MANUFACTURER-RETAILER CONTRACTING UNDER AN UNKNOWN DEMAND DISTRIBUTION

Martin A. Lariviere

Fuqua School of Business

Duke University

Evan L. Porteus

Graduate School of Business

Stanford University

Draft

December, 1995

Abstract

We consider a manufacturer introducing a new product into a distribution channel and examine what wholesale price should be charged. The setting in many ways is simple. The channel is abbreviated with the manufacturer selling directly to the retailer. The contract is also simple, merely a flat wholesale price. Demand is stochastic but independent and identically distributed in each period. Complications arise from two additional assumptions. First, neither party knows some parameter of the demand distribution. The system evolves informationally as the channel has more experience with, and information about, the product. Second, we assume unmet demand is both lost and unobserved, so only sales data are available. The autonomous retailer's stocking level consequently dictates the rate at which the channel acquires information. The manufacturer's pricing policy, in turn, influences the retailer's actions.

We explore how the wholesale price evolves as beliefs are updated in a Bayesian fashion. Pricing is driven by the precision of information and not the size of the market. In particular, we show that the manufacturer charges a lower price following a stockout than after an exact observation. That is, she prices more aggressively following a signal of relatively weak demand (unsold stock) than after a signal of strong demand (empty shelves). The apparent anomaly is explained by relating the precision of information to the number of observed stockouts and the elasticity of retailer orders to the precision of information; stockouts are less informative, and an uncertain retailer is relatively price sensitive.

In many markets, manufacturers do not sell directly to consumers but distribute products through decentralized channels; wholesalers and distributors take possession of the goods before retailers offer them to customers. When the links of the distribution chain are independent firms, a manufacturer must bundle her product with a contract governing channel transactions. Contracting within a distribution channel has received considerable attention within economics (*e.g.*, Spengler, 1950; Mathewson and Winter, 1984; Tirole, 1988, provides a survey) and marketing (*e.g.*, Jeuland and Shugan, 1983; Moorthy, 1987). Both fields emphasize the impact of the contractual form on total channel profits and their division. Both generally presume a known deterministic demand curve. Issues of inventory or stochastic demand are ignored.

In the operations management literature, researchers have examined schemes that lower a supplier's ordering and inventory costs in an economic order quantity setting. Monahan (1984) and Lee and Rosenblatt (1986) suggest quantity discounts to alter buyer's ordering patterns. Weng (1995) shows that quantity discounts alone will not maximize channel profits. Pasternack (1985) considers both inventory and stochastic demand; his retailer faces a classic, single period newsvendor problem for a known demand distribution. He proves that channel profits are maximized when the manufacturer offers a returns policy.

Here, we study a manufacturer who launches a new product in an uncertain environment and then dynamically adjusts the wholesale price as information is revealed. The manufacturer deals directly with a self-interested retailer who holds stock in anticipation of demand and sells the product at a fixed retail price. Demand is stochastic with all realizations independently and identically distributed, but neither party knows some parameter of the demand distribution. Both gain information about that parameter by observing sales. Thus, despite stationary demand, the demand distribution appears different in each period because estimates of the unknown parameter change. The system evolves informationally as the channel gains experience with the product. Further, we assume that unmet demand is both lost and unobserved. Beliefs about the *demand* distribution are updated based on observed *sales*. As sales data provide only a lower bound on realized demand when the retailer stocks out, the autonomous retailer's stocking decision dictates the rate at which the channel acquires information. The manufacturer's pricing policy, in turn, influences the retailer's actions.

We explore how the manufacturer's wholesale price schedule evolves over a finite horizon as beliefs about the demand distribution are updated in a Bayesian fashion. In our model, the pricing decision is driven by the precision of information and not the size of the market. In particular, we show that the manufacturer charges a lower price following a stockout than after an exact observation. That is, she prices more aggressively following a signal of relatively weak demand (unsold stock) than after a signal of strong demand (empty shelves). The apparent anomaly is explained by relating the precision of information to the number of observed stockouts and the elasticity of retailer orders to the precision of information; stockouts are less informative, and an uncertain retailer is relatively price sensitive.

The next section gives the basic assumptions and structure of the model. We also present the retailer's optimal stocking policy for a given price schedule. To maintain the tractability of the retailer's problem, we must restrict the manufacturer to a limited set of admissible price schedules. In section 2, we first show that if the manufacturer uses such a schedule her profits have a very simple form that allows one to find the optimal admissible schedule easily. We then argue that admissible schemes are all the manufacturer need consider; she may well prefer following some alternative schedule but is unable to commit to its use. Section 3 characterizes the behavior of the optimal price schedule for the special case of exponential demand and perishable inventory. It is here that we discuss the relation between the wholesale price and the state of information. Section 4 considers a simple example that allows us to examine the behavior of prices over time and the impact of retailer inventory.

1. MODEL BASICS AND THE RETAILER'S PROBLEM

Over a finite horizon of length N, demand realizations are independent and identically distributed draws from some "true" distribution. The family of the demand distribution is known, but some parameter ω of its density $\psi(\cdot|\omega)$ is not. There is a prior on the unknown parameter with density $g(\omega)$. One also has the predictive or marginal density, $\phi(\xi) = \int_{-\infty}^{\infty} \psi(\xi|\omega) g(\omega) d\omega$. $\phi(\xi)$ is the updated demand density, the best estimate of the demand density given the prior. The standard approach is to assume that $g(\omega)$ is from the conjugate family of the demand distribution (DeGroot, 1970). The posterior of the unknown parameter

(which is the prior for the next period) would then come from the same family as $g(\omega)$, and the subsequent updated demand distribution would be from the same family as $\phi(\xi)$. Additionally, one would have a fixed-dimensional sufficient statistic that would completely characterize the prior. By updating the statistic following an observation, the updated sufficient statistic would similarly fully characterize the subsequent posterior. One would not need to retain each observation explicitly.

The difficulty is that sales and not demand are observed. Although demand realizations are identically distributed, sales observations are not. Sales data contain both exact and censored observations, and a conjugate prior does not necessarily exist even if there is one for fully observed demand. Braden and Freimer (1991) characterize "newsboy" distributions for which a conjugate prior and a fixed-dimensional sufficient statistic exists even with censored observations. The newsboy family has the form $\Psi(\xi|\omega) = 1 - e^{-\omega d(\xi)}$ with $d(\xi) \geq 0$. The gamma distribution with shape parameter a and scale parameter S is the conjugate prior for all newsboy distributions. The updated gamma parameters (a_n, S_n) are the sufficient statistics for period n. Let m_n be the number of exact observations by the start of period n and s_i the i-th sales observation. (a_n, S_n) are calculated as:

$$a_n = a_1 + m_n \qquad\qquad S_n = S_1 + \sum_{i=1}^{n-1} d(s_i)$$

In updating the scale parameter, exact and censored observations are treated equally: For either, $d(\xi)$ is evaluated at the observed value and added to the sum. The shape parameter a_n, on the other hand, changes only with exact observations. It can (and will) be interpreted as a measure of the precision, or state, of the system's information.

