity to risk affects both sides of the manufacturer-retailer relationship under various scenarios of strategic power, and how these dynamics are altered by a return policy. A key finding is that the penalty for ignoring risk sensitivity can be substantial. This will suggest an informational motive affecting the use of return policies, a consequence of the potential difficulty of inferring another party's risk sensitivity and the positive incentive for deception. © 2002 by New York University. All rights reserved.

Keywords: Risk; Manufacturer return policies; Channels of distribution; Supply chain management

Introduction

Risk is nearly always regarded as a key concern in the structuring of business relationships. This certainly applies to vertical supply/distribution agreements between independent firms (whose parties are referred to as a *manufacturer* and a *retailer*). In this context, manufacturer return policies have been identified by both the academic and practitioner literatures as a popular vehicle for addressing the risk engendered by market demand uncertainty in a variety of industries (e.g., books, magazines, newspapers, recorded music, computer hardware and software, greeting cards, and pharmaceuticals) (cf. Pasternack, 1985; Padmanabhan and Png, 1995; Kandel, 1996). One key question remains unanswered, as noted by Padmanabhan and Png (1995, p.66): "Using a returns policy as insurance does not address the important implementation issue of who should offer the returns policy. A manufacturer must consider whether it or the retailers can better absorb the risk of excess inventory. Even if the manufacturer accepts returns from retailers, the risk does not disappear, but merely shifts up the distribution channels." These authors reason that firms with greater depth/breadth of assets and activities tend to be less con-

cerned with uncertainty. But the undeniable message is that firms do care about risk, and different firms may care to differing extents. This realization is essential to the proper management of interfirm relationships.

However, most existing formal models of this setting remain silent on this issue because they either contain no uncertainty at all, or assume risk-neutrality of all decision makers. In the latter case, expected profit (or cost) is the exclusive currency used to measure individual preferences, and transfers of this currency between parties are described sometimes as a "sharing" of risk¹. But any logic that fails to differentiate between certain and uncertain payoffs is fundamentally at odds with the notion of sensitivity to risk, and therefore may offer spurious recommendations.

The objective of this research is to address this shortcoming by analyzing formally how sensitivity to risk affects both sides of a manufacturer-retailer relationship under various scenarios of relative strategic power, and how these dynamics are altered by the introduction of a manufacturer return policy. At a number of junctures the analysis will articulate how risk sensitivity leads to behaviors that are qualitatively different from those predicted by risk-neutral analysis, and show that the penalty for ignoring risk sensitivity in channel policy design can be substantial. This will suggest an informational motive affecting the use of return policies, a consequence of the potential difficulty of infer-

* Tel.: +1-408-554-4561; fax +1-408-554-5157.

E-mail address: atsay@stanfordalumni.org (A.A. Tsay).

ring another party's risk sensitivity and the positive incentive for deception.

To model a channel facing stochastic demand and comprising multiple independent decision-makers, we extend a formulation of Padmanabhan and Png (1997), which we will refer to as PP². This research generalizes PP's in two key ways. The first is in considering the case in which the retailer is the strategic leader in the channel, in addition to PP's premise that the manufacturer dictates the distribution policy. The second is in the explicit acknowledgment of sensitivity to risk. Whereas PP's manufacturer and retailer each maximize individual expected profit, we assume each to assess any random financial outcome Z via a value function of the form $\{E[Z] - k \text{StdDev}[Z]\}$. We will refer to this as a *mean-standard deviation (MS) value function* and k as the *risk aversion parameter*³. Bar-Shira and Finkelshtain (1999) argue that using value functions that increase in mean and decrease in standard deviation is more robust than approaches based on expected utility. Tsang (1972) and Adar et al., (1977) also use this concept, and Saha (1997) suggests that such forms are more analytically tractable. The MS value function is a special case of this, and has been utilized by numerous researchers, including Johnson and Simik (1971) and Lau (1980)⁴.

§2 reviews the relevant literature. §3 outlines notation and modeling assumptions, derives the equilibrium behaviors and outcomes, and then performs a number of lines of analysis. The main issues addressed are (1) the effect on channel behavior of sensitivity to risk, (2) the effect of the distribution policy on channel behavior and the outcomes for each party in the presence of risk sensitivity, (3) the importance of explicitly acknowledging sensitivity to risk, and (4) the influence on these issues of the balance of power in the channel. §4 concludes with a summary of managerial implications and areas for future research. All proofs and derivations are presented in the Appendix.

Positioning in the literature

In the literature of manufacturer return policies, two general formulations have been used to model a channel facing stochastic demand. They are differentiated mainly by the representation of the retailer's decision problem. The predominant approach treats the retailer as a classical "newsvendor" facing a random demand in a single period (cf. Pasternack, 1985; Kandel, 1996; Emmons and Gilbert, 1998; Donohue, 2000; Webster and Weng, 2000; Tsay, 2001). The retailer orders product (and possibly sets the retail price) before the resolution of that demand, and salvages any overstock at some value less than the procurement cost. The other approach is that of PP, which uses a price-sensitive demand curve with a random shift parameter. Here, the retailer orders, observes demand, and then sets the retail price at which all units sell. This representation of events is a modeling abstraction whose intent is to capture

the fact that over-ordering is detrimental to the retailer, a central property of the newsvendor framework as well⁵. Under both regimes, the retailer's aversion to overstock scenarios depresses the quantity that the manufacturer can sell.

A return policy is one mechanism by which a manufacturer can increase the retailer's initial order. In the process, the manufacturer's payoff converts from certain to uncertain. Because ostensibly the retailer benefits at the same time that the manufacturer accepts exposure to risk, popular vernacular tends to label this as a "sharing" or "transfer" of risk. Existing models attempt to formalize this and determine when mutual benefit is possible. Yet, as noted earlier, they tend to do this without acknowledging the difference between certain and uncertain outcomes.

We are aware of only two studies about distribution channels that have considered multiple interdependent decision-makers with individual agendas and attitudes towards risk. Webster and Weng (2000) addressed sensitivity to risk on the manufacturer's side only. Rather than quantifying explicitly the manufacturer's preferences towards uncertainty, these researchers pursued the premise that the manufacturer unequivocally will prefer offering full returns if every *resulting realization* of profit will be at least as great as the certain profit attainable with no returns. However, this may be too extreme a requirement, as in reality a manufacturer might willingly endure the prospect of a poor profit outcome if the full set of potential realizations is sufficiently attractive on average (or in some other aggregate sense). Webster and Weng's framework provides no way to evaluate this case. Spulber (1985) assumed the manufacturer and retailer to each maximize the expected value of a concave (hence, risk-averse) utility function, and determined which party should bear the demand risk. When appropriate, this responsibility can be transferred to the manufacturer via a consignment contract, the economic equivalent of a return policy. However, Spulber's formulation ordains that the manufacturer directly controls the retail inventory decision (which need only guarantee the retailer a certain reservation utility), and hence is not a true multi-player analysis.

Risk has been incorporated into other types of decision models in a number of ways, including the various "Safety First" objectives of Roy (1952), Telser (1955), and Kataoka (1963)⁶. Robison and Barry (1987) provide extensive background on these. Others (e.g., Harlow, 1991; Tse et al., 1993) have advocated building decision models around measures of downside risk, such as lower partial moments. These formulations are more tractable in finance and economics applications since the distribution of the stochastic profit often is a simple transformation of the distributions of individual asset prices. Unfortunately, generally this is not so for stochastic inventory models, as Lau (1980) demonstrated⁷. Furthermore, the complexity of these alternatives may become prohibitive in settings with multiple decision-makers.

Table 1

Components of the model

(Note: A tilde denotes a random variable, and superscripts and subscripts identify the channel member (R or M, for retailer or manufacturer, respectively), distribution policy, and demand scenario, as necessary.)

