Journal of Economic Theory **105**, 503–517 (2002)
doi:10.1006/jeth.2001.2890

Equilibrium Price Dispersion with Consumer Inventories*

Pilky Hong

SAP Korea; and Department of Management, KonKuk University, Seoul, Korea
pil.ky.hong@sap.com

R. Preston McAfee

Department of Economics, University of Texas, Austin, Texas 78712

and

Ashish Nayyar

PM Industrial Economics, Inc., 88 Kearny, Suite 1300, San Francisco, California 94108
anayyar@pm-industrial.com

Received June 15, 1998; final version received August 8, 2001;
published online June 11, 2002

A model with two types of consumers, shoppers and captives, is constructed that leads to an equilibrium price dispersion. Shoppers may hold inventories of the good; the level of consumer inventories leads to state-dependent price dispersions. It is shown that prices and quantities display negative serial correlation. The model is tested using grocery store data, which display the predicted correlations. *Journal of Economic Literature* Classification Numbers: D43, L13, D83, L81.

Key Words: price dispersion; inventories; sales; consumer search; grocery stores.

INTRODUCTION

In his seminal 1961 article, George Stigler [9] noted that the law of one price is empirically false; it is frequently the case that nearby stores sell identical, brand-name items for significantly different prices. Distinct prices arise for items that have no service component, and even when the shopping ambience is similar or irrelevant, as in delivered items. Stores with

* The authors thank Dan Slesnick, Fred Thum, and an anonymous referee for comments.

0022-0531/02 $35.00

consistently higher prices than nearby stores do not experience zero demand, even for identical items.

A phenomenon of equal or greater economic significance is sale prices, where prices are marked down for a short period of time. Frequently such sales may reduce prices of standardized items, such as brand name paper towels, soft drinks, or canned goods, by 50% or more. Various models exist that account for such sales, including those by Butters [2], Varian [10], Burdett and Judd [1], Fishman and Rob [4], Robert and Stahl [7], and McAfee [5, 6]. These models differ in the particular distributions of consumer information and in advertising technology, but share two characteristics. The models generate price dispersion by a dispersion of information among consumers that leads to an equilibrium randomization of price by firms, and the models are essentially static models.

None of the existing models account for two important attributes of price dispersion for many goods. First, price dispersions exist on storable goods such as paper towels. Presumably, consumer inventories should lead to a reduction in the ability of firms to randomize over time, since consumers will stock up when prices are low, leading to higher inventories. Second, casual observation suggests that there is a negative correlation of prices over time, with sales followed by relatively high prices. We document this negative serial correlation empirically. The existing price dispersion models do not account for the observed correlation of prices, for in the absence of inventories, the price distribution will remain constant over time.[1]

It is intuitive that inventories will induce a negative serial correlation of prices. Low prices today induce consumers to stock up on frequently used items. Thus, the most price-sensitive consumers are out of the market following a sale, unless the subsequent prices are very low. As a consequence, after a sale, firms will price high to extract maximum revenue from the price-insensitive consumers. The key assumption driving this result is that the "comparison shoppers," or consumers who buy from the cheapest store, are also the most likely to acquire inventories when prices are low. This assumption makes economic sense in that the most price-sensitive consumers will use both comparison shopping and inventories as a way of lowering their costs of consumption, while other consumers with a high value of time will tend to use neither method of lowering costs.

[1] Conlisk *et al.* [3] produce a negative serial correlation of prices through the Coase conjecture in a model with a durable good, which is economically similar to the case of a storable good. This model features a monopolist and produces price cycles due to the entry of new consumers but does not produce random prices, which in the present study arise only through competition.

The model follows Varian [10] and divides consumers into two types: shoppers and captives. Both types of consumers desire to consume one unit of the good per time period. As in Varian, shoppers buy from the cheapest store and captives buy from a fixed store, in both cases provided the price is not too high. Shoppers may hold an inventory, while captives are assumed not to hold inventories. Since all shoppers are assumed identical, a state variable is created by the level of shoppers' inventories. It is shown that a high level of inventories tends to lead to high prices, since it is relatively more profitable to charge monopoly prices to captives rather than seek to sell to the inventory-rich shoppers. Conversely, a low level of inventories makes attraction of more shoppers desirable, and firms will price more competitively in this case, since total sales will include not just selling the level of consumption of shoppers, but also selling for their inventories.

