Vertical Control of Price and Inventory

By Harish Krishnan and Ralph A. Winter*

This paper offers a simple approach to the theory of decentralizing inventory and pricing decisions along a supply chain. We consider an upstream manufacturer selling to two outlets, which compete as differentiated duopolists and face uncertain demand. Demand spillovers between the outlets arise in the event of stockouts. The price mechanism, in which each outlet pays a two-part price and chooses price and inventory, virtually never coordinates incentives efficiently. Contracts that can elicit first-best decisions include resale price floors or buy-back policies (retailer-held options to sell inventory back to the manufacturers). (*JEL* D21, L13, L14, M11)

The choice of inventory for a firm facing uncertain demand is a classic problem of economics (Kenneth J. Arrow, Theodore Harris, and Jacob Marshcak 1951) and a central problem of management science. A recent literature recognizes that the optimal inventory is not just a single-agent decision problem, but rather involves the alignment of incentives all along a supply chain. Moreover, as Raymond Deneckere, Howard P. Marvel, and James Peck (1996, 1997) have shown, the vertical control of inventory decisions is tightly linked to control of pricing.

This paper introduces a framework that synthesizes and extends the theory of decentralizing inventory and pricing decisions. Can a manufacturer set a wholesale price for its product, relying on the distributors of its product to set optimal pricing and inventory levels? We isolate the sources of incentive incompatibility in simple price-mediated exchange and then characterize the contracts that do elicit the right incentives. Our framework is based on two principles. First, an organization faces an incentive problem when an agent within the organization does not appropriate the full collective benefits of actions taken. Second, the incentive problem is resolved when some of the agent's actions are constrained at the optimal levels, and prices or reward systems internalize the externalities imposed by the remaining actions on other agents within the organization.

We apply this framework to shed light on a puzzling set of cases. When vertical price floors (resale price maintenance) were struck down as illegal in the early 1970s, manufacturers were suddenly constrained in their design of distribution systems.[1] Retailer inventories collapsed for some products and the distribution of the products suffered. We aim to understand why the price system itself is unable to convey the right incentives for pricing and inventory decisions—and why the required incentive contracts take the form of vertical price restraints.

The framework can be applied to contracts beyond vertical restraints in multi-task, multi-agent vertical supply chains. We take buy-back policies as an example. A buy-back policy is the right to sell unused inventory back to the manufacturer or an agreement to give credit for any

* Krishnan: Sauder School of Business, University of British Columbia, Vancouver, BC V6T 1Z2 (e-mail: harish.krishnan@sauder.ubc.ca); Winter: Sauder School of Business, University of British Columbia, Vancouver BC V6T 1Z2 (e-mail: ralph.winter@sauder.ubc.ca). We gratefully acknowledge support from the Social Sciences and Humanities Research Council and the Natural Sciences and Engineering Research Council. We thank Xuan Zhao for very helpful discussions on this topic and the participants of numerous seminars for insightful comments. We are especially grateful to the editor and three referees for their detailed and constructive comments.

[1] Until *Leegin Creative Leather Products, Inc. v. PSKS, Inc.*, S. Ct. No. 06-480 in June 2007, resale price maintenance (RPM) had been per se illegal under *Dr. Miles Medical Co. v. John D. Park and Sons,* 220 US 373 (1911). The exception had been a period of "fair trade" state laws that had allowed the practice. When these laws were repealed in the 1970s, resale price maintenance once again became illegal. (Even then, however, a manufacturer could unilaterally adopt a plan to establish suggested resale prices in advance and lawfully terminate retailers who fail to adhere to those prices (*United States v. Colgate & Co.*, 250 US 300, 39 Sup. Ct. 465, 7 A.L.R. 443).)

unsold inventory. Bookstores, for example, send the covers of paperback books and magazines to publishers to obtain credit.

This paper contributes to three related literatures. The economics literature on vertical restraints (e.g., G. Frank Mathewson and Ralph A. Winter 1984; Michael L. Katz 1989; Deneckere, Marvel, and Peck 1996, 1997; David A. Butz 1997; James D. Dana, Jr., and Kathryn E. Spier 2001) has mainly a positive motivation in explaining observed contracts. The management science and supply chain literature has a prescriptive motivation in analyzing the optimal means of managing inventory in a multi-stage distribution system. Gerard P. Cachon's (2003) survey of supply chain coordination cites more than 150 articles and the area has been active since the survey was written. An antitrust literature addresses the normative issue of whether particular vertical restraints should be allowed (Richard A. Posner 1981; Frank H. Easterbrook 1984).[2]

We begin by outlining the classic newsvendor inventory model, as well as the general theory of vertical restraints that we will apply to the inventory problem. Section II develops the theory of pricing and inventory decisions in the simplest cases of a purely vertical supply chain (only one retailer) and competing retailers without demand spillovers from a stocked-out retailer. This framework serves as well to connect our approach to the literature. Section III presents the full model. Section IV applies the theory to a set of cases.

[2] The research papers that lie closest to our model are Deneckere, Marvel, and Peck (1997) and Fernando Bernstein and Awi Federgruen (2005). Bernstein and Federgruen's model is highly complex and assumes away the possibility of demand spillovers from a stocked-out retailer to other retailers. We suggest that this type of externality is central and can be incorporated in a simple framework. Deneckere, Marvel, and Peck offer a fundamental contribution in linking resale price maintenance to inventory management in a model with market power upstream and perfect competition downstream. Our model incorporates market power downstream in the form of a differentiated duopoly. This has the advantage of revealing the sources of coordination failure of simple price-mediated exchange and the role of contracts beyond price restraints in resolving the coordination failures.

I. Background

A. *The Optimal Inventory Problem for a Single Firm*

The classic newsvendor inventory problem considers a firm sourcing a good at a constant per unit cost, c, and facing an exogenous price, p. The firm must order an amount y of the good prior to the realization of demand, and any product not sold is worth nothing. Denoting the distribution of uncertain demand by G, the firm's expected profit is $p\int_0^y x\, dG(x) + py[1 - G(y)] - cy$. Maximizing this with respect to y yields the fractile solution $1 - G(y^*) = c/p$. Equivalently, $G(y^*) = (p - c)/p$. Note that if an upstream seller is providing the product to a downstream retailer at a uniform price w, then the retailer's first-order condition substitutes w for c, and the retailer therefore orders too little inventory. The standard double-marginalization or vertical externality distorts the retailer's decision: the retailer ignores the upstream margin $(w - c)$ that accrues to the upstream manufacturer with each additional unit of y ordered.

The simple newsvendor problem is concave and the first-order conditions are sufficient as well, as is necessary for the solution. This is not always true in extensions that (a) endogenize price (e.g., Nicholas C. Petruzzi and Maqbool Dada 1999); (b) involve multiple outlets with demand that spills over from one firm to the other in the event of a stockout (e.g., Serguei Netessine and Nils Rudi 2003); or (c) involve a vertical structure in which an upstream firm relies on inventory decisions by downstream firms (e.g., Martin A. Lariviere and Evan L. Porteus 2001). Since our model will incorporate all three of these features, we recognize the possibility of nonconcavities. In this article, however, we take a first-order approach to the incentive problem (i.e., we take agents' first-order conditions as sufficient, not just necessary, for the agent optimum), restricting consideration to payoffs that are quasi-concave.[3]

[3] We illustrate in a numerical example that there is a range of parameters where the quasi-concavity assumption is satisfied, and in footnote 10 we remark on the effects of relaxing the assumption.

B. *The Simple Analytics of Vertical Price Restraints*

The structure we will consider throughout is a manufacturer upstream producing a product at unit cost, c, and selling to two differentiated retailers downstream, who in turn sell to consumers. While our aim is to understand contracts that coordinate decisions on pricing and inventory, it is helpful to set out in advance a structure in which demand is certain and depends on price and another retailer action such as sales effort or service. This framework is developed in more detail in Ralph A. Winter (1993).

