8 SUPPLY CHAIN CONTRACTING AND COORDINATION WITH STOCHASTIC DEMAND

Martin A. Lariviere

Fuqua School of Business
Duke University
Durham, NC 27708

S. Tayur et al. (eds.), *Quantitative Models for Supply Chain Management*

Recent years have seen a growing interest among both academics and practitioners in the field of supply chain management. With that has come a growing body of work on supply chain contracts. Few firms are so large and few products so simple that one organization can manage the entire provision of the good. Rather, most supply chains require the coordination of independently managed entities who seek to maximize their own profits. Issues of who controls what decisions and how parties will be compensated become critical. An understanding of contractual forms and their economic implications is therefore an important part of evaluating supply chain performance.

The literature on supply chain contracts can be roughly split into two classes. The first takes a particular contract and determines what optimal actions are assuming that the contract terms are fixed. Some attention may be given to the impact of contract parameters on agents' profits or costs, but there is generally no attempt to establish systematically whether the contract allows the decentralized system to perform as well as a centralized one (i.e., whether the contract can *coordinate* the system). Eppen and Iyer (1997) and Brown and Lee (1997) are two recent examples that fall within this class.

The second class takes agents' optimal policies under a contract as given and considers whether the terms of trade can be adjusted to at least improve, if not coordinate, the supply chain. Generally, one supposes that a player may propose the specific terms within a contractual form and asks what contract she would offer. This line of work is closely related to the work on vertical restraints in the economics literature (Mathewson and Winter, 1984) and channel coordination in marketing (Jeuland and Shugan, 1983; Moorthy, 1987). In this chapter, we essentially ignore the first class of papers and focus on a subset of the second. Our basic model is a one period setting in which a manufacturer sells to a retailer facing a newsvendor problem.

Most stochastic inventory models rest in part on intuition gained from the single period newsvendor problem. It is our contention that the same holds true for supply chain contracting under stochastic demand. Under the contracts we consider here, the fundamental inventory problem remains sufficiently simple that one can characterize its solution with some precision. Because the solution to the inventory problem is well understood, we can develop a detailed analysis of the economic incentives the contracts provide. In particular, we will show that contracting on excess inventory in the form of returns policies can greatly improve channel performance. Returns policies allow for the parties to place "bets" on how realized demand compares to the chosen stocking level. Improved channel performance follows from the manipulation of these wagers.

In what follows, we first present the basic assumptions of the model and evaluate the performance of the integrated channel. Section 8.2 presents the simplest terms of trade, a price-only contract. Subsequent sections examine richer contractual forms that allow the channel to be coordinated. Section 8.3 considers buy back contracts while section 8.4 analyzes quantity flexibility contracts. Section 8.5 briefly discusses other coordination schemes found in the literature. Section 8.6 offers concluding remarks and directions for future research.

8.1 MODEL ASSUMPTION AND INTEGRATED CHANNEL PERFORMANCE

We consider a one-period setting in which a manufacturer sells to a retailer facing demand from consumers. If the channel fails to provide sufficient stock, unmet demand is lost. The retail price is fixed at r per unit regardless of the terms the manufacturer offers. The manufacturer can produce the good at a constant marginal cost c. Any unsold inventory can be salvaged at a value v per unit. To avoid trivial problems, assume $r > c > v$. Obviously, the setting can be characterized as a newsvendor problem.

We assume that demand ξ has a continuous distribution $F(\xi)$ on the non-negative reals with density $f(\xi)$. In addition it will be convenient to assume that $F(\xi)$ is invertible and that $f(\xi)$ has a continuous derivative $f'(\xi)$. Also let $\bar{F}(\xi) = 1 - F(\xi)$. The demand distribution, as well as all cost and revenue information, are common knowledge.

It is straightforward to verify that the profits of an integrated firm (i.e., one that controls both manufacturing and sales to the public) for stocking level y are:

$$\Pi_I(y) = (r-c)\,y - (r-v)\int_0^y F(\xi)\,d\xi. \tag{8.1}$$

The problem is concave in the stocking level, and the optimal solution is given by:

$$y_I = F^{-1}\left(\frac{r-c}{r-v}\right)$$

where F^{-1} is the inverse of the cumulative distribution. Denote the maximum system profits by $\Pi_I^* = \Pi_I(y_I)$. Note that system profits are completely determined by the stocking level. Looking ahead to when we consider contracting between independent parties below, we can prove that terms coordinate the channel if we can show that they induce the choice of the centralized system's optimal stocking level, y_I.

As we consider contracting in this environment, we assume the manufacturer acts as a Stackelberg leader. She offers the terms of trade as a take-it-or-leave-it proposition to the retailer, which he can only accept or reject. We assume he accepts the terms if they allow him to earn a non-negative return (perhaps net of some opportunity cost). Clearly, this is a gross simplification; it would be more natural to assume that the parties negotiate over the terms. Unfortunately, properly modeling bargaining is sufficiently complex that no consensus exists regarding the appropriate equilibrium concept (Salanié, 1997). We skirt the issue by assigning all the bargaining power to one player. Where appropriate, we will comment on how the results would differ if the retailer were able to offer the terms of trade.

8.2 PRICE-ONLY CONTRACTS

Here we consider price-only contracts; the manufacturer offers the good at a per unit wholesale price w, and the retailer retains possession of any excess stock. We first examine the retailer's problem under a price-only contract before turning to the manufacturer's problem. We present sufficient conditions for the manufacturer's profits to be unimodal and characterize the optimal wholesale price (where "optimal" should be interpreted as maximizing the manufacturer's profits). We show that whenever the manufacturer proposes the terms of trade, a price-only contract fails to coordinate the channel. We close the section with a brief discussion of standard remedies from economics for improving channel performance.

8.2.1 The retailer's problem

The retailer faces a problem analogous to that of the integrated channel given in (8.1). The principal difference is that the retailer must buy stock at price w instead of producing it at cost c. If we assume that the retailer has the same salvage opportunities as the integrated channel, we then have:

$$\Pi_R(y) = (r-w)\,y - (r-v)\int_0^y F(\xi)\,d\xi.$$

Clearly, $\Pi_R(0) = 0$ and $\Pi'_R(0) > 0$ if $w < r$. Thus if we normalize the retailer's opportunity cost to zero, he will accept any contract such that the wholesale price is less than the retail price and order a positive quantity from the manufacturer.[1] The retailer's problem is concave and the optimal solution is given by:

$$y(w) = F^{-1}\left(\frac{r-w}{r-v}\right).$$

8.2.2 The manufacturer's problem

Acting as a Stackelberg leader, the manufacturer correctly anticipates how the retailer will order for any wholesale price. She therefore anticipates facing a demand curve $y(w)$, yielding a profit function of:

$$\begin{aligned}\Pi_M(w) &= (w-c)y(w) \\ &= (w-c)F^{-1}\left(\frac{r-w}{r-v}\right).\end{aligned} \tag{8.2}$$

There are a few points worthy of mention from (8.2). First, under a price-only contract, the manufacturer's profits are deterministic. She knows exactly what the retailer will order at every wholesale price and bears no responsibility for

[1] This will not necessarily hold if there is a positive stock out penalty. See Lariviere and Porteus, 1997.

the goods once the retailer takes possession. All uncertainty regarding channel profits is foisted onto the retailer. The richer contracts we consider in subsequent sections differ from price-only contracts by allowing the manufacturer to assume some of the risk arising from random demand.

A second observation is that (8.2) is simply not the most convenient expression with which to work. We would like to develop conditions for the manufacturer's profit function to be well behaved. When, for example, is it unimodal? Answering such questions could be exceedingly difficult since we lack general statements about the inverse of the cumulative density function. We would be forced to evaluate (8.2) distribution by distribution.

To allow for broader statements, we follow Lariviere and Porteus (1998) to develop an alternative expression for manufacturer profits. First, instead of working with the demand curve $y(w)$, we work with the inverse demand curve $w(y)$, where

$$w(y) = (r - v)\bar{F}(y) + v.$$

The change may be interpreted as follows. Previously, we had assumed the manufacturer chose her wholesale price while anticipating selling the most the retailer would freely take at that price. Now, we are assuming that manufacturer chooses how much she wants to sell anticipating receiving the most per unit that the retailer would freely pay to take all of the offered stock. In a competitive setting, choosing quantities instead of price can lead to markedly different outcomes. Here, however, the approaches are equivalent since the manufacturer holds a monopoly position.

The second change we make is to alter the costing conventions. Let $\hat{w} = w - v$ and $\hat{r} = r - v$. The inverse demand curve can now be written as:

$$\hat{w}(y) = \hat{r}\bar{F}(y). \tag{8.3}$$

Under the revised accounting structure, the retailer immediately credits himself for salvaging the goods upon receipt, thus lowering the effective wholesale price. To keep himself from booking deceptively high profits, he must also lower the retail price.[2] While in what follows we will refer to $\hat{w}$ and $\hat{r}$ as the wholesale and retail prices, respectively, one should keep in mind that they are, in fact, mark ups over the salvage value.

An immediate observation following from (8.3) is that the "market-clearing" wholesale price $\hat{w}(y)$ is proportional to the retail price $\hat{r}$. Another consequence is a simpler expression for manufacturer profits:

$$\begin{aligned} \Pi_M(y) &= y(\hat{w}(y) - \hat{c}) \\ &= y\left(\hat{r}\bar{F}(y) - \hat{c}\right) \end{aligned} \tag{8.4}$$

where $\hat{c} = c - v$. Expressing the marginal cost of production as a mark up over the salvage value assures that (8.2) and (8.4) yield the same results for a given

[2]The interpretation requires an obvious modification if v is negative and the retailer pays to dispose of stock.

$(y, \hat{w}(y))$ pair (or, equivalently, a $(\hat{w}, y(\hat{w}))$ pair). We also have an alternative statement of the manufacturer's problem: Choose y to maximize $\Pi_M(y)$.