Variable	Description
i	index for demand scenario; $i \in \{l, h\}$ for “low” and “high”, respectively
j	index for distribution policy; $j \in \{nr, fr\}$ for “no return” and “full return” (at full price), respectively
$\tilde{\alpha}$	primary demand, $= \begin{cases} \alpha_l \text{ (low) with probability } \lambda \\ \alpha_h \text{ (high) with probability } (1 - \lambda) \end{cases}$ with $\alpha_h > \alpha_l > 0$ and $0 < \lambda < 1$
$p_{j,i}$	retail price under distribution policy j in demand scenario i
β	market sensitivity to retail price; $\beta > 0$
$q_{j,i}$	market demand under distribution policy j in demand scenario i ; $\equiv \alpha - \beta p_{j,i}$
s_j	retailer's order size under distribution policy j
w_j	unit wholesale price charged by the manufacturer under distribution policy j
c	manufacturer's unit production cost
$\tilde{\Pi}_j^R$	retailer's profit under distribution policy j
$\tilde{\Pi}_j^M$	manufacturer's profit under distribution policy j
$\pi_{j,i}^R$	realization of $\tilde{\Pi}_j^R$ in demand scenario i
$\pi_{j,i}^M$	realization of $\tilde{\Pi}_j^M$ in demand scenario i
k_M	manufacturer's sensitivity to risk, with $k_M \geq 0$
k_R	retailer's sensitivity to risk, with $k_R \geq 0$
V_j^M	manufacturer's MS value function; $\equiv E[\tilde{\Pi}_j^M] - k_M \cdot \text{StdDev}[\tilde{\Pi}_j^M]$
V_j^R	retailer's MS value function; $\equiv E[\tilde{\Pi}_j^R] - k_R \cdot \text{StdDev}[\tilde{\Pi}_j^R]$

As noted, the two primary ways to model the inventory decision problem of a retailer facing uncertain demand are the newsvendor model and PP's formulation. Because attempts to incorporate risk into the newsvendor framework (e.g., Lau, 1980; Eeckhoudt et al., 1995; Agrawal and Seshadri, 2000) generally have not yielded results conducive to investigation of multiparty interactions, we use a channel representation similar to PP's. In this context, MS objective functions are a reasonable and mathematically tractable representation of risk preferences. We use this combination to advance an area in which apparently there is little existing research, and generate empirically testable hypotheses.

The Model

The following discussion recapitulates the PP analysis in the course of generalizing their formulation to handle risk. Basic constructs of the model are presented in Table 1.

Assuming the following decision structure facilitates direct comparison to PP's findings:

Stage 1. The channel leader declares the distribution policy j (w and any return policy).

Stage 2. The retailer chooses s_j , while $\tilde{\alpha}$ is unknown.

Stage 3. Uncertainty about $\tilde{\alpha}$ is resolved, and then the retailer selects $p_{j,i}$, sells the minimum of $q_{j,i}$ and s_j at that price, and executes on any relevant terms of the distribution policy as appropriate (e.g., returning any overstock if so allowed).

The decision structure represents relative strategic power in the channel. PP's work and nearly every other analysis of return policies across a variety of literatures assume the manufacturer to be the channel leader. This would be reasonable when a large manufacturer with a powerful brand deals with a small to midsized retail firm. Except where explicitly stated, this assumption will be in effect. In addition, we will consider an alternative in which the retailer dominates the channel, hence is able to dictate the distribution policy.

Under the assumption that the manufacturer charges a per-unit price⁸, we study the two policies considered by PP: *no returns*, and *full returns for full credit*. With all appropriate caveats, we follow PP in assuming common knowledge of all parameters so as to focus on the channel dynamics. We will address the significance of this assumption with respect to the risk sensitivity parameters.

The complete equilibrium for each policy is obtainable by reverse induction, as detailed in the Appendix. These are summarized in Table 2, using the additional notation $\Lambda_R \equiv \lambda + k_R \sqrt{\lambda(1-\lambda)}$, $\Lambda_M \equiv \lambda + k_M \sqrt{\lambda(1-\lambda)}$ and $\bar{\alpha} \equiv \Lambda_M \alpha_l + (1 - \Lambda_M) \alpha_h$. This generalizes PP's Table 3 and provides new results, particularly for the retailer's value function.

These results will be used in the following subsections to explore the implications of risk sensitivity for distribution policy design. Since the main contribution of this research is the formalization of the channel partners' sensitivity to risk, attention is first directed towards how such sensitivity affects the behavior of each party under both policies. This is discussed in subsection 3.1, providing a basis for interpreting subsequent findings. Subsection 3.2 will then provide a direct comparison between the equilibrium decisions and outcomes under the different policies, culminating in an assessment of when each party will prefer one policy over the other. Subsection 3.3 will quantify the manufacturer's penalty for erroneously assuming the retailer to be risk-neutral, thus underscoring the importance of acknowledging risk sensitivity. This will suggest an informational motive for the use of return policies. Finally, subsection 3.4 will reveal how the manufacturer-retailer balance of power influences channel behavior in the presence of sensitivity to risk by allowing the retailer to dictate the terms of trade.

This progression will offer some answers to the major questions motivating this research: (1) When is each party better off facing demand risk directly, rather than seeking to shift this burden to the other party? (2) How does each party's relative sensitivity to risk influence any compensation given or received as part of a transfer of risk?

Table 2

Equilibrium for each distribution policy

(Note: $\Lambda_R \equiv \lambda + k_R \sqrt{\lambda(1-\lambda)}$, $\Lambda_M \equiv \lambda + k_M \sqrt{\lambda(1-\lambda)}$, and $\bar{\alpha} \equiv \Lambda_M \alpha_l + (1 - \Lambda_M) \alpha_h$)

Variable	No Returns ($j = nr$)	Full Returns ($j = fr$)
wholesale price: w_j	$\frac{(1 - \Lambda_R)\alpha_h + \beta c}{2\beta}$	$\frac{\bar{\alpha} + \beta c}{2\beta}$
retail order quantity: s_j	$\frac{(1 - \Lambda_R)\alpha_h - \beta c}{4(1 - \Lambda_R)}$	$\frac{2\alpha_h - \bar{\alpha} - \beta c}{4}$
price in low demand scenario: $p_{j,l}$	$\frac{\alpha_l}{2\beta}$	$\frac{2\alpha_l + \bar{\alpha} + \beta c}{4\beta}$
price in high demand scenario: $p_{j,h}$	$\frac{3(1 - \Lambda_R)\alpha_h + \beta c}{4\beta(1 - \Lambda_R)}$	$\frac{2\alpha_h + \bar{\alpha} + \beta c}{4\beta}$
demand in low demand scenario: $q_{j,l}$	$\frac{\alpha_l}{2}$	$\frac{2\alpha_l - \bar{\alpha} - \beta c}{4}$
demand in high demand scenario: $q_{j,h}$	$\frac{(1 - \Lambda_R)\alpha_h - \beta c}{4(1 - \Lambda_R)}$	$\frac{2\alpha_h - \bar{\alpha} - \beta c}{4}$
manufacturer value function: V_j^M	$\frac{[(1 - \Lambda_R)\alpha_h - \beta c]^2}{8\beta(1 - \Lambda_R)}$	$\frac{(\bar{\alpha} + \beta c)^2 - 4\beta c \beta_h}{8\beta}$
retailer value function: V_j^R	$\frac{\Lambda_R(\alpha_l)^2 + \frac{[(1 - \Lambda_R)\alpha_h - \beta c]^2}{4(1 - \Lambda_R)}}{4\beta}$	$\frac{\Lambda_R(2\alpha_l - \beta c)^2 + (1 - \Lambda_R)(2\alpha_h - \beta c)^2}{16\beta}$

How sensitivity to risk affects channel behavior

We can gain better understanding of the dynamics of how risk sensitivity influences the behaviors of the channel members under each policy by examining the comparative statics for the equilibria. These results are documented in Table 3, and follow from Table 2.

In the case of no returns, only the retailer's risk sensitivity (k_R) matters; the manufacturer's risk sensitivity (k_M) has no effect because the manufacturer encounters no uncertainty. On becoming more risk averse the retailer orders more conservatively (s decreases). The manufacturer lowers the wholesale price to counteract this reduction in its sales, but is unable to do so enough to avoid a net reduction in its value function.

Interestingly, though, the retailer is not necessarily worse off for having a larger k_R . If the high outcome for primary demand is not "too much higher" than the low primary demand, then a more risk-averse retailer actually does bet-

ter. (The specific condition is $(\alpha_h)^2 < 4(\alpha_l)^2 + \left(\frac{\beta c}{1 - \Lambda_R}\right)^2$. This is true, for example, when α_h is less than double α_l .) This is true even though the retailer is a strategic follower here, and can be explained as follows. As with any reasonable measure of risk aversion, increasing k_R has a directly negative effect on the retailer's value function when all else is equal. However, as noted, this also leads the manufacturer to reduce w , benefiting the retailer. The net effect depends on the relative magnitudes of these two countervailing forces. The closer α_h is to α_l , the smaller will be the variability in retailer profit, mitigating the direct impact of k_R on the retailer's objective. The stated condition defines when this is outweighed by the wholesale price reduction.