Assume that the demands for the product downstream at the two retailers are symmetric, with the demand at retailer 1 given by $q_1(p_1, s_1; p_2, s_2)$, where p_i and s_i are the price and sales effort or service—measured in units of the dollar cost of the effort—provided at the two retailers. Retailers bear costs given by the wholesale price, w, paid to the upstream manufacturer, a fixed fee paid to the manufacturer for the right to carry the product, and expenditure s_i on sales effort. The profit for retailer 1 gross of the fixed fee is denoted by π_1, and the total profit, for the manufacturers and both retailers, is denoted by Π. These profit functions are given by

$$(1) \quad \pi_1(p_1, s_1; p_2, s_2) = q_1(p_1, s_1; p_2, s_2)(p_1 - w) - s_1;$$

$$(2) \quad \Pi_1(p_1, s_1; p_2, s_2) = q_1(p_1, s_1; p_2, s_2)(p_1 - c) + q_2(p_1, s_1; p_2, s_2)(p_2 - c) - s_1 - s_2.$$

Assume that Π is maximized at a symmetric set of prices and effort levels, and denote this optimum by (p^*, s^*). Furthermore, for purposes of this background, assume that the profit functions are concave. Can the symmetric (p^*, s^*) be elicited with a single instrument, w, or are more complex contracts required? The key to understanding any incentive distortions in retailer decisions is to isolate and decompose the difference between the marginal gain in *individual* profit and the marginal gain in *total* profit from a change in either p_i or s_i. From equations (1) and (2), it follows that, at a symmetric configuration $(p_1 = p_2 \equiv p$ and $s_1 = s_2 \equiv s)$, this difference can be expressed for retailer 1 (and similarly for retailer 2) as follows:

$$(3) \quad \frac{\partial \pi_1}{\partial p_1} = \frac{\partial \Pi}{\partial p_1} - \overbrace{\frac{\partial q_1}{\partial p_1}(w - c)}^{\text{vertical externality}} - \overbrace{\frac{\partial q_2}{\partial p_1}(p - c)}^{\text{horizontal externality}};$$

$$(4) \quad \frac{\partial \pi_1}{\partial s_1} = \frac{\partial \Pi}{\partial s_1} - \overbrace{\frac{\partial q_1}{\partial s_1}(w - c)}^{\text{vertical externality}} - \overbrace{\frac{\partial q_2}{\partial s_1}(p - c)}^{\text{horizontal externality}}.$$

The individual retailer's private optimum in setting p_1 is distorted from the collective optimum by two externalities: when p_1 is raised, the manufacturer collects the wholesale markup, $(w - c)$, on a smaller demand through retailer 1. This is the *vertical* externality. The effect of this externality is to distort the retailer's price upward. The second externality operates through the cross-elasticity of demand between the two retailers. When p_1 is raised, the competing retailer collects the retail markup $(p - w)$ on $\partial q_2/\partial p_1$ additional units and the manufacturer collects the wholesale markup $(w - c)$ on the same additional units; these add up to the term labelled *horizontal* externality in equation (3). The effect of this externality is to distort the retailer's price downward. The same two externalities distort the sales effort decision. For each instrument, the vertical and horizontal externalities act to distort the decentralized decision in opposite directions.

When will the right choice of w alone elicit the optimum p^* and s^* at both outlets? That is, when does the price system elicit (privately) efficient incentives in this model of a distribution system? This efficiency property will hold when the value of $w \equiv \hat{w}$ that renders the sum of the last two terms of (3) to zero at the optimum $(p^*, s^*; p^*, s^*)$ also renders the sum of the last two terms of (4) to zero. The externalities must balance in both first-order conditions at the same value of w.[4] If they do not, then the price system

[4] The externality-balancing condition can be expressed as $\varepsilon_p^r/\varepsilon_s^r = \varepsilon_p^m/\varepsilon_s^m$, where ε_p^r and ε_s^r are retailer elasticities with respect to p and s, and the right-hand-side terms are market elasticities. This condition for the efficiency of the

fails. The elicitation of correct incentives in both p and s can be achieved, however, with a restraint on p. For example, if the sum of the externalities in (4) is negative at $\hat{w}$, then there is a bias toward too much price competition. A price *floor* at p^* achieves the first-best because it protects the retailer margin as w is lowered below $\hat{w}$ so as to elicit the optimal choice of s^*. If the sum of the externalities in (4) is positive at $\hat{w}$, then there is a bias toward too little price competition, and a price *ceiling* at p^* is necessary.

Of course, this framework does not explain *why* downstream outlets should be biased one direction or the other. Why, for example, were price floors observed much more often than price ceilings during periods when these restraints were both legal? The externality-balancing argument merely provides a structure, which we apply below to a setting in which the downstream decisions are price and inventory.

II. Price and Inventory Decisions in Simple Cases

We begin our analysis of the decentralization problem with the characterization of efficient contracts in two simple cases: a single retailer; and competing retailers without spillovers. The contracts that we consider include two-part tariffs, vertical price restraints, and buy-back policies. Throughout the paper we set aside the complete contract that specifies both downstream agent decisions, price and inventory, since we are interested in the feasibility of decentralizing these decisions.

A. *Purely Vertical Setting: One Retailer*

The simplest model of decentralized pricing and inventory decisions involves a single manufacturer selling to a single retailer. Suppose that a manufacturer produces a product at constant unit cost c and then sells to a retailer, who in turn sells to consumers. The final quantity demanded $q(p;\theta)$ depends upon retail price p and a random variable θ. The manufacturer sells the product to the retailer at a wholesale price w, possibly with a fixed fee F, or offers an alternative contract. The retailer accepts or rejects the contract and then must choose the level of inventory y and set the price p before knowing the realization of θ. Our model allows for the consideration of retailer fixed costs, and we comment on this after Proposition 1 below, but we omit fixed costs from the development of the model.

We denote the number of transactions by $t(p,y;\theta) = \min[q(p;\theta),y]$ and the number of "overstocks" by the retailer as $O(p,y;\theta) = [y - q(p;\theta)]^+$.[5] We let $s(p,y) = \Pr(q(p;\theta) > y) = \Pr(O(p,y;\theta) > 0)$ denote the stockout probability. The profit function for the retailer is $\pi(p,y;w,\theta) = pt(p,y;\theta) - wy - F$.[6]

The manufacturer and the retailer are both risk-neutral. The manufacturer offers a contract that maximizes profits subject to the participation constraint that the contract be acceptable to the retailer and the incentive compatibility constraint that the inventory and price decisions be in the retailer's best interest. We are interested in characterizing the set of contracts that can achieve the first-best expected profits for the manufacturer. A first-best contract maximizes total expected profits for the system (the manufacturer and the retailer) and leaves the retailer with zero expected profits, i.e., leaves the participation constraint binding. The total profits are given by $\Pi(p,y;\theta) = pt(p,y;\theta) - cy$. Assuming differentiability, maximizing $E\Pi(p,y;\theta)$ with respect to p and y yields the following first-order conditions:

$$(5) \qquad Et(p,y) + \frac{\partial Et(p,y)}{\partial p} \cdot p = 0;$$

$$(6) \qquad ps - c = 0,$$

where the arguments of $s(p,y)$ are suppressed. We denote by (p^*,y^*) the first-best (p,y) that solve these conditions, and let $s^* = s(p^*,y^*)$.

A contract achieves the first-best if it can elicit (p^*,y^*) and incorporates an incentive-neutral

price system can also be derived as a direct implication of the classic theorem by Robert Dorfman and Peter O. Steiner (1954) that optimal choice of p and s leads to equality of the elasticity ratio with the ratio of revenue to s.

[5] Throughout, we use the notation $[X]^+ \equiv \max(X,0)$.

[6] Throughout the paper, we suppress the argument F in the function π_i.

means of transferring rents between the manufacturer and the retailer. Under a uniform pricing contract (i.e., a wholesale price only), the retailer will maximize $E\pi(p,y;\theta)$ with respect to p and y. The retailer's first-order conditions are

$$(7) \qquad Et(p,y) + \frac{\partial Et(p,y)}{\partial p} \cdot p = 0;$$

$$(8) \qquad ps - w = 0.$$

From a comparison of (8) and (6), it is immediate that the uniform pricing contract fails to achieve the first-best outcome. In our terminology, $w > c$ creates a positive "vertical externality," flowing upstream through $(w - c)$, in the agent's choice of inventory y. The usual double-markup effect on price is *not* at work, however: (7) is identical to (5). This is because, conditional on inventory, both the retailer and the collective interest are served by setting the price to maximize expected revenue. The cost of inventory is sunk when price is set. The absence of any *direct* incentive distortion on price plays a role in our predictions in the full model of Section III.

The incentive problem can be resolved with a standard residual-claimancy contract. Where the manufacturer can charge a fixed fee, it would set $w = c$ and collect all profits through the fixed fee, effectively selling the firm to the downstream agent. All decisions are then optimal because the problem is internalized.

Consider as an alternative the set of instruments consisting of w and a buy-back policy under which the outlet has the option to sell back the inventory ex post to the manufacturer for a "buy-back price" b. (This set of instruments has no fixed fee.) Under a buy-back policy, the retailer's profit function is $\pi(p,y;w,b,\theta) = pt(p,y;\theta) - wy + bO(p,y;\theta)$. When the manufacturer adopts these instruments, the agent's first-order conditions are

$$(9) \qquad Et(p,y) + \frac{\partial Et(p,y)}{\partial p} \cdot (p - b) = 0$$

and

$$(p - w)s - (w - b)(1 - s) = 0,$$

or, equivalently,

$$(10) \qquad ps - w + b(1 - s) = 0.$$

Comparing (9) and (10) with (5) and (6), we can get (10) to match (6) at (p^*, y^*) by setting $b = (w - c)/(1 - s^*)$. However, because $w > c$ (to collect rents), this requires $b > 0$ and at p^* the retailer therefore has an incentive to raise price above p^* (from a comparison of (9) and (5)). In our terminology, the positive vertical externality on inventory must be offset by an inducement $b > 0$ to hold more inventory. But $b > 0$ then *creates* a vertical externality in price. An agent's action $dp > 0$ and implied $d(Et) < 0$ result in an opportunity cost to the manufacturer $bd(Et)$ that must be paid in each non-stockout realization of demand. The retailer ignores this cost. The use of a buy-back policy, therefore, necessitates a price restraint. Assuming quasi-concavity, the appropriate restraint is a price *ceiling* because the retailer has an incentive to raise prices.