8.2.3 Characterizing the optimal solution

In examining the manufacturer's problem, a few points are obvious. First, she will never sell more than y_I since that would necessitate a wholesale price below the marginal cost of production. Second, if $F(\xi) > 0$ for all $\xi > 0$, she will choose a sales quantity strictly between 0 and y_I which will result in a wholesale price between $\hat{c}$ and $\hat{r}$. The optimal sales quantity will be an interior solution, and first order conditions must hold.[3]

But are first order conditions sufficient? Is it possible for the manufacturer's problem to have multiple local maxima? Alternatively, is it possible that there is no solution to the manufacturer's problem? Both are possible; the manufacturer's first order conditions may have multiple or no solution. To see the former, consider the following density defined for $\xi \in (0,1)$:

$$f(\xi) = Cos(32\pi\xi) + 1. \tag{8.5}$$

The corresponding profit function exhibits multiple local maxima and minima. While first order conditions must hold at the optimal solution, they may hold at multiple points, some of which may be local minima.

For an example of a distribution for which first order conditions never hold, consider the Pareto distribution with $f(\xi) = k\theta^k/\xi^{k+1}$ for $\xi \geq \theta$. The corresponding inverse demand curve,

$$\hat{w}(y) = \hat{r}\bar{F}(y) = \hat{r}\theta^k y^{-k},$$

is isoelastic with elasticity $1/k$. For the manufacturer's profits to be concave we require that the elasticity be greater than one or, equivalently, that $k < 1$. However, for $k < 1$, the mean of the Pareto, and hence the retailer's problem, is undefined. One can show that for any problem with finite expected demand the manufacturer's problem is convex; she will charge $\hat{w} = \hat{r}$ and sell the minimum quantity θ.

The density given in (8.5) is admittedly perverse and an unrealistic model of demand. The Pareto example is a little more troubling; while not commonly used as a demand distribution, it is not a completely absurd choice for a modeler. It forces one to question what other distributions fail to yield sensible solutions to the manufacturer's problem.

---

[3] There are two exceptions to note. If $F(\xi) = 0$ for all $\xi \in [0, \bar{\xi}]$, the manufacturer may sell $\bar{\xi}$ units at a wholesale price of r. We cannot rule out the possibility that she prefers the corner solution of extracting all profits on the certain sales of $\bar{\xi}$. In addition, if the retailer has a positive opportunity cost that exceeds his profits under the manufacturer's optimal retail price, the manufacturer will deviate from the price found from the first order conditions. It is straightforward to show that the manufacturer prefers cutting her wholesale price to charging the unconstrained price and paying a lump sum to gain participation (Lariviere and Padmanabhan, 1997).

In order to assure that the manufacturer's first order conditions have a unique solution, we must place limits on the demand distribution that rule out the likes of (8.5) and the Pareto. To that end, define $h(\xi)$ as the failure, or hazard, rate function of the demand distribution.

$$h(\xi) = \frac{f(\xi)}{\bar{F}(\xi)}.$$

In our setting, $h(\xi)\,d\xi$ may be interpreted as the probability that demand will lie in the interval $[\xi, \xi + d\xi]$ given that demand is at least ξ. A distribution is said to have an increasing failure rate (IFR) if $h'(\xi) > 0$ for all ξ.

Define $g(\xi) = \xi h(\xi)$ as the *generalized failure rate.* We say that a distribution has an increasing generalized failure rate (IGFR) if $g'(\xi) > 0$ for all ξ. Clearly, a distribution that is IFR is also IGFR, but the distinction is not vacuous. Consider the Weibull distribution,

$$f(\xi) = \theta k \xi^{k-1} e^{-\theta \xi^k},$$

for $k > 0$ and $\theta > 0$, or the gamma distribution,

$$f(\xi) = \theta^k \xi^{k-1} e^{-\theta\xi} / \Gamma(k),$$

for $k > 0$ and $\theta > 0$. Both are IFR for a restricted set of parameters (i.e., $k > 1$) but are IGFR for all parameters. IGFR distributions offer additional flexibility over IFR ones that are particularly relevant in our setting. For example, IFR distributions must have a coefficient of variation of less than one (Barlow and Proschan, 1965) while IGFR distributions are not similarly restricted.

The following theorem shows that an IGFR demand distribution is also sufficient for manufacturer profits to be well-behaved.[4]

Theorem 1 *Let $\nu(\xi)$ denote the own-price elasticity of the retailer's orders to the manufacturer, and let y^1 be the smallest value of ξ such that $g(\xi) = 1$. If no such ξ exists, let $y^1 = \infty$.*

1. *The elasticity of retailer's orders is given by*

$$\nu(\xi) = 1/g(\xi). \tag{8.6}$$

2. *The manufacturer's first order conditions may be written as*

$$\hat{w}(y)\left(1 - 1/\nu(y)\right) = c. \tag{8.7}$$

[4]This and most of the results in the current section are from Lariviere and Porteus, 1998.

3. *If the demand distribution is IGFR, then the manufacturer's profits are unimodal on* $[0,\infty)$*, concave on* $[0,y^1]$*, and strictly decreasing on* $[y^1,\infty)$*. Any solution to (8.7) is a unique global maximum and must lie in the interval* $[0,y^1]$*.*

Proof: For the first part of the theorem, note that $\nu(\xi)$ is defined as $\nu(\xi) = -\frac{1}{\hat{w}'(\xi)}\frac{\hat{w}(\xi)}{\xi}$. The result is then immediate from the definition of $\hat{w}(\xi)$. The second part follows from a standard microeconomic result (Kreps, 1990). Differentiating revenue, $y\hat{w}(y)$, yields marginal revenue, $MR(y)$:

$$MR(y) = \hat{r}\bar{F}(y) - y\hat{r}f(y) = \hat{r}\bar{F}(y)(1-g(y)).$$

The condition (8.7) thus reduces to setting marginal revenue equal to marginal cost. For the final part of the theorem, note that IGFR implies that $\bar{F}(y)$ and $1-g(y)$ are decreasing on $[0,\infty)$. Both are positive on $[0,y^1]$ so marginal revenue is positive and decreasing there. Hence revenue and profits are concave on $[0,y^1]$. As $MR(y^1)=0$, any solution to (8.7) must lie in $[0,y^1]$ and from concavity must be unique. Profits are falling on $[y^1,\infty)$ because revenue is strictly decreasing while costs are strictly increasing over this region. □

Note that $\nu(\xi)$ is the elasticity of the orders the retailer places with the manufacturer, not the elasticity of end consumer demand. We have not modeled consumer behavior and, indeed, have assumed a fixed retail price. On the other hand, by modeling the retailer's ordering policy, we have defined an induced demand curve (and the equivalent induced inverse demand curve) that the manufacturer faces. $\nu(\xi)$ measures the percent change in the retailer's orders for a percent change in the wholesale price. Thus, if $\nu(\xi) > 1$, a wholesale price cut is offset by such an increase in volume that total revenue increases. Conversely, if $\nu(\xi) < 1$, a price increase will boost total revenue. An immediate consequence of the theorem is the standard result that the optimal sales quantity will lie on the elastic portion of the demand curve (i.e., $\nu(\xi) > 1$).

The relationship between the elasticity of orders and the generalized failure rate captured by (8.6) offers some intuition for why IGFR is the correct sufficient condition for a well-behaved profit function. IGFR implies that the elasticity of orders falls monotonically with the sales quantity. If the elasticity of orders falls monotonically, then marginal revenue falls monotonically and can equal the marginal cost of production at only one point.

The two poorly-behaved examples we considered above fail to be IGFR. The generalized failure rate for the density in (8.5) is not monotonic since the density itself drops to zero periodically. The Pareto, on the other hand, has a constant generalized failure rate. For meaningful parameter values, the generalized failure rate is greater than one and marginal revenue is everywhere negative.

While the theorem limits the distributions that one can be assured are well-behaved, the IGFR class is broad enough to capture most of the distribution a modeler would choose to employ. In addition to the distributions discussed above, the uniform and the normal are both IFR and hence IGFR. In addition to assuring a well behaved profit function, IGFR gives us some additional attractive properties.

Theorem 2 *If the demand distribution is IGFR, then the optimal sales quantity y^* is increasing $\hat{r}$ and decreasing in $\hat{c}$. Additionally, $\hat{w}(y^*) \geq \hat{w}(y^1)$.*

Proof: The first part follows immediately from noting that IGFR is sufficient for marginal revenue to be monotonically decreasing on $[0, y^1]$. The second part follows from the fact that y^* must lie in $[0, y^1]$. □

Thus an IGFR distribution assures that the optimal sales quantity changes as one would expect with respect to cost and revenue parameters. The lower bound on the manufacturer's profit maximizing wholesale price $\hat{w}(y^*)$ given by her revenue maximizing price $\hat{w}(y^1)$ is useful since y^1 may be found simply from the generalized failure rate. For example, the generalized failure rate for the Weibull is $g(\xi) = \theta k \xi^k$. Solving for the sales quantity such that $\nu(y) = 1$ yields $y^1 = \left(\frac{1}{\theta k}\right)^{1/k}$. The lower bound on the wholesale price is then

$$\hat{w}(y^1) = \hat{r}e^{-1/k}. \tag{8.8}$$

We are consequently able to conclude that the optimal wholesale price for the Weibull goes to the retail price as the parameter k goes to infinity.

8.2.4 Special cases

The preceding results give a general characterization of the optimal solution. Further assumptions regarding the demand distribution allow some additional analysis. First we require some definitions. We say a distribution is from a "scaled" family if the distribution depends on a parameter θ and there exists an increasing, positive function $\tau(\theta)$ such that

$$F(\xi|\theta) = F\left(\frac{\xi}{\tau(\theta)}|1\right).$$

We immediately have that for scaled families $f(\xi|\theta) = \frac{1}{\tau(\theta)} f\left(\frac{\xi}{\tau(\theta)}|1\right)$, $h(\xi|\theta) = \frac{1}{\tau(\theta)} h\left(\frac{\xi}{\tau(\theta)}|1\right)$, and $g(\xi|\theta) = g\left(\frac{\xi}{\tau(\theta)}|1\right)$. The most obvious examples of scaled families are uniform $[0, \theta]$ random variables or exponential random variables with mean θ. Less obvious examples are the Weibull and gamma distributions.

We say that a family of distributions is "shifted" if, for some parameter $\theta \geq 0$,

$$F(\xi|\theta) = F(\xi - \theta|0).$$

This implies that $f(\xi|\theta) = f(\xi - \theta|0)$, $h(\xi|\theta) = h(\xi - \theta|0)$, and $g(\xi|\theta) \geq g(\xi - \theta|0)$. The normal distribution with mean θ and fixed standard deviation σ is an example of a shifted family.