In what follows, we restrict $\psi(\xi|\omega)$ to be a newsboy distribution and assume a gamma prior. Additionally, we wish to use the state space reduction technique of Azoury (1985), which imposes further requirements. Of the newsboy distributions presented in Braden and Freimer (1991), only the Weibull distribution satisfies the conditions of Azoury (1985). We therefore assume the true demand has a Weibull distribution with known k and unknown ω, which yields:

$$\psi\left(\xi|\omega\right) = \omega k \xi^{k-1} e^{-\omega \xi^k} \qquad g\left(\omega|a,S\right) = \frac{S^a}{\Gamma(a)} \omega^{a-1} e^{-S\omega} \qquad \phi\left(\omega|a,S\right) = \frac{ak\xi^{k-1}S^a}{\left(S+\xi^k\right)^{a+1}}$$

The Retailer's Problem

We now review the solution to the retailer's stocking problem when beliefs about the unknown parameter are updated in a Bayesian fashion and unmet demand is lost and unobserved (hereafter, we simply say lost sales). A full analysis is presented in Lariviere and Porteus (1995). All costs and revenues are linear. The product is sold at a fixed retail price *r*. The retailer is charged *h* for each unit that remains unsold and *p* for each unit of unsatisfied demand. Future profits are discounted by a factor α, and *x* denotes the retailer's inventory level before ordering.

The manufacturer follows a linear price schedule, posting a wholesale price for each period and supplying as much as the retailer wants at that price. Further, we restrict the manufacturer to a set of "admissible" linear price schedules in which the posted wholesale price depends on the period, the current state of information, and the normalized retailer inventory. (The inventory normalization is explained below.) An admissible schedule is independent of the scale parameter *S* or the exact history of sales. We denote the wholesale price for the current period by *w*, suppressing its dependence on other parameters for simplicity. The manufacturer cannot contract on future prices although the retailer is assumed to anticipate correctly the price for each possible future state. We return to the importance of this assumption later.

Limiting the manufacturer to admissible schedules is obviously restrictive. We argue below that such schemes are all she need consider because they are the only linear price schedules that can be self-enforcing. Given the manufacturer's inability to commit to future prices, an admissible schedule is the only linear schedule whose use the retailer rationally anticipates.

Lariviere and Porteus (1995) assume a fixed wholesale price is charged over the horizon but point out that their *Theorem 1* (as stated below) generalizes to the case in which wholesale prices follow one of our admissible schedules. Thus, letting $f_n\left(x|a,S\right)$ denote the retailer's maximum expected discounted profit over periods *n, n+1, …, N* when starting period *n* with *x*

units of stock before ordering and sufficient statistics a and S, and letting $f_n(x|a) = f_n(x|a,1)$ be the solution of the *normalized* problem with $S = 1$, we obtain the following result.

> *Theorem 1. Assume that the manufacturer uses an admissible schedule, that the underlying demand distribution is Weibull with known k and unknown ω, and that the prior on ω is gamma with parameters a and S such that $ak > 1$. Then letting $q(S) = S^{\frac{1}{k}}$, the following hold:*
>
> *(a) The optimal values of the original problem can be found by scaling the normalized values: $f_n(x|a, S) = q(S)f_n\left(\frac{x}{q(S)}|a\right)$.*
>
> *(b) The optimal level, $y_n(x, a, S)$, of inventory after ordering in period n for the original problem can be found by scaling the optimal level, $y_n\left(\frac{x}{q(S)}, a\right)$, of inventory after ordering in period n for the normalized problem: $y_n(x, a, S) = q(S)y_n\left(\frac{x}{q(S)}, a\right)$*

Henceforth, we assume that the conditions of the theorem hold unless otherwise stated. The ordering level for the original problem depends on the critical numbers of the normalized one. It is straightforward to show that the same fraction of demand is met in either system. The optimal level of service depends only on the cost structure and the state of information as measured by a. The amount of stock needed to provide that service depends on the size or scale of the market as measured by $q(S)$. Whether the retailer expects a profit or loss similarly depends on the cost structure, normalized inventory, and the state of information; the magnitude of the profit or loss depends on the scale of the market. The function $q(S)$ is a natural measure of market size and is used to normalize retailer inventory (*i.e.,* $\frac{x}{q(s)}$) in applying an admissible price schedule. Intuitively, a given inventory level may be significant in a small market but trivial in a large one.

2. THE MANUFACTURER'S PRICING PROBLEM

We now turn to the determination of the optimal price schedule. We tackle the problem in stages, solving a dynamic program by first limiting the policies that may be used and then arguing that the optimal policy must lie in the restricted set. We begin by considering the

manufacturer's profits from using an admissible schedule as defined above. As these have a simple form, determining the optimal admissible schedule is straightforward. We then show that the optimal admissible schedule is the only linear schedule to which the manufacturer may credibly commit.

Channel Structure

Before analyzing the pricing problem, we must formalize the manufacturer-retailer relationship. The channel members are informationally equivalent; both know the family of the demand distribution and have the same prior on the unknown parameter ω, and both observe retail sales. Admittedly, these are dramatic simplifications. There is no compelling reason to believe that the parties would have an identical assessment of a new product's prospects even if they had access to the same information, and a manufacturer can rarely peer directly into a retailer's business. However, the assumptions force updated estimates of the unknown parameter to coincide. Everyone thus has identical values of a_n and S_n and derives the same updated demand distribution. Further, assuming that retail sales are visible to the manufacturer prevents the retailer from manipulating information to gain a more favorable price.

The firms are differentiated by their respective decision rights. The retailer controls his shelf space, ordering to maximize his own profits given the posted price. He refuses the product if stocking it would be unprofitable. The manufacturer designs and offers the terms of trade. Moving first, she anticipates the stocking policy and acts as a Stackelberg leader in setting the price. The "optimal" wholesale price is thus the one that is best for the manufacturer and not the contract that maximizes overall channel profits as considered by some (*e.g.*, Jeuland and Shugan, 1983). The terms of trade are presented as a take-it-or-leave-it offer with no opportunity for counteroffers or bargaining. The retailer accepts any deal offering a non-negative return.

Expected Manufacturer Profits from an Admissible Schedule

The manufacturer has a constant marginal cost of production c and discounts future profits by a factor β per period. For now, assume the retailer never drops the product. The manufacturer follows an arbitrary admissible schedule W, and the current period's wholesale

price is denoted simply by w. The retailer correctly anticipates the use of W and orders optimally. For simplicity, we suppress the order quantity's dependence on any parameter and write y^* for $y_n(x,a,S)$ and γ^* for $y_n(x,a)$.

Let $\pi_n(x|a, S, W)$ be the manufacturer's expected profits at the start of period n from following the price schedule W for the remainder of the horizon given current shape and scale parameters and retailer inventory. The only uncertainty the manufacturer faces is the evolution of the system. Her current period profits, unlike the retailer's, are deterministic; she knows exactly what the retailer will order at any wholesale price. Expected profits are then:

$$\pi_n(x|a, S, W) = (w - c)(y^* - x) + \beta \int_0^{y^*} \pi_{n+1}(y^* - \xi|a + 1, S + \xi^k, W)\phi(\xi|a, S)d\xi$$

$$+ \beta \int_{y^*}^{\infty} \pi_{n+1}(0|a, S + y^{*k}, W)\phi(\xi|a, S)d\xi$$

where $\pi_{N+1}(x|a, S, W) \equiv 0$ for all x, a, S, and W. As with the retailer's problem, we define the normalized case as having $S = 1$ and write $\pi_n(x|a, W)$ for $\pi_n(x|a, 1, W)$. Similarly, let $\phi(\xi|a) = \phi(\xi|a, 1)$ be the normalized updated demand distribution.