When the retailer's price decision is constrained with a price ceiling, the buy-back allows an incentive-neutral transfer of profits to the manufacturer. From (9) and (5), raising w and b in the ratio $(w - c)/b = (1 - s^*)$ leaves the retailer's inventory incentives unchanged. Raising w and b while maintaining this ratio increases the flow of rents to the manufacturer. Thus w and b can be raised sufficiently to collect all rents from the retailer without affecting incentives.[7]

We summarize the observations above as the following proposition:

PROPOSITION 1: *In the model with a single retailer downstream, the following contracts elicit first-best profits for the manufacturer:*

(a) A wholesale price $w = c$ and a fixed fee F;

(b) Assuming quasi-concavity of the profit function in price: a wholesale price $w > c$, a buy-back price $b > 0$, and a price ceiling at p^*.

[7] Our argument here parallels the seminal contribution of Barry A. Pasternack (1985), with the additional observation regarding the requirement of a price ceiling. This result (Proposition 1(b) below) extends to all models analyzed in this article. In the remainder of this article, however, we focus on the case where a fixed fee F is feasible.

With positive retailer fixed costs, or retailer bargaining power, Proposition 1(b) demonstrates the ability of a buy-back contract to substitute for the simple residual-claimancy contract in a purely vertical setting. As retailer fixed costs or retailer bargaining power approach zero, however, the contract under Proposition 1(b) is degenerate: in this limit, the manufacturer simply extracts all rents by combining the price ceiling at p^* with $w = p^*$ and $b = p^*$.

In this setting with a single multi-task agent, the efficient contracts are simple: either a residual-claimancy contract or (if retailer fixed costs are zero) a degenerate buy-back policy. In the *multi-agent multi-task* settings below, the residual-claimancy contract is not enough to resolve incentive problems; and the buy-back policy plays a wider role in correcting not only inventory incentive distortions as above but also pricing incentive distortions.[8]

B. *Competing Retailers with No Spillovers*

In this special case, starting with conventional Marshallian demands, $q_i(p_1,p_2;\theta)$, $i = 1,2$, we must specify a process generating transactions in the event that at least one inventory constraint is binding. We impose the assumption that each consumer goes to the store at which she would shop if there were no chance of a stockout; and if there is a stockout then she does not buy. This assumption allows us to go directly from Marshallian demands to transactions: $t_1(p_1,p_2,y_1;\theta) = \max(q_1(p_1,p_2;\theta),y_1)$. Let $s_i(p_1,p_2;y_1) = \Pr(q_i(p_1,p_2;\theta) > y_i)$ denote the stockout probability at outlet i.

Assuming differentiability and a symmetric optimum, the first-best (p^*,y^*) in this case (expressed below for outlet 1) satisfies

$$(11)\quad Et_1(\mathbf{p},\mathbf{y}) + \frac{\partial Et_1(\mathbf{p},\mathbf{y})}{\partial p_1}\cdot p + \frac{\partial Et_2(\mathbf{p},\mathbf{y})}{\partial p_1}\cdot p = 0;$$

$$(12)\qquad ps - c = 0,$$

where $\mathbf{p} \equiv (p_1,p_2)$ and $\mathbf{y} \equiv (y_1,y_2)$.

The second equation is the same as in the purely vertical case, but equation (11), the first-order condition for price, changes because the first-best now incorporates the impact of a price change on demand at both outlets.

One resolution to the problem is a two-part tariff $\{w = c;F\}$ combined with a price floor. To see this, compare equations (11) and (12) (the efficient solution) with the first-order conditions of a retailer facing $w = c$ and no restraints. Since $w = c$, equation (8) matches equation (12). Equation (7) is missing the last term of (11), which is positive, so the agent has the incentive to price too low. A price floor, along with $w = c$, therefore elicits the first-best when a fixed fee F is available to transfer rents.[9]

Another resolution to the problem, which does not involve price restraints, is a two-part tariff $\{w > c;F\}$ combined with an inventory buy-back. In this case, the retailer's conditions, (9) and (10), will match the efficiency conditions (11) and (12), provided that b and w are set so as to solve the following two equations (which are linear in b and w):

$$b = -p^*\cdot\left.\frac{\partial Et_2(\mathbf{p},\mathbf{y})/\partial p_1}{\partial Et_1(\mathbf{p},\mathbf{y})/\partial p_1}\right|_{(\mathbf{p}^*,\mathbf{y}^*)};$$

and

$$w - c = b(1 - s^*).$$

[8] Pricing decisions are distorted in the multi-agent settings below for the same reason that they are *not* distorted in the single-agent setting above: there is no externality on price in a newsvendor model.

[9] Recall that we are adopting a first-order approach to the incentive problem, assuming quasi-concavity of the retailer profit function. When the general objective functions are non-quasi-concave, a *dictated price* on the part of the manufacturer plays the same role in solving incentive problems (without the prediction of whether this price restraint is a floor or a ceiling). The competitive inventory problem then collapses to one with exogenous prices and is nicely concave. We are interested, however, in whether price ceilings or floors are optimal because we observe both of these contracts. They are opposite instruments, and from a policy perspective, the law is completely different on each restraint, so analysis should distinguish the two. Accordingly, in this paper we adopt the common approach in applied models of assuming a range of parameters where payoffs are quasi-concave and therefore equilibrium exists. A similar issue arises regarding the symmetry of optimal inventory levels in that this symmetry assumption also requires a restriction of parameter values (see James D. Dana 2001).

Thus, providing that a fixed fee is available, incentive incompatibilities can be resolved simply with prices, w and b, i.e., without the need for restraints of any kind. The following proposition summarizes the contracts derived in this section.

PROPOSITION 2: *With competing retailers and no spillovers, the following contracts elicit first-best profits for the manufacturer:*

(a) A wholesale price $w = c$, a fixed fee F, and a vertical price floor at p^*;

(b) A wholesale price $w > c$, a buy-back price b, and a fixed fee F (and no price restraints).

Recall that with only one agent, and the *two tasks* of setting price and inventory, the simple residual-claimancy contract $\{w > c; F\}$ solved the incentive problem easily. No restraints were necessary in this case. Note, as well, that in models where there is one task (price or inventory) downstream, but *two agents*, there is also a contract $\{w > c; F\}$ that solves the incentive problem easily: w is set high enough that the horizontal distortion and vertical distortion are exactly offsetting. In our case of two agents *and* two tasks, two-part pricing does not work. The multi-task/multi-agent aspect of the setting is at the heart of the failure of two-part pricing.

C. *Connection to the Literature*

In inventory and pricing decisions along a supply chain, there are four externalities that potentially distort an agent's decisions: a vertical externality on price, a horizontal externality on price, and the same two externalities on the inventory decision. The two special cases in this section restrict the set of externalities a priori. The one-retailer (purely vertical) model rules out both horizontal externalities, and the second case considered, competing retailers without spillovers, rules out the horizontal externality on inventory. These simple models have direct counterparts in the often complex supply chain literature. Pasternack (1985) considers a single downstream retailer in a model with exogenous prices, thus assuming away both horizontal externalities and the vertical price externality, and incorporating a single incentive distortion—the vertical externality on inventory. Pasternack's result is that buy-backs can correct the distortion in inventory incentives. We find that buy-backs resolve the distortions in both inventory and pricing decisions, even allowing horizontal competitive externalities. In fact, as will become clearer in the next section, the most direct role of the buy-back is to correct a distortion in pricing, not inventory decisions.

Steven A. Lippman and Kevin F. McCardle (1997) consider two competing newsvendors with exogenous pricing, again restricting attention to a single distortion: the horizontal externality in inventory. The result of their complex model, that the newsvendors choose too much inventory (in terms of the collective interest), can be seen as an immediate implication of the single externality at work in their model. Ravi Anupindi and Yehuda Bassok (1999) add a vertical inventory externality to the horizontal externality of Lippman and McCardle.

Several other papers consider similarly restricted models. Howard P. Marvel and Peck (1995), Hamilton Emmons and Stephen M. Gilbert (1998), and Yuyue Song, Saibal Ray, and Shanling Li (2006) all consider a model with a single retailer and endogenous price and inventory decisions, thus incorporating two externalities (the two vertical externalities). Krishnan, Roman Kapuscinski, and Butz (2004) develop a model with two vertical externalities on inventory and sales effort (but not on prices), as does Terry A. Taylor (2002). Bernstein and Federgruen (2005) develop a model with competing retailers, in which three externalities (vertical inventory, vertical price, and horizontal price) are operative but spillovers are excluded.[10]

[10] Our model cannot be considered a generalization of all of these theories because we restrict attention to first-best (collective profit-maximizing) contracts. Much of the management science literature rules out fixed fees, an assumption that is ad hoc in the sense of not being imposed by the informational or arbitrage conditions of the model. We adopt fixed fees in some contracts, which allows the objective to be maximization of collective profits. These fixed fees are feasible under the assumptions of our model; and, in reality, nonlinear pricing of many types is observed and could substitute for fixed fees in our contracts. (The realistic assumption that a manufacturer must be concerned about retailers' profitability of carrying its product is an alternative justification for the focus on total efficiency.) In short, we restrict attention to contracts that are optimal under the assumptions of our model.