Theorem 3 *Suppose that the demand distribution $F(\xi|\theta)$ is IGFR for all values of θ.*

1. *If $F(\xi|\theta)$ is from a scaled family, the optimal order quantity is proportional to $\tau(\theta)$ and the resulting wholesale price is independent of θ. That is:*

$$y^*(\theta) = \tau(\theta)\, y^*(1) \quad \textit{and} \quad \hat{w}^*(\theta) = \hat{w}^*(1)$$

where $y^*(\theta)$ and $\hat{w}^*(\theta)$ are the optimal values of y and $\hat{w}$ given θ.

2. *If $F(\xi|\theta)$ is a shifted family, then for $\Delta > 0$:*

$$y^*(\theta+\Delta) < y^*(\theta) + \Delta \quad \textit{and} \quad \hat{w}^*(\theta) < \hat{w}^*(\theta+\Delta).$$

If in addition $F(\xi|\theta)$ is IFR, then $y^*(\theta) < y^*(\theta+\Delta)$.

Proof: For the first part, the manufacturer's first order conditions for an arbitrary value of θ can then be written as:

$$\begin{aligned} \hat{r}\bar{F}(y^*(\theta)|\theta)(1-g(y^*(\theta)|\theta)) &= \\ \hat{r}\bar{F}(y^*(\theta)/\tau(\theta)|1)(1-g(y^*(\theta)/\tau(\theta)|1)) &= c. \end{aligned} \tag{8.9}$$

However, for $\theta = 1$, first order conditions are:

$$\hat{r}\bar{F}(y^*(1)|1)(1-g(y^*(1)|1)) = c.$$

Thus (8.9) must be satisfied by $y^*(\theta) = \tau(\theta)\, y^*(1)$. The corresponding wholesale price is found from:

$$\hat{w}^*(\theta) = \hat{r}\bar{F}(y^*(\theta)|\theta) = \hat{r}\bar{F}\left(\frac{\tau(\theta)\, y^*(1)}{\tau(\theta)}|1\right) = \hat{w}^*(1).$$

For the second part, we first show that $y^*(\theta+\Delta) < y^*(\theta) + \Delta$. Since F is a shifted family:

$$\begin{aligned} \bar{F}(y^*(\theta)+\Delta|\theta+\Delta) &= \bar{F}(y^*(\theta)-\theta|0) = \bar{F}(y^*(\theta)|\theta) \\ h(y^*(\theta)+\Delta|\theta+\Delta) &= h(y^*(\theta)-\theta|0) = h(y^*(\theta)|\theta) \end{aligned}$$

which gives that $g(y^*(\theta)+\Delta|\theta+\Delta) > g(y^*(\theta)|\theta)$. We then have:

$$\begin{aligned} &\hat{r}\bar{F}(y^*(\theta)+\Delta|\theta+\Delta)(1-g(y^*(\theta)+\Delta|\theta+\Delta)) \\ < \;&\hat{r}\bar{F}(y^*(\theta)|\theta)(1-g(y^*(\theta)|\theta)) = \hat{c}. \end{aligned}$$

The final equality follows from the definition of $y^*(\theta)$. That $y^*(\theta+\Delta) < y^*(\theta)+\Delta$ follows from the IGFR property which implies that marginal revenue is a decreasing function. That $\hat{w}^*(\theta) < \hat{w}^*(\theta+\Delta)$ immediately follows from $y^*(\theta+\Delta) < y^*(\theta)+\Delta$.

For $y^*(\theta) < y^*(\theta+\Delta)$, note that

$$\begin{aligned} \bar{F}(y^*(\theta)|\theta+\Delta) &= \bar{F}(y^*(\theta)-(\theta+\Delta)|0) > \bar{F}(y^*(\theta)-\theta|0) = \bar{F}(y^*(\theta)|\theta) \\ h(y^*(\theta)|\theta+\Delta) &= h(y^*(\theta)-(\theta+\Delta)|0) < h(y^*(\theta)-\theta|0) = h(y^*(\theta)|\theta), \end{aligned}$$

The second inequality follows from the IFR assumption and implies that

$$g\left(y^*\left(\theta\right)|\theta+\Delta\right) < g\left(y^*\left(\theta\right)|\theta\right).$$

The remainder of the proof is essentially the reverse of the previous case. □

The first part of the theorem states that for many distributions the optimal wholesale price never depends on some distribution parameter. For example, the same wholesale price is optimal for all markets with exponential demand. More generally, one can show that a scaled family has a constant coefficient of variation, and it is the coefficient of variation that determines the optimal price when demand is gamma or Weibull.

An alternative interpretation of the result is that the manufacturer induces the retailer to serve the same fraction of demand for all distributions in the family. Recall that the elasticity of retailer $\nu(\xi)$ gives the percentage change in retailer's orders for a percentage change in the wholesale price. With a scaled family we have $\nu(\xi|\theta) = \nu\left(\frac{\xi}{\tau(\theta)}|1\right)$. The responsiveness of retailer orders to a price change at a given fractile of the demand distribution is the same for all members of the family. Thus the optimum is always found at the same fractile.

The second part of the theorem relates the optimal sales quantity and wholesale price to an additive parameter. As with a scaled distribution, a higher θ for a shifted family corresponds to a larger market, and the manufacturer responds to an increased market size by increasing her sales quantity (although the increase is less than the increase in the market size). Unlike a scaled family, a shifted distribution leads to the manufacturer charging a higher price. These observations immediately give the following.

Theorem 4 *If the demand distribution $F(\xi|\theta)$ is from either a shifted or a scaled family, then the manufacturer's profits are increasing in θ.*

Note that for a shifted family, varying θ changes the mean but has no impact on the variance of the distribution. Thus while scaled families have constant coefficients of variations, shifted families have coefficients of variations that decrease as θ increases. Further, the optimal price is higher for lower coefficients of variation. This logic yields the following:

Theorem 5 *Suppose demand is normally distributed. The optimal wholesale price is determined by the coefficient of variation and the manufacturer charges more the smaller the coefficient of variation is.*

Proof: We first show that the normal can be expressed as a scaled family. Let θ denote the mean, γ denote the coefficient of variation (giving a standard deviation of $\theta\gamma$), and $f(\xi|\theta,\theta\gamma)$ denote the density. We have:

$$\begin{aligned} f(\xi|\theta,\theta\gamma) &= \frac{1}{\theta\gamma\sqrt{2\pi}}\exp\left[-\left(\xi-\theta\right)^2/2\gamma^2\theta^2\right] \\ &= \frac{1}{\theta\gamma\sqrt{2\pi}}\exp\left[-\left(\xi/\theta-1\right)^2/2\gamma^2\right] \\ &= \frac{1}{\theta}f(\xi/\theta|1,\gamma). \end{aligned}$$

The normal is therefore a scaled family, and the optimal wholesale price depends only on the coefficient of variation. To see that the wholesale price must decrease in the coefficient of variation, consider two markets, one with mean one and coefficient of variation γ_1 and the other with mean γ_1/γ_2 and coefficient of variation γ_2. Assume $\gamma_1 > \gamma_2$. The demand distributions are now from a shifted family since they have the same standard deviation but different means. As $\gamma_1/\gamma_2 > 1$, market two must have the higher price, but it is also the market with the smaller coefficient of variation. □

We derived a similar result for the Weibull using the lower bound on the optimal wholesale price given in (8.8) above. The coefficient of variation for the Weibull is decreasing in the parameter k while the lower bound increases in k. To gain some intuition to support these results, suppose the coefficient of variation were zero and demand were known with certainty. The retailer's order is then completely inelastic for any wholesale price less than or equal to the retail price. The manufacturer responds by pushing the wholesale price up to $\hat{r}$, capturing all channel profits for herself. This observation suggests if the optimal wholesale price is monotonic in the coefficient of variation (as it is for the normal), it must fall as variability increases.

8.2.5 Supply chain performance

While we have determined the best the manufacturer can do for herself, we must acknowledge that it will not be the best outcome for the supply chain as a whole. The manufacturer will always choose to sell a quantity less than y_I (or, equivalently, will always charge a wholesale price above the marginal cost of production). Total profits in the decentralized system will thus be less than the profits of a centralized system.

This phenomenon of double marginalization has long been known in the study of industrial organization (Spengler, 1950). The standard presentation has a manufacturer selling to a retailer facing a known, downward-sloping demand curve. By selling at a wholesale price above her marginal cost of production, the manufacturer induces the retailer to set a retail price above what an integrated firm would charge. Sales, and system profits, are thus below what an integrated channel could achieve. In our setting, the retail price is fixed. Double marginalization consequently appears as inadequate stocking levels. Insufficient inventories, not overly high retail prices, are the source of the problem. In an applied setting, this issue is explored by Pasternack (1980).

Economists have suggested a number of possible remedies to improve supply chain performance. These generally fall under the heading of "vertical restraints." The most common suggestions are franchising, quantity forcing, resale price maintenance (RPM) and exclusive territories. Tirole (1988) provides an introduction to the topic while Mathewson and Winter (1984) evaluate their relative merits in a number of environments.

RPM (i.e., allowing the manufacturer to dictate the retail price) and exclusive territories provide little help in our setting. We have explicitly assumed that the retail price is fixed regardless of the terms offered and have implicitly assumed that the retailer already enjoys a local monopoly. The impact of RPM when the market size is uncertain and the retailer has control over the retail price is considered in Deneckere, Marvel, and Peck (1996 and 1997) and Butz (1997). A comparison of RPM and exclusive territories in such a setting is offered by Rey and Tirole (1986).

On the other hand, franchising and quantity forcing are capable of coordinating our system. That is, allowing the decentralized system to earn the same profits as a centralized one. Further, both allow for an arbitrary split of the system's profits. Under franchising, the manufacturer charges an upfront fee A to carry the product (regardless of the stocking level) and then sells the product at a wholesale price of w.[5] Since the size of the lump sum payment is independent of the order quantity, the retailer will willingly pay any quantity such that

$$A \leq \Pi_R(w) - \kappa,$$

where $\Pi_R(w) = \Pi_R(y(w))$ and $\kappa \geq 0$ is the retailer's opportunity cost for carrying the product. It is straightforward to show that $\Pi_R(w)$ is decreasing in w. The largest fee the manufacturer can extract without subsidizing the retailer (i.e., selling the product at a loss) is

$$A = \Pi_R(c) - \kappa = \Pi_I^* - \kappa.$$

She can thus capture all channel profits except for the minimum amount the retailer requires to participate in the system. Note that system profits are independent of A. It is transferring at marginal cost that eliminates double marginalization; the franchise fee serves only to redistribute profits and consequently enables an arbitrary split of profits.