The full return policy reflected in the last two columns of Table 3 provides another point of reference. Under this regime the manufacturer accepts full exposure to market

Table 3

Effect of risk sensitivity on the equilibria

Variable	No Returns ($j = nr$)		Full Returns ($j = fr$)	
	Increasing k_R	Increasing k_M	Increasing k_R	Increasing k_M
wholesale price: w_j	↓	no effect	no effect	↓
retail order quantity: s_j	↓	no effect	no effect	↑
price in low demand scenario: $p_{j,l}$	no effect	no effect	no effect	↓
price in high demand scenario: $p_{j,h}$	↑	no effect	no effect	↓
demand in low demand scenario: $q_{j,l}$	no effect	no effect	no effect	↑
demand in high demand scenario: $q_{j,h}$	↓	no effect	no effect	↑
manufacturer value function: V_j^M	↓	no effect	no effect	↓
retailer value function: V_j^R	↑ for sufficiently small α_h , ↓ otherwise	no effect	↓	↑

Table 4

Effect of the distribution policy on the system equilibrium
(Note: $\Lambda_R \equiv \lambda + k_R \lambda (1 - \lambda)$ and $\Lambda_M \equiv \lambda + k_M \lambda (1 - \lambda)$)

Variable	Directional Impact Of Allowing Full Returns (= $\text{sign}[\Delta_{nr \rightarrow fr}(\cdot)]$)
wholesale price: w_j	$\text{sign} \left[\frac{\Lambda_R}{\Lambda_M} - \left(1 - \frac{\alpha_l}{\alpha_h} \right) \right]$
retail order quantity: s_j	positive
price in low demand scenario: $p_{j,l}$	positive
price in high demand scenario: $p_{j,h}$	negative
demand in low demand scenario: $q_{j,l}$	negative
demand in high demand scenario: $q_{j,h}$	positive

risk. Retailer risk sensitivity does not affect the retailer's decisions since in effect these need not be determined until after the resolution of uncertainty (although k_R does impact the valuation of the yet uncertain payoff). A risk-averse manufacturer cuts w to increase the retailer stock level, depressing the retail price in both demand scenarios. This closes the gap between the manufacturer's possible profit outcomes, reducing the standard deviation⁹. But the *mean* profit also drops, reducing the total manufacturer value. Concessions made by the manufacturer due to its own concern with risk turn out to benefit the retailer.

Overall, even in spite of controlling the distribution policy, the manufacturer derives no benefit from any form of risk aversion in the channel under either policy. However, the retailer can benefit from *manufacturer* risk aversion when the policy exposes the manufacturer to risk, and from its *own* risk aversion if the variability in demand is sufficiently small. A managerial implication is that all else equal, a retailer should seek manufacturers that are more sensitive to risk if they offer return privileges. Of course, such manufacturers may be less willing to offer such policies, as we will see in Proposition 1.

How the distribution policy affects channel behavior with risk sensitivity, and who benefits

In this section we perform a direct comparison between the equilibrium decisions and outcomes under the different policies, which PP did not fully pursue for the risk-neutral setting. Since PP's model is a special instance of ours, this section contains new results for their setting as well.

For any variable x_j , define the incremental change in the equilibrium value of that variable on installing a return policy to be $\Delta_{nr \rightarrow fr}(x_j) \equiv x_{fr} - x_{nr}$. The signs of $\Delta_{nr \rightarrow fr}(x_j)$ for all decision variables are characterized in Table 4.

Table 4 reports that the installation of a return policy affects all retail decisions in a direction that is invariant to specific risk attitudes anywhere in the channel. The retailer will unequivocally order more in the pursuit of sales volume in the high demand scenario (albeit at a lower retail price), due to the support provided for a higher price when demand

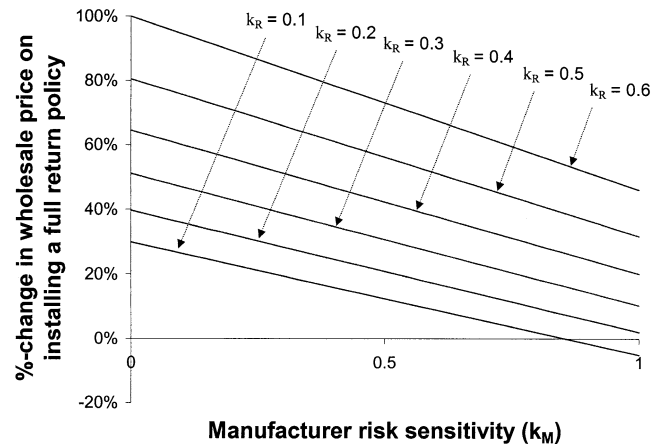


Fig. 1. Percentage change in wholesale price on installing a full return policy. (Note: k_R is the retailer's risk sensitivity.)

is low (since returning excess to the manufacturer avoids over-depressing the inventory-clearing price). Since $\Delta_{nr \rightarrow fr}(p_{j,l}) > 0$ and $\Delta_{nr \rightarrow fr}(p_{j,h}) < 0$, clearly the return policy reduces the dispersion of retail prices across demand states, as PP found (Padmanabhan and Png, 1997, p.90).

The impact on the wholesale price *does* depend on the relative risk attitudes, which conflicts with PP's conclusion that the manufacturer will *always* require a higher w as compensation for the risk accompanying a return policy. Their result applies when the manufacturer is sufficiently insensitive to risk (small k_M). The actual condition depends on the retailer's risk attitude as well, since this determines the magnitude of the retailer's reaction to the return policy, hence the manufacturer's prospects¹⁰. Indeed, in the extreme case of a risk-neutral manufacturer ($k_M = 0$), undeniably the return policy increases the wholesale price (since then $\Lambda_R/\Lambda_M > 1 > 1 - \alpha_l/\alpha_h$, and the more risk sensitive the retailer the bigger the increase. However, when the manufacturer is sufficiently risk sensitive (i.e., k_M is large and therefore Λ_R/Λ_M is small), the manufacturer's pricing will be more greatly influenced by the uncertainty brought on by the return policy. As noted earlier, the standard deviation of the manufacturer's profit increases with w . Hence, if pressed to offer a return policy, the manufacturer might also cut the wholesale price as a defensive response to *variation* in profit, even if this would compromise the *mean* profit.

These properties are illustrated in Fig. 1 for a representative example with $\alpha_h = 12$, $\alpha_l = 6$, $\beta = 1$, $\lambda = 0.3$, and $c = 0$. This reports the percentage change in the equilibrium wholesale price (formally, this is $\Delta_{nr \rightarrow fr}(w_j)/w_{nr}$) that will occur when a full return policy is installed, highlighting the influence of each party's sensitivity to risk. The $k_R = 0.1$ curve illustrates that a sufficiently risk sensitive manufacturer will reduce the wholesale price when introducing a return policy. (This occurs here once k_M exceeds approximately 0.85)

We next comment on when the manufacturer will favor the return policy.

The manufacturer's preferences

The manufacturer's policy preferences are determined by the sign of $\Delta_{nr \rightarrow fr}(V_j^M)$. As seen in the Appendix (see proof of Proposition 1), this is a complicated expression without an obvious interpretation. For insight we follow PP in considering the special case of $c = 0$, as stated in Proposition 1.

PROPOSITION 1. *Manufacturer preferences regarding the distribution policies, when $c = 0$: (a) The manufacturer prefers to allow full returns for full credit if and only if*

$$\frac{\alpha_h}{\alpha_l} < \frac{\Lambda_M}{\sqrt{1 - \Lambda_R - (1 - \Lambda_M)}} \equiv T$$

where $\Lambda_R \equiv \lambda + k_R \sqrt{\lambda(1 - \lambda)}$ and $\Lambda_M \equiv \lambda + k_M \sqrt{\lambda(1 - \lambda)}$.

(b) *The more risk sensitive the retailer, the larger the set of conditions in which the manufacturer will prefer to allow returns.*

(c) *The more risk sensitive the manufacturer, the smaller the set of conditions in which the manufacturer will prefer to allow returns.*

(d) *Regardless of the relative risk sensitivities, the manufacturer will benefit from allowing returns only by increasing the wholesale price relative to the case of no returns.*

The manufacturer's policy selection is simply a matter of choosing one's poison, as market uncertainty creates pain one way or another. The manufacturer can bear the consequences *indirectly* through the retailer's reactions (by disallowing returns), or opt for the *direct* exposure (by offering return privileges). Part (a) of Proposition 1 generalizes PP's Proposition 2 and states that the relative magnitudes of the two evils are determined by the range of market uncertainty, expressed through the ratio α_h/α_l . In particular, the manufacturer is better off directly facing market uncertainty that is sufficiently small, as defined by values of α_h/α_l that do not exceed the threshold T . Part (b) suggests that the retailer's risk sensitivity magnifies the former evil, so that a return policy becomes more attractive. Meanwhile, part (c) reports that the greater the manufacturer's aversion to the uncertain, the weaker the manufacturer's desire to face that uncertainty *directly*. The importance of this progression is in framing the motivation for the use of return policies in terms of risk attitudes. Part (d) formalizes the mechanism by which a manufacturer is compensated for exposure to risk, and rules out the possibility that a return policy can be viable solely due to the resulting increase in sales volume (which applies to PP's setting as well). It also confirms to the retailer that although return privileges may seem desirable at face value, there is no free lunch.