III. The Full Model: Competing Newsvendors in a Vertical Setting

Our analysis to this point has hinged on comparison of first-order conditions for the manufacturer (who maximizes joint profits) and first-order conditions for the decentralized retailer. In our discussion of the full model in this section, the first-order conditions are more involved. We adopt the approach outlined in Section I by decomposing the *difference* in first-order conditions between the manufacturer and the retailer into vertical and horizontal externalities.

A. *Assumptions*

In the full model, a single manufacturer sells the product through two competing retailers with the possibility of demand spillovers in the event of a stockout. The demands for the product at the two outlets are $q_1(p_1,p_2;\theta)$ and $q_2(p_1,p_2;\theta)$, where θ is a random variable. The retailers choose inventory levels, y_1 and y_2, and set prices before the realization of uncertainty. We define $\lambda_1(p_1,p_2;\theta)$ as the proportion of any excess demand from outlet 1 that would spill over to outlet 2 in the event of a stockout at outlet 1, and similarly for λ_2. (This imposes a restriction that λ_i not depend on y_i.) Finally, we assume that the joint distribution $H(q_1,q_2|p_1,p_2)$ on the random demands is differentiable and symmetric in the sense that the marginal distributions of $q_1(p_1,p_2;\theta)$ and $q_2(p_1,p_2;\theta)$ are identical. Asymmetric realizations of demand are, of course, possible. For fixed $\mathbf{p} \equiv (p_1,p_2)$, the support of the distribution of $q_i(p_1,p_2;\theta)$ is assumed to be an interval.

The following is an example that fits this structure. The two retailers are located at the ends of a unit line segment, with consumers located along the line between two firms. Each consumer purchases zero units or one unit of the product. A consumer bears a cost t per unit distance to travel to a store, so that the location, l, of a consumer on the line captures the consumer's relative preference for purchasing the product from one store versus the other. In addition, consumers vary in their absolute value of the product: consumers' reservation prices (including travel costs) for the product vary over an interval $[\underline{r},\bar{r}]$. Thus, a consumer type is a point in $[0,1] \times [\underline{r},\bar{r}]$. A density of consumers on this space of types (the "demand density") is a mapping from this consumer space to $[0,\infty)$. The demand density itself is random, depending on the random variable θ. The random density is denoted by $g(l,r;\theta)$. A realized demand density will in general be asymmetric, favoring one store or another. (To anticipate the discussion below, a store with a particularly favorable realization of demand may end up being stocked out. In this case, the demand from those affected consumers who would gain net surplus from purchasing at the other store represents a demand spillover.) However, the ex ante density of demand, $h(l,r) \equiv \int g(l,r;\theta)\,d\theta$, is assumed to be symmetric.

The timing of the game is as follows. The manufacturer offers contracts; the two retailers observe the contract offers and simultaneously choose to accept or not; the two retailers simultaneously choose price and inventory levels from some intervals $[0,\bar{p}]$ and $[0,\bar{y}]$; uncertain demand is realized; and, finally, consumers make purchase decisions.

Given decisions $(p_1,y_1;p_2,y_2) \equiv (\mathbf{p},\mathbf{y})$ and the realization θ, the demand at outlet i can be expressed as $D_i(\mathbf{p},y_j;\theta) = q_i(\mathbf{p};\theta) + \lambda_j(\mathbf{p};\theta)(q_j(\mathbf{p};\theta) - y_j)^+$. The transactions (sales) realized by outlet i are given by

$$
\begin{aligned}
(13) \qquad t_i(\mathbf{p},\mathbf{y};\theta) &= \min\{y_i, D_i(\mathbf{p},y_j;\theta)\} \\
&= y_i - (y_i - D_i(\mathbf{p},y_j;\theta))^+ \\
&= y_i - O_i(\mathbf{p},\mathbf{y};\theta),
\end{aligned}
$$

where $O_i(\mathbf{p},\mathbf{y};\theta)$ is the number of "overstocks" at outlet i.

The aggregate realized profit of the supply chain can be expressed as

$$
\begin{aligned}
(14) \qquad \Pi(\mathbf{p},\mathbf{y};\theta) = {} & p_1 t_1(\mathbf{p},\mathbf{y};\theta) - cy_1 \\
& + p_2 t_2(\mathbf{p},\mathbf{y};\theta) - cy_2 .
\end{aligned}
$$

The realized profit of outlet 1 can be expressed as

$$(15) \qquad \pi_1(\mathbf{p},\mathbf{y};w,\theta) = p_1 t_1(\mathbf{p},\mathbf{y};\theta) - wy_1 - F.$$

In the example outlined earlier, a symmetric equilibrium in prices and inventory levels leads to a partition of the set of consumer types, as

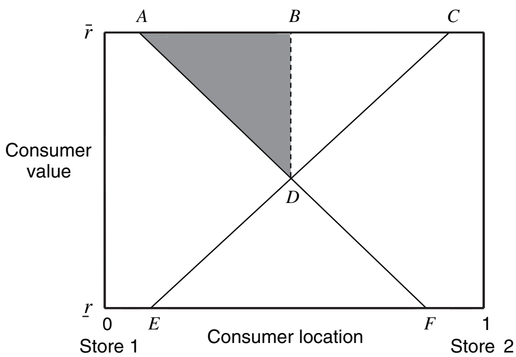

Figure 1. Equilibrium Partition of Consumer Types
(The shaded area represents those consumer types who attempt to purchase from store 1 and then, if order is unfilled, from store 2)

illustrated in Figure 1. The set of consumer types who would obtain a nonnegative surplus from purchasing from store 1 are $\{(l, r) | r - p - tl \geq 0\}$, who are the consumers to the upper-left of the line 0-E-C in the diagram. Similarly, consumers to the upper-right of the line 1-F-A would obtain nonnegative surplus purchasing from store 2. In equilibrium, consumers to the left of E-D-B *plan* to purchase from store 1; consumers to the right of F-D-B plan to purchase from store 2. Consumers in E-D-F stay home. Consumers in A-B-D are those who will spill over to store 2 if unable to purchase from store 1, their preferred outlet. This partition is established as soon as prices are announced, prior to the realization of demand, and does not depend upon inventory decisions.

With a strong asymmetric realization of demand favoring, say, store 1, this store will be stocked out. If store 1 can fill, say, only two-thirds of its orders, then one-third of its potential consumers go without purchasing from it. We assume, to complete the example, that the probability of getting the product (the "fill rate") is independent of a consumer's type. Thus, one-third of the consumers in the region A-B-D who would have purchased from store 1, but also gain positive surplus from store 2, are added to the demand at store 2. That is, consumers who value the product sufficiently highly, and who are not inconvenienced too severely by purchasing from store 2, spill over to that store if their order is unfilled at store 1.

B. *Failure of the Price System to Coordinate Incentives: The Missing Externality*

From (14) and (15) we can derive the first-order conditions and compare the individual incentives and collective efficiency in price and inventory decisions. The difference in first-order conditions can be decomposed as follows:

$$(16) \quad \frac{\partial E\pi_1}{\partial y_1} = \frac{\partial E\Pi}{\partial y_1} - \overbrace{(w - c)}^{\text{vertical externality}} - \overbrace{p_2 \frac{\partial Et_2}{\partial y_1}}^{\text{horizontal externality}};$$

$$(17) \quad \frac{\partial E\pi_1}{\partial p_1} = \frac{\partial E\Pi}{\partial p_1} - \overbrace{p_2 \frac{\partial Et_2}{\partial p_1}}^{\text{horizontal externality}}.$$

The terms labelled in these equations describe the externalities parallel to those outlined in equations (3) and (4) of Section I. The last equation captures what we think is a fundamental feature of the decentralization of price and inventory decisions in the newsvendor model: controlling for the level of inventory, *pricing decisions are not subject to a vertical externality.* Conditional upon the inventory choice of an outlet, the manufacturer has no direct (vertical) interest at all in the price at which the inventory is resold: the wholesale revenue, costs, and hence profits are completely determined by the inventory purchase. This does not mean that the manufacturer is indifferent to the pricing decision. To the contrary, the downstream pricing decision is virtually *always* distorted because the manufacturer cares about maximizing total profits.