Theorem 2. Letting $U(t) = (1 + t^k)^{\frac{1}{k}}$, *the following hold:*

(a) $\pi_n(x|a, W) = (w - c)(\gamma^* - x) + \beta \int_0^{\gamma^*} U(t)\pi_{n+1}\left(\frac{\gamma^* - t}{U(t)}\middle|a + 1, W\right)\phi(t|a)dt$ (1)

$$+ \beta \int_{\gamma^*}^{\infty} U(\gamma^*)\pi_{n+1}(0|a, W)\phi(\xi|a)dt$$

(b) The expected value of the original problem can be found by scaling the normalized value: $\pi_n(x|a, S, W) = q(S)\pi_n\left(\frac{x}{q(S)}\middle|a, W\right)$.

The proof appears (as do all others) in the *Appendix*. From the first part of the theorem, expected manufacturer profits from following schedule W in the normalized case can be found from a simplified system in which future rewards are expressed in terms of normalized profits.

By the second part, the corresponding quantity for the original system is determined simply by the appropriate scaling of the normalized profits. As the conditions of *Theorem 1* hold, the results are not surprising for a single period problem; the retailer orders $y^* - x = q(S)\left(\gamma^* - \frac{x}{q(S)}\right)$. The manufacturer's sales, and hence her profits, are scaled by the normalizing function. The proposition establishes that future profits are similarly scaled when the manufacturer uses an admissible schedule.

The Optimal Admissible Price Schedule

Because the expression for normalized manufacturer profits (1) is independent of the scale parameter, a price schedule that maximizes expected normalized profits must be an admissible schedule. Determining the optimal admissible schedule is then relatively straightforward. The process is further simplified by noting that in both the original and normalized system, the manufacturer's current period profits depend on what she has charged earlier only through past retailer orders and past demand realizations. What has been ordered and what is known of past demand, however, are completely captured by the current sufficient statistics (a, S) and the retailer's inventory x. The optimal admissible schedule may consequently be developed recursively.

Let $\Pi_n(x|a, S)$ be the expected manufacturer's profits from using the optimal admissible price schedule for the original system from period n to the end of the horizon. For $n = 1, \ldots, N$, it must satisfy the optimality equations:

$$\Pi_n(x|a, S) = \max_{w \geq 0} \left\{ (w - c)(y^* - x) + \beta \int_0^{y^*} \Pi_{n+1}(y^* - \xi | a + 1, S + \xi^k) \phi(\xi | a, S) d\xi \right.$$

$$\left. + \beta \int_{y^*}^{\infty} \Pi_{n+1}(0 | a, S + y^{*k}) \phi(\xi | a, S) d\xi \right\}$$

where $\Pi_{N+1}(x|a, S) \equiv 0$ for all x, a, and S. Let $\Pi_n(x|a)$ be the corresponding quantity for the normalized system.

8

Corollary 1. Assume there is a unique optimal wholesale price for each period, information state, and inventory level triplet in the normalized case, and denote the optimal admissible schedule for that case as W^*.

(a) W^* may be found by solving the following optimality equations:

$$
\Pi_n(x|a) = \max_{w \geq 0} \left\{ (w - c)(\gamma^* - x) + \beta \int_0^{\gamma^*} U(t)\Pi_{n+1}\left(\tfrac{\gamma^*-t}{U(t)}|a + 1\right)\phi(t|a)dt \right.
$$

$$
\left. + \beta \int_{\gamma^*}^{\infty} U(\gamma^*)\Pi_{n+1}(0|a)\phi(t|a)dt \right\}
$$

(2)

for $n = 1, \ldots, N$, where $\Pi_{N+1}(x|a) \equiv 0$ for all x, and a.

(b) The optimal values of the original problem can be found by scaling the normalized values: $\Pi_n(x|a, S) = q(S)\Pi_n\left(\tfrac{x}{q(S)}|a\right)$.

(c) W^* is optimal for both systems.

The proof follows directly from *Theorem 2*. Since the manufacturer's profits from the full system are just a scalar multiple of her profits from the normalized system, the same admissible schedule is optimal for both. We henceforth assume that the conditions of the corollary hold and that a unique, optimal admissible schedule exist.

Gaining Retailer Participation

To this point, we have assumed that the retailer never drops the product although nothing in the formulation of (2) guarantees the retailer a satisfactory return. The optimal admissible schedule W^* may price the product out of the market, charging so much that the retailer refuses to stock it. Recalling *Theorem 1*, the problem is easily resolved.

Corollary 2. Solving (2) subject to gaining retailer participation yields an admissible price schedule.

The retailer's participation constraint in period n reduces to $f_n\left(\frac{x}{q(S)}\middle|a\right) \geq 0$. Thus, even if the constraint binds, the resulting wholesale price will depend only on the period, state of information, and normalized inventory. As the retailer can credibly threaten to drop any product that does not allow him to break even, the manufacturer will always choose to accommodate him and deviate from the unconstrained price schedule. The optimal admissible price schedule should consequently be understood to be the optimal schedule subject to gaining retailer participation.

Credible Price Schedules

We have so far shown that the manufacturer's profits from admissible schedules have a simple form and that the structure of her profits may be exploited to find the optimal admissible schedule. The analysis has been predicated on restricting the manufacturer to schedules driven by the current period, information state, and normalized retailer inventory. Prices cannot be functions of the sales level or the scale parameter although the manufacturer could clearly do no worse with richer contracts. However, she will be unable to implement more detailed schemes. She is limited by her inability to commit to future prices. The manufacturer is assumed bound only by the price offered for the current period, so she may post any price at the start of the next period regardless of earlier promises. Only price schedules that form subgame perfect equilibria (Selten, 1975) are then credible, and only the optimal admissible schedule meets that criterion.

> *Theorem 3. The optimal admissible schedule determined subject to retailer participation is the only linear schedule that forms a subgame perfect equilibrium.*

The result follows from the uniqueness of the optimal wholesale price. For each state of information and normalized inventory level, the manufacturer has a unique optimal price in the final period. Any prior claim to deviate from those prices is not credible. The manufacturer can consequently only commit to the optimal admissible schedule in the final period. Given that she will use the optimal admissible schedule in the last period, she has a unique optimal price for each state in the penultimate period, and a similar argument then iterates.

Table 1
A Simple Inadmissible Schedule

	The Optimal Admissible Schedule	An Inadmissible Schedule
Period 1	*4.29*	*4.29*
Period 2 ($s_1 < y_1$)	*4.98*	*4.98*
Period 2 ($s_1 = y_1$)	*4.59*	*4.59 if $s_1 \geq 14$* *9.50 if $s_1 = 14$*

y_1 is the retailer's Period 1 order, and s_1 observed Period 1 sales.
$p = 0$, $h = 0.1$, $r = 10$, $c = 2$, $a_1 = 1.1$, $S_1 = 10$

We illustrate the proposition by an example. *Table 1* presents the optimal admissible schedule and a simple inadmissible schedule for a two period horizon with exponential demand and perishable inventory. The inadmissible schedule forms a Nash equilibrium but fails to be subgame perfect. Under the optimal admissible schedule, the manufacturer charges *4.29* in the first period and either *4.59* (following a stockout) or *4.98* (following an exact observation) in the second. The retailer would order *12.7* units in the first period. The inadmissible schedule also charges *4.29* in the first period and *4.98* following an exact observation but only imposes *4.59* following stockouts at or above *14* units. After lower stockouts, *9.50* is charged. The intent is clear: To coerce a larger initial order. The manufacturer would sell more in the first period but charge her preferred prices in the second. The retailer's second period margin is imperiled unless he takes *14* in the first. An order of *14* is thus part of a Nash equilibrium. If the retailer believes the manufacturer will follow the announced policy, he orders *14*. Given that the retailer orders *14*, the manufacturer follows the posted schedule.