These properties are illustrated in Fig. 2 for the same

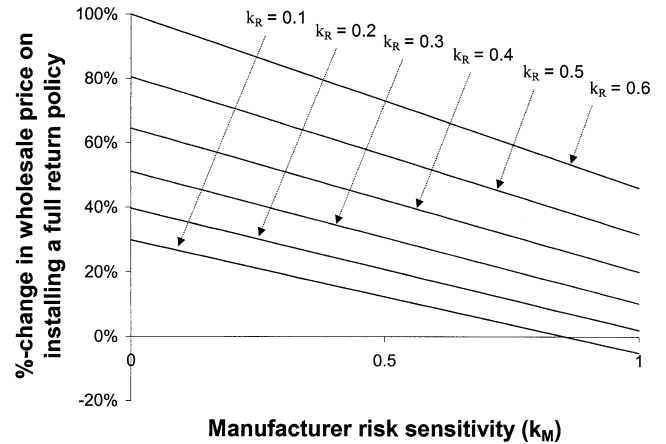


Fig. 2. Percentage change in manufacturer's value function on installing a full return policy. (Note: k_R is the retailer's risk sensitivity.)

example considered in Fig. 1. This reports the percentage change in the manufacturer's value function (formally, this is $\Delta_{nr \rightarrow fr}(V_j^M)/V_{nr}^M$) that will occur when a full return policy is installed. Portions of the curves above the horizontal axis correspond to circumstances under which the return policy benefits the manufacturer. When each curve intersects this axis, the condition in part (a) of Proposition 1 holds with equality. We next explore the retailer's perspective

The retailer's preferences

A retailer without channel power simply decides whether to do business under the terms proposed by the manufacturer. Table 2 indicates that the retailer's value function will be positive under either policy, so the retailer will participate willingly. Nevertheless, it is worth investigating how the retailer is affected by a return policy since this will determine whether the retailer might seek out other channel partners or propose alternative policies in the long run. These results also provide a benchmark for the case of a powerful retailer, which we will analyze later. Proposition 2 reports conditions under which a manufacturer-specified return policy will benefit the retailer.

PROPOSITION 2. *The retailer's preferences regarding distribution policies, when $c = 0$:*

(a) *The retailer prefers full returns for full credit if and only if $\alpha_h/\alpha_l > T$, for T as defined in part (a) of Proposition 1.*

(b) *The more risk sensitive the retailer, the smaller the set of conditions in which the retailer will prefer return privileges.*

(c) *The more risk sensitive the manufacturer, the larger the set of conditions in which the retailer will prefer return privileges.*

Part (a) of Proposition 2 is significant for a number of reasons. While for a fixed wholesale price a retailer always prefers a return policy for insurance value, this is not necessarily so when the manufacturer's adjustment of w is

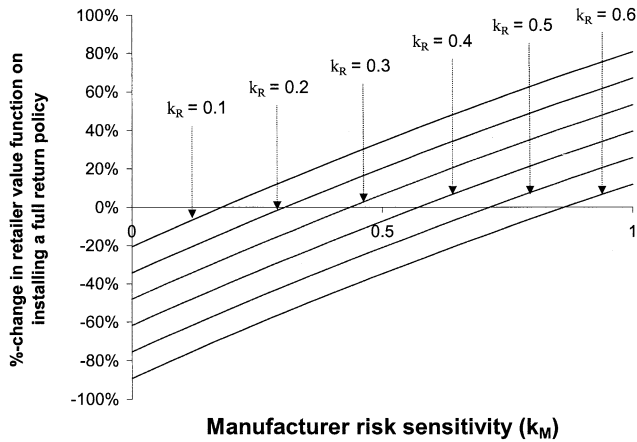


Fig. 3. Percentage change in retailer's value function on installing a full return policy. (Note: k_R is the retailer's risk sensitivity.)

taken into consideration. In fact, we have identified circumstances under which the seemingly counterintuitive can occur: the manufacturer may wish to offer full returns for full credit, but the retailer objects. Comparing part (a) to Proposition 1 indicates that the two parties' preferences are in conflict. That is, while the environmental parameters determine which party will support the return policy, there is no "win-win" outcome. While the strictness of the result is likely an artifact of the modeling assumptions, this suggests why return policies are sometimes implemented in intermediate forms, such as partial returns for full credit or full returns at partial credit. Part (b) reflects the fact that the more sensitive the retailer is to risk, the more the manufacturer will charge for the return privilege (via w). However, as the manufacturer becomes more risk sensitive, it will reduce the wholesale price attached to the return policy so as to control the variability in its own profit. This benefits the retailer, as part (c) reports.

These properties are illustrated in Fig. 3 for the same example considered in the previous two figures. This reports the percentage change in the retailer's value function (formally, this is $\Delta_{nr \rightarrow fr}(V_j^R)/V_{nr}^R$ that will occur when a full return policy is installed. Portions of the curves above the horizontal axis correspond to circumstances under which the return policy benefits the retailer. When each curve intersects this axis, the condition in part (a) of Proposition 2 holds with equality. Note that the crossover point is the same in both Figs. 2 and 3 for each k_R value, although the crossings occur in opposite directions. This depicts the strict conflict between the channel partners' preferences.

The manufacturer's penalty for ignoring risk sensitivity

In this section we underscore the importance of acknowledging risk sensitivity by quantifying the manufacturer's penalty for erroneously assuming the retailer to be risk-neutral. This will suggest a motive for the use of return policies that apparently has not previously been considered.

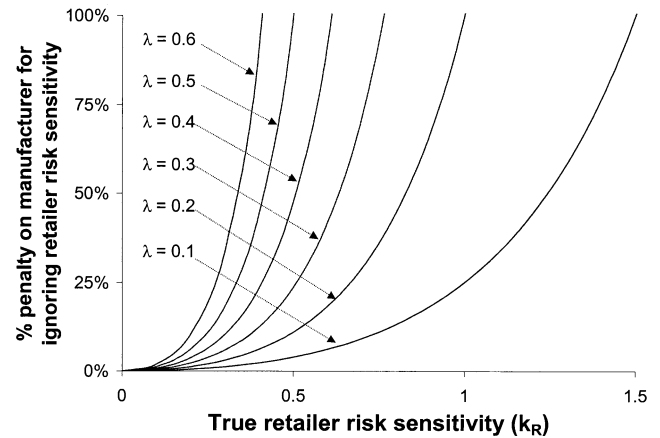


Fig. 4. Lower bound on manufacturer's percentage penalty for ignoring retailer risk sensitivity (no returns). (Note: λ is the probability of the low demand scenario.)

When returns are disallowed, the manufacturer's wholesale price should properly reflect the retailer's risk sensitivity as delineated in Table 2. Suppose instead that the manufacturer uses a strategy suitable for a risk-neutral retailer ($k_R = 0$), but that the retailer's actions reflect the true, strictly positive k_R . The decision impacting the manufacturer is the retailer's order, that is, $[(1 - \Lambda_R)\alpha_h - \beta w]/[2(1 - \Lambda_R)]$. The manufacturer will choose a wholesale price, denoted as $\bar{w}_{nr} = [(1 - \lambda)\alpha_h + \beta c]/[2\beta]$, that is inappropriately high. The resulting value to the manufacturer will be

$$V_{nr}^M(\bar{w}_{nr}) = \frac{[(1 - \Lambda_R)\alpha_h - \beta c]^2}{8\beta(1 - \Lambda_R)} - \frac{(\alpha_h k_R)^2 \lambda (1 - \lambda)}{8\beta(1 - \Lambda_R)}$$

where the first term is what would result from correct wholesale pricing and the second term is the penalty for ignoring the retailer's risk sensitivity. Naturally the latter vanishes when the retailer is truly risk-neutral ($k_R = 0$). The penalty can be expressed in percentage terms as

$$[\% \text{ penalty}] = \left(\frac{\alpha_h k_R}{(1 - \Lambda_R)\alpha_h - \beta c} \right)^2 \lambda (1 - \lambda).$$

While this can be computed for any specific set of parameters, for illustration we plot this in Fig. 4 as a function of the true k_R when $c = 0$ (which provides a lower bound on $[\% \text{ penalty}]$)¹¹.

While a manufacturer is unlikely to be completely ignorant of its channel partner's concern for risk, and learning will occur over time, this analysis shows that estimation error has a nontrivial impact¹². The issue is problematic since inferring another firm's risk attitudes could be challenging. While this would be a concern even if the problem were merely one of estimating an unknown, our analysis suggests that the retailer has incentive to feign risk sensitivity by whatever means possible. (We learned earlier that the manufacturer's proper response to retailer risk sensitivity is to reduce the wholesale price.)

Table 5
Effect of risk sensitivities when the retailer is channel leader

Variable	No Returns ($j = nr$)		Full Returns ($j = fr$)	
	Increasing k_R	Increasing k_M	Increasing k_R	Increasing k_M
wholesale price: w_j	no effect	no effect	no effect	↑
manufacturer value function: V_j^M	no effect	no effect	no effect	no effect
retailer value function: V_j^R	↓	no effect	↓	↓

One managerial response might be to pursue strategies that de-emphasize the need for such information. Allowing retailer returns can achieve this. This is apparent from the rightmost column of Table 2, in which no operational decisions depend on k_R at all. The return policy renders all retail decisions risk-invariant, so that the manufacturer needs only comprehend its own risk attitude¹³. This issue should therefore increase the set of circumstances under which the manufacturer will prefer to offer returns (cf. Proposition 1). Naturally, the existing literature has not addressed the possibility that return policies can have lower informational needs since estimation of risk attitudes is not an issue under the premise of risk-neutral decision-making.