A significant prediction of the model is that, when inventories are high, there is a positive probability that the monopoly price will be charged (this probability may be one for some parameters). This prediction appears to be in accord with empirical fact. Most grocery store items have a "regular" price that is charged frequently and may be viewed as a monopoly price for the price-insensitive. When sale prices are charged, the sale prices vary; this is in accord with casual empiricism concerning grocery store prices. Most other price dispersion models have a smooth distribution of prices and do not exhibit a positive probability of a particular price.[2]

The next section sets out the model's assumptions and derives the economically significant conclusions. The third section summarizes an empirical analysis of grocery store pricing. The paper ends with an evaluation of the robustness of the conclusions. All proofs are relegated to Appendix 2.

THE MODEL

There are n identical firms selling a homogeneous good in each time period $t = 0, 1, 2, \ldots$. Each firm has a mass c of captive[3] consumers that will buy one unit each period from that firm, provided the price does not exceed p^m. In addition, there is a mass of s shoppers, who buy from the lowest priced firm, if they buy at all. Shoppers also value a unit of

[2] Exceptions include Robert and Stahl [7] and McAfee [5], where the mass point arises for very different reasons.

[3] While consumers who visit only a single store but stock up when prices are low may constitute an important class, the introduction of such a class complicates the model significantly. In the interests of simplicity we ignore this class of consumers.

consumption per time period at p^m. Shoppers can keep up to one unit in inventory. All agents use the discount factor δ. Production costs are set to zero for simplicity.

Formally, there are three types of agents in the model: captives, shoppers, and firms. Captives can buy or not buy in each period. Let p_t represent the price paid in period t and $\tilde{q}_t \in \{0, 1\}$ be the amount purchased by a captive; a captive's payoff is, then,

$$\sum_{t=0}^{\infty} \delta^t (p^m - p_t)\, q_t.$$

Let a shopper's inventory at time t be represented by $\imath_t$. Let y_t be the lowest price offered in period t. A shopper wishes to maximize

$$\sum_{t=0}^{\infty} \delta^t [p^m q_t - y_t(\imath_t - \imath_{t-1} - q_t)]$$

subject to $\tilde{\imath}_t \in \{0, 1\}$, $\tilde{q}_t \in \{0, 1, 2\}$, and $\imath_t \geqslant \imath_{t-1} + q_t$. The last condition ensures that the inventory holder does not attempt to resell to the firm.

Fix a firm i and let the lowest price offered by the other firms be y_t. Firm i can choose the price p_t in each time period t. A strategy for firm i is a history-dependent distribution of prices.

Let 1_A be the characteristic function of the set A. Let w_t be the average purchases of shoppers; that is, w_t is the average of $\imath_t - \imath_{t-1} - q_t$. Measurability problems are potentially created by such a formulation, because such an average need not be well defined. There are only two possible types of shoppers—those with an inventory and those without. There are at most three actions. Measurability problems can be avoided by making current inventory a sufficient statistic for a shopper's type and keeping track of inventories.

The payoff to a firm, then, is

$$\sum_{t=0}^{\infty} \delta^t p_t [c + s(1_{p_t < y_t} + \kappa 1_{p_t = y_t})].$$

If $p_t = y_t$, the firm shares the market with others; κ is intended to represent the share accruing to a firm, which of course depends in the obvious way on how many firms are sharing the market.

We look for a subgame perfect Markov equilibrium that satisfies several conditions. First, as noted above, we assume that current inventory is a sufficient statistic for a shopper's behavior. As a practical matter, we will construct an equilibrium where all shoppers do the same thing (purchase for inventory or not). Second, the state variable in the Markov equilibrium is the level of shopper inventories. Thus, the behavior of firms is predicated only on the level of shopper inventories and not on any other historically determined variable. Thus, a strategy for firm i in a Markov equilibrium is a distribution $F_\imath$ which gives the cumulative distribution of prices, given inventory $\imath$.