In which direction is the retailer pricing decision distorted? From (17), this depends on the sign of $\partial Et_2/\partial p_1$. The random derivative $\partial t_2/\partial p_1$ varies across states θ depending on which store (if either) is stocked out. From the expression $D_2(\mathbf{p}, y_1; \theta) = q_2(\mathbf{p}; \theta) + \lambda_1(\mathbf{p}; \theta)(q_1(\mathbf{p}; \theta) - y_1)^+$, we can construct the matrix in Table 1 indicating the value of $\partial t_2/\partial p_1$ in various states. In states where store 2 is stocked out, its transactions are given by 0 and $\partial t_2/\partial p_1 = 0$ in Table 1. When neither store is stocked out, $\partial t_2/\partial p_1 = \partial q_2/\partial p_1 > 0$ and we obtain the conventional

TABLE 1—VALUE OF $\partial t_2/\partial p_1$ ACROSS DIFFERENT STATES OF θ

	Store 1 stocked out	Store 1 not stocked out
Store 2 stocked out	0	0
Store 2 not stocked out	$\frac{\partial q_2}{\partial p_1} + \lambda_1(p;\theta) \cdot \frac{\partial q_1}{\partial p_1} + q_1 \cdot \frac{\partial \lambda_1}{\partial p_1}$	$\frac{\partial q_2}{\partial p_1}$

result, that an increase in p_1 has a positive effect on transactions at store 2. Call $\partial q_2/\partial p_1 > 0$ the "direct effect" of an increase in p_1 on t_2. When only store 2 is stocked out, however, there is an offsetting effect of dp_1 on t_2, which has two components captured by the last two terms of the bottom-left cell of Table 1. To understand the offsetting effect, which we refer to as the "fill-rate effect," consider the example illustrated in Figure 1. The demand at store 1 decreases with an increase in p_1 (as the line *CE* shifts to the left) and therefore fewer customers will spill over to store 2. This is captured by the negative term $\lambda_1(p;\theta) \cdot \partial q_1/\partial p_1$ in Table 1. The final term $q_1 \cdot \partial\lambda_1/\partial p_1$ captures the fact that the proportion of excess demand that spills over from store 1 to store 2, which is the ratio of consumers in the triangle A-B-D to consumers in the trapezoid 0-E-D-B, also changes as p_1 increases. For example, the reduced demand by consumers very near store 1 may free up stock for consumers who would otherwise spill over to store 2. Thus, the second component of the fill rate effect may also be negative.

It is not hard to construct extreme examples of densities of demand for which the fill-rate effect dominates the direct effect, and the cross-derivative of p_1 on Et_2 is actually negative. In Figure 1, suppose that there is a (smooth and symmetric) ex ante density $h(l,r)$ that is very large along the borders ED and FD in equilibrium, but very close to zero along the line BD. Then, the direct effect, which is given by the additional consumers captured by store 2 when the line BD shifts left, is very small. The fill-rate effect, which depends on the density of consumers along ED and FD, as well as the size of A-B-D, is large. We cannot formally rule out the possibility that the fill-rate effect dominates since we do not incorporate into our model evidence on the distribution of demand. But simulations of the model indicate that for the widest reasonable range of parameters, the direct effect dominates the fill-rate effect. In sum:

PROPOSITION 3: *In the full model, unconstrained outlets generically*[11] *fail to achieve the efficient outcome* (p^*, y^*). *If the direct effect of* dp_1 *on* Et_2 *dominates the fill-rate effect, then at* (p^*, y^*), *each outlet would profit by unilaterally reducing price.*

Note that Proposition 3 requires no concavity assumptions whatsoever—not even assumptions that guarantee the existence of a pure strategy equilibrium. It relies only on the necessary first-order conditions for the aggregate optimum. Economists often view incentive distortions in terms of externalities, or missing markets. Ironically, the source of the inefficiency in downstream decisions here is a missing externality.

C. *Contracts That Achieve Coordination*

We consider, as earlier, the following contract terms: a linear wholesale price w and a fixed fee F; vertical price restraints; and a buy-back policy with a buy-back price b. As before, denote by s^* the probability of a stockout at either outlet at the first-best $(\mathbf{p}^*, \mathbf{y}^*)$. Denote the own-price elasticity of expected transactions and the cross-price elasticity of expected transactions, evaluated at $(\mathbf{p}^*, \mathbf{y}^*)$ by $\varepsilon_p^t \equiv d\ln Et_1/d\ln p_1$ and $\varepsilon_{pc}^t \equiv d\ln Et_2/d\ln p_1$. The following is proved in the Appendix.

PROPOSITION 4: *If profit functions are quasi-concave, the efficient* (p^*, y^*) *can be elicited with:*

[11] "Generically" means that the statement is true for "essentially all" model parameters. Specifically, if, at $(\mathbf{p}^*, \mathbf{y}^*)$, $\partial Et_2/\partial p_1 = 0$ (by coincidence) in equilibrium, a small perturbation in the model parameters will render the derivative non-zero.

(a) A price floor at p^*, a fixed fee F, and a linear wholesale price $w^* = p^*s^*$, if the direct effect of dp_1 on Et_2 dominates the fill-rate effect;

(b) A wholesale price $w > p^*s^*$, a fixed fee F, and a buy-back price given by

$$b^* = -\frac{\varepsilon^t_{pc}}{\varepsilon^t_p}p^*.$$

The role of the price floor in Proposition 4(a) is clear. The missing externality in the outlet's first-order condition on price leaves the outlet with the incentive to drop price below p^* (when the direct effect of dp_1 on Et_2 dominates the fill-rate effect). The floor constrains the retailer against pricing below the first-best. Setting $w^* = p^*s^*$ then leaves the outlet, which faces a newsvendor problem, with efficient incentives: its optimum is achieved by setting inventory such that the probability of stockout equal to w^*/p^* (from the well-known solution to this problem), and this ratio is selected by the manufacturer to equal s^*. The only way the outlet can achieve a probability of stockout equal to s^* is by choosing the first-best efficient y^*, so the inventory incentive problem is resolved as well.

An optimal buy-back policy in Proposition 4(b) resolves the incentive problem by *creating* a vertical externality of precisely the right magnitude to fill in the missing externality. This contract achieves the efficient solution purely through the use of price instruments to elicit the right incentives; no restraint on retailer actions is necessary. The key is that the wholesale price, w, is free for use purely as an incentive device because the fixed fee is available to redistribute profits.[12]

Proposition 4 has characterized contracts in the case where a fixed fee, F, is feasible. The result (Proposition 1(b)) that a buy-back policy combined with a price ceiling can, in the absence of fixed fees, achieve the first-best also extends to the current setting. Buy-backs substitute for price floors but are complementary with price ceilings in our model.

The predictive content of our model can be summarized as follows. First, the price system is (generically) unable to convey the right incentives for pricing and inventory decisions and will be dominated by vertical restraint contracts and other contracts. Price floors are supported as the most plausible form of price restraint. Our model as it stands predicts that RPM and buy-back policies are perfect substitutes. A prohibition of one instrument would simply lead to the adoption of the other.

D. *Numerical Example*

To illustrate the results presented in this section and to explore the welfare implications of price restraints in the model, we simulate a special case of the example depicted in Figure 1. We adopt the following structure for the distribution of densities of demand: a random "mass" of consumers at each store location—these consumers incur an infinite cost of shopping at the rival store and are called "loyal" consumers—as well as a uniform density of "common" or comparison shoppers on the line segment joining the two stores. The travel cost per unit distance is t for the comparison shoppers. We assume that the density of the comparison shoppers is itself random, as are the masses of loyal consumers for each outlet.[13] This spatial model is illustrated in Figure 2.[14]

[12] The role of buy-backs in creating a vertical externality, which dampens price competition, is in a sense the opposite of the role that buy-backs would take in a model with (a) no fixed fees (or nonlinear pricing); (b) no uncertainty; and (c) complete price flexibility, i.e., ex post price setting. Consider the David M. Kreps and José A. Scheinkman (1983) theorem that ex ante quantity setting with ex post Bertrand competition leads to a Cournot outcome. A manufacturer with no means of collecting profits via a fixed fee or high inframarginal prices could desire more intense retailer competition than retailer Cournot behavior. Buy-backs would be a profitable means of "undoing" the Kreps-Scheinkman effect through a weakening of the commitment by retailers to quantities, bringing the retail market equilibrium from Cournot closer to the (more competitive) Bertrand equilibrium (V. Padmanabhan and I. P. L. Png 1997). In our setting, buy-backs do not play this role because of limited price flexibility (ex ante pricing) and the availability of fixed fees.

[13] This is the simplest representation of random demand in a spatially differentiated duopoly in which the density of demand is symmetric across the two outlets but any particular realization of demand is (in general) not symmetric.

[14] V.G. Narayanan, Ananth Raman, and Jasjit Singh (2005) consider a similar spatial model but set aside the possibility of asymmetric demand realizations, and therefore

Each of the comparison shoppers buys either zero or one unit, and all of these customers have the same reservation price r. This set-up generates a linear demand curve for these consumers, with a choke price of r. We assume that the demand curve generated by loyal customers is also linear and, for simplicity, also has a choke price of r: $LD_i(p_i;\theta_i) \equiv \theta_i(r - p_i)$. The density of comparison shoppers is given by a random variable θ_c, and the size of loyal customer demand at i, θ_i, is also random. The demand distribution is given by the joint probability distribution $G(\boldsymbol{\theta})$, where $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_c)$.