Policy preferences when the retailer controls the channel

Thus far we have assumed the manufacturer to dictate the distribution policy, the only case considered by PP and many others in the related literature. Conversely, a powerful retailer might drive the terms of trade. Juxtaposing these alternatives will suggest how the manufacturer-retailer balance of power influences channel behavior in the presence of sensitivity to risk.

In addition to choosing the order quantity and retail price as before, the retailer will dictate the wholesale price to maximize its own value function. Hence the retailer controls all decisions, with the only constraint being that the manufacturer's value function must achieve some threshold that assures participation. Without loss of generality, we assume that threshold to be zero. While the change in control structure may alter the retailer's channel policy preferences relative to the case of the powerful manufacturer (cf. subsection 3.2, in particular Proposition 2), clearly this can never make the retailer worse off since a decision-maker given control over more decision variables always has the option to just maintain their previous values.

Proposition 3 describes the wholesale price that will prevail under each policy and the impact on preferences.

PROPOSITION 3. *When the retailer is the channel leader with respect to the distribution policy:*

(a) *When no returns are allowed, the wholesale price will be exactly c .*

(b) *When full returns for full credit are allowed, the wholesale price will be*

$$\frac{(\bar{\alpha} + \beta c) - \sqrt{(\bar{\alpha} + \beta c)^2 - 4\beta\alpha_h c}}{2\beta}$$

which is strictly greater than c .

(c) *Risk sensitivities affect the wholesale price and value functions as detailed in Table 5.*

A comparison of parts (a) and (b) to Table 2 confirms that channel power is beneficial. The dominant retailer's advantage is manifested in wholesale prices that are lower under each policy than what the manufacturer would prefer. As evident from comparing part (c) to Table 3, the dominant retailer no longer benefits from the risk sensitivity of manufacturers who offer return policies. As channel leader the retailer has the responsibility for insuring a certain outcome for the follower. Here this entails paying a positive margin to compensate a manufacturer for the risk exposure associated with accepting returns (as reported in part (b)), and raising this margin for the more risk sensitive manufacturers. A property that is common across power structures is that greater risk sensitivity at the channel leader's partner is detrimental to the leader's well being.

We next consider how the reversal in the balance of power affects preferences towards return policies. The weak manufacturer will receive zero value under either policy, so we focus only on the retailer's prospects. We do this by studying the sign and magnitude of the percentage change in the retailer's value function due to implementing a return policy, that is, $\Delta_{nr \rightarrow fr}(V_j^R)/V_{nr}^R$.

The equilibrium value functions are sufficiently complex that no simple answers are available analytically. However, certain properties become evident under additional assumptions. To simplify demand uncertainty we assume α_l and α_h to be equally likely (i.e., $\lambda = 1/2$). The midpoint between α_l and α_h is denoted as $\bar{\alpha}$, and we use δ to measure each point's deviation from $\bar{\alpha}$ (i.e., $\delta = \alpha_h - \bar{\alpha} = \bar{\alpha} - \alpha_l$). Varying δ for a fixed $\bar{\alpha}$ isolates the effect of demand variability. An extensive set of numerical experiments was performed under these conditions. Representative findings are illustrated in Fig. 5, which assumes $k_R = 0.3$, $\bar{\alpha} = 9$, $\beta = 1$, and $c = 1$.

Fig. 5 suggests two conditions that lead the dominant retailer to prefer the return policy. (1) *A very uncertain market, corresponding to a large δ (indicating a large spread between α_h and α_l).* All else equal, the more uncertain the demand, the more likely the retailer will want to offload the risk on another party. At the same time, Proposition 1(a) indicates that when the manufacturer has channel power, a return policy will be less likely to occur when α_h is much larger than α_l . Hence the final outcome will depend

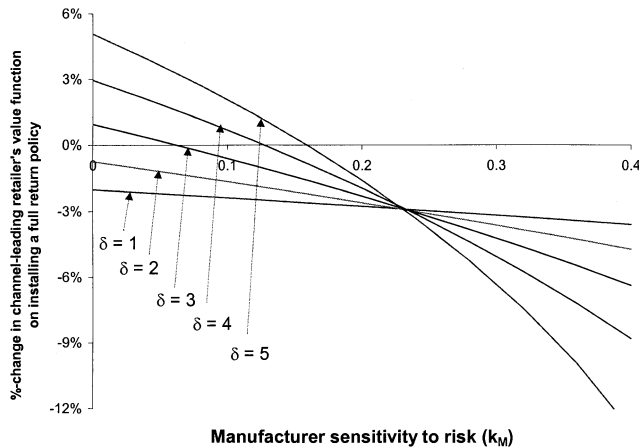


Fig. 5. Channel-leading retailer's percentage benefit from installing a return policy. (Note: δ is a measure of demand uncertainty.)

strongly on the balance of strategic power. (2) A manufacturer that is relatively insensitive to risk (small k_M). Such a manufacturer requires less compensation for the risk created by allowing returns. This property is unaffected by the channel power structure (see Proposition 1(c)).

As would be expected, the larger the k_M , the larger the uncertainty (δ) required for the dominant retailer to prefer returns. Furthermore, as δ increases, the effect of k_M becomes more pronounced.

The analysis thus far has treated channel power and risk sensitivity as independent. However, in reality it is likely that the strategic factors that allow a party to dominate a channel (e.g., depth/breadth of activities) will render that party less risk sensitive than its channel partners. Fig. 5 would then suggest that a powerful retailer is unlikely to insist on return policies, since this appears to be the outcome when $k_M > k_R$. Under similar reasoning, Proposition 1 implies that a dominant manufacturer will tend to accept returns. These conclusions are not general, though, since they ignore the effects of the other parameters.

The preceding investigation provides two major findings that defy conventional wisdom. The first is that a powerful channel member does not automatically prefer to offload risk onto its trading partners. In fact, even when the retailer controls all channel decisions, it may still do better bearing all risk itself. The second is that the retailer does not always prefer the unfettered right to return excess product, and the manufacturer does not always oppose this. While this is true when w is held fixed, it is not so in general.

Finally, revisiting the possibility that other parties' risk sensitivity might not be perfectly observable again reveals an informational issue that is salient to the channel leader's choice of distribution policy. It is apparent from Proposition 3(c) that the manufacturer has positive incentive to feign risk aversion when the retailer demands a return policy, since this will increase the corresponding wholesale price. (A chart similar to Fig. 4 would illustrate the retailer's penalty for overestimating the manufacturer's risk sensitiv-

ity, but we omit this for space considerations.) As in the case of the powerful manufacturer, this suggests an informational motive for the channel leader to bear the risk itself. The difference here is that the policy that accomplishes this is one that disallows returns.

Conclusion

This research was motivated by the thesis that any understanding of the risk-sharing rationale for multiparty business agreements is incomplete if all affected parties are presumed to be risk-neutral. This shortcoming has been addressed here in the context of distribution channel relationships, identifying a number of issues that can guide future empirical investigation.

The specific approach was to formalize how sensitivity to risk (in the context of MS preferences) affects behaviors and outcomes on both sides of a manufacturer-retailer supply relationship, and how these dynamics are altered by a manufacturer return policy. This investigation studied not only a manufacturer-dominated channel, as PP and others have considered, but also one defined by a powerful retailer. Juxtaposing these has generated new insights about how channel power interacts with risk preferences.

The analysis has articulated how each party will act, and circumstances under which each will prefer the return policy. These findings are interpretable in terms of the relative risk sensitivities of the parties and the market conditions. The explicit consideration of risk sensitivity has yielded a number of meaningful managerial insights, including the following:

- The penalty for errors in estimating a channel partner's sensitivity to risk can be substantial. The issue is problematic since inferring another firm's risk attitudes could be challenging. While this would be a concern even if the problem were merely one of estimating an unknown, our analysis suggests that the party with less leverage in choosing the channel policy has incentive to feign risk sensitivity by whatever means possible. This can favor the policy that avoids this informational issue.
- Risk sensitivity leads to behaviors that can differ qualitatively from those predicted by risk-neutral analysis. For example, if a channel-leading manufacturer is sufficiently averse to risk, the unit wholesale price charged might actually be lower under a return policy than when no returns are allowed. This is a risk sensitive manufacturer's response to the profit variability induced by retailer returns.
- Being risk sensitive does not always make a party worse off. Net gain might result from how the other party compensates for its partner's reaction to risk.
- A powerful channel member does not automatically prefer to offload channel risk on its trading partners.

In particular, a risk sensitive retailer does not always prefer the right to return excess product for full credit, and a risk sensitive manufacturer does not always oppose this. While this is true when the wholesale price is held fixed, it is not true in general.

