

UC Berkeley

Postprints from Department of Agricultural and Resource Economics, UCB

Title

Sales: Tests of Theories on Causality and Timing

Permalink

<https://escholarship.org/uc/item/0076k5r6>

Journal

International Journal of Industrial Organization, 26(6)

Authors

Villas-Boas, Sofia B

Perloff, Jeff

Berck, Peter

et al.

Publication Date

2008

Peer reviewed



Sales: Tests of theories on causality and timing[☆]

Peter Berck¹, Jennifer Brown, Jeffrey M. Perloff^{*,1}, Sofia Berto Villas-Boas¹

Department of Agricultural & Resource Economics, University of California, Berkeley, United States

Received 27 June 2007; received in revised form 19 December 2007; accepted 20 December 2007

Available online 10 January 2008

Abstract

Modern theories of sales make conflicting predictions about the temporal pattern of sales, which we test using grocery scanner data. We examine both frozen orange juice, which consumers can store, and refrigerated orange juice, which is more perishable, to determine what role—if any—durability plays in the pattern of sales. We examine (1) whether sales are more frequent for national brands than for private label brands; (2) whether manufacturers, retailers, or others determine sales; (3) the distribution of prices; (4) whether sales and their effects differ for durable (frozen) or nondurable (refrigerated) products; and (5) the temporal ordering of sales. We use correlations, runs tests, probit regressions, and vector-autoregressive analyses to test our hypotheses. Some of our more striking findings are that retailers rather than manufacturers determine sales, private labels have sales as often as national brands, and that a sale of a national brand is more likely to “cause” sales of other products than is a sale of a private label product.

© 2008 Elsevier B.V. All rights reserved.

JEL classification: L11; L81; M31

Keywords: Sales; Price promotions; Causality

Does a sale—a temporary price reduction—of one product lead to a similar sale for rival products? Is the

observed pattern of sales consistent with the predictions of current theories of sales? Are sales determined by manufacturers or retailers? Do sales patterns differ between private labels and name brands or between durable and nondurable products? We answer these questions and others using price data for refrigerated and frozen orange juice products in U.S. grocery stores.

Theories of sales make conflicting predictions about the temporal pattern of sales. The few existing empirical studies of sales (e.g., Lal, 1990; Warren and Barsky, 1995; Aguirregabiria, 1999; Pesendorfer, 2002; Hosken and Reiffen, 2004; Hendel and Nevo, 2006a,b; Nakamura and Steinsson, 2007) have found mixed support for various

[☆] Perloff thanks the United States Department of Agriculture, Economic Research Service, a National Research Initiative grant, and the Giannini Foundation for support. We are grateful to Rui Huang, Jenny Lui, and Kristin Kiesel for help with our data set. We thank an anonymous referee and the editor for useful advice.

* Corresponding author. Department of Agricultural and Resource Economics, MC 3310, 207 Giannini Hall, Berkeley, University of California Berkeley CA 94720-3310, United States.

E-mail address: perloff@are.berkeley.edu (J.M. Perloff).

¹ Member of the Giannini Foundation.

theories of sales. Because we use a different type of data set than did these earlier studies, we are able to address several important issues that they could not.

To our knowledge, this paper is the first detailed empirical study of the temporal pattern of sales between brands and between retailers, and the first study of the role of product durability in determining sales patterns.² We contrast the patterns of sales for frozen orange juice, which can be stored, to that of refrigerated orange juice, which is much more perishable, to determine what role—if any—durability plays in pricing patterns. Other than [Pesendorfer \(2002\)](#), our study is the only one to examine the roles of retailers, chains, and manufacturers in determining patterns of sales. Our paper is also unusual in empirically examining hypotheses about sales from most of the existing theoretical literature. Most previous empirical papers examined only a single theory.

In Section 1, we derive testable hypotheses from the various theories of sales. In Section 2, we briefly describe our grocery scanner data. In Section 3, we use summary statistics, correlations of prices across stores, and histograms to examine three hypotheses: are private labels less likely to hold sales than national brand; is the timing of sales determined by manufacturers, chains, or retail stores; and are the observed price distributions consistent with the theoretical models of sales. In Section 4, we examine two sets of hypotheses concerning the timing and temporal ordering of sales. We start with a simple probit model of whether a sale occurs this period given the length of time since previous sales and other variables that are proxies for whether national manufacturers or local retailers determine prices. We then turn to a vector autocorrelation model to forecast a brand's price conditional on past prices of that brand and those of its rivals. We use Granger tests of temporal ordering (“causality tests”) to determine whether the sale of one brand is followed in a predictable way by the sale of another brand or its own later sales. We also simulate impulse responses to determine the magnitude of these time-series effects. We summarize our results in the last section and draw conclusions.