In this framework, shoppers will purchase a unit for inventory provided the price p does not exceed a critical level p^c. There are two possible states of the system defined by the inventory levels of shoppers and denoted by $\imath = 0, 1$ for inventories of zero or one, respectively. Buying for inventory in state zero means purchasing two units: one for consumption and the other for inventory. As noted above, we only consider Markov equilibria, where each agent's behavior depends only on the state of shopper inventories. Let $\mu_\imath$ be the mean price in state $\imath$; it will turn out that $\mu_1 \geqslant \mu_0$ since the competition to sell two units to shoppers leads to lower prices than the competition to sell a single unit to shoppers.

The equilibrium condition on p^c is

$$\delta\mu_0 \leqslant p^c \leqslant \delta\mu_1. \tag{1}$$

That inequality (1) leads to optimal shopper behavior is straightforward. If $p^c \geqslant p$, then a given shopper conjectures that all other shoppers will purchase the good for inventory, and thus that the subsequent inventory level will be one. Thus, the subsequent price will, on average, be μ_1. Since $\delta\mu_1 \geqslant p^c \geqslant p$, it pays the consumer to buy a unit for inventory now, rather than wait and pay μ_1 on average one period later. Similarly, if $p > p^c$, a shopper conjectures that the other shoppers will not purchase a unit for inventory, and thus that the subsequent state will be zero. Since $p > p^c \geqslant \delta\mu_0$, it does not pay the shopper to buy an additional unit for inventory.

Inequality (1) reveals two other facts about the model. First, there is an externality; the decisions of other shoppers influence the level of inventories, which in turn influence the subsequent period's prices. Second, this inequality generally results in a continuum of equilibria. Although the equilibrium levels of both μ_0 and μ_1 depend on p^c, they do so in a continuous way, and thus as p^c is varied, inequality (1) is satisfied for an interval of p^c values. Although these equilibria are qualitatively similar in their properties, they result in distinct levels of utility for consumers and firms, and thus they are quantitatively distinct.[4] Note that the initial state does

[4] There are other Markov equilibria created by allowing p^c to vary by state. However, letting p^c vary by state appears unduly complex when compared to the simple rule for consumer behavior dictated by a cutoff, or reservation, value for purchase for inventory, and these other equilibria are not considered further.

not affect the nature of the equilibria. Regardless of whether none, some, or all of the shoppers hold inventories in the first period of the game, by the second period all shoppers will have the same level of inventories. For instance, if the initial state is such that some shoppers hold inventories while others do not, either the shoppers with inventories will run down inventories while the remaining shoppers only purchase one unit, or the shoppers with inventories will purchase one unit while the remaining shoppers purchase two units.[5]

Firms in this model will randomize their prices, because it always pays to break ties, in the event that shoppers will purchase. Thus, in particular, in state 0, it pays to undercut another firm slightly; such undercutting will not result in price equal to marginal cost because it is more profitable to charge p^m to captives rather than charge marginal cost. Firms will not choose prices in excess of p^m because such prices produce zero sales.

We denote by F_ι the cumulative distribution of prices in state ι and denote the density, where defined, by f_ι. The present value of a firm's profits could, in principle, depend on the state, and we denote this present value by V_ι. It is straightforward to see that

$$V_0 = \begin{cases} pc + ps(1-F_0(p))^{n-1} + \delta V_0(1-F_0(p^c))^{n-1} & \\ \qquad + \delta V_1(1-(1-F_0(p^c))^{n-1}) & \text{if} \quad p > p^c \\ pc + 2ps(1-F_0(p))^{n-1} + \delta V_1 & \text{if} \quad p \leqslant p^c. \end{cases} \tag{2}$$

In state 0, if a firm chooses a price $p > p^c$, it will sell to its captives, earning pc, and sell one unit to shoppers if no other firm posts a price less than p, which occurs with probability $(1-F_0(p))^{n-1}$. The firm's future value is either δV_0 or δV_1, depending on whether some other firm posts a price less than p^c. If the firm prices below p^c, the subsequent state will be 1 for sure. In this case, the firm sells two units to shoppers when the firm has the lowest price (one for consumption and one for inventory).