Although Nash, the equilibrium is not subgame perfect. It requires the retailer to believe that for some states of the world (*i.e.*, stockouts below *14*), the manufacturer will not charge her profit maximizing price. If the retailer were to call her bluff (say, by ordering *12.7*), the

manufacturer would renege on the second period prices and revert to the optimal admissible schedule. Unless the manufacturer is endowed with some commitment mechanism, the retailer should ignore any claim to follow another policy and expect the optimal admissible schedule in the final period. The logic carries over to longer horizons and more intricate pricing schemes. Unless the manufacturer constructs a richer contract limiting future pricing decisions (which is beyond the scope of our model), the only linear price schedule the retailer can rationally anticipate is the profit-maximizing admissible price schedule. Any other announced plan lacks credibility.

3. BEHAVIOR OF THE OPTIMAL PRICE SCHEDULE

We now characterize the dependence of the optimal price schedule on model parameters. We do this through a simple example that presumes true demand is exponential and that all inventory is perishable. With perishable inventory, the retailer's problem becomes a trivial series of newsvendor problems if either the demand distribution is known or Bayesian updating is done under fully observed demand. With Bayesian updating and lost sales, however, the retailer's problem remains challenging because the stocking level still determines the rate at which information is acquired. Lariviere and Porteus (1995) show that the optimal (normalized) retailer order for period n is:

$$ \gamma^* = \left(\frac{r + p + h + \alpha\big(\rho_{n+1}(a+1) - \rho_{n+1}(a)\big)}{w + h} \right)^{\frac{1}{a}} - 1 $$

where $\rho_n(a) = (a-1)f_n(0|a)$ and $\rho_{N+1}(a) \equiv 0$ for all a.

Perishable inventory also simplifies the manufacturer's problem. From the retailer's ordering policy, the manufacturer faces a deterministic demand curve. When the retailer carries inventory into the period, the manufacturer confronts a residual demand curve; the retailer buys only the difference between his order-up-to quantity and his stock on hand. It is immediate to show that the manufacturer's marginal revenue falls, which favors a lower wholesale price. In the perishable case, the manufacturer will never be constrained by retailer inventory. The perishable inventory analogue of (2) is:

$$\Pi_n(a) = \max_{w \geq 0} \left\{ (w-c)\gamma^* + \beta \int_0^{\gamma^*} U(t)\Pi_{n+1}(a+1)\phi(t|a)dt + \beta \int_{\gamma^*}^{\infty} U(\gamma^*)\Pi_{n+1}(a)\phi(t|a)dt \right\}$$

for $n = 1,\dots,N$, where $\Pi_{N+1}(a_{N+1}) \equiv 0$ for all a. The optimal wholesale price for period n and state of information a must satisfy:

$$\left[1 - \frac{w-c}{a(w+h)} - \frac{\beta\big(\sigma_{n+1}(a+1) - \sigma_{n+1}(a)\big)}{a\delta(a)} \right]^a = \frac{w+h}{\delta(a)} \tag{3}$$

where $\sigma_n(a) = (a-1)\Pi_n(a)$ and $\delta(a) = r + p + h + \alpha\big(\rho_{n+1}(a+1) - \rho_{n+1}(a)\big)$.

The Informational Dynamics of Pricing in the Single Period Problem

Clearly, there is generally no explicit solution for (3). However, the corresponding condition for the single period case is substantially simpler:

$$\left(1 - \frac{w-c}{a(w+h)} \right)^a = \frac{w+h}{r+p+h} \tag{4}$$

and allows us to examine how the optimal price for the one period horizon changes with various parameters.

> *Theorem 4. Consider a one period problem, and assume the retailer's participation constraint does not bind.*
>
> *(a) The optimal wholesale price w is increasing in the retail price r, the shortage penalty p, and the marginal cost of production c.*
>
> *(b) $\frac{\partial w}{\partial r} = \frac{\partial w}{\partial p} < 1$, and $\frac{\partial w}{\partial c} < 1$ if $a > \frac{w-c}{c+h}$.*
>
> *(c) The wholesale price also increases with the state of information a.*

The first part of the theorem is hardly surprising. It may seem odd to take the wholesale price as a function of the retail price as conventionally one has the reverse. In our case, the retail price is fixed, making it a piece of data for the manufacturer's problem. To the manufacturer,

increases in the retail price and the stockout penalty are equivalent because they have identical impacts on retailer behavior (despite markedly different impacts on retailer profits as discussed in Lariviere and Porteus, 1995). The increase in w following an increase in either is less than one-for-one; gains from a higher retail price are split between the parties. Similarly, higher production costs are passed on to the retailer. If the state of information is sufficiently high, the increase in the wholesale price is less than the increase in marginal cost. Otherwise, the wholesale price will increase more than the unit cost.

The more interesting result is the final part of the theorem. The manufacturer charges more as the precision of channel information improves. Consequently, the manufacturer posts a higher price if there were an exact observation in the prior period than if not. Stated differently, the manufacturer charges more following a signal of weak demand (unsold stock) than following a signal of strong demand (empty shelves). The outcome contradicts standard intuition; one would generally expect a higher price in a larger market. Here, the reverse holds, and the manufacturer demands a higher price in a smaller market.

The Elasticity of Retailer Orders and the State of Information

We identify two factors that underlie this apparent anomaly. First, by *Corollary 1*, the optimal price is independent of the market scale. What exact and censored observations portend for potential market size is immaterial. The type of observation only matters to the extent that it affects the state of information. Second, the retailer's ordering policy shifts systematically with the state of information. We now show that his sensitivity to price variation depends on the shape parameter and that he becomes increasingly less sensitive to price changes as the channel has more information. Consider the inverse demand curve the manufacturer faces for an arbitrary demand distribution $F(y)$:

$$w(y) = (p + r + h)(1 - F(y)) - h$$

We now substitute $\hat{w} = w + h$ and $\hat{r} = r + h$. The change represents an alternative costing convention. The retailer now charges himself for holding the product at its acquisition but rebates the charge when selling a unit. As we have implicitly assumed that costs are

14

denominated in start-of-the-period units, the analysis is unaffected by the accounting change. With the new convention, the inverse demand curve is $\hat{w}(y) = (p + \hat{r})(1 - F(y))$. Letting $f(y)$ be the density of the demand distribution, the corresponding own-price elasticity is:

$$\hat{v}(y) = \frac{1}{y}\left(\frac{1 - F(y)}{f(y)}\right) \tag{5}$$

$\hat{v}(y)$ is the elasticity of demand the retailer places on the manufacturer, not the elasticity of consumer demand faced by the retailer. We have not explicitly modeled the consumer purchase decision and have assumed a constant retail price. On the other hand, by modeling the retailer's ordering policy, we have defined an induced demand curve (and the equivalent inverse demand curve) that the manufacturer faces. $\hat{v}(y)$ is the elasticity of retailer orders, measuring the percentage change in the quantity the retailer demands for a percentage change in the wholesale price. It is inversely related to the hazard rate of demand whenever the retailer faces a single period newsvendor problem.