- In choosing from among potential channel partners a firm should consider their relative risk attitudes, to the extent that these attributes may be ascertained. For instance, a weak retailer should seek manufacturers that are more sensitive to risk if they offer return privileges. Such manufacturers will tend to be more accommodating on the wholesale price.

While the intent was to use as simple a model as possible to highlight an issue of import, this work is obviously limited by the particular assumptions applied. This also suggests the following areas for future research:

- (1) The presence of multiple retailers, possibly differing in risk sensitivity or other attributes, would lead to an interesting but much more complex modeling challenge¹⁴. Per Robinson-Patman considerations, the manufacturer would seek to specify a single distribution policy that would apply to all the retailers, or possibly allow the retailers to choose from the same menu of options. In the latter case the optimal design of the policy would require orchestrating which of the retailers would opt for return privileges. Another salient issue would be risk pooling that might mitigate the manufacturer's cost of accepting returns, although the general effect would depend on how the retailers competitively interact, as PP's work would suggest.
- (2) Multiple manufacturers might bid for the business of a retailer using a return policy in addition to pricing terms. This setting could also give rise to a more sophisticated type of risk-management strategy in which the retailer would allocate its total purchase across a portfolio of manufacturers, with return policies in effect for only some of the purchases.
- (3) Because the results in this report are specific to the MS value function, consideration of alternative representations remains an open area.
- (4) These results assume the specific form of demand suggested by PP. They also rely on common knowledge of all parameters. In fact, the manufacturer and retailer might have different information/beliefs on a variety of factors, especially the market uncertainty. This would create additional uncertainties for risk sensitive entities to consider and attempt to manage, especially where incentive for deliberate deception exists (as we determined to be true of risk sensitivities).
- (5) Analysis was restricted to the two specific policies that PP contrasted. As discussed, an intermediate policy (e.g., full returns at partial credit, or partial

returns at full credit) might provide mutual benefit to the channel partners.

- (6) More complex wholesale pricing schemes, such as quantity discounts or two-part tariffs, could conceivably interact with risk preferences in meaningful ways.
- (7) This framework could potentially be applied to other distribution policies intended to share risk, such as price protection plans (e.g., Lee et al., 2000) or slotting allowances (e.g., Lariviere and Padmanabhan, 1997; Desiraju, 2001).

Notes

1. Expected profit maximization is a useful theoretical construct that lends tractability to many behavioral models. This is usually rationalized as being an appropriate representation for firms, whose assets, diversification, and long-term perspectives remove the risk-sensitivity that usually appear in individuals. However, there is a long history of evidence suggesting that even large firms often do not conduct themselves in a strictly risk-neutral fashion (see, for instance, Lanzilotti, 1958 and Swalm, 1966). One cause of this might be that accountability for a firm's decisions still resides with individual decision makers, whose behavior can easily be colored by the resulting uncertainty for their personal livelihoods and professional futures. Another is that the performance of firms (especially those that are publicly traded) is increasingly judged through a myopic lens that does not forgive unpleasant surprises.
2. PP provide two models. In the first, which studies a manufacturer return policy in the presence of retail competition, the absence of uncertainty precludes discussion of risk. The second (in their §6) makes demand uncertainty explicit while suppressing the feature of retail competition. This research builds upon the latter model.
3. This interpretation of k is advocated by the framework of Meyer (1987), in which this function displays "constant risk aversion." This representation "dollarizes" a party's aversion to risk, and the second term can be viewed as a risk premium since it penalizes based on the magnitude of uncertainty. $k = 0$ restores risk-neutrality.
4. A related form with similar interpretation is $\{E[Z] - k \text{Var}[Z]\}$. This has been used extensively in the literatures of finance (cf. Markowitz, 1952; 1959) and economics (cf. Robison and Barry, 1987). The MS form is used in this model for mathematical tractability.
5. If one were to calculate an average "sell-through" price in the newsvendor model (averaged across the

units sold at full price and the units salvaged), this would be seen to drop with the amount over ordered.

6. Roy's model: *minimize* $\text{Prob}\{\text{Profit} \leq T\}$ for a given T . *Telser's model*: *maximize* $E[\text{Profit}]$ such that $\text{Prob}\{\text{Profit} \leq T\} \leq \alpha$ for given T and α . *Kataoka's model*: *maximize* T such that $\text{Prob}\{\text{Profit} \leq T\} = \alpha$ for a given α .
7. Lau (1980) looks at how a newsvendor behaves for a variety of objective functions other than expected profit. For the MS value function one can articulate the equation that implicitly defines the optimal order Q or a random demand D . However, this requires evaluation of the mean and standard deviation of the random variable representing actual sales, that is, $A \equiv \min[Q, D]$, and closed forms for these moments are available only for a specific exotic class of demand distributions. Even then, obtaining the optimal Q requires numerical solution of a complex equation. Lau does show that the optimal Q is lower for this value function than would be preferred by the risk-neutral newsvendor. (Magee, 1975 also comes to this conclusion.) He achieves similar results with a polynomial utility function.
8. While more complex pricing schemes may be more profitable in theory, practical considerations lead many firms to use as simple of forms as possible. According to Ingene and Parry (1995) "(1) Complex schedules involve relatively greater administration, bargaining, and contract development costs; (2) ascertaining optimal quantity levels imposes nontrivial information acquisition costs on manufacturers, especially as the number of retailers increases; and (3) complex schedules may generate negative goodwill and, in extreme cases, lead to lawsuits."
9. For fixed wholesale price w the standard deviation of manufacturer profit is $\sqrt{\lambda(1-\lambda)}(\pi_{fr,h}^M(w) - \pi_{fr,l}^M(w))$ (cf. Lemma 1 in the Appendix), and $\pi_{fr,h}^M(w) - \pi_{fr,l}^M(w) = (\alpha_h - \alpha_l)w/2$. So the standard deviation moves with the wholesale price.
10. On their p.93, PP conjecture that retailer risk aversion would increase the w associated with a return policy, which is consistent with our result. However, they do not entertain the possibility of risk sensitivity on the manufacturer's part.
11. These curves provide worst-case analysis for *underestimating* the true k_R . One could generate similar graphs for the penalty for *overestimating* k_R , and analogous general properties would result.
12. As is shown graphically and is demonstrable analytically, the percentage penalty for ignoring the retailer's risk-sensitivity increases dramatically with the true k_R for any λ . In fact, the penalty becomes arbitrarily large (100% or even greater) as k_R approaches its maximum allowable value of $\sqrt{(1-\lambda)/\lambda}$. Also, increasing the value of λ raises the penalty curve. This is because the low demand outcome, whose probability is λ , is the scenario that creates risk for the retailer and in turn depresses the amount purchased from the manufacturer. Therefore, an increase in λ will broaden the discrepancy between the behavior of a risk-neutral retailer and a risk-sensitive one, exacerbating the manufacturer's penalty for failing to adjust for risk-sensitivity.
13. This model does not consider partial return policies, but one could speculate that these might at least partially mitigate the problem. However, this is not definite. The retailer makes decisions not to minimize risk, but to optimally trade off risk and return. Thus, there is no guarantee that the retailer will necessarily take on less absolute risk under a partial return policy than a full return policy, especially since the manufacturer's wholesale price will also differ between the cases. Since the retailer's risk sensitivity still influences retail-level decisions as long as less than the full initial purchase can be returned, the estimation issue will persist.
14. No existing work on return policies in a stochastic environment (using either the PP or newsvendor paradigm) has explicitly handled multiple retailers, even without consideration of risk-sensitivity. The technical difficulty is that the demand distribution faced by the manufacturer is substantially more complex when driven by multiple retail markets, even if they do not overlap. Even with a simple two-point market demand uncertainty, two scenarios for one retailer become up to 2^n scenarios when treating multiple retailers. Using continuous distributions for the retail demand does not help much since the demand distribution generally becomes intractable when filtered through the retailer's inventory policy. This is a well-known finding in the multiechelon inventory literature.

Acknowledgments

I have benefited from many conversations with Karen Donohue on this subject. Marty Lariviere, Paddy Padmanabhan, and Terry Taylor graciously provided detailed commentary on an early draft of this manuscript. The guidance provided by all individuals involved with the editorial process, Editor L.P. Bucklin in particular, has been invaluable.

Appendix

DERIVATION OF Table 2. The equilibria detailed in Table 2 are obtained by reverse induction analysis similar to

that used by PP. The following Lemma provides a property of MS value functions that is central to this analysis.