Franchising works by eliminating distorted incentives that affect the retailer's choice. Quantity forcing works by simply eliminating the retailer's choice. The manufacturer offers the product a wholesale price of w and insists that the retailer take some quantity Q. Clearly, if $Q = y_I$, inventory in the system is equal to the integrated channel's quantity, so system profits must equal Π_I^*. It is now the wholesale price that serves no purpose in coordination and is freed to redistribute profits.[6] The retailer's profits are

$$\Pi_I^* - (w - c)\, y_I.$$

Hence he accepts any contract such that

$$w \leq \frac{\Pi_I^* - \kappa}{y_I} + c.$$

Quantity forcing is therefore also able to support an arbitrary split of profits and, in particular, can be structured to drive the retailer to indifference.

[5]Note that here and for the rest of the chapter we return to our original costing convention.
[6]See Cachon and Lariviere (1997) for ways in which this pricing flexibility can be exploited.

8.3 RETURNS POLICIES: BUY BACK CONTRACTS

The nostrums economists favor for improving supply chain performance have generally been developed for settings in which demand is deterministic. Consequently, no sale is ever lost and no unit ever goes unsold. That is not so in our setting, and that change allows for the consideration of some additional contracts. In this and the next section we focus on returns policies. See Padmanabhan and Png (1995) for a managerial discussion of such programs.

Given a stocking level, some demand realizations will result in excess stock. We show that by judiciously assigning responsibility for that surplus the system can be coordinated. Accountability can be adjusted through a price mechanism (buy backs) or a quantity scheme (quantity flexibility contracts). Further, either approach can support an arbitrary split of profits without any ancillary payments. Surprisingly, a player's profits are increasing in the burden she bears for excess stock. Either method is consequently capable of performing as well as franchising or quantity forcing without being as heavy-handed as those schemes.

Here we consider buy back contracts, postponing quantity flexibility until the next section. We first present the basics of the contract, the retailer's optimal policy under the contract and some properties of the resulting manufacturer's profits. We then determine the coordinating contract and demonstrate the flexibility of buy backs by considering the various settings in which they have been applied.

8.3.1 Contract basics

We now assume that in addition to posting a wholesale price w, the manufacturer stands ready to buy back any unsold stock from the retailer at a per unit rate $b < w$. If realized demand ξ is less than the order quantity y, the retailer receives $b\,(y-\xi)$. Restricting b to be less than w ensures that the manufacturer does not create an arbitrage opportunity for the retailer, allowing him to buy stock in order to return it for a profit. Additionally, for the deal to be attractive to the retailer, b must be greater than what the retailer can achieve by salvaging the stock himself, which raises an important issue: To evaluate the effectiveness of such contracts, we must make some assumption about the salvage opportunities open to the players. It will obviously be impossible to replicate the outcome of the integrated channel if our buy back scheme results in some stock being salvaged at unfavorable terms.

There are a number of assumptions that will prevent this outcome. First we can assume that the manufacturer can salvage the product at the integrated channel rate v while the retailer can only salvage the product for an amount less than or equal to v. Offering a buy back of b to the retailer costs the manufacturer $b - v$ per unit.[7] If, on the other hand, only the retailer can salvage at the integrated channel level, we can assume that the manufacturer can verify the retailer's excess inventory. The manufacturer can then pay the retailer $b - v$ per unit, leaving the retailer with an effective salvage rate of b. We shall see that the power of returns policies is that they allow payments that are conditional on both the chosen stocking level and realized demand. This second assumption recognizes that these payments can be divorced from the actual movement of goods as long as the parties can verify that the payments are warranted.

Both alternatives must be recognized as simplifications of reality. Moreover, they both rest on a larger assumption: Imposing a returns policy on the decentralized channel introduces no additional cost beyond that incurred by the centralized system. Thus in the first case, returning stock to the manufacturer requires no transportation or handling costs that the centralized system could avoid. In the second, verifying the retailer's excess inventory must be costless.

8.3.2 The retailer's problem

With these assumptions, we can write the retailer's objective as:

$$\Pi_R(y) = (r - w)\, y - (r - b) \int_0^y F(\xi)\, d\xi.$$

The retailer still faces a newsvendor problem. The optimal solution is now

$$y(w, b) = F^{-1}\left(\frac{r - w}{r - b}\right).$$

It is straightforward to verify that the retailer's optimal order quantity and profits for a fixed wholesale price are increasing in b. Holding the manufacturer's wholesale price constant, he prefers a generous returns policy.

8.3.3 The manufacturer's problem

Under price only contracts, the manufacturer's profits were deterministic. For any wholesale price, she knew exactly what she would sell, and the retailer's presence sheltered her from market uncertainty. With a buy back contract, she is now exposed to the possibility of a poor demand outcome. Her profits are given by:

$$\Pi_M(w, b) = (w - c)\, y(w, b) - (b - v) \int_0^{y(w,b)} F(\xi)\, d\xi. \tag{8.10}$$

[7] Note that b does not have to be positive. For example, the retailer may pay the manufacutrer to take excess stock off his hands as long as b is less costly than his next-best disposal opportunity.

For price-only contracts, we could develop sufficient conditions on the demand distribution for the manufacturer's problem to be well-behaved. As the following shows, buy backs are more complicated, and her objective is never well-behaved.

Theorem 6 *The Hessian of* $\Pi_M(w,b)$ *is nowhere negative definite. Thus second order conditions for a local maximum are never satisfied.*

Proof: For second order conditions to be satisfied, the Hessian must be negative definite, which requires that its determinant be positive. Its determinant, however, is

$$-\frac{\left(r\left(w-c\right)+v\left(r-w\right)-b\left(r-c\right)\right)^2}{\left(r-b\right)^6 f\left(y\left(w,b\right)\right)^2},$$

which is always negative. □

For price only contracts we could place limits on the demand distribution that assured that first order conditions were sufficient to determine the optimal contract. While intuitively appealing, that approach will not work here. We require an alternative approach to determine the equilibrium contract.

8.3.4 Coordinating contracts

While we cannot identify the optimal contract through standard optimization approaches, some form of returns policy is beneficial to the system. Pasternack (1985) shows that having the manufacturer accept no returns leads to a suboptimal outcome for the system. Similarly a policy of full returns (i.e., fully refunding the wholesale price on all unsold items) is also suboptimal. An intermediary policy, however, results in channel coordination, as the following theorem (from Pasternack, 1985) shows.

Theorem 7 *Suppose that the manufacturer offers a contract* $(w(\varepsilon), b(\varepsilon))$ *for* $\varepsilon \in (0, r-c]$ *where*

$$w(\varepsilon) = r - \varepsilon \quad \text{and} \quad b(\varepsilon) = r - \frac{\varepsilon\left(r-v\right)}{r-c}.$$

1. *The retailer orders the integrated channel quantity (i.e.,* $y(w(\varepsilon), b(\varepsilon)) = y_I$*) and system profits are equal to the integrated channel profits.*
2. *Retailer profits are increasing in* ε*. Specifically,* $\Pi_R\left(w(\varepsilon), b(\varepsilon)\right) = \frac{\varepsilon}{r-c}\Pi_I^*$*.*
3. *Manufacturer profits are decreasing in* ε*. Specifically,* $\Pi_M\left(w(\varepsilon), b(\varepsilon)\right) = \left(1-\frac{\varepsilon}{r-c}\right)\Pi_I^*$*.*

Proof: Note that for all allowed ε, $r > w\left(\varepsilon\right) > b\left(\varepsilon\right) \geq v$. So the contract is feasible even if the retailer may salvage excess items at rate v. Next observe that for all ε:

$$\frac{r-w\left(\varepsilon\right)}{r-b\left(\varepsilon\right)} = \frac{r-c}{r-v}.$$

The retailer faces the same critical fractile as the integrated channel and thus orders the same amount. As channel profits only depend on the total stock in the system, the decentralized system's profits must equal integrated system profits. To determine retailer profits, we have:

$$\begin{aligned}
\Pi_R\left(w(\varepsilon), b(\varepsilon)\right) &= (r - w(\varepsilon))y^* - (r - b(\varepsilon)) \int_0^{y^*} F(x)dx \\
&= \frac{\varepsilon}{r-c}\left((r-c)\,y^* - (r-v)\int_0^{y^*} F(x)dx\right) \\
&= \frac{\varepsilon}{r-c}\Pi_I^*.
\end{aligned}$$

System profits are fixed, so the manufacturer earns $\left(1 - \frac{\varepsilon}{r-c}\right)\Pi_I^*$. □

The coordinating contract is therefore not unique. Rather a continuum of contracts exists. Possible contracts differ in how they divide the channel profits. A few observations add insight to this outcome. First, if the manufacturer wishes to sell at a price above the marginal cost of production, the retailer faces a higher acquisition cost than the integrated channel does. Double marginalization suggests that the retailer will order less than the integrated channel unless the manufacturer can simultaneously raise the retailer's marginal revenue. Granting a higher pay out in low demand states accomplishes this. Not surprisingly, the higher the wholesale price, the more the manufacturer must compensate the retailer in the event of low demand.

Second, what the retailer orders depends on a single value, the critical fractile. Under price only contracts, the manufacturer has a single contract parameter that affects the fractile, and the only way to coordinate the channel is to transfer at marginal cost. With a buy back contract, on the other hand, she may manipulate two parameters to have one value come out correctly. The system is under-determined, and a continuum of solutions exists.

This observation together with the third part of the theorem suggests why Theorem 6 holds. The manufacturer would prefer to push ε as low as possible and capture all channel profits. However, she cannot set ε to zero since this would result in a full returns policy which Pasternack (1985) has shown to be suboptimal. She consequently does not have an optimal contract; for any feasible, coordinating contract she will prefer to cut ε even further to grab a larger share of the profits.