The surplus that a comparison shopper at location x experiences buying from outlet 1 is $r - p_1 - tx$ and from outlet 2 is $r - p_2 - t(d - x)$; the customer will buy from the outlet where she obtains the higher net surplus. There are three marginal consumer types: the consumer M_0 who is indifferent between purchasing at each outlet; the consumer M_1 who (if outlet 2 is stocked out) is indifferent between purchasing at outlet 1 and not purchasing; and M_2, defined similarly. These marginal consumers are given by $M_1 = \min\{(r - p_1)/t, d\}$; $M_2 = \max\{d - (r - p_2)/t, 0\}$ and (if $M_1 > M_2$), $M_0 = (M_1 + M_2)/2$.

Comparison customers first attempt to obtain the product from their preferred (higher surplus) outlet; if their preferred outlet is stocked out, then those customers who are willing to shop at the other outlet will spill over to that outlet. We assume "proportional rationing:" all customers who shop at a particular outlet have the same probability of getting the product whether they buy directly or on rebound from the stocked-out other outlet.[15] The "common" demand at outlet 1 (the number of customers on the line segment who will make their first purchase attempt at outlet 1) is given by $CD_1 \equiv \theta_c(\min\{(M_1 + M_2)/2, M_1\})$ and at outlet 2 is $CD_2 \equiv \theta_c(d - \max\{(M_1 + M_2)/2, M_2\})$. The proportion of customers who spill over from outlet i to outlet j is given by $\lambda_{ij} = [(M_1 - M_2)^+/2]/(LD_i + CD_i)$.

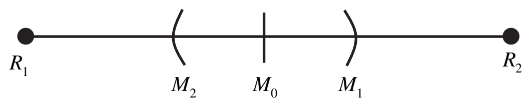

FIGURE 2. SPATIAL MODEL: THE THREE MARGINAL CONSUMERS

Using this setup, we numerically simulate and compare the centralized and decentralized (i.e., with no restraints) solutions. Each outlet chooses price p_i and inventory y_i prior to observing demand. In the centralized case, a single decision maker chooses (p_1, y_1, p_2, y_2) to maximize collective profits; it is this centralized solution that is implemented with vertical restraints. In the decentralized case, the outlets independently and simultaneously choose (p_i, y_i).

In comparing the centralized solution with the decentralized equilibrium, we restrict attention to symmetric outcomes. For very high travel costs, each outlet is independent of the other (their markets do not overlap), and a unique, symmetric solution will result.[16] For intermediate values of travel cost, the two outlets have overlapping markets and have a symmetric solution. (For sufficiently small travel costs, a symmetric solution does not exist.) Where multiple symmetric equilibria exist, which occurs for some parameters because of the nonconcavities in payoffs as we have discussed, we restrict attention to the equilibrium that maximizes joint profits.

In addition to the positive comparison of the impact of the movement from the purely decentralized to the centralized solution, we address the obvious normative question. Consider, for example, the use of price floors (as in Proposition 4(a)), applied to the numerical example. In adopting a price floor in this model, the manufacturer trades off a higher retail price and a consequent drop in quantity demanded in exchange for greater inventory and a resulting increase in the expected number of transactions. Does the manufacturer's

spillover demand, at the two outlets. In our model, the existence of "loyal" customers with uncertain demand allows for the possibility of asymmetric demand realizations and demand spillovers, even under symmetric ex ante densities of demands. Since Hal R. Varian's (1980) classic article, models of retail markets incorporating both shoppers and loyal consumers have been common.

[15] A specific assumption on the mechanism rationing demand in the event that *both* firms are stocked out is necessary for welfare analysis, but not for positive analysis (since the transactions are simply y_1 and y_2 in this event).

[16] The uniqueness follows from the specific assumptions on demand that we use in the numerical example.

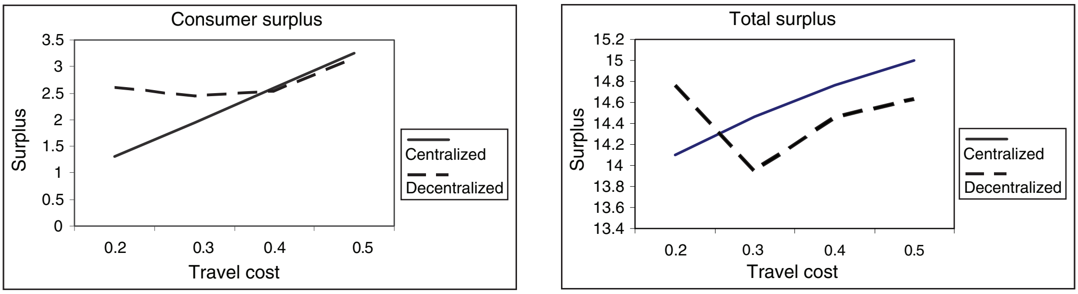

FIGURE 3. WELFARE CALCULATIONS FOR SPATIAL MODEL WITH $r = 2$, $d = 2$, $c = 0.1$, $\xi_1, \xi_2 \sim Uniform\{1,2\}$, $\xi_c \sim Uniform\{3,4\}$, t VARYING

willingness to make this tradeoff signal an increase in total welfare as well? That is, do vertical restraints increase total welfare?

A sample of the simulation results is presented in Figure 3. We numerically verify that centralization can increase total welfare and even consumer surplus alone. In addition, in each of the cases where centralization raises welfare in this example, the centralized outcome can be implemented with a price *floor.* This example supports the Supreme Court departure from per se illegality of resale price maintenance in *Leegin Creative Leather Products, Inc. v. PSKS, Inc.*

IV. Application

In Section I of this paper, we reviewed a framework that allowed us to pose the following question: are decentralized retailers biased toward excessive price competition and inadequate inventories, in which case we should observe vertical price floors, or toward high prices and excessive inventories, which would lead to vertical price ceilings? Our model shows that price restraints can resolve incentive distortions in pricing *and* inventory decisions, and suggests that in markets where inventory is important, the bias is toward excessive price competition. The optimal price restraint is traced to a comparison of the direct effect versus the fill-rate effect in the cross-elasticity of prices on rival's expected transactions and, except in the extreme case where the fill-rate dominates, a price floor is optimal.

The model developed in this paper is, of course, not an exact match with any real market, and when we bring the theory to cases, we must recognize its simplicity. The relaxation of the model's assumptions in three directions is particularly important. First, the results generalize to the case of an arbitrary number of retailers. Second, because retailers incur fixed costs in the decision to carry a product, if a retailer's gain to carrying inventory is distorted sufficiently, then the retailer will not carry the product at all. In other words, the addition of fixed costs of stocking a product yields the "outlets hypothesis"—that manufacturers use RPM to encourage more outlets to carry their products or to carry a broader line of their products (J. R. Gould and L. E. Preston 1965).[17] Third, our model predicts perfect substitutability between instruments. While the implication of substitutability is falsifiable, available data do not allow a test. Instead, we note that in reality instruments vary in their implementation costs and effectiveness. The implication is then that a legal restriction on instruments—buy-backs or price floors—will disrupt efficient distribution.

Almost every retail product is sold from inventory—from chewing gum to farm implements to diamond rings—and our prediction is consistent with the fact that vertical price floors were indeed more popular than price ceilings when these restraints were legal.[18] Individual

[17] Note that our model does not assume that the unrestrained retail market equilibrium is symmetric, and in this sense potentially explains cases in which a manufacturer uses vertical contracts to ensure that its product is not carried by discounters.

[18] When resale price maintenance was permitted, it was used in a wide variety of retail markets, including many lines of clothing (jeans, shoes, socks, underwear, shirts), jewelry, sports equipment, candy, biscuits, automobiles,

case studies provide more direct evidence of the role of RPM in correcting distortions in inventory incentives. The "fair trade" state laws that allowed RPM were repealed in the early 1970s, allowing economists to study the impact of a sudden elimination of RPM from the feasible set of contracts. Anthony P. Hourihan and Jesse W. Markham (1974) conducted case studies of nine manufacturing companies that had been using the restraint. In five of the nine cases studied, prices were unaffected, consistent with the price restraint not binding retailer decisions. In three of the four cases where retail prices did decline, the availability of the product to consumers dropped because of a drop in inventory, a drop in the selection of items to carry within the product line, and, in two cases, a drop in the number of outlets carrying the product. The products carried by the firms with decreases in inventory were housewares and tableware, both of which are products with a strong seasonal demand, so our assumptions of demand uncertainty and perishability are plausible.[19]

In other cases as well, the termination of resale price maintenance led to fewer outlets, to the detriment of the manufacturer. Corning Glass Works used this restraint from 1937 until it was prevented from doing so in a case brought by the Federal Trade Commission in 1975.[20] In interviews ten years after the case, Corning executives indicated that one of the most important effects of the case was the loss of many of its smaller outlets.[21] In another example, after legislation had ended an earlier era of Fair Trade Pricing, the number of dealers selling Schick shavers fell from 35,000 to 7,000 in one year (P. W. S. Andrews and Frank A. Friday 1960).[22]