We defined $\hat{v}(y)$ for any distribution. If $F(y)$ is now the appropriate updated demand distribution, we can relate the elasticity of retailer orders to the state of information.

Theorem 5. If $F(y)$ is the updated demand distribution for a newsboy distribution with a gamma prior, own-price elasticity $\hat{v}(y)$ decreases with the state of information if:

a) the horizon is one period and the true demand distribution is any newsboy distribution, or

b) the horizon is N and the true demand distribution is exponential.

Demand as perceived by the manufacturer becomes less elastic as the shape parameter increases. When information is imprecise, the increase in volume after a price cut is large, and the manufacturer favors a lower price. When the channel has more refined information, the drop in sales following a price hike is less severe. The manufacturer exploits increased certainty by raising her price. Indeed, if demand were deterministic, the manufacturer would charge arbitrarily close to the retail price and capture all channel profits, leaving the retailer with the

15

smallest acceptable margin possible. She would succeed because retailer orders would be completely inelastic for any wholesale price below the retail price. At the other extreme, the retailer is very uncertain and hence extremely price sensitive. The manufacturer responds with a relatively low levy that increases as information becomes more precise.

The deterministic example highlights the result's dependence on the assumed power structure. In the deterministic case, the manufacturer can exploit price insensitivity because she presents the terms of trade as a take-it-or-leave-it offer. We cannot characterize the outcome of bargaining over price. The same is true in our setting. The manufacturer's leadership position allows her to take advantage of the growing inelasticity of retailer orders through a higher price.

4. PRICES OVER TIME AND NON-PERISHABLE INVENTORY

Our analysis of the optimal schedule has so far been limited to changes within a given period. We have examined how the optimal price varies with the state of information holding the remaining horizon fixed. Equally interesting is how the wholesale price changes with the remaining horizon holding the state of information fixed. We would like to know whether the manufacturer follows a penetration (*i.e.*, $w_n(a) \le w_{n+1}(a)$) or a skimming (*i.e.*, $w_n(a) > w_{n+1}(a)$) strategy. Also, our characterization to date has been limited to perishable inventory. It is unclear how the results of the previous section would change if retailer inventory was not perishable.

We address these points through an example for which we have an explicit solution for the one period pricing problem. If inventory is perishable, there is an explicit solution for longer horizons. We continue to assume that the true demand distribution is exponential and work with the "shifted" costing convention introduced above: $\hat{w} = w + h$ and $\hat{r} = r + h$. Additionally, we have $\hat{c} = c + h$ to make manufacturer's costs comparable. Equation (4) becomes:

$$\left(1 - \frac{\hat{w} - \hat{c}}{a\,\hat{w}}\right)^a = \frac{\hat{w}}{\hat{r} + p} \tag{6}$$

We now assume that $\hat{c} = 0$. Unsold product is salvaged at its marginal cost of production. The product would be riskless for an integrated firm that owned both the manufacturer and retailer. The integrated channel would stock an infinite amount to serve *100%*

16

of demand. While clearly unrealistic, the example offers insights because the story is not as simple in a decentralized channel. A profit maximizing manufacturer prices above her marginal cost, so the retailer will not meet all of demand. The outcome is a variation of the standard double marginalization argument (Spengler, 1950). A wholesale price above marginal cost does not affect the retail price but does induce the retailer to hold less stock than the integrated firm.

Prices over Time

Clearly, if $\hat{c} = 0$, the optimal wholesale price for the one period problem with perishable inventory is $\hat{w}(a) = (\hat{r} + p)\left(1 - \frac{1}{a}\right)^a$, which induces an order of $\gamma^* = \frac{1}{a-1}$; the retailer stocks to meet average demand. It is not coincidental that $\hat{v}\left(\frac{1}{a-1}\right) = 1$. As production is costless, the manufacturer maximizes revenue in the last period, and this is accomplished where elasticity is one. The corresponding manufacturer and retailer profits for a one period horizon are:

$$\sigma(a) = (\hat{r} + p)\left(1 - \frac{1}{a}\right)^a \qquad \rho(a) = \hat{r} - (\hat{r} + p)\left(1 - \frac{1}{a}\right)^a\left(1 + \frac{a}{a-1}\right)$$

The ratio of decentralized to centralized profits is $1 - \frac{\hat{r}+p}{\hat{r}}\left(1 - \frac{1}{a}\right)^{a-1}$, and the decentralized channel under-performs a centralized one. The relative inefficiency increases with the stockout penalty. The integrated firm never incurs any shortage penalty because it satisfies all demand. Both manufacturer and retailer profits increase with the state of information. This goes beyond the results of Lariviere and Porteus (1995). They show that retailer profits are increasing in the state of information for a fixed wholesale price. Here, retailer profits are increasing in *a* even as the manufacturer optimally adjust the wholesale price upwards.

To consider a two period problem, we first define:

$$\overline{\sigma}_2(a) = \frac{\sigma_2(a+1) - \sigma_2(a)}{\hat{r} + p} \qquad \overline{\rho}_2(a) = \frac{\rho_2(a+1) - \rho_2(a)}{\hat{r} + p}$$

Lariviere (1995) shows that neither $\overline{\sigma}_2(a)$ nor $\overline{\rho}_2(a)$ depends on the cost parameters $\hat{r}$ and p and that this result can be extended to an arbitrary horizon. From (3), the first period wholesale price is:

$$\hat{w}_1(a) = (\hat{r}+p)\left(1+\alpha\overline{\rho}_2(a)\right)\left[1-\frac{1}{a}\left(1+\frac{\beta\overline{\sigma}_2(a)}{1+\alpha\overline{\rho}_2(a)}\right)\right]^a$$

Prices are constant across the horizon only if both the manufacturer and retailer totally discount the future (*i.e.*, $\alpha = \beta = 0$). If the manufacturer values the future at all (*i.e.*, $\beta > 0$), she prices below the level that maximizes her current profits. She sacrifices current earnings to induce a higher stocking level, increasing the likelihood of reaching a more profitable future state. Conversely, if the retailer values the future (*i.e.*, $\alpha > 0$), the manufacturer charges more. Lariviere and Porteus (1995) show that a forward-looking retailer holds excess stock to increase the rate of information acquisition. The retailer buys more at every wholesale price, expanding the manufacturer's market. Part of that growth is taken through a higher wholesale price.

As the first period price depends on both the manufacturer and retailer discount factors, the price path over the horizon does as well. From manipulating the prices, we have:

Theorem 6. In a two period problem with $\hat{c} = 0$, the manufacturer uses penetration pricing (i.e., $\hat{w}_2(a) \geq \hat{w}_1(a)$) if :

$$\beta \geq \beta^* = \frac{(a-1)\left(1+\alpha\overline{\rho}_2(a)\right)\left(1-\left(1+\alpha\overline{\rho}_2(a)\right)^{\frac{-1}{a}}\right)}{\overline{\sigma}_2(a)}$$

The critical value β^* is increasing in the retailer's discount factor α. The first period price $\hat{w}_1(a)$ increases with α while $\hat{w}_2(a)$ is unchanged. Consequently, the more the retailer looks to the future, the greater the range of β for which the manufacturer uses a skimming strategy.