Expression (2) holds with equality at all points in the support of a firm's randomization, that is, in the support of F_0. Otherwise, when p is not in the support of F_0, V_0 is at least as large as the right-hand side of (2). This construction builds in both the Bellman maximization, which sets V_0 at the maximum of the right-hand side of (2), and the fact that, to induce randomization, the firm must be indifferent to choosing all values in the support of the distribution.

[5] It may be possible to induce shoppers to randomize by choosing p^c so that they are made indifferent, which is made rational by the firms' subsequent reaction to partial inventories.

There are no mass points, except possibly at p^c and p^m. No firm would choose a price just above a mass point, because a slight price cut produces a discrete gain in the expected quantity of consumers. Thus, there is an interval above the mass point that is never selected. But this means choosing the mass point was unprofitable relative to increasing price slightly, unless a slight increase in price eliminates demand, which only happens at p^c and p^m.

The value function in state 1 is derived similarly to (2)

$$V_1 = \begin{cases} pc + \delta V_0 (1 - F_1(p^c))^{n-1} + \delta V_1 (1 - (1 - F_1(p^c))^{n-1}) & \text{if} \quad p > p^c \\ pc + ps(1 - F_1(p))^{n-1} + \delta V_1 & \text{if} \quad p \leqslant p^c. \end{cases} \tag{3}$$

The upper bound on prices is p^m, which is the most any consumer will pay. In addition, p^m is the supremum of the support of both price distributions F_0 and F_1. The reason is that, if the upper bound on F_ι were $\bar{p} < p^m$, it would be more profitable for a firm charging sufficiently close to $\bar{p}$ to instead charge p^m, since it almost surely fails to sell to the shoppers at $\bar{p}$, and makes a strictly higher profit on the captives at p^m.

Were $\bar{p}$ the maximum, a firm charging close to $\bar{p}$ has the probability of selling to shoppers given by $1 - F_0(p) \to 0$ as $p \to \bar{p}$. In other words, if $\bar{p} < p^m$, the profit function is strictly increasing on $(\bar{p}, p^m]$, which contradicts a support less than $\bar{p}$. Therefore, in both states 0 and 1, the maximum of the support of the price distributions is p^m. From this, we have

LEMMA 1. $V_0 = V_1 = \frac{cp^m}{1-\delta}$.

An immediate consequence of (3) and Lemma 1 is that in state 1, no firm charges a price in the interval (p^c, p^m). Moreover, there is a strictly positive probability of charging p^m in state 1. These observations are intuitive. In state 1, shoppers only buy if price is less than p^c. Thus, if a firm is going to charge more than p^c, there is no loss from raising price to p^m, the monopoly price. Note that while the profits are the same from charging the monopoly price to captives, the equilibrium entails prices below the monopoly price with positive probability for firms competing for the shoppers.

Lemma 1 permits the direct calculation of the price distributions in the two states. In state 1, the support of prices is $[L_1, p^c] \cup \{p^m\}$ where

$$L_1 = \frac{c}{c+s} p^m,$$

and the c.d.f. of prices, F_1, satisfies

$$(1-F_1(p))^{n-1}=\frac{c(p^m-p)}{p^s} \qquad \text{for} \quad L_1 \leqslant p \leqslant p^c. \tag{4}$$

There is a mass point at p^m, with the probability of price p^m being $1-F_1(p^c)$. The distribution F_0 is somewhat more complicated. Although it contains no mass points, it also contains a gap, since prices slightly above p^c must bring lower profit than charging p^c. Intuitively, if a firm is contemplating a price above and close to p^c, the firm is better off lowering the price to p^c to double sales to the shoppers in the event that it is the low price firm. Let M_0 be the right endpoint of this unprofitable interval, so that the probability of a state 0 price in (p^c, M_0) is zero. Then

$$(1-F_0(p))^{n-1}=\begin{cases} \dfrac{c(p^m-p)}{2ps} & \text{if} \quad L_0 \leqslant p \leqslant p^c \\ \dfrac{c(p^m-p)}{ps} & \text{if} \quad M_0 \leqslant p \leqslant p^m. \end{cases} \tag{5}$$

The value of M_0 satisfies $F_0(M_0)=F_0(p^c)$, which gives

$$M_0=\frac{2p^mp^c}{p^m+p^c}. \tag{6}$$

From $F(L_0)=0$, we obtain

$$L_0=\frac{c}{c+2s}p^m<\frac{c}{c+s}p^m=L_1.$$

Equations (4) and (5) provide explicit solutions for F_0 and F_1, which makes the following lemma a routine calculation.