Markets where inventory incentives should be important are those with substantial uncertainty in demand and products with a limited shelf life (perishability). Four factors may account for limited shelf life: seasonality in demand (greeting cards, toys and other holiday gifts, sports equipment); physical depreciation (pharmaceuticals); product obsolescence (magazines, newspapers); and fashion goods with a limited period of popularity (clothing, books, recorded music). These are indeed industries where RPM has been popular.[23]

Beyond RPM, our simple model captures the role of contractual resolutions of incentive distortions in inventory and pricing decisions. Buy-backs have been used for books, cosmetics, greeting cards, and a wide variety of products with uncertain demand and perishability or a short selling season.[24] The instrument is often adopted in the more general form of "returns policies," such as subsidizing unsold inventory, as in the return for credit of the front covers of paperback books or magazines.[25]

V. Conclusion

An organization faces an incentive problem when an agent within the organization does not appropriate the full net benefits to the organization of the agent's decisions. The organization can respond to an incentive problem by altering the net benefits through internal prices that

gasoline, and small and large appliances (stereos, shavers, washing machines). See Mathewson and Winter (1998). Estimates of the proportion of retail sales subject to RPM in the United States during the 1950s run from 4 percent to 10 percent (Frederic M. Scherer and David Ross 1990, 549; Thomas R. Overstreet 1983, 6). In both the United Kingdom and Canada, the practice was even more popular than in the United States: in 1960, some 25 percent of goods and services were subject to RPM in the United Kingdom; and in Canada, before the law prohibiting RPM was enacted in 1951, an estimated 20 per cent of goods sold through grocery stores and 60 per cent sold through drugstores were fair-traded (Overstreet 1983, 153, 155).

[19] Of course, the assumption of *complete* perishability in our model is a simplification.

[20] *In the Matter of Corning Glass Works*, 85 FTC 1061 (1975), modifying 82 FTC 1675 (1973), aff'd 509 F.2d 293 (7th Cir. 1975).

[21] Pauline M. Ippolito and Overstreet (1996, 325).

[22] Similar results are reported for three other companies; see Andrews and Friday (1960, 27–29). See also Overstreet (1983) for an excellent discussion of general evidence in this regard.

[23] See, for example, Ippolito (1991).

[24] Padmanabhan and Png (1995), for example, report that book publishers have accepted returns starting with Viking Press in 1932, and that college textbook stores in the United States returned about 40 percent of all books as of the mid-1990s.

[25] Note that a contract with a linear price w and a royalty r per unit sold is pay-off equivalent to a contract with a linear price $w + r$ and a buy-back price r. In both cases, a downstream outlet pays w for all units not sold, and $w + r$ for units sold. All results that use buy-back policies extend to per-unit royalties. Empirically, buy-back prices appear to dominate per-unit royalties perhaps because of the relative ease of monitoring.

internalize the nonappropriated returns. Or it can constrain the agent's decisions in some dimensions and ensure that the externalities are internalized in the other dimensions.

We apply these simple principles to the coordination of price and inventory decisions in a distribution system. The strategies take the form of reward parameters (the wholesale price and buy-backs) and vertical restraints on downstream prices. An incentive problem leads to the failure of the price system and to the need for more complex contracts whenever incentive distortions, or externalities, do not "balance out." We show that in the decentralization of price and inventory decisions, externalities virtually never balance out. The problem, ironically, is a missing externality.

Price floors and price ceilings are used to address the opposite types of incentive problems. In Section I, we show that the need for restraints—price floors, for example—hinges not on whether outlets price too low, but on whether the optimal *mix* of decisions or strategies at the retail level mirrors the efficient mix. Price floors are optimal in the newsvendor model to counter the missing externality, except in the extreme case where the fill-rate dominates. Other contracts can also internalize externalities in downstream outlets' decisions.[26] Two instruments, a buy-back price and a wholesale price, elicit the efficient pair of targets, p^* and y^*, without the need for any restraints at all, when these instruments are set at levels that create exactly offsetting externalities on price and inventory decisions. The most direct role of the buy-back price is to correct for the missing-externality distortion in the *pricing* decision—contrary to what one might expect, in that buy-backs would seem directed toward inventory incentives. Our model is the simplest possible framework within which the full range of vertical and horizontal interactions of price and inventory decisions are at work.

[26] The supply chain literature has recently studied a wide range of contracts. Examples include *quantity-flexibility contracts,* in which a retailer has some flexibility in modifying the order size ex post; *price discount sharing schemes,* which are contractual clauses that provide a functional link between wholesale prices and retail prices; trade promotions such as *bill-backs, count-recounts,* and *markdown allowances,* which all refer to schemes by which the costs of downstream promotions are shared by the upstream firm; *sales rebates,* whereby the retailer gets a rebate if a sales threshold is reached; *revenue sharing contracts,* where the retailer pays a royalty on the total revenue; and so on. The externality-balancing framework presented in this paper can help explain the role that each of these instruments plays in aligning incentives along a supply chain.

Appendix

PROOF OF PROPOSITION 4:

To prove (a), first assume that the outlets are constrained to charge p^*. They then play a game in inventory decisions y_1 and y_2, with payoff functions given by $\pi_i(\mathbf{p}^*, y; \theta) = p^* \min\{y_i, q_i(\mathbf{p}^*; \theta) + \lambda_j(\mathbf{p}^*; \theta)[q_j(\mathbf{p}^*; \theta) - y_j]^+\} - wy_i - F$. Note that $\pi_i(p^*, y; \theta)$ is concave in y_i and therefore so is $E\pi_i(p^*, y; \theta)$. This guarantees a pure strategy equilibrium in (y_1, y_2). The assumption that the support of $q_i(\mathbf{p}^*; \theta)$ is an interval of the real line can be used to show that the payoffs are differentiable. Let $\bar{y}$ be the maximum possible $q_i(\mathbf{p}^*; \theta)$ at $\mathbf{p}^*$. Within the range $[0, \bar{y}]$, the reaction function of each firm is differentiable, and has a strictly negative slope.[27] Note also that the optimal response to 0 is positive and that the optimal response to $\bar{y}$ is less that $\bar{y}$. It follows that a *unique* symmetric pure strategy exists.[28]

To prove that $w^* = p^* s^*$, note that the optimal inventory y^* is characterized by the first-order condition (following from (14))

$$\left[p^* \frac{\partial Et_1}{\partial y_1} - c + p^* \frac{\partial Et_2}{\partial y_1} \right] \Bigg|_{(p^*, y^*)} = 0.$$

From this equation, setting $w = p^* \dfrac{\partial Et_1}{\partial y_1}\Bigg|_{(p^*, y^*)}$ ensures that

$$w - c + p^* \frac{\partial Et_2}{\partial y_1}\Bigg|_{(p^*, y^*)} = 0,$$

which then implies that the last two terms of (16) sum to zero, ensuring that the individual and collective first-order conditions for y coincide.

[27] This can be shown by taking the derivative of $E\pi_i$ w.r.t. y_j. The derivative and the expectation operation can be switched by the dominated convergence theorem. We can then differentiate (piecewise) π_i. There will be a range of values of θ where $\partial\pi_i/\partial y_j$ is negative, and a range where it is 0. Taking expectations will give us a strictly negative value—assuming that $\partial\pi_i/\partial y_j < 0$ on a set of positive measure.

[28] A similar argument shows that reaction curves are downward-sloping for the buy-back payoff functions below.

Note that $(\partial Et_1/\partial y_1)|_{(\mathbf{p}^*,\mathbf{y}^*)} = s^*$, proving that $w^* = p^* s^*$.[29]

To complete the proof of part (a), we need to show that a price floor at p^*, if imposed, is indeed binding. To prove this, note that equation (17) shows that $(\partial E\pi_i/\partial p_i)|_{(p_i = p^*)} < 0$ and under the quasi-concavity assumption on $E\pi_i$ the price floor is binding; and the first-order condition on inventory is sufficient for the optimum, implying that y^* is elicited.