The critical value's dependence on the state of information is not as simple. From *Figure 1*, one sees that β^* is not monotonic in the shape parameter; the range of manufacturer discount factors in which a skimming strategy is used at first increases but then falls as more information is available. The graphed example has α equal to one and is consequently an upper bound on β^* because the critical value is increasing in α: If the manufacturer would set $\hat{w}_2(a) \geq \hat{w}_1(a)$ for α equal to one, she would do the same for any retailer discount factor. As the maximum

Figure 1
β^* *for* $\alpha = 1$
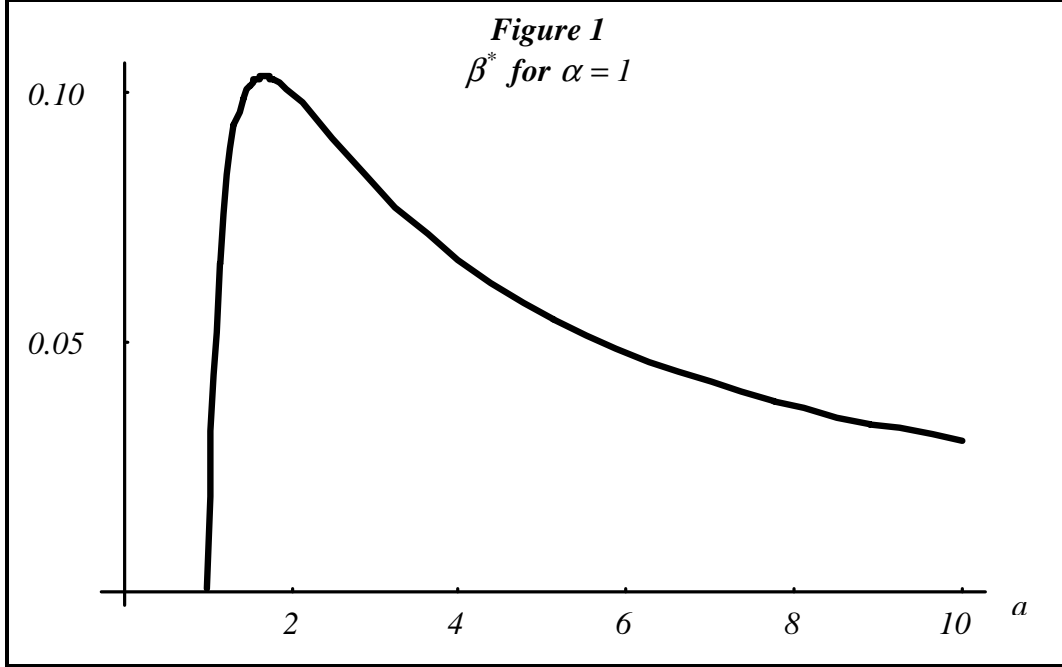
value on the upper bound for this special case is just greater than *0.10*, the manufacturer would follow a penetration strategy for all but extremely low manufacturer discount factors.

Non-Perishable Inventory

We now consider how retailer inventory influences the manufacturer's pricing decision. As discussed above, when the retailer carries inventory into the period, the manufacturer faces a residual demand curve making her marginal revenue lower at every wholesale price. She consequently charges less than she would if inventory were perishable. For a one period problem in which the retailer has normalized inventory x, the analogue of (6) is:

$$\left(\frac{1}{x+1}\left(1 - \frac{\hat{w} - \hat{c}}{a\hat{w}}\right)\right)^a = \frac{\hat{w}}{\hat{r} + p} \tag{7}$$

If the retailer's inventory is $x^*(a) = \left(\frac{\hat{r}+p}{\hat{c}}\right)^{\frac{1}{a}} - 1$, the stocking level of the integrated channel, (7) has a unique solution: $\hat{w} = \hat{c}$. If the retailer has inventory above $x^*(a)$, the manufacturer would have to price below cost to induce the retailer to order more stock. For retailer inventory below

19

the integrated stocking level, the manufacturer can always find a price above cost that would induce the retailer to take more stock.

The integrated channel's stocking level is decreasing in the state of information. One can therefore construct examples in which the optimal wholesale price falls as the shape parameter increases for a fixed level of inventory. Suppose the retailer's inventory is $x^*(a + 1)$. If the state of information is a, the manufacturer offers the product at some price above marginal cost. On the other hand, if the state of information increases by one, the price drops to $\hat{c}$, and the wholesale price would be decreasing in the state of information contrary to *Theorem 4*.

One does not have to go to such an extreme case for the wholesale price to be decreasing in the shape parameter. We return to our example with $\hat{c} = 0$. The solution to (7) is now:

$$\hat{w} = (\hat{r} + p)\left(\frac{a - 1}{a(x + 1)}\right)^a$$

As the integrated firm would hold an infinite amount of stock, the manufacturer in the decentralized channel can always induce the retailer to order more at a price above marginal cost. The price, however, is decreasing in the retailer inventory level. *Figure 2* graphs the optimal price as a function of the shape parameter for three inventory levels (*0, 0.2,* and *0.8*). One sees that the wholesale price is increasing for low states of information but eventually falls with the state of information when inventory is positive. The shape parameter at which the price begins to decline decreases with the inventory level. For this simple case, we are able to determine the inventory level below which the wholesale price increases with the state of information.

20

Figure 2
The Optimal Wholesale Price as a Function of
the Shape Parameter for Three Inventory Levels
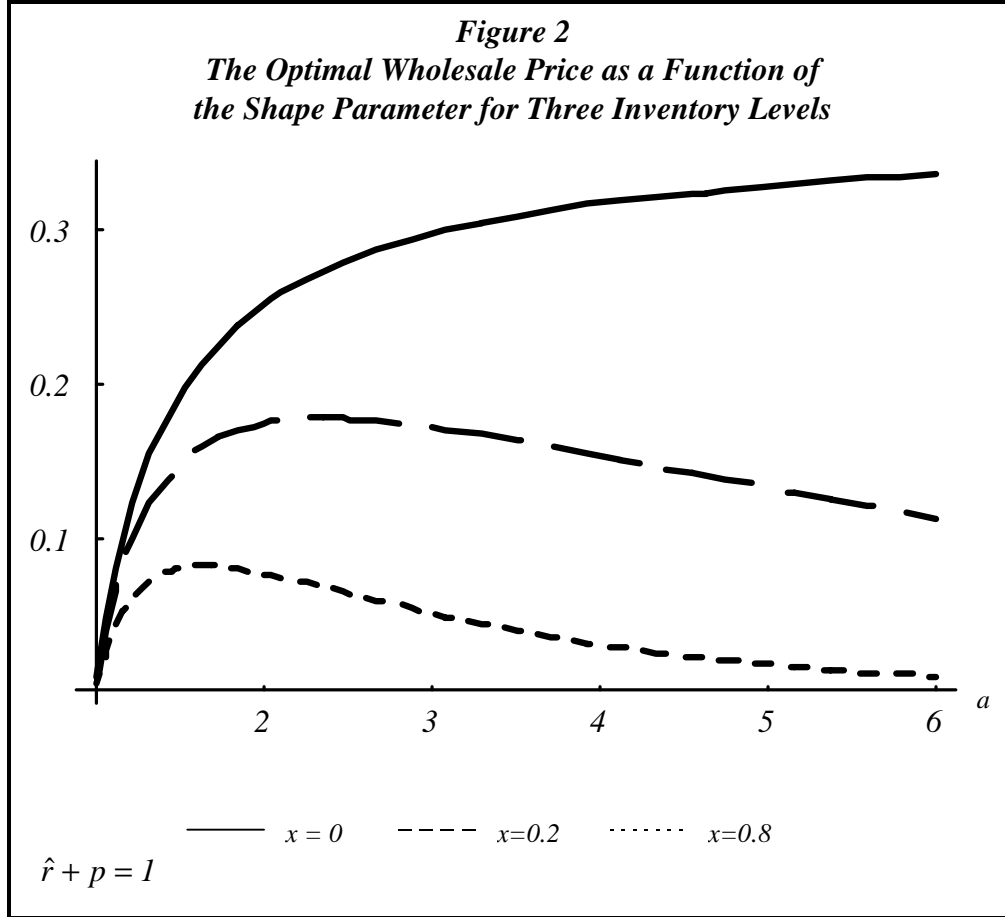
Theorem 7. If $\hat{c} = 0$, the wholesale price increases with the state of information if:

$$x \leq \frac{(a-1)}{a} e^{\frac{1}{a-1}} - 1.$$

Theorem 4 must then be qualified. If retailer inventory is sufficiently low, the optimal wholesale price will be increasing in the state of information. If retailer inventory is not sufficiently low, the price will be decreasing in the state of information. The caveat reflects two countervailing forces that affect the relationship between the wholesale price and the state of information. On the one hand, the retailer becomes less price sensitive as the state of information increases, and the manufacturer consequently wants to raise the wholesale price. On the other, the retailer's order-up-to level for a fixed wholesale price falls with the state of information; he can satisfy a greater fraction of his needs with stock already on hand. The manufacturer compensates by cutting the wholesale price. Which force prevails depends on the problem.

5. CONCLUSION

We have examined a manufacturer launching a new product into an existing retail channel. At the time of introduction, neither the manufacturer nor the retailer knows some parameter of the demand distribution. Further, unmet demand is both lost and unobserved; beliefs about the demand distribution must be updated based on sales observations. While the retailer's inventory thus dictates the rate at which the channel acquires information, the retailer sets the stocking level to maximize his own profits because he is independent of the manufacturer. Indeed, he may even drop the product if he cannot earn a satisfactory return. The manufacturer must then consider the retailer's stocking policy when setting the wholesale price for each period of the finite horizon. She acts as a Stackelberg leader, anticipating the retailer's ordering policy and consequently facing a deterministic, induced demand curve.

Building on Lariviere and Porteus (1995), we show that the manufacturer's optimal linear schedule subject to gaining retailer participation is driven by the precision of available information, normalized retailer inventory, and the period. It is not dependent on the size of the market or the exact history of sales. The optimal schedule can therefore be determined prior to launching the product. For the case of perishable inventory with exponential demand, we show that the optimal wholesale price is increasing in the state of information. As a consequence, a higher price is charged following a signal of weak demand (product left unsold) than following a signal of strong demand (a stockout). The result follows from the induced demand curve the manufacturer faces becoming increasingly inelastic as the precision of information improves. As he has better information regarding the unknown parameter, the retailer becomes less price sensitive. The manufacturer exploits his growing insensitivity by raising the wholesale price. We also show through a simple example that the manufacturer will follow a penetration strategy raising prices over the horizon if her discount rate is sufficiently high.

Our results are subject to limitations. We demonstrate, for example, that the wholesale price may not increase with the state of information if inventory is not perishable. If retailer inventory is large, the optimal wholesale price may fall with the state of information. In our example, the retailer's optimal stocking level (all else being equal) decreases with the state of information. The manufacturer's potential market shrinks putting downward pressure on the

wholesale price. Increasing retailer price insensitivity but decreasing overall demand pull the optimal wholesale price in opposite directions.

The optimal schedule's independence from the exact history of sales and the size of the market is due to our distributional assumptions and preclusion of longer term contracts. While we give an example in which the manufacturer could do better with an inadmissible linear schedule if she could commit to its use, analyzing the retailer stocking policy for a general inadmissible schedule would be hopelessly complex. Similarly, the assumption of Weibull demand is based on the desire to use a newsboy distribution. As Braden and Friemer (1991) note, a modeler dealing with censored data but unsatisfied with any of the distribution they suggest should expect "serious analytic and computational problems." (p. 1390) If we forgo assuming that unmet demand is unobserved, a greater range of families become analytically feasible. We conjecture that the optimal linear price schedule with backlogged sales would be independent of the scale of the market for the settings considered by Azoury (1985) because the retailer's ordering policy would be still given by a scaling of a normalized problem.

Finally, a fixed retail price is problematic. On the one hand, the retailer may order in a self-interested manner, but on the other, he has cannot adjust the retail price. Again, tractability necessitates the assumption although it is more defensible than, say, the assumption of Weibull demand. We have implicitly assumed that the channel members have some experience in the product's category on which to base their priors. The product is not completely new and must fit into an existing market. The retail price is then set to target a particular segment. The channel, however, is uncertain of the size of the segment.

Proof of Theorem 2

We begin with a lemma that follows from simple algebraic manipulations:

Lemma A1. Given $q(S) = S^{\frac{1}{k}}$, $U(t) = (1 + t^k)^{\frac{1}{k}}$, and $\phi(\xi|a) = \phi(\xi|a,1)$, the following hold:

(a) $\phi(\xi|a, S) = \dfrac{1}{q(S)} \phi\left(\frac{\xi}{q(S)}\middle|a\right)$.

(b) $q(S + \xi^k) = q(S) U\left(\frac{\xi}{q(S)}\right)$.

We prove the theorem by a backward induction on n. It clearly holds for $i = N$. We suppose it is true for $i = n + 1$ and verify that it holds for $i = n$ using the inductive hypothesis, *Lemma A1*, and a change of variables $t = \frac{\xi}{q(S)}$.

$$
\pi_n(x|a, S, W) = (w - c)(y^* - x) + \beta \int_0^{y^*} q(S + \xi^k) \pi_{n+1}\left(\frac{y^* - \xi}{q(S + \xi^k)}\middle|a + 1, W\right) \phi(\xi|a, S) d\xi
$$

$$
+ \beta \int_{y^*}^{\infty} q(S + y^{*k}) \pi_{n+1}(0|a, W) \phi(\xi|a, S) d\xi
$$

$$
= (w - c)(y^* - x) + \beta \int_0^{y^*} U\left(\frac{\xi}{q(S)}\right) \pi_{n+1}\left(\frac{y^* - \xi}{U\left(\frac{\xi}{q(S)}\right) q(S)}\middle|a + 1, W\right) \phi\left(\frac{\xi}{q(S)}\middle|a\right) d\xi
$$

$$
+ \beta \int_{y^*}^{\infty} U\left(\frac{y^*}{q(S)}\right) \pi_{n+1}(0|a, W) \phi\left(\frac{\xi}{q(S)}\middle|a\right) d\xi
$$

$$
= q(S)\Bigg[(w - c)\left(\gamma^* - \frac{x}{q(S)}\right) + \beta \int_0^{\gamma^*} U(t) \pi_{n+1}\left(\frac{\gamma^* - t}{U(t)}\middle|a + 1, W\right) \phi(t|a) dt
$$

$$
+ \beta \int_{\gamma^*}^{\infty} U(\gamma^*) \pi_{n+1}(0|a, W) \phi(t|a) dt \Bigg]
$$

$$
= q(S) \pi_n\left(\frac{x}{q(S)}\middle|a, W\right)
$$

Proof of Theorem 3

Define the state of the system in period *n* as the prevailing shape parameter and normalized inventory level. We make the following observation:

> *Lemma A2. It is optimal to use the optimal admissible price schedule in the current period for the prevailing state if the manufacturer plans on using the optimal admissible price schedule in all possible future periods and states.*

The lemma follows immediately from *Corollary 1* and the assumption of a unique profit maximizing price: Given that the optimal schedule will always be used in the future, there is a unique, optimal price for the current period and state. Because the optimal schedule is determined recursively, that unique price must be part of the optimal schedule.