LEMMA 2. *F_0 is stochastically dominated by F_1. Thus, in particular, $\mu_0<\mu_1$.*

That zero inventories produce lower prices than positive inventories is intuitive. With zero inventories, demand is more elastic—a price cut leads to larger demand, at least for price less than p^c. Thus, the firms compete harder to obtain shoppers when the shoppers have no inventory, since in this circumstance, shoppers buy more. This effect appears to be relatively robust.

The exact characterization of the properties of equilibria with dispersed prices is rather complicated to state, because there is a continuum of equilibria, which can have different properties concerning the possibility of sales to inventory in states 0 and 1. The following result, however, summarizes some of the restrictions that equilibrium behavior implies.

THEOREM 1. *There exists at least one equilibrium.*

(i) *There exist equilibria with* $F_1(p^c) > 0$, *that is, a positive probability of sales to inventory in state* 1, *if and only if* $\delta > c/(c+s)$.

(ii) *Sales to inventory in state* 0, *i.e.,* $F_0(p^c) > 0$, *can occur if and only if* $\delta > c/(c+2s)$.

If δ is too small, it is not profitable for firms to price to sell for inventory. This is because consumers demand too low a price in order to make it worthwhile to buy a unit for inventory, given their high discount rate. The critical level of discounting depends on the state. When consumers possess an inventory, the value of price-cutting is reduced relative to the case where consumers do not have an inventory. If $\delta \leqslant c/(c+2s)$ the equilibrium is unique and the state is always 0, as no sales to inventory are made. Were the state to be 1, all firms would charge p^m and the state returns and stays at 0. In state 0, no firm ever prices below p^c, so no consumer buys for inventory.

If $(c/(c+2s)) < \delta \leqslant (c/(c+s))$ in state 1 all firms charge p^m, and thus state 1 invariably leads to state 0. In state 0, however, there are equilibria where firms do sell for inventory, returning the state to 1 with positive probability.

A variety of testable predictions emerge from this model. In particular, the model predicts a negative serial correlation of prices and quantities, over time. This is intuitive, since low prices in the current period will lead to high consumer inventories, which in turn lead to high prices in the subsequent period. Moreover, this negative serial correlation is true for both individual firm prices and for the best market price.

THEOREM 2. *Individual firm prices are negatively serially correlated. Individual firm sales are negatively correlated. The quantity at a given time is positively correlated with the previous period's price, holding constant the current price.*

The model makes a variety of other predictions that are unlikely to be robust, as they depend on the two-state nature of the environment, the two types of consumers, or the fixed quantity demanded. In particular, the support of the price distribution is independent of the number of firms, and the shape of the distribution depends on the number of firms only through the simple formula embodied in (4) and (5). In addition, the discount factor influences the price distribution only through its effects on p^c, and the effect of the discount factor on p^c occurs only through the inequality (1). While these comparative statics are predictions of the model, they seem unlikely to generalize beyond the two-type model. However, the predictions summarized in Theorem 2 seem likely to be robust to changes in the model's structure, for the main requirement is that the consumers who are price-sensitive also are the consumers most likely to use inventories. Thus, when inventories are high, prices are high since the price-sensitive consumers are out of the market.