The coordinating contracts have a number of additional properties. First, it is good to be responsible. Either party would prefer to take greater responsibility for excess stock as this leads to higher profits. This holds despite the fact that for any fixed wholesale price, the retailer would prefer the manufacturer offer an easy returns policy. But there is, of course, a difference between being easy and being cheap. Under coordinating contracts, generous return rates are paired with high wholesale prices. The price hikes are more than enough to offset any gains from the higher buy back rate. The manufacturer is in essence bundling insurance with the physical good. The more generous the insurance is, the higher the bundle is priced. The retailer, however, is risk neutral and would not pay for the indemnity if it were offered separately. It is the bargaining power we have assigned to the manufacturer that allows her to force the retailer to purchase the bundle.

Another interesting feature is that the coordinating contracts $(w(\varepsilon), b(\varepsilon))$ are independent of the demand distribution. If the manufacturer were to contract with a second retailer who faced a different demand distribution but had an identical cost structure, she could offer the same contract and capture the same fraction of total channel profits. While this independence is generally useful, one must be careful in interpreting the result. For example, one might be tempted to infer that the manufacturer does not need to know the demand distribution to design a coordinating contract; as long as the retailer knows the demand distribution, the correct actions will be taken. Mathematically, this is true, but it makes little economic sense. If the manufacturer does not know the demand distribution, she cannot know Π_I^*. We must then suppose that a rational agent is willing to enter into a transaction even though she cannot estimate her expected profits from the deal. Further, if the retailer were willing to take any contract the manufacturer might propose, we must infer that he has no opportunity cost. If more realistically the retailer has a positive opportunity cost κ, the manufacturer would want to offer a contract that assure acceptance without leaving any excess rents for the retailer. That is, she would choose ε such that:

$$\varepsilon = \frac{\kappa\,(r-c)}{\Pi_I^*}.$$

Clearly to define an acceptable contract, she must have enough information to calculate Π_I^*. Note that the presence of a positive retailer opportunity cost resolves the indeterminacy inherent in Theorem 7. Where before we said that for any positive ε the manufacturer could find a better contract, we now have a specific value of ε defining the best she can do for herself.

A final property to note is that different contracts result in the parties having wildly different variability in their profits. If ε equals $r - c$, the manufacturer transfers at cost and has profits that are identically zero. As ε falls to zero, both her profits and the variability of her earnings increase. The outcomes are reversed for the retailer. Again, the result should be evaluated with care. Both players are assumed to be risk neutral; they are concerned only with their mean earnings and not the variance of those earnings. A buy back contract is not a risk-sharing device in the economic sense. Risk sharing is usually interpreted as a means of maximizing some measure of the total utility of a group of risk averse agents. Common results are, for example, that if one agent is risk neutral he should bear all the risk (see Kreps, 1990, or Salanié, 1997). Here, both players are risk neutral and we seek to maximize the sum of their profits. From a social welfare perspective, all of the coordinating contracts lie on the Pareto frontier and are equally effective. Shifting the variability in profits is just a second-order consequence of using the terms of trade to both assure the stocking of the system-optimal quantity and allocate profits among the players.

8.3.5 Applications of buy backs

The coordinating contracts of Pasternack (1985) are surprisingly general and robust. The result can be interpreted as saying that if a setting can be manipulated to look like a newsvendor problem, it can be successfully decentralized through a system of linear prices. The usefulness of this approach can be seen by the number of researchers who have applied or extended the result. In general, a returns policy can have a significant impact on system performance even if it fails to replicate the outcome of the integrated firm.

Kandel (1996) covers much of the same ground as Pasternack (1985) but from an economist's perspective. In particular, he emphasizes the incentive for a manufacturer to implement a consignment policy. He also notes that if the demand distribution depends on the retail price, coordination cannot be achieved through buy backs unless the manufacturer can impose resale price maintenance. Although he does not formally model them, he does consider risk aversion and the provision of promotional effort and customer service.

Donohue (1996) considers two settings. In the first, the manufacturer has two production technologies and the retailer learns additional information prior to the demand realization. The first production technology has a low marginal cost but a long lead time while the second has a high marginal cost but short lead time. The retailer's additional information allows for a better estimate of realized demand but is learned at point at which only the high cost technology is feasible. She shows that the system can be coordinated by allowing the retailer to place a supplemental order after observing the information. The wholesale price for items purchased in the supplemental order is higher than for those purchased in the initial order, but any leftover stock may be returned at a single return rate. Again, a continuum of coordinating contracts exists and is independent of the demand distribution.

In her second model, there is a single production opportunity which may be initiated after the retailer has gained his information but production requires a specific component with a long lead time; there is no substitute for the component available after the retailer learns his refined information. Thus the total production quantity is limited by the amount of the component purchased. The coordinating contract now includes an option. The retailer pays the manufacturer o per unit to reserve the production of a finished good but must pay w per unit when exercising the option. The manufacturer again offers a buy back rate b. Once more we have a continuum of contracts that is independent of the demand distribution that will coordinate the system.

Note that in both models, there are two decisions that must be coordinated. In the first, it is the production quantities for the two technologies. In the second, it is the amount of the component to purchase and the final production quantity. The contracts consequently have three parameters. The extra degree of flexibility allows for the coordination of both actions.

Emmons and Gilbert (1998) also have a system in which two decisions must be made. Their model is similar to what we have considered here except that the retailer sets both the stocking level and the retail price. The demand distribution is a function of the retail price. Its density $f(\xi|r)$ satisfies

$$f(\xi|r) = \frac{1}{D(r)}\hat{f}\left(\frac{\xi}{D(r)}\right),$$

where $D(r)$ is a decreasing function of the retail price and $\hat{f}(\xi)$ is a probability density with mean one. While buy back contracts cannot coordinate the channel, there does exist a range of wholesale prices such that for any price in that range, the manufacturer and the retailer are both better off when the manufacturer offers a positive buy back rate.

In Ha (1998), a supplier proposes a contract to a manufacturer for the supply of a component. The manufacturer incurs a cost to turn the component into a finished good and must choose the number of units to produce as well as the retail price. The retail price affects demand in an additive fashion. For a given retail price r, demand is given by $D(r)+\xi$, where $D(r)$ is a decreasing function and ξ is a mean zero random variable. If the supplier knows the manufacturer's cost to complete production, the system can be coordinated through franchising, quantity forcing, or resale price maintenance combined with a buy back policy. Obviously, the latter reduces to our setting. If there is asymmetric information regarding the manufacturer's cost, the supplier can design a menu of incentive compatible contracts. For each possible cost realization, the manufacturer chooses a unique contract. The menu relies on resale price maintenance and a non-linear price schedule. The supplier is worse offer under asymmetric information since she must leave higher returns for manufacturers with lower costs to induce self revelation.

Cachon and Lariviere (1997) also consider alternative information structures. Like Ha (1998), they have a manufacturer contracting with an upstream supplier except now it is the manufacturer who proposes the contract. They show that the effectiveness of different contractual forms depends on the system's enforcement mechanism and information structure. If the supplier is obliged to cover the manufacturer's order and there is symmetric information, then offering to pay a termination fee for each item in the original order not taken allows the manufacturer to coordinate the system. It is the quantity forcing aspect of the enforcement mechanism that produces the result, not the marginal incentives of the price and termination fee. If enforcement is not possible, the manufacturer never offers a termination fee under full information. However, if the manufacturer is privately informed about the demand distribution, a manufacturer expecting a large market offers a termination fee to signal her information to the supplier. The fee plays only an information role. A high wholesale price is still required to induce the supplier to provide adequate capacity.

Narayanan and Raman (1997) note that most retailers carry multiple products in each category. Thus while a manufacturer loses sales following a stock out, the retailer might capture sales of a competing good. They consider a manufacturer selling to a retailer who also carries another product (say, a private label) that is never out of stock. If the manufacturer's good stocks out, a fraction of its customer p switch to the other good. In addition, the manufacturer can exert effort to increase demand for the product. They consider three channel management structures: retailer managed inventory (similar to our price-only contacts); vendor managed inventory (in which the manufacturer sets the stocking level of her product); and contracting on unsold inventory (similar to our buy back contracts). They show that double marginalization will result in too little stock and too little effort relative to the integrated channel when retail managed inventory is used. Vendor managed inventory results in the integrated channel effort level but too much stock since the vendor does not internalize the sales of the other product. Which management structure leads to better channel performance depends on the parameters of the problem. When the manufacturer implements a buy back, the channel outperforms simple retailer managed inventory although it does not perfectly coordinate the channel. How buy backs stack up to vendor managed inventory depends on the parameters of the problem.

While Narayanan and Raman (1997) have one retailer carrying two competing products, Padmanabhan and Png (1997) have two competing retailers carrying one product that must be purchased from the same manufacturer. The retailers face a linear demand curve with an uncertain intercept. The sequence of play has the manufacturer posting the terms of trade and the retailers ordering. Only after ordering do the retailers observe the realized intercept and determine which price to charge. Absent a returns policy, the retailers will have so much stock that they will engage in intense competition when the demand realization is low. Fearing a disastrous outcome in a small market, they will order relatively little stock and not be able to take full advantage of large market realizations. By offering a returns policy, the manufacturer mitigates "excessive" competition in small markets and encourages higher stocking levels that allow the channel to take advantage of large market outcomes.[8] The authors limit the analysis to full returns policies, so in some instances the manufacturer may prefer a price-only contract. Intuitively, allowing for partial returns should make returns more attractive.

Finally, a pair of papers suggest decentralization schemes related to buy backs for multiechelon inventory systems similar to Clark and Scarf (1960). Chen (1997) supposes that the supply chain is owned by a principal who hires a manager to run each echelon. The echelons are run as cost centers evaluated on long run average costs. If the costs charged to each echelon are based on accounting inventory (which measures what net inventory would have been if the upstream supply had been perfectly reliable), each manager's problem reduces to setting a base stock level to a critical fractile. By setting the holding and back order cost for each echelon appropriately (the solution is not unique), the principal can induce each manager to set the system optimal base stock level for his echelon.

---

[8]Deneckere, Marvel, and Peck (1996 and 1997) present related arguments for resale price maintenance.

Cachon and Zipkin (1997) consider a two echelon supply chain in which the upstream manufacturer and the downstream retailer set their base stock levels in a competitive fashion. They consider the difference in outcome if the players choose local or echelon base stock levels and show that the system optimal solution is rarely a Nash equilibrium. They also show that a simple linear payment scheme can result in the competitive system matching the performance of the centralized. Under this system, the manufacturer subsidizes the retailer's holding of inventory and pays a penalty when she is back ordered. The retailer, on the other, pays a penalty to the manufacturer when he is back ordered. Again, the coordinating contract is not unique. As in Donohue (1996) two decisions must be coordinated, and the contract must involve three parameters.