² In an interesting study, [Hosken and Reiffen \(2004\)](#) investigate sale patterns for a large variety of product categories using Consumer Price Index (CPI) data. However, because they use CPI data, their monthly price observations may not correspond to the same product/store combination over time, which masks differences in sales behavior across stores. We have detailed, weekly brand-specific information by store and by market, so we can examine many questions that they cannot.

1. Sales theories

We start by reviewing the best-known theories of sales and use these theories to derive testable hypotheses. Although many theories offer explanations for price changes over time, few are clearly labeled as theories of sales. One literature describes price discrimination over time with consumers varying in tastes or knowledge about prices ([Salop, 1977](#); [Salop and Stiglitz, 1982](#); [Conlisk et al., 1984](#)). In another sales literature, firms use mixed strategies to set prices ([Shilony, 1977](#); [Varian, 1980](#); [Lal, 1990](#); [Lal and Villas-Boas, 1998](#)). Most of the models in these literatures assume that there are two (or more) types of consumers with varying search costs—for example, informed consumers with no cost of search and other consumers with relatively high costs of search. Other sales theories are less easy to categorize.³

We summarize the predictions from these literatures into five sets of hypotheses. After stating each hypothesis or set of hypotheses, we briefly describe the underlying theories.

Hypothesis 1. Private label products are less likely to go on sale than are national brands.

Many people expect that national brands conduct sales more frequently than do private labels. Two theories support such a prediction, while a third theory draws the opposite conclusion.

Sales may be used to discourage purchases of a private label brand. For example, suppose that a market has a single national, brand-name product and a private label. Some customers are loyal to the brand and others are price-sensitive and will buy whichever is least expensive. If the product can be stored for n periods, the national brand could hold a sale every n th period, thereby largely eliminating the private label (depending on how often price-sensitive customers shop) without losing much revenue from its loyal customers ([Perloff and Wu, 2007](#)). [Lal \(1990\)](#) shows a similar result for duopoly national brands. This result turns on the product being storable.

An alternative explanation that does not require that the product be storable is that the private label is a near generic and is sold at a price close to marginal cost. It is

³ We do not discuss theories that cannot be tested with our weekly grocery store scanner price data. [Hendel and Nevo \(2006a,b\)](#) examine consumer inventory behavior in response to sales. [Warren and Barsky \(1995\)](#) show that sales are more likely to occur on weekends and holidays, which is consistent with their hypothesis that sales are a result of shocks to consumers shopping intensity. [Aguirregabiria \(1999\)](#) presents reduced-form evidence consistent with a model of firm joint price and inventory decisions and structural evidence to show that the frequencies of sales are positively related to cost of placing inventory orders.

not profit-maximizing for a “competitive” firm to hold sales. This argument turns on the idea that the private label is not a major brand that competes directly with the national brands, which is incorrect in many cases (e.g., Safeway’s national Select private label brand). [Raju et al. \(1990\)](#) comes to the opposite conclusion. A brand with strong brand loyalty competes with one that has less brand loyalty. If the strong brand holds a sale to attract its rival’s customers, it loses relatively more revenue from its large group of loyal consumers than would the weaker brand. Consequently, the stronger brand conducts sales less frequently than does the weak brand in equilibrium. Presumably, private labels have less brand loyalty than do national brands; hence this theory predicts private labels should hold sales more often.

Hypothesis 2. Sales are determined by manufacturers rather than by chains or retailers.

Most of the price discrimination and game-theoretic sales models explicitly or implicitly assume that manufacturers determine when sale occur. For example, the story in [Lal \(1990\)](#) that the national brands hold sales to kill off the retailers’ private labels only makes sense if the manufacturers determine when sales occur. However, in recent years, grocery retailing has become much more concentrated and hence retailers apparently have gained price-setting power relative to manufacturers. [Villas-Boas \(2007\)](#) provides evidence that retailers are more likely to determine average retail prices than are manufacturers. [Blattberg and Neslin \(1989\)](#), [Lal \(1990\)](#), [Blattberg, Briesch, and Fox \(1995\)](#), and [Pesendorfer \(2002\)](#) note that manufacturers’ trade promotions provide weak incentives for retailers to conduct sales, which retailers may ignore. Thus, we predict that chains or stores are more likely to determine sales than are manufacturers.

Hypothesis 3. Distribution of prices over an extended period of time:

A. Closed-form (possibly bell-shaped), continuous frequency,

B. Smooth histogram with a mass point at the highest price,

C