LEMMA 1. *Let X be some random variable that takes values x_l with probability λ and x_h with probability $(1 - \lambda)$, where $x_h > x_l$ and let $V(X) \equiv E[X] - k \cdot \text{StdDev}[X]$. Then $V(X) = \Lambda x_l + (1 - \Lambda)x_h$ where $\Lambda \equiv \lambda + k\sqrt{\lambda(1 - \lambda)}$.*

PROOF OF LEMMA 1. Note that $E[X] = \lambda x_l + (1 - \lambda)x_h$ and $E[X^2] = \lambda(x_l)^2 + (1 - \lambda)(x_h)^2$, so that $\text{Var}[X] = E[X^2] - (E[X])^2 = \lambda(1 - \lambda)(x_h - x_l)^2$. Thus $V(X) \equiv E[X] - k\sqrt{\text{Var}[X]} = (\lambda + k\sqrt{\lambda(1 - \lambda)})x_l + (1 - (\lambda + k\sqrt{\lambda(1 - \lambda)}))x_h$.

An implication of Lemma 1 is that the value to an MS decision-maker of a stochastic payoff with a two-point distribution is mathematically similar to an expectation of that payoff. However, the true probability of the less desirable outcome (λ) is replaced by a larger “risk-adjusted probability” that we denote as Λ . That is, risk aversion has an impact comparable to increasing the weight a risk-neutral decision-maker would place on the lower outcome. Consideration is restricted to MS value functions for which $0 \leq k < \sqrt{(1 - \lambda)/\lambda}$, so that $\Lambda \in (0, 1)$, that is, Λ and $(1 - \Lambda)$ can be properly viewed as nontrivial probabilities.

The scenario with no returns will be examined first, followed by the scenario of full returns for full credit. In each case the analysis works backwards from the retailer’s pricing decision.

Independent retailer, with no return policy

- **Stage 3: The retailer sets the retail price given the available stock s .** Since the retail price is chosen after demand uncertainty is resolved and the inventory is procured, the retailer simply maximizes revenue. Hence the analysis of PP continues to apply exactly. When primary demand is low, the retailer will leave some stock unsold, with $p_{nr,l}(s) = \alpha_l/(2\beta)$ so that $q_{nr,l}(s) = \alpha_l/2$. When primary demand is high, the retailer will price to sell all stock. So $p_{nr,h}(s) = (\alpha_h - s)/\beta$ and then $q_{nr,h}(s) = s$. Thus $\pi_{nr,l}^R(s) = (\alpha_l)^2/(4\beta) - w_{nr}s$ and $\pi_{nr,h}^R(s) = (\alpha_h - s)s/\beta - w_{nr}s$. (In Remark 2 on their p.88, PP note that this strategy will be rational if $(\alpha_h - \alpha_l) \geq \beta c/(1 - \lambda)$, that is, high primary demand is sufficiently greater than low primary demand, which they assume.)
- **Stage 2: The retailer chooses s .** By Lemma 1, the retailer’s objective can be written as $V_{nr}^R(s) = \Lambda_R \pi_{nr,l}^R(s) + (1 - \Lambda_R) \pi_{nr,h}^R(s)$, where $\Lambda_R \equiv \lambda + k_R \sqrt{\lambda(1 - \lambda)}$. $V_{nr}^R(s)$ will be maximized with an order size of $[(1 - \Lambda_R)\alpha_h - \beta w]/[2(1 - \Lambda_R)]$, which we denote as $\hat{s}_{nr}(w)$. (It is straightforward to show that $d\hat{s}_{nr}(w)/dk_R < 0$ for any wholesale price, i.e., any increase in the retailer’s risk aversion decreases the retailer’s order.)

- **Stage 1: The manufacturer chooses w .** With no return policy, the manufacturer’s profit is deterministic, so that manufacturer risk sensitivity is immaterial. The manufacturer’s objective is simply $V_{nr}^M(w) = (w - c)\hat{s}_{nr}(w)$. This is clearly concave in w , and is maximized at $w_{nr} = [(1 - \Lambda_R)\alpha_h + \beta c]/(2\beta)$, which completely specifies the equilibrium.

Independent retailer, with full returns at full price

- **Stage 3: The retailer sets the retail price given the available stock s .** The return option enables the retailer to obtain exactly the ideal inventory for each realization of the primary demand since s can be chosen so as never to be a constraint. This means the revenue-maximizing level of sales can be supported in each demand scenario, that is, $q_{fr,h}(w) = (\alpha_h - \beta w)/2$ and $q_{fr,l}(w) = (\alpha_l - \beta w)/2$. (When the primary demand turns out low, all excess stock will simply be returned to the manufacturer.) The corresponding retail prices are $p_{fr,h}(w) = (\alpha_h + \beta w)/(2\beta)$ and $p_{fr,l}(w) = (\alpha_l + \beta w)/(2\beta)$. Since this is equivalent to deferring the retailer’s decisions until after all uncertainty is resolved, risk sensitivity does not affect the retailer’s ordering behavior (although it will certainly influence the *ex ante* evaluation of the value function, since the retailer continues to face an uncertain prospect). The possible profit realizations for a given w are $\pi_{fr,l}^R = (p_{fr,l}(w) - w)q_{fr,l}(w)$ and $\pi_{fr,h}^R = (p_{fr,h}(w) - w)q_{fr,h}(w)$.
- **Stage 2: The retailer chooses s .** The retailer will choose $s_{fr}(w) = q_{fr,h}(w) = (\alpha_h - \beta w)/2$ so as to have on hand exactly the optimal inventory for a high primary demand.
- **Stage 1: The manufacturer chooses w .** The manufacturer’s profit under the two demand scenarios will be $\pi_{fr,l}^M(w) = q_{fr,l}(w) \cdot w - s_{fr}(w) \cdot c = (\alpha_l - \beta w)w/2 - (\alpha_h - \beta w)c/2$ and $\pi_{fr,h}^M(w) = q_{fr,h}(w) \cdot w - s_{fr}(w) \cdot c = (\alpha_h - \beta w)(w - c)/2$. By Lemma 1, the manufacturer’s objective is

$$V_{fr}^M(w) = \Lambda_M \pi_{fr,l}^M(w) + (1 - \Lambda_M) \pi_{fr,h}^M(w) \\ = \frac{(\bar{\alpha} - \beta w)w - (\alpha_h - \beta w)c}{2}$$

where $\Lambda_M \equiv \lambda + k_M \sqrt{\lambda(1 - \lambda)}$ and $\bar{\alpha} \equiv \Lambda_M \alpha_l + (1 - \Lambda_M) \alpha_h$. V_{fr}^M is concave in w , and is maximized at $w_{fr} = (\bar{\alpha} + \beta c)/(2\beta)$, which completely specifies the equilibrium.

DERIVATION OF TABLE 4. The explicit values of the decision variables are detailed in Table 2. The appropriate differences are cataloged in Table 6, and the signs follow directly.

PROOF OF PROPOSITION 1. Installing the full return policy has a net effect on the manufacturer’s value function of the magnitude

Table 6

Values of $\Delta_{nr \rightarrow fr}(\cdot)$ for all variables(Note: $\Lambda_R \equiv \lambda + k_R \sqrt{\lambda(1-\lambda)}$, $\Lambda_M \equiv \lambda + k_M \sqrt{\lambda(1-\lambda)}$, and $\bar{\alpha} \equiv \Lambda_M \alpha_l + (1 - \Lambda_M) \lambda_h$)

Variable	Change In Value On Allowing Full Returns ($\Delta_{nr \rightarrow fr}(\cdot)$)
wholesale price: w_j	$\frac{\Lambda_M \alpha_l + (\Lambda_R - \Lambda_M) \alpha_h}{2\beta}$
retail order quantity: s_j	$\frac{\Lambda_M(\alpha_h - \alpha_l) + \beta c(\Lambda_R/(1 - \Lambda_R))}{4}$
price in low demand scenario: $p_{j,l}$	$\frac{\bar{\alpha} + \beta c}{4\beta}$
price in high demand scenario: $p_{j,h}$	$-\frac{\Lambda_M(\alpha_h - \alpha_l) + \beta c(\Lambda_R/(1 - \Lambda_R))}{4\beta}$
demand in low demand scenario: $q_{j,l}$	$-\frac{\bar{\alpha} + \beta c}{4}$
demand in high demand scenario: $q_{j,h}$	$\frac{\Lambda_M(\alpha_h - \alpha_l) + \beta c(\Lambda_R/(1 - \Lambda_R))}{4}$

$$\Delta_{nr \rightarrow fr}(V_j^M) = \frac{1}{8\beta} \left[(\bar{\alpha} + \beta c)^2 - 4\alpha_h \beta c - \frac{[(1 - \Lambda_R)\alpha_h - \beta c]^2}{1 - \Lambda_R} \right].$$

When $c = 0$, this becomes

$$\Delta_{nr \rightarrow fr}(V_j^M) = \frac{(\alpha_h)^2}{8\beta} [(\Lambda_M)^2(\alpha_l)^2 + 2\Lambda_M(1 - \Lambda_M)\alpha_l + [(1 - \Lambda_M)^2 - (1 - \Lambda_R)]]$$

This is an upward-facing parabola in α_l , and will be positive outside the two zeros, which are $\alpha_h[-(1 - \Lambda_M) \pm \sqrt{1 - \Lambda_R}]/\Lambda_M$. The negative root is immaterial since α_l is positive by definition, so the manufacturer will prefer a return policy if $\alpha_l > \alpha_h[\sqrt{1 - \Lambda_R} - (1 - \Lambda_M)]/\Lambda_M$, or, equivalently $\alpha_h/\alpha_l < \Lambda_M[\sqrt{1 - \Lambda_R} - (1 - \Lambda_M)] \equiv T$. Likewise, the manufacturer will prefer no returns if $\alpha_h/\alpha_l > T$. This proves (a). Parts (b) and (c) are evident since $dT/dk_R > 0$ and $dT/dk_M < 0$. (This follows directly from the fact that $dT/d\Lambda_R > 0$ and $dT/d\Lambda_M < 0$.)

For part (d), note that $T \equiv \Lambda_M[\sqrt{1 - \Lambda_R} - (1 - \Lambda_M)] < \Lambda_M[1 - \Lambda_R - (1 - \Lambda_M)] = \Lambda_M[\Lambda_M - \Lambda_R]$, where the inequality is true because $1 - \Lambda_R < 1$. So the willingness of the manufacturer to offer a return policy implies $\alpha_h/\alpha_l < \Lambda_M[\Lambda_M - \Lambda_R]$, which is equivalent to $\Lambda_R/\Lambda_M > 1 - \alpha_l/\alpha_h$. As noted in Table 4, this condition indicates that $\Delta_{fr,nr}(w_j) > 0$.

PROOF OF PROPOSITION 2. Installing the full return policy has a net effect on the retailer's value function of the magnitude

$$\Delta_{nr \rightarrow fr}(V_j^R) = \frac{1}{16\beta} \times \left[\frac{(3\alpha_h(1 - \Lambda_R) - \beta c)(\alpha_h(1 - \Lambda_R) + \beta c)}{1 - \Lambda_R} - (3\bar{\alpha} - \beta c)(\bar{\alpha} + \beta c) \right]$$

for part (a), note that when $c = 0$,

$$\Delta_{nr \rightarrow fr}(V_j^R) = \frac{3}{16\beta} [(1 - \Lambda_R)(\alpha_h)^2 - (\bar{\alpha})^2]$$

Routine algebra shows this to be positive if and only if $\alpha_h/\alpha_l > T$. Parts (b) and (c) use logic from the analogous parts of Proposition 1.

PROOF OF PROPOSITION 3. The equilibria under this new power structure can be obtained by retracing the development that led to Table 2, except with different conditions on the determination of the wholesale prices.

(a) With no returns, the retailer's value function as a function of the wholesale price is

$$V_{nr}^R(w) = \frac{1}{4\beta} \left[\Lambda_R(\alpha_l)^2 + \frac{[(1 - \Lambda_R)\alpha_h - \beta w]^2}{1 - \Lambda_R} \right]$$

Clearly the retailer will seek as low a w as possible, as would be expected. The constraint on this is that the manufacturer's value function, which has the value

$$V_{nr}^M(w) = (w - c) \frac{(1 - \Lambda_R)\alpha_h - \beta w}{2(1 - \Lambda_R)},$$

must remain non-negative. The lowest possible w that achieves this clearly has the value c , giving the manufacturer a deterministic outcome of exactly 0.

(b) With full returns, the retailer's value function becomes

$$V_{fr}^R(w) = \frac{\Lambda_R(\alpha_l - \beta w)^2 + (1 - \Lambda_R)(\alpha_h - \beta w)^2}{4\beta}$$

while the manufacturer's value function becomes

$$V_{fr}^M(w) = \frac{(\bar{\alpha} - \beta w)w - (\alpha_h - \beta w)c}{2}$$

Clearly the retailer still seeks as low a wholesale price as possible, and solving the quadratic equation $V_{fr}^M(w) = 0$ provides the lowest allowable. That this value strictly exceeds c is obvious mathematically since setting $w \leq c$ renders V_{fr}^M strictly negative. This is because there is no profit margin to be earned on any unit supplied, yet the outcome is uncertain. (Existence of the desired w is guaranteed since there clearly exists $w \geq c$ for which V_{fr}^M is strictly positive.)

The remainder of the equilibria under this new power

structure can be obtained by retracing the development that led to Table 2, except with these new wholesale prices.

(c) How risk sensitivity impacts the wholesale prices and value functions can be established by routine differentiation.

References

- Adar, Z., Barnea, A., & Lev, B. (1977). A comprehensive cost volume profit analysis under uncertainty. *Accounting Review*, 52(1), 137–149.
- Agrawal, V., & Seshadri, S. (2000). Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing and Service Operations Management*, 2(4), 410–423.
- Bar-Shira, Z., & Finkelshtain, I. (1999). Two-moments decision models and utility-representable preferences. *Journal of Economic Behavior and Organization*, 38(2), 237–244.
- Donohue, K. L. (2000). Efficient supply contracts for fashion goods with forecast updating and two production modes. *Management Science*, 46(11), 1397–1411.
- Desiraju, R. (2001). New product introductions, slotting allowances, and retailer discretion. *Journal of Retailing*, 77(3), 335–358.
- Eeckhoudt, L., Gollier, C., & Schlesinger, H. (1995). The risk-averse (and prudent) newsboy. *Management Science*, 41(5), 786–794.
- Emmons, H., & Gilbert, S. M. (1998). Note: the role of returns policies in pricing and inventory decisions for catalogue goods. *Management Science*, 44(2), 276–283.
- Harlow, W. V. (1991). Asset allocation in a downside risk framework. *Financial Analysts Journal*, 47(5), 28–40.
- Ingene, C. A., & Parry, M. E. (1995). Coordination and manufacturer profit maximization: the multiple retailer channel. *Journal of Retailing*, 71(2), 129–151.
- Johnson, G., & Simik, S. (1971). Multiproduct CVP analysis under uncertainty. *Journal of Accounting Research*, 9(Autumn), 278–286.
- Kandel, E. (1996). The right to return. *Journal of Law and Economics*, 39(1), 329–356.
- Kataoka, S. (1963). A stochastic programming model. *Econometrica*, 31(1–2), 181–196.
- Lanzilotti, R. F. (1958). Pricing objectives in large companies. *American Economic Review*, 48, 921–940.
- Lariviere, M., & Padmanabhan, V. (1997). Slotting allowances and new product introductions. *Marketing Science*, 16(2), 112–128.
- Lau, H. (1980). The newsboy problem under alternative optimization objectives. *Journal of the Operational Research Society*, 31(6), 525–535.
- Lee, H. L., Padmanabhan, V., Taylor, T. A., & Whang, S. (2000). Price protection in the personal computer industry. *Management Science*, 46(4), 467–482.
- Magee, R. P. (1975). Cost-volume-profit analysis, uncertainty and capital market equilibrium. *Journal of Accounting Research*, 13(2), 257–266.
- Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance*, 7(March), 77–91.
- Markowitz, H. M. (1959). *Portfolio selection: efficient diversification of investments*. New York, NY: John Wiley.
- Meyer, J. (1987). Two-moment decision models and expected utility maximization. *American Economic Review*, 77(3), 421–430.
- Padmanabhan, V., & Png, I. P. L. (1995). Returns policies: make money by making good. *Sloan Management Review*, 37(1), 65–72.
- Padmanabhan, V., & Png, I. P. L. (1997). Manufacturer's returns policies and retail competition. *Marketing Science*, 16(1), 81–94.
- Pasternack, B. A. (1985). Optimal pricing and returns policies for perishable commodities. *Marketing Science*, 4(2), 166–176.
- Robison, L. J., & Barry, P. J. (1987). *The competitive firm's response to risk*. New York, NY: Macmillan.
- Roy, A. D. (1952). Safety-first and the holding of assets. *Econometrica*, 20(3), 431–449.
- Saha, A. (1997). Risk preference estimation in the nonlinear mean standard deviation approach. *Economic Inquiry*, 35(4), 770–782.
- Spulber, D. F. (1985). Risk sharing and inventories. *Journal of Economic Behavior and Organization*, 6(1), 55–68.
- Swalm, R. O. (1966). Utility theory-insights into risk taking. *Harvard Business Review*, 44(November–December), 123–136.
- Telser, L. (1955–56). Safety-first and hedging. *Review of Economic Studies*, 23(Spring), 1–16.
- Tsay, A. A. (2001). Managing retail channel overstock: markdown money and return policies. *Journal of Retailing*, 77(4), 457–492.
- Tse, K. S. M., Uppal, J., & White, M. A. (1993). Downside risk and investment choice. *The Financial Review*, 28(4), 585–605.
- Tsiang, S. C. (1972). The rationale of the mean-standard deviation analysis: skewness, preference, and the demand for money. *American Economic Review*, 62(June), 354–371.
- Webster, S., & Weng, Z. K. (2000). A risk-free perishable items returns policy. *Manufacturing and Service Operations Management*, 2(1), 100–106.