Together, these last two papers show the power of buy back contracts. In comparison to other decentralized schemes that have been suggested for multiechelon systems (e.g., Lee and Whang, 1996), these are simple and intuitive. The Chen approach is somewhat more robust. The resulting equilibrium under his scheme is supported by iterated dominance; for the manager of the n-th echelon, picking the system optimal is a dominant strategy assuming that all managers below him follow the optimal policy. In contrast, Cachon and Zipkin cannot rule out that the existence of multiple Nash equilibria under a coordinating contract. The robustness of Chen's scheme, however, comes at a significant cost: All managers must contract individually with a principal, which essentially limits the approach to intrafirm transactions. In Cachon and Zipkin, the parties contract directly with each other, which potentially allows for greater applicability.

8.4 RETURNS POLICIES: QUANTITY FLEXIBILITY CONTRACTS

As the previous section has shown, a returns policy in the form of a buy back rate can coordinate a system in which a manufacturer sells to a retailer facing a newsvendor problem. We now examine an alternative way of implementing returns, quantity flexibility (QF) contracts. QF contracts are frequently used for components in the electronics and computer industry and have occasionally been used in the automotive industry. They also bear some resemblance to "take-or-pay" contracts used in natural resource markets. Tsay and Lovejoy (1998) provides a discussion of these points and a detailed analysis of the contract in a multiperiod setting.

Here we limit the analysis to a single period, so our newsvendor model applies. A QF contract is specified by three parameters: a wholesale price w, a downward adjustment parameter $d \in [0, 1)$, and an upward adjustment parameter $u \geq 0$. The sequence of play has the manufacturer offering the terms of trade and the retailer placing an initial order y. The manufacturer commits to providing $y(1+u)$ units to the system. If realized demand ξ is between $y(1-d)$ and $y(1+u)$, the retailer buys ξ units at price w from the manufacturer to sell to the market. The manufacturer thus provides an "upside" coverage to the retailer of u-percent above his initial order. The QF contract also requires a "downside" commitment from the retailer in the form of a minimum purchase requirement; he may cancel d-percent of his order but must take the remainder.[9] If realized demand is below $y\,(1-d)$, the retailer must take $y\,(1-d)$ units at a wholesale price of w, filling demand and salvaging the rest.

It will be convenient to define $\chi = \frac{1-d}{1+u}$. Since total channel stock will be $y\,(1+u)$, χ represents the fraction of channel stock for which the supplier is responsible and is thus a measure of the flexibility offered the retailer with lower values of χ reflecting greater flexibility.

To see that relationship between QF contracts and buy backs, suppose that under a QF contract the retailer had to pay for all $y\,(1+u)$ units up front but could still cancel his commitment down to $y\,(1-d)$ units. For each unit canceled, the manufacturer would refund the wholesale price. Interpreted this way, a QF contract allows the retailer to cancel or return a fraction of his order for a full refund of the wholesale price. In contrast, a buy back contract allows the retailer to cancel his full order for a fractional refund of the wholesale price.

In the remainder of this section, we first present the retailer's problem and then consider coordinating contracts. We close with a comparison of QF contracts with buy backs.

8.4.1 The retailer's problem

The retailer's objective under a QF contract is given by:

$$\begin{aligned} \Pi_R(y) &= (r-w)y(1+u) - (r-w)\int_{y(1-d)}^{y(1+u)} F(x)dx \qquad (8.11) \\ &\quad -(r-v)\int_0^{y(1-d)} F(x)dx. \end{aligned}$$

The first term represents the retailer's profits if he were to sell everything. The subsequent terms represent adjustments for lower demand realizations. If demand falls between $y(1-d)$ and $y(1+u)$, he loses revenue of r but saves w. However, when demand is below $y(1-d)$, he only saves v since demand is below his minimum purchase commitment and he is obliged to salvage some units. Note that for a fixed y and w, the retailer's profits are increasing in d and u. That is, if at a given order quantity and fixed wholesale price the manufacturer offers greater upside coverage or demands less downside commitment, the retailer is better off. The following theorem (largely from Tsay, 1997), establishes some properties of the $\Pi_R\,(y)$ and the optimal retailer order y^*.

Theorem 8 *Suppose that the manufacturer offers a QF contract with parameters (w, d, u).*

---

[9]Many of the result of this section go through if the total stock in the system is given by $y + \hat{u}$ and the retailer's minimum purchase is $y - \hat{d}$. Such terms would allow the retailer a fixed range $(\hat{d} + \hat{u})$ of flexibility instead of a fixed percentage of flexibility. The QF formulation is preferable since a percentage commitment will always be positive while an additive commitment will not if $y < \hat{d}$.

1. *The retailer's profits $\Pi_R(y)$ are concave in y.*

2. *The retailer's optimal order y^* unique and is implicitly defined by*

$$\begin{aligned} \Pi'_R(y^*) = (r-w)(1+u)\bar{F}(y^*(1+u)) & \\ -(w-v)(1-d)F(y^*(1-d)) &= 0. \end{aligned} \tag{8.12}$$

3. *The retailer's optimal order y^* is decreasing in the wholesale price and increasing in the downward adjustment parameter d. It is increasing in the upward adjustment parameter u if*

$$g\left(y^*\left(1+u\right)\right) < 1, \tag{8.13}$$

where $g(\xi)$ is the generalized hazard rate of the demand distribution.

4. *The total amount of stock in the system $y^*(1+u)$ is increasing in u.*

Proof: Differentiating $\Pi_R(y)$ twice yields

$$\Pi''_R(y) = -(r-w)(1+u)^2 f(y(1+u)) - (w-v)(1-d)^2 f(y(1-d)),$$

which is always negative. As profits are concave, first order conditions are sufficient for a unique solution. The third part of the theorem follows from applying the implicit function theorem to (8.12). In particular we have:

$$\begin{aligned} \frac{\partial y^*}{\partial u} &= -(r-w)\left(\bar{F}(y^*(1+u)) - (1+u)y^* f\left(y^*(1+u)\right)\right)/\Pi''_R(y^*) \\ &= -(r-w)\bar{F}(y^*(1+u))\left(1-g\left(y^*\left(1+u\right)\right)\right)/\Pi''_R(y^*). \end{aligned}$$

As $\Pi''_R(y^*)$ is negative, we require that $1-g\left(y^*\left(1+u\right)\right)$ be positive for y^* to be increasing in u. The final part of the theorem follows from differentiating $y^*(1+u)$ and using the expression for $\frac{\partial y^*}{\partial u}$. □

Note that if $d=u=0$, than the retailer's problem reduces to a newsvendor problem. Otherwise it will not be generally possible to determine an explicit solution to (8.12) even if $F(\xi)$ is easily invertible. In the standard newsvendor problem, one balances the chance of being over and under realized demand at one point. Under a QF contract, that balance must be struck at two distinct points. One point corresponds to the manufacturer's upside commitment and the other to the retailer's downside minimum purchase. An explicit solution is hence not readily possible unless either $F(\xi)$ or $\bar{F}(\xi)$ is a homogeneous function so that either $F(\lambda\xi) = \lambda^\rho F(\xi)$ or $\bar{F}(\lambda\xi) = \lambda^\rho \bar{F}(\xi)$ for $\lambda > 0$ and some ρ. For example, if $F(\xi)$ is homogeneous, the solution to (8.12) has a modified critical fractile structure:

$$F\left(y^*\right) = \frac{(r-w)(1+u)}{(r-w)(1+u)^{\rho+1} + (w-v)(1-d)^{\rho+1}}.$$

The alternative case has similar results. It is straightforward to show that $F(\xi)$ is homogeneous for the uniform and that $\bar{F}(\xi)$ is homogeneous for the Pareto distribution.

The comparative statics are largely as one would expect. The retailer orders less when he has a slimmer margin. Increasing d lowers both the absolute level of the minimum purchase requirement and the marginal rate at which the minimum purchase obligation is increasing. Hence he orders more. The case for the upward adjustment is not as clear. On the one hand as u increases, the retailer has greater upside coverage for any order level. He could reduce his current order, lower his minimum purchase requirement, and still enjoy a larger upside coverage. On the other hand, the rate at which an increase in his order extends his upside coverage has increased – arguing for a boost in his order.

The outcome depends on how likely an incremental increase of upside coverage is going to be used. That, in turn, depends on the generalized hazard rate. Note that (8.13) corresponds to the elasticity of orders under a price-only contract being greater than one, which will be so under the optimal price-only contract. Consequently, if a manufacturer moves from the optimal price-only contract to a QF contract with upside flexibility without changing the wholesale price, the retailer's order will increase. If $F(\xi)$ is IGFR, however, the retailer's order will at some point be decreasing in u. The final part of the theorem establishes that even if the retailer reduces his order in response to an increase in the upward adjustment parameter, he does not reduce it so much that the total amount of stock in the system falls.

8.4.2 Coordinating contracts

We now seek coordinating contracts similar to those we had for buy backs. To begin, note that if we want the system to stock y_I units, we must induce the retailer to order $\frac{y_I}{1+u}$. Following Tsay (1997), we substitute $\frac{y_I}{1+u}$ for y in (8.12) and perform a few manipulations to yield:

$$\begin{aligned} w^*(d,u) &= v + \frac{(r-v)(1+u)\frac{c-v}{r-v}}{(1+u)\frac{c-v}{r-v} + (1-d)F\left(y_I\frac{1-d}{1+u}\right)} \\ &= v + \frac{c-v}{\frac{c-v}{r-v} + \chi F(\chi y_I)}. \end{aligned} \tag{8.14}$$

Note that for the allowed ranges of d and u, $0 < \chi F(\chi y_I) \leq \frac{r-c}{r-v}$, which in turn implies that $c \leq w^*(d,u) < r$. The contract is thus economically sensible. The optimal wholesale price depends on the adjustment parameters only through the ratio χ, the fraction of the total system inventory for which the retailer is ultimately responsible. Although a QF contract generally has three parameters, a coordinating QF contract in reality only has two: the wholesale price and the flexibility offered the retailer, χ. For simplicity we will write the coordinating wholesale price simply as a function of χ, i.e., $w^*(\chi)$.