To prove (b), note that outlet 1's realized profit function with a buy-back price is given by

$$\pi_1^b(\mathbf{p},\mathbf{y};\theta) = p_1 t_1(\mathbf{p},\mathbf{y};\theta) + bO_1(\mathbf{p},\mathbf{y};\theta) - wy_1 - F.$$

Note that

$$\pi_1^b(\mathbf{p},\mathbf{y};\theta) = \pi_1(\mathbf{p},\mathbf{y};\theta) + b(y_1 - t_1(\mathbf{p},\mathbf{y};\theta));$$

$$E\pi_1^b(\mathbf{p},\mathbf{y};\theta) = E\pi_1(\mathbf{p},\mathbf{y};\theta) + b(y_1 - Et_1(\mathbf{p},\mathbf{y};\theta)).$$

The now-familiar comparison of individual versus collective incentives at (p^*, y^*), with the contractual parameter b, is given by the following:

$$\text{(A1)} \quad \frac{\partial E\pi_1^b}{\partial y_1} = \frac{\partial E\Pi}{\partial y_1} - \overbrace{\left[(w-c) - b\left(1 - \frac{\partial Et_1}{\partial y_1}\right)\right]}^{\text{vertical externality}} - \overbrace{p_2\frac{\partial Et_2}{\partial y_1}}^{\text{horizontal externality}};$$

$$\text{(A2)} \quad \frac{\partial E\pi_1^b}{\partial p_1} = \frac{\partial E\Pi}{\partial p_1} - \overbrace{b\frac{\partial Et_1}{\partial p_1}}^{\text{vertical externality}} - \overbrace{p_2\frac{\partial Et_2}{\partial p_1}}^{\text{horizontal externality}}.$$

[29] Note that $Et_1 = y_1 - \int_{\{\theta | y_1 > D_1(\mathbf{p},y_2;\theta)\}} (y_1 - D_1(\mathbf{p},y_2;\theta))\,dG(\theta)$ and $D_1(\mathbf{p},y_2;\theta)$ is a random variable, which we can denote by $\tilde{D}_1$. Define $\tilde{G}(x)$ to be the distribution of $\tilde{D}_1$, i.e., $\tilde{G}(x) = \Pr(\tilde{D}_1 \le x)$. Applying the Leibnitz rule, we get that $\partial Et_1/\partial y_1 = 1 - \int_0^{y_1} d\tilde{F}(\tilde{D}_1) = 1 - \tilde{G}(y_1)$, which is the probability of a stockout at outlet 1.

Note that setting $b = -p_2(\partial Et_2/\partial p_1)/(\partial Et_1/\partial p_1)$ makes the last two terms of (A2) sum to zero. Multiplying the numerator and denominator of the fraction by p^*/Et^* yields the expression for b^*/p^* in the proposition, where t^* is the transactions function evaluated at the optimum. Setting $w = c + b[1 - (\partial Et_1/\partial y_1)] - p_2(\partial Et_2/\partial y_1)$ makes the last two terms of (A1) equal to zero. Since, from the proof of part (a) above we know that $c - p^*(\partial Et_1/\partial y_1)|_{(p^*,y^*)} = p^* s^*$, this also shows that $w^* > p^* s^*$ (because $(\partial Et_1/\partial y_1) < 1$).

REFERENCES

Arrow, Kenneth J., Theodore Harris, and Jacob Marschak. 1951. "Optimal Inventory Policy." *Econometrica*, 19(3): 250–72.

Andrews, P. W. S., and Frank A. Friday. 1960. *Fair Trade: Resale Price Maintenance Reexamined.* London: Macmillan and Co.

Anupindi, Ravi, and Yehuda Bassok. 1999. "Centralization of Stocks: Retailers vs. Manufacturer." *Management Science,* 45(2): 178–91.

Bernstein, Fernando, and Awi Federgruen. 2005. "Decentralized Supply Chains with Competing Retailers under Demand Uncertainty." *Management Science*, 51(1): 18–29.

Butz, David A. 1997. "Vertical Price Controls with Uncertain Demand." *Journal of Law and Economics*, 40(2): 433–59.

Cachon, Gerard P. 2003. "Supply Chain Coordination with Contracts." In *Supply Chain Management: Design, Coordination and Operation*. ed. A. G. de Kok and Stephen C. Graves, 229–339. Amsterdam: Elsevier Science, North-Holland.

Dana, James D., Jr. 2001. "Monopoly Price Dispersion under Demand Uncertainty." *International Economic Review*, 42(3): 649–70.

Dana, James D., Jr., and Nicholas C. Petruzzi. 2001. "Note: The Newsvendor Model with Endogenous Demand." *Management Science,* 47(11): 1488–97.

Dana, James D., Jr., and Kathryn E. Spier. 2001. "Revenue Sharing and Vertical Control in the Video Rental Industry." *Journal of Industrial Economics*, 49(3): 223–45.

Deneckere, Raymond, Howard P. Marvel, and James Peck. 1996. "Demand Uncertainty, Inventories, and Resale Price Maintenance." *Quarterly Journal of Economics*, 111(3): 885–913.

Deneckere, Raymond, Howard P. Marvel, and James Peck. 1997. "Demand Uncertainty and Price Maintenance: Markdowns as Destructive Competition." *American Economic Review*, 87(4): 619–41.

Dorfman, Robert, and Peter O. Steiner. 1954. "Optimal Advertising and Optimal Quality." *American Economic Review,* 44(5): 826–36.

Easterbrook, Frank H. 1984. "Vertical Arrangement and the Rule of Reason." *Antitrust Law Journal*, 53(2): 135–74.

Emmons, Hamilton, and Stephen M. Gilbert. 1998. "Note. The Role of Returns Policies in Pricing and Inventory Decisions for Catalogue Goods." *Management Science*, 44(2): 276–83.

Gould, J. R., and L. E. Preston. 1965. "Resale Price Maintenance and Retail Outlets." *Economica*, 32(127): 302–12.

Hourihan, A. P., and J. W. Markham. 1974. *The Effects of Fair Trade Repeal: The Case of Rhode Island*. Cambridge, MA: Marketing Science Institute and Center for Economic Studies.

Ippolito, Pauline M. 1991. "Resale Price Maintenance: Empirical Evidence from Litigation." *Journal of Law and Economics,* 34(2): 263–94.

Ippolito, Pauline M., and Thomas R. Overstreet, Jr. 1996. "Resale Price Maintenance: An Economic Assessment of the Federal Trade Commission's Case against the Corning Glass Works." *Journal of Law and Economics*, 39(1): 285–328.

Katz, Michael L. 1989. "Vertical Contractual Relations." In *Handbook of Industrial Organization Volume 1,* ed. Richard Schmalensee and Robert D. Willig, 655–721. Amsterdam: Elsevier Science, North-Holland.

Kreps, David M., and José A. Scheinkman. 1983. "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes." *Bell Journal of Economics*, 14(2): 326–37.

Krishnan, Harish, Roman Kapuscinski, and David A. Butz. 2004. "Coordinating Contracts for Decentralized Supply Chains with Retailer Promotional Effort." *Management Science*, 50(1): 48–63.

Laraiviere, M., and E. Porteus. 2001. "Selling to the Newsvendor: An Analysis of Price-Only Contracts." *Manufacturing & Service Operations Management,* 3(4): 293–305.

Lippman, Steven A., and Kevin F. McCardle 1997. "The Competitive Newsboy." *Operations Research*, 45(1): 54–65.

Marvel, Howard P., and James Peck. 1995. "Demand Uncertainty and Returns Policies." *International Economic Review*, 36(3): 691–714.

Mathewson, Frank, and Ralph A. Winter. 1984. "An Economic Theory of Vertical Restraints." *RAND Journal of Economics*, 15(1): 27–38.

Mathewson, Frank, and Ralph A. Winter. 1998. "The Law and Economics of Resale Price Maintenance." *Review of Industrial Organization*, 13(1–2): 57–84.

Mathewson, Frank, and Ralph A. Winter. 1984. "An Economic Theory of Vertical Restraints." *RAND Journal of Economics*, 15(1): 27–38.

Narayanan, V. G., Ananth Raman, and Jasjit Singh. 2005. "Agency Costs in a Supply Chain with Demand Uncertainty and Price Competition." *Management Science*, 51(1): 120–32.

Netessine, S., and N. Rudi. 2003. "Centralized and Competitive Inventory Models with Demand Substitution." *Operations Research*, 51(2): 329–35.

Overstreet, Thomas R., Jr. 1983. *Resale Price Maintenance: Economic Theories and Empirical Evidence.* Washington, DC: Federal Trade Commission.

Padmanabhan, V., and I. P. L. Png. 1995. "Returns Policies: Make Money by Making Good." *Sloan Management Review,* 37(1): 65–72.

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

Pasternack, Barry Alan. 1985. "Optimal Pricing and Return Policies for Perishable Commodities." *Marketing Science,* 4(2): 166–76.

Petruzzi, Nicholas C., and Maqbool Dada. 1999. "Pricing and the Newsvendor Problem: A Review with Extensions." *Operations Research*, 47(2): 183–94.

Posner, Richard A. 1981. "The Next Step in the Antitrust Treatment of Restricted Distribution: Per Se Legality." *University of Chicago Law Review*, 48(1): 6–26.

Scherer, Frederic M., and David Ross. 1990. *Industrial Market Structure and Economic Performance.* Chicago: Rand McNally.

Song, Yuyue, Saibal Ray, and Shanling Li. Forthcoming. "Structural Properties of Buyback Contracts for Price-Setting Newsvendors."

Manufacturing & Service Operations Management.

Taylor, Terry A. 2002. "Supply Chain Coordination Under Channel Rebates with Sales Effort Effects." *Management Science*, 48(8): 992–1007.

Varian, Hal R. 1980. "A Model of Sales." *American Economic Review*, 70(4): 651–59.

Winter, Ralph A. 1993. "Vertical Control and Price versus Nonprice Competition." *Quarterly Journal of Economics*, 108(1): 61–76.