The model prohibits captives from holding inventories. Were captives able to hold inventories, n state variables can be introduced, associated with the inventories of each firm's captives. Moreover, captives will use a critical value for determining when to augment inventories which is different than shoppers. Due to the major complexity of such an analysis, we have not undertaken it here. However, we feel confident that the major qualitative results would be obtained in a model with captives holding inventories. The reason for the serial correlation in prices is that, if prices were sufficiently low last period, inventories are high this period, and thus the incentives for low prices (to attract large quantity purchases from shoppers) are reduced, leading to high prices this period. This logic is not overturned by the ability of captives to hold inventories: the same incentives with regard to shoppers hold. However, a new cost of low prices is added, because the expected payment by captives falls as the captives take greater advantage of the lower prices offered to attract shoppers.

The model is not capable of providing sensible comparative statics on welfare for three reasons. First, the quantities have been held constant, so that the dispersed pricing has no deadweight loss associated with it. Second, the model has abstracted from costs; as a consequence, the purchases for inventory are associated with zero efficiency loss. In general, purchasing for inventory is inefficient in that the costs are incurred earlier than they need be. As a practical matter, trips to the store are free in the model, but costly in reality. (Although for a single product, such costs may be realistically ignored, especially if the product under consideration is a negligible portion of the consumer's overall budget, and thus does not determine the frequency of store visits. Such an assumption is useful in that it justifies the quasilinear utility posited.) Finally, multiple equilibria present a not inconsiderable challenge to reasonable welfare comparisons, although in principle one could examine predictions for the set of equilibria. We now turn to an empirical test of Theorem 2 using grocery store data.

THE TEMPORAL PATTERN OF PRICES

We test some of the predictions of the model using data on weekly prices and quantities for 10 branded grocery products.[6] We have 49 weeks of data for both prices and quantities. Appendix 1 presents summary statistics.

We set the period to be a week not because the data come in this form but because a week seems to be the natural period, in that grocery stores advertise prices weekly, and these sale prices tend to be available for a week until the next newspaper advertising supplement comes out, often on Wednesday.

We examine four predictions of the model with these data.[7]

1. The maximum price occurs with positive probability.[8]
2. Negative serial correlation of prices.
3. Negative serial correlation of quantities.
4. Current sales depend positively on the previous period's price, holding current price constant.

All four of these predictions appear to be robust to variations in the model's structure, and they appear to depend only on the importance of consumer inventories in pricing. Moreover, each is easily testable.

Table I reports the results of our tests of the four predictions. The second column reports the proportion of the maximum price observed. The third column reports the estimated correlation coefficients and the respective t statistics for prices. The fourth column does the same for quantities. The fifth column reports the results of a regression of current quantity on current and lagged price.

[6] The products include soft drinks, fruit juice, paper towels and canned food. The products were sold by a Randalls store in Houston, Texas between 09/17/96 and 08/20/97. We thank Randalls Food Markets Inc. for generously providing the data.

[7] Our data are for products sold at a single store. Thus, we are unable to test the model's predictions with respect to the pattern of prices across firms. However, the model provides predictions with respect to the pricing of products by a single store as well.

[8] Since prices are discrete, it could be that all prices that are observed are observed a significant proportion of the time.

TABLE I

Price and Quantity Correlations

		$P_t = a + bP_{t-1} + u_t$		$Q_t = c + dQ_{t-1} + u_t$		$Q_t = e + fP_t + gP_{t-1} + u_t$	
Product	Proportion maximum price	Estimated b	t statistic	Estimated d	t statistic	Estimated g	t statistic
1	0.300	−0.174	−0.259	−0.388	−1.670	15.400	0.223
2	0.270	−0.716	−7.820	−0.653	−7.210	6.770	2.264
3	0.450	−0.428	−3.300	−0.487	−3.900	8.840	0.778
4	0.450	−0.107	−0.426	−0.577	−4.810	153.800	0.407
5	0.060	−0.161	−0.002	−0.025	0.000	56.700	2.027
6	0.040	−0.649	−6.710	−0.145	−0.077	47.300	2.741
7	0.450	−0.459	−3.800	−0.420	−3.550	263.900	2.708
8	0.430	−0.330	−2.740	−0.622	−6.250	33.600	2.156
9	0.410	−0.524	−6.070	−0.522	−5.310	85.500	3.446
10	0.450	−0.617	−8.110	−0.534	−5.870	69.200	2.202