Given our restrictions on d and u, it must be that $0 < \chi \leq 1$. We again have that a continuum of coordinating contracts exists, and they differ in the flexibility they offer the retailer. As before, if the retailer is given no flexibility, the product must be transferred at marginal cost ($w^*(1) = c$). At the other extreme, the retailer cannot be offered complete flexibility. He must remain responsible for some fraction (even if trivially small) of the system's inventory. It is simple to show that the coordinating wholesale price is decreasing in the flexibility parameter χ and that the limiting price as χ goes to zero is the retail price r.

The remaining issue is how system profits are divided. We again have that both parties would prefer to bear greater responsibility for excess stock.

Theorem 9 *If the manufacturer offers a coordinating contract* $(w^*(\chi), \chi)$*, the retailer's profits* $\Pi_R(w^*(\chi), \chi)$ *are increasing in* χ.

Proof: Under a coordinating contract, the derivative of the retailer's profits with respect to χ can be written as:

$$\begin{aligned}
\frac{\partial \Pi_R(w^*(\chi), \chi)}{\partial \chi} &= -w^{*\prime}(\chi)\left(y_I - \int_{\chi y_I}^{y_I} F(\xi)\, d\xi\right) - y_I F(\chi y_I)(w^*(\chi) - v) \\
&= \frac{(w^*(\chi) - v)}{\frac{c-v}{r-v} + \chi F(\chi y_I)} ((F(\chi y_I) + \chi y_I f(\chi y_I)) \\
&\quad \left(y_I - \int_{\chi y_I}^{y_I} F(\xi)\, d\xi\right) - y_I F(\chi y_I)\left(\frac{c-v}{r-v} + \chi F(\chi y_I)\right)\Bigg),
\end{aligned}$$

where the second equality follows from noting that:

$$w^{*\prime}(\chi) = -\frac{(w^*(\chi) - v)((F(\chi y_I) + \chi y_I f(\chi y_I)))}{\left(\frac{c-v}{r-v} + \chi F(\chi y_I)\right)}.$$

To prove the result, it is sufficient to show that

$$F(\chi y_I)\left(y_I - \int_{\chi y_I}^{y_I} F(\xi)\, d\xi\right) > y_I F(\chi y_I)\left(\frac{c-v}{r-v} + \chi F(\chi y_I)\right).$$

This follows from noting that

$$y_I - \int_{\chi y_I}^{y_I} F(\xi)\, d\xi = y_I\left(\frac{c-v}{r-v} + \chi F(\chi y_I)\right) + y_I \int_{\chi y_I}^{y_I} f(\xi)\, d\xi,$$

where we use the fact that $F(y_I) = \frac{r-c}{r-v}$. □

With the theorem, we have that QF contracts are almost as flexible as buy backs. Under either, a range of coordinating contracts exists, and the more flexibility the manufacturer offers, the higher the wholesale price and manufacturer profits are. The only reason to qualify the comparison is that the under QF contracts, the coordinating wholesale price depends on the demand distribution. The manufacturer can no longer rely on one contract to coordinate all markets. This is hardly surprising since the retailer's optimal order balances overage and underage costs at two distinct points of the demand distribution. It is also as we argued above not that dramatic a shortcoming since it often makes no economic sense to assume that the manufacturer is uninformed about the demand distribution. Nevertheless, one can show that for some demand distributions the coordinating contract is independent of some parameter.

Theorem 10 *If $F(\xi|\theta)$ is from a scaled family, the coordinating wholesale price $w^*(\chi|\theta)$ is independent of θ for all values of χ. That is:*

$$w^*(\chi|\theta) = w^*(\chi|1).$$

Proof: For a scaled family, we have that $y_I(\theta) = \tau(\theta) y_I(1)$ and thus that $F(\chi y_I(\theta)|\theta) = F(\chi y_I(1)|1)$. Therefore, $w^*(\chi|\theta) = w^*(\chi|1)$. □

The theorem immediately suggests that the same coordinating contract can be used in multiple markets if the demand distribution for each market is from the same scaled family. In particular, if demand in each market is normal with the same coefficient of variation, then one coordinating QF contract suffices.

8.4.3 Quantity flexibility vs. buy backs

The remaining issue to consider is whether a manufacturer offering a returns policy has any reason to favor a QF contract over a buy back. Relatively little research has been done on QF contracts, but there are reasons to believe that they will generally prove as useful as buy backs. For example, Tsay (1997) has extended their application to a setting similar to Donohue (1996). Further, it is simple to show that QF contracts, like buy backs, will fail to coordinate the channel when the demand distribution depends on the retail price as in Emmons and Gilbert (1998).

Hence, the sole reason we have established for favoring buy backs is that coordinating buy back contracts do not depend on the demand distribution. Even this property is arguably over-valued, and the QF dependence on the distribution can be relaxed for some distributional families. Indeed, one may argue that the distributional dependence benefits QF contracts since the manufacturer cannot induce self selection through coordinating buy back contracts.

To see this point, consider a manufacturer selling to two retailers who face different demand distributions and have positive opportunity costs. The manufacturer knows that each market is characterized by one of two demand distributions but cannot verify a particular market's prevailing distribution. Ideally, the manufacturer would like to offer a menu of contracts that induces retailers with different demand distributions to choose different contracts for themselves. Further, she would like to avoid sacrificing any channel profits and grab as large a share of profits as possible. Coordinating buy backs simply cannot do this. Because the coordinating contracts are independent of the demand distribution, all retailers will prefer the one that offers the least flexibility regardless of their demand distribution. Self selection necessarily requires deviating from the coordinating contract and sacrificing channel profits. In any given setting, self selection through coordinating QF contracts might also require sacrificing coordination, but there is at least a possibility of achieving both goals.

8.5 ALTERNATIVE CONTRACTS

Returns policies are an effective means of improving system performance, but other possibilities exists. Here we consider two. The first approach can be classified as penalty schemes. In lieu of offering a safety net to the retailer, the manufacturer brandishes a stick. The second contract is the standard-setting scheme of Atkinson (1979). Unlike returns policies, both penalties and standards are best suited for intrafirm coordination.

8.5.1 Penalty methods

Suppose that the manufacturer offers the good to the retailer at a wholesale price w with no return policy but with a demand for a payment of p per unit for any missed sales. The retailer's objective becomes:

$$\Pi_R(y) = (r-w)\,y - (r-v)\int_0^y F(\xi)\,d\xi - p\int_y^\infty \bar{F}(\xi)\,d\xi,$$

and his optimal order is:

$$y(w,p) = F^{-1}\left(\frac{r+p-w}{r+p-v}\right).$$

It is straightforward to show that if the manufacturer imposes a penalty of

$$p(w) = \frac{(w-c)(r-v)}{c-v}$$

with a wholesale price of w then the retailer orders the integrated channel quantity (i.e., $y(w,p(w)) = y_I$).

It is straightforward to see that $p(c) = 0$ and that $p(w)$ is increasing in w. We also have that the retailer's profits are decreasing in w since this decreases his expected revenue and increases his expected penalty payments. Consequently, the manufacturer can capture all channel profits at a wholesale price strictly below the retail price. If she were to push the wholesale price to the retail price (as she can under returns policies), the retailer would rationally refuse the contract since it would saddle him with an uncompensated loss.

Where a returns policy works by manipulating the consequence of having excess stock, a penalty method alters the consequences of being short. As such, its implementation may be difficult. In particular, it must be possible to observe lost sales. This can be relaxed slightly. Suppose that instead of imposing a linear payment rate, the manufacturer charges a lump sum in the event of a stock out and offers the retailer a menu of contracts in which the lump sum depends on the retailer's initial order. In particular suppose that for a given wholesale price w, the penalty payment is

$$P(y|w) = p(w)\,\frac{\int_y^\infty \bar{F}(\xi)\,d\xi}{\bar{F}(y)}.$$

For a stocking level of y, the retailer incurs the penalty with probability $\bar{F}(y)$, and his objective becomes equivalent to that under the linear payment scheme.

The quantity $\int_y^\infty \bar{F}(\xi)\,d\xi/\bar{F}(y)$ is known as the mean residual life in the reliability literature (Barlow and Proschan, 1965) and may be interpreted here as the expected number of items short given that demand is greater than the stocking level. Instead of basing the charges on the realized shortfall, the manufacturer exploits the retailer's risk neutrality and bases the charges on the expected shortfall given the stocking level. The information requirements have been significantly reduced. The manufacturer does not have to observe how many sales were lost, only that a stock out occurred.

Implementation may still be difficult. Under a returns policy, the retailer has an incentive to cooperate and allow the manufacturer to audit his sales since demand information is tied to his receiving compensation in a down market. Here, demand information is tied to being penalized and the retailer has every reason to make the auditing process difficult. Consequently, a penalty scheme may be more suitable for use within a firm instead of between firms since verifying information would hopefully be simpler in an intrafirm setting.

Intrafirm coordination through an alternative penalty scheme is considered by Celikbas, Shanthikumar, and Swaminathan (1997). Here, marketing provides a forecast of demand to manufacturing who must then produce the good prior to the realization of demand. The two functions are run by separate managers. They propose coordinating the system by evaluating manufacturing as a cost center and having the manufacturing manager pay a per unit penalty to the principal if the production quantity is less than both realized demand and marketing's forecast. The penalty induces manufacturing to internalize the cost of a lost sale. The marketing manager is sold the rights to the market but must pay a per unit penalty to the principal for each unit that his forecast exceeds the demand. As excess inventory remains on the books of manufacturing, the penalty forces the marketing manager to internalize the cost of excess production.

The scheme is somewhat more complicated than establishing an internal pricing or returns policy between the functions since it requires the parties to contract with the principal instead of each other. Also note that the data point marketing provides is not what one would usually term a "forecast." For the scheme to work, the complete demand distribution must be common knowledge so the data transferred from marketing does not lead anyone to alter their beliefs regarding the market. The forecast plays no informational role but is a useful contracting device for manipulating incentives.

8.5.2 Standard setting

Atkinson (1979) also considers intrafirm coordination but has a more classical principal-agent format. A risk neutral principal hires a risk averse agent to set the stocking level in a newsvendor problem. He shows that if the agent and the principal have the same evaluation of the demand distribution, the agent will set the stocking level below what the principal would choose.