Turning to the theorem itself, suppose there were a linear pricing schedule $\tilde{W}$ that formed a subgame perfect equilibrium but deviated from the optimal admissible price schedule. As we are considering inadmissible schedules, the state now includes all available information (*e.g.* the exact history of sales or retail orders) in addition to the shape parameter and normalized inventory level. The assumed order of play has the manufacturer posting a price at the start of each period that is binding for that period. A schedule for future states may be announced but is not binding. Each period-state pair thus forms a subgame.

Let $\tilde{n}$ be the latest period for which $\tilde{W}$ deviates from the optimal admissible schedule for some state. We have two possibilities, $\tilde{n} = N$ or $\tilde{n} < N$. For the former, the manufacturer has a unique profit maximizing price that depends on past observations only through the shape parameter and is part of the admissible schedule. The proposed schedule, however, differs from the optimal admissible schedule in period *N* for some state. Thus, for at least one subgame, the manufacturer would have an incentive to unilaterally deviate from $\tilde{W}$. $\tilde{W}$ cannot be part of a subgame equilibrium. The second case is similar. By *Lemma A2*, there is a unique optimal wholesale price for each subgame that begins in period $\tilde{n}$ that only depends on the shape parameter and normalized inventory level, but $\tilde{W}$ calls for a different price in at least one such subgame. Hence, $\tilde{W}$ cannot be subgame perfect.

Proof of Theorem 4

We first show that for a one period problem with exponential demand, the manufacturer's profits are concave in the wholesale price and thus any w solving (4) in the text is thus the unique, optimal wholesale price.

Lemma A3. Let $\pi(w|a) = (w - c)\gamma^*$. *$\pi(w|a)$ is concave in w.*

Define $T = \frac{r + p + h}{w + h}$. Then $\gamma^* = T^{\frac{1}{a}} - 1$ and:

$$\frac{\partial \pi(w|a)}{\partial w} = T^{\frac{1}{a}}\left(1 - \frac{w - c}{a(w + h)}\right) - 1$$

$$\frac{\partial^2 \pi(w|a)}{\partial w^2} = -T^{\frac{1}{a}}\left(\frac{w(a - 1) + c(a + 1) + 2ah}{a^2(w + h)^2}\right)$$

Recalling from *Theorem 1* that $ak > 1$ completes the proof. (4) follows setting $\frac{\partial \pi(w|a)}{\partial w} = 0$. To prove the theorem, we apply the implicit function theorem to the first order conditions to yield:

$$\frac{\partial w}{\partial r} = \frac{\partial w}{\partial p} = \frac{T^{-1}\left((a - 1)w + ah + c\right)}{\left((a - 1)w + (a + 1)c + 2ah\right)}$$

$$\frac{\partial w}{\partial c} = \frac{a(h + w)}{\left((a - 1)w + (a + 1)c + 2ah\right)}$$

$$\frac{\partial w}{\partial a} = \frac{(h + w)\left(a(w - c) - Log(T)\left(a(w + h) - (w - c)\right)\right)}{a\left((a - 1)w + (a + 1)c + 2ah\right)}$$

As $a > 1$, the first two expressions are clearly positive, proving the first part of the proposition. The second part follows by noting that $T > 1$. The final part will hold if:

$$\frac{w - c}{a(w + h) - (w - c)} > \frac{1}{a}Log(T)$$

which from the first order conditions is equivalent to $T^{\frac{1}{a}} > Log\left(T^{\frac{1}{a}}\right)$ and thus always true.

Proof of Theorem 5

For (a), the updated demand distribution for a newsboy distribution with a gamma prior is

$F(y) = 1 - \left(\frac{S}{S+d(y)}\right)^a$. Denoting the first derivative of $d(y)$ as $d'(y)$, (5) becomes:

$$\hat{v}(y) = \left(\frac{S+d(y)}{ay\, d'(y)}\right)$$

which is decreasing in the state of information a. For exponential demand, $F(y) = 1 - \left(\frac{S}{S+y}\right)^a$ and

$$y^* = S\left(\left(\frac{r+p+h+\alpha(\rho_{n+1}(a+1)-\rho_{n+1}(a))}{w+h}\right)^{\frac{1}{a}} - 1\right).$$ The inverse demand curve for the manufacturer in period

n is then:

$$\hat{w}_n(y) = \left(\hat{r} + p + \alpha\big(\rho_{n+1}(a+1) - \rho_{n+1}(a)\big)\right)\left(1 - F(y)\right)$$

Its elasticity reduces to $\hat{v}(y) = \frac{S+y}{ay}$, which is again decreasing in a.

Proof of Theorem 7

Differentiating $\hat{w}$ with respect to a, one has:

$$\frac{\partial \hat{w}}{\partial a} = (\hat{r} + p)\left(\frac{a-1}{a(x+1)}\right)^a \left(Log\left(\frac{a-1}{a(x+1)}\right) + \frac{1}{a-1}\right)$$

which is positive if $Log\left(\frac{a-1}{a(x+1)}\right) \geq \frac{-1}{a-1}$ or, equivalently, $x \leq \frac{a-1}{a}e^{\frac{1}{a-1}} - 1$.

27

REFERENCES

Azoury, Katy S., "Bayes Solution to Dynamic Inventory Models under Unknown Demand Distribution," *Management Science*, 31, 9 (September, 1985), 1150-1160.

Braden, David J. and Marshall Freimer, "Informational Dynamics of Censored Observations," *Management Science*, 37, 11 (November, 1991), 1390-1404.

DeGroot, Morris H., *Optimal Statistical Decisions*, McGraw-Hill, New York, New York, 1970.

Jeuland, Abel P. and Steven M. Shugan, "Managing Channel Profits," *Marketing Science*, 2,3 (Summer, 1983), 239-272.

Lariviere, Martin A., *Optimal Manufacturer-Retailer Contracting under an Unknown Demand Distribution*, Unpublished Doctoral Dissertation, Stanford University, 1995.

Lariviere, Martin A. and Evan L. Porteus, "Stalking Information: Bayesian Inventory Management with Unobserved Lost Sales, working paper, Duke University.

Lee, Hau L. and Meir J. Rosenblatt, "A Generalized Quantity Discount Pricing Model to Increase Supplier's Profits," *Management Science*, 32, 9 (September, 1986), 1177-1185.

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

Monahan, James P., "A Quantity Discount Profit Model to Increase Vendor Profits," *Management Science*, 30, 6 (June, 1984), 720-726.

Moorthy, K. Sridhar, "Managing Channel Profits: Comment," *Marketing Science*, 6, 4 (Fall, 1987), 375-379.

Pasternack, Barry Alan, "Optimal Pricing and Return Policies for Perishable Commodities," *Marketing Science*, 4, 2 (Spring, 1985), 166-176.

Selten, R. "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4 (1975), 22-55.

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

Tirole, Jean, *The Theory of Industrial Organization*, The MIT Press, Cambridge, Massachusetts, 1988.

Weng, Z. Kevin, "Channel Coordination and Quantity Discounts," *Management Science*, 41, 9 (September, 1995), 1509-1522.