There are several reasons why these results should be interpreted with care. First, the sample size of 49 potentially limits the power of the test. Second, the model abstracts from the supply side. However, for the products examined, constant marginal cost is not an unreasonable assumption at the retailer level. Third, there are several factors that the model does not account for. The presence of consumers who buy only from a single store but hold inventories is one such factor. In addition, the model considers only a single good. Thus, the effects of the prices of close substitutes and complements are neglected. Finally, consumers typically buy more than one product in a single trip to the grocery store. Hence, they may be concerned about the average price of a basket of goods rather than the price of any single product. This, however, does not preclude negative serial correlation of prices (and hence quantities).

CONCLUSION

The model provides an analysis of the effects of consumer inventories on price dispersion. The model formally demonstrates what casual empiricism suggests: there is a temporal dependence of prices on past outcomes even in an environment with constant marginal cost. One application of the model is monopoly pricing. In the model of Robert and Stahl [7] the monopolist advertises heavily when prices are low. Our analysis suggests that the demand-inducing effect of advertising may be overestimated because part of the increased demand, when prices are low, may be for inventory in the anticipation of subsequently higher prices. A useful extension of the model would be to consider the more general case of downward-sloping demand. The key is the elasticity of demand. The ability to hold inventories means that elasticity of demand will be higher for shoppers. This should lead to negative serial correlation of prices and quantities. Thus, the model's predictions about the behavior of prices and quantities appear to be quite robust.

APPENDIX 1

TABLE II

Summary Statistics

Product	Variable	Mean	Standard deviation	Minimum	Maximum
1	Quantity	100.45	71.12	28	429
	Price	0.65	0.11	0.35	0.72
2	Quantity	19.61	11.33	4	52
	Price	1.71	0.31	0.99	1.99
3	Quantity	67.33	28.98	24	176
	Price	1.57	0.13	0.99	1.69
4	Quantity	374.37	229.53	130	1170
	Price	0.26	0.03	0.17	0.29
5	Quantity	45.22	37.29	5	233
	Price	1.24	0.13	0.8	1.59
6	Quantity	39.24	26.76	10	191
	Price	1.26	0.11	0.99	1.59
7	Quantity	70.43	34.25	23	168
	Price	0.46	0.04	0.41	0.5
8	Quantity	233.41	114.79	58	523
	Price	3.12	0.82	1.99	3.99
9	Quantity	346.02	242.225	61	945
	Price	3.05	0.83	1.99	3.99
10	Quantity	356.16	211.46	98	922
	Price	3.13	0.83	1.99	3.99

APPENDIX 2

Proof of Lemma 1. For $p = p^m$, we obtain from (3) and (4)

$$V_0 = cp^m + \delta V_0 (1 - F_0(p^c))^{n-1} + \delta V_1 [l - (1 - F_0(p^c))^{n-1}],$$

and

$$V_1 = cp^m + \delta V_0 (1 - F_1(p^c))^{n-1} + \delta V_1 [1 - (1 - F_1(p^c))^{n-1}].$$

Subtracting,

$$V_0 - V_1 = \delta(V_0 - V_1)[(1 - F_0(p^c))^{n-1} - (1 - F_1(p^c))^{n-1}],$$

which yields $V_0 = V_1 = cp^m + \delta V_0$. ∎

Proof of Lemma 2. Since $f_1(p) = 0$ for $p \in (p^c, p^m)$, we need to show

$$F_1(p) \leqslant F_0(p), \qquad \text{for} \quad p \in (p^c, p^m).$$

But this is equivalent to $(1 - F_1(p))^{n-1} \geqslant (1 - F_0(p))^{n-1}$ for $p \in (p^c, p^m)$, which is immediate from (5) and (6). ∎

Proof of Theorem 1. Consider the function $\delta\mu_0 - p^c$, where μ_0 is viewed as a function of p^c. We have

$$\mu_0 = \int_{\frac{cp^m}{c+2s}}^{p^c} pn(1 - F_0(p))^{n-1} f_0(p)\, dp + \int_{\frac{2p^m p^c}{p^m + p^c}}^{p^m} pn(1 - F_0(p))^{n-1} f_0(p)\, dp.$$

$\delta\mu_0 - p^c$ is continuous in p^c and is positive at $p^c = 0$ and negative at $p^c = p^m$. Thus, there exists a solution to $\delta\mu_0 - p^c = 0$. This solution is an equilibrium by (1) and Lemma 2, since $\mu_1 \geqslant \mu_0$. $F_1(p^c) > 0$ if, and only if, $p^c > L_1$. This inequality is satisfiable if and only if

$$0 < \delta\mu_1 - p^c|_{p^c = L_1} = \delta p^m - \frac{cp^m}{c+s},$$

which gives the desired condition. Finally, a solution with $F_0(p^c) > 0$ exists if and only if

$$0 < \delta\mu_1 - p^c|_{p^c = L_0} = \delta p^m - \frac{cp^m}{c+2s},$$

which completes the proof. ∎

Proof of Theorem 2. Let p_t^j denote the price of firm j at time t, $\mu_\imath^j$ denote the mean price of the firm in state $\imath$ at time $t+1$, and μ^j denote the unconditional mean price of an individual firm at time $t+1$. Note, by Lemma 2, that $\mu_1^j \geqslant \mu^j \geqslant \mu_0^j$. Finally, let Y_t denote $\min_j\{p_t^j\}$. Then,

$$\begin{aligned}
&\mathrm{COV}(p_t^j, p_{t+1}^j)\\
&\quad = E[p_t^j(p_{t+1}^j - E[p_{t+1}^j])\\
&\quad = E[p_t^j(p_{t+1}^j - \mu^j) \mid Y_t \leqslant p^c]\,\Pr[Y_t \leqslant p^c]\\
&\qquad + E[p_t^j(p_{t+1}^j - \mu^j) \mid Y_t > p^c]\,\Pr[Y_t > p^c]\\
&\quad = E[p_t^j \mid Y_t \leqslant p^c]\,\Pr[Y_t \leqslant p^c](\mu_1^j - \mu^j)\\
&\qquad + E[p_t^j \mid Y_t > p^c]\,\Pr[Y_t > p^c](\mu_0^j - \mu^j)\\
&\quad \leqslant E[p_t^j \mid Y_t > p^c]\,\Pr[Y_t \leqslant p^c](\mu_1^j - \mu^j)\\
&\qquad + E[p_t^j \mid Y_t > p^c]\,\Pr[Y_t > p^c](\mu_0^j - \mu^j)\\
&\quad = E[p_t^j \mid Y_t > p^c][\Pr[Y_t \leqslant p^c](\mu_1^j - \mu^j) + \Pr[Y_t > p^c](\mu_0^j - \mu^j)] = 0.
\end{aligned}$$

The other cases are similar. ∎

REFERENCES

1. K. Burdett and K. L. Judd, Equilibrium price dispersion, *Econometrica* **51** (1983), 362–376.
2. G. R. Butters, Equilibrium distribution of sales and advertising prices, *Rev. Econ. Stud.* **44** (1977), 465–491.
3. J. Conlisk, E. Gerstner, and J. Sobel, Cyclic pricing by a durable goods monopolist, *Quart. J. Econ.* **99** (1984), 489–505.
4. A. Fishman and R. Rob, The durability of information, market efficiency and the size of firms, *Int. Econ. Rev.* **36** (1995), 19–36.
5. R. P. McAfee, Endogenous availability, cartels, and merger in an equilibrium price dispersion, *J. Econ. Theory* **62** (1994), 24–47.
6. R. P. McAfee, Multiproduct equilibrium price dispersion, *J. Econ. Theory* **67** (1995), 83–105.
7. J. Robert and D. O. Stahl, Informative price advertising in a sequential search model, *Econometrica* **61** (1993), 657–686.
8. J. Sobel, Durable goods monopoly with entry of new consumers, *Econometrica* **59** (1991), 1455–1485.
9. G. J. Stigler, The economics of information, *J. Polit. Econ.* **69** (1961), 213–225.
10. H. R. Varian, A model of sales, *Amer. Econ. Rev.* **70** (1980), 651–659.