He further argues that the manager through greater involvement in the daily operations of the business may have information about market demand unavailable to the principal. That is, between the time the parties enter into their contract and the time at which the stocking decision must be made, the manager is able to revise his estimate of the demand distribution but the principal is not. This leads to a standards-based contract. The principal offers a fixed wage ω and a gains sharing parameter ϕ (assumed to be between zero and one) while posting a standard y_0. Letting $\pi(y)$ represent realized system profits when the initial stocking level is y, the retailer's total compensation can be written as

$$\omega + \phi(\pi(y) - \pi(y_0))$$

The agent receives ω for certain and shares in a fraction of the gains or loss that result from deviating from the principal's standard.

Suppose the principal sets the standard optimally given her initial information (i.e., so that y_0 is the critical fractile of the initial distribution). Atkinson shows that the agent only deviates from the standard if his revised information leads him to believe the critical fractile has shifted. Further, the stock level moves in the correct direction; the agent's choice always lies between the standard and what the risk neutral principal would have chosen given the agent's revised beliefs. If the players start with a common prior and update beliefs in a Bayesian fashion, the principal is better off implementing the standard based scheme than paying a fixed wage and insisting on implementing y_0.

The standard scheme is intended to be used within a firm and is not necessarily optimal. Dealing with a risk averse agent is going to impose some loss on the system, and there is no guarantee that this scheme minimizes that loss. Nonetheless, the paper is of interest. Contracting with a risk averse manager is an important topic to consider, and the proposed contract is simple and intuitive. It manages to balance the need for risk sharing (by providing a fixed wage) with the need to induce the manager to act on his revised information.

8.6 SUMMARY AND DIRECTIONS FOR FUTURE RESEARCH

We have reviewed a series of results related to perhaps the simplest supply chain contracting model with stochastic demand: a manufacturer selling to a newsvendor. We have presented conditions for the manufacturer's problem under price-only contracts to be well-behaved and shown that the optimal wholesale price is closely related to a generalization of the failure rate of the demand distribution. We have also shown that returns policies of various forms are a powerful tool for improving supply chain performance. Returns policies allow for payments that are conditional on how realized demand compares to the chosen stocking level. As such, they alter the marginal incentive to hold inventory and can reduce – or eliminate – the impact of double marginalization.

Given the apparent power of returns policies, it is not surprising that they are common in industries such as publishing. Indeed, one may wonder why they are not even more common. Relatively little work has examined this issue, but the fact that returns policies are not ubiquitous suggests that in some environments one or more of our assumptions fail to hold. Two obvious considerations are enforceability and cost. To begin with the latter, we assumed that offering a returns policy imposed no additional costs on the system relative to the integrated channel. That is obviously a simplification; merely accounting for the amount returned and providing compensation adds some costs. If system performance under a price-only contract is sufficiently good, it may not be worth incurring the additional cost of running a returns policy. Lariviere and Porteus (1998) offer some evidence that performance under price-only contracts improves as the coefficient of variation falls. It may be that demand in some markets is not sufficiently variable to warrant a returns policy.

The relevance of contract enforceability is best illustrated by quantity flexibility contracts. A manufacturer who offers a QF contract leaves herself facing a newsvendor with demand distributed between $y\,(1-d)$ and $y\,(1+u)$. Building sufficient stock to cover all demand is not necessarily optimal if the manufacturer can freely choose her production quantity. If the retailer anticipates that the manufacturer is not committed to providing the full upside coverage, his optimal order will change. Thus our analysis of QF contracts implicitly assumes an unmodeled enforcement mechanism that compels the manufacturer to fulfill her contractual obligation. Cachon and Lariviere (1997) show that absent such a mechanism returns policies can collapse under symmetric information.

An alternative explanation for not employing returns is offered by Marvel and Peck (1995). They note that there are two forms of market uncertainty: how many customers will show up and how a typical customer will value the product. The problem we have considered here is an extreme representation with the number of arrivals unknown but every arrival's valuation known to be r. As we have seen, returns are valuable in this setting. At the other extreme with the number of arrivals certain but the price they are willing to pay unknown, returns can hurt the manufacturer; the retailer sets a higher price in the presence of a returns policy and sales fall. In most markets, the manufacturer must balance both concerns.

While exploring the limits on the applicability of returns policy is one possible research direction, another fruitful avenue is to extend the lessons learned here to other settings. We argued in the introduction that the basic model we have studied provides an intuitive foundation for designing contracts to improve the performance of more complicated supply chains. This contention is supported by work such as Donohue (1996) and especially Chen (1997) and Cachon and Zipkin (1997), which develop schemes for multiechelon inventory systems.

An important extension to consider is the coordination of multiple actions. As we noted above, a returns policy cannot coordinate the system when the retailer controls the retail price. Finding a modification to returns policies to coordinate such a setting would be very useful. Similar statements could be made for actions such as manufacturer or retailer promotional effort. One might also want to consider the role of forecasting. Atkinson (1979), Donohue (1996), and Tsay (1997) all partially address this issue by considering systems in which some informational refinement occurs over time, but none of these authors explicitly models how that revision occurs. In particular, they do not consider that gathering information might be costly for the retailer and thus that the contract terms might affect the quality of information available to the supply chain. All of these would be worth pursuing, and likely involve some variant of the schemes considered here.

Acknowledgments

I am grateful for the helpful comments of Gerard Cachon, Evan Porteus, Andy Tsay, and Paul Zipkin.

References

Atkinson, A. A., "Incentives, Uncertainty, and Risk in the Newsboy Problem," *Decision Sciences*, 10 (1979), 341-353.

Barlow, R. E. and F. Proschan, *Mathematical Theory of Reliability* Cambridge: Princeton University Press, 1965.

Butz, D., "Vertical Price Controls with Uncertain Demand," *Journal of Law and Economics*, 40 (1997), 433-459.

Brown, A. O. and H. L. Lee, "Optimal Pay to Delay Capacity Reservation with Applications to the Semiconductor Industry," Stanford University working paper, 1997.

Cachon, G. P. and M. A. Lariviere, "Contracting to Assure Supply, or What did the Supplier Know and When did He Know It?," Duke University working paper, 1997.

Cachon, G. P. and P. H. Zipkin, "Competitive and Cooperative Inventory Policies in a Two-Stage Supply Chain," Duke University working paper, 1997.

Celikbas, M., J. G. Shanthikumar, and J. M. Swaminathan, "Coordinating Production Quantities and Marketing Efforts through Penalty Schemes," University of California at Berkeley working paper, 1997.

Chen, F., "Decentralized Supply Chains Subject to Information Delays," Columbia University working paper, 1997.

Clark, A.J. and H.E. Scarf, "Optimal Policies for a Multi-Echelon Inventory Problem," *Management Science*, 6 (1960), 475-490.

Deneckere, R., H. Marvel, and J. Peck, "Demand Uncertainty, Inventories, and Resale Price Maintenance," *Quarterly Journal of Economics*, 111 (1996), 885-913.

Deneckere, R., H. Marvel, and J. Peck, "Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition," *American Economic Review*, 87, 4 (1997), 619-641.

Donohue, K., "Supply Contracts for Fashion Goods: Optimizing Channel Profits," Wharton School working paper, 1996.

Emmons, H., and S. Gilbert, "Returns Policies in Pricing and Inventory Decisions for Catalogue Goods," *Management Science*, 44, 2 (1998), 276-283.

Eppen, G. and A. Iyer, "Backup Agreements in Fashion Buying - The Value of Upstream Flexibility," *Management Science*, 43, 11 (1997), 1469-1484.

Ha, A., "Supply Contract for a Short-Life-Cycle Product with Demand Uncertainty and Asymmetric Cost Information," Yale School of Management working paper, 1998.

Jeuland, A. and S. Shugan, "Managing Channel Profits," *Marketing Science*, 2 (1983), 239-272.

Kandel, E., "The Right to Return," *Journal of Law and Economics*, 39 (1996), 329-56.

Lariviere, M. A. and V. Padmanabhan, "Slotting Allowances and New Product Introductions," *Marketing Science*, 16, 2 (1997), 112-128.

Lariviere, M. A. and E. L. Porteus, "Stalking Information: Bayesian Inventory Management with Unobserved Lost Sales," Stanford University working Paper, 1997.

Lariviere, M. A., and E. L. Porteus, "Selling to the Newsvendor," Duke University working paper, 1998.

Lee, H. and S. Whang, "Decentralized Multi-Echelon Supply Chains: Incentives and Information," Stanford University working paper, 1996.

Kreps, D. M., *A Course in Microeconomic Theory* Cambridge: Princeton University Press, 1990.

Marvel, H. P. and J. Peck, "Demand Uncertainty and Returns Policies," *International Economics Review*, 36, 3 (1995), 691-714.

Mathewson, G. F. and R. A. Winter, "An Economic Theory of Vertical Restraints," *The Rand Journal of Economics*, 15 (1984), 27-38.

Moorthy, K. S., "Managing Channel Profits: Comment," *Marketing Science*, 6 (1987), 375-379.

Narayanan, V. G. and A. Raman, "Assignment of Stocking Decision Rights under Incomplete Contracting," Harvard University working paper, 1997.

Padmanabhan, V. and I.P.L. Png, "Returns Policies: Make Money by Making Good," *Sloan Management Review*, Fall (1995), 65-72.

Padmanabhan, V. and I.P.L. Png, "Manufacturer's Returns Policy and Retail Competition," *Marketing Science*, 16, 1 (1997), 81-94.

Pasternack, B. A., "Filling Out the Doughnuts: The Single Period Inventory Model in Corporate Pricing Policy," *Interfaces*, 10, 5 (1980), 96-100.

Pasternack, B. A., "Optimal Pricing and Returns Policies for Perishable Commodities," *Marketing Science*, 4, 2 (1985), 166-176.

Rey, P. and J. Tirole, "The Logic of Vertical Restraints," *The American Economics Review*, 5, 76 (1986), 921-939.

Salanié, B., *The Economics of Contracts: A Primer* Cambridge: MIT Press, 1997.

Spengler, J. J., "Vertical Integration and Antitrust Policy," *Journal of Political Economy*, 58 (1950), 347-352.

Tirole, J., *The Theory of Industrial Organization* Cambridge: MIT Press, 1988.

Tsay, A. A., "The Quantity Flexibility Contract and Supplier-Customer Incentives," Santa Clara University working paper, 1997.

Tsay, A. A. and W. S. Lovejoy, "Quantity Flexibility Contracts and Supply Chain Performance," Santa Clara University working paper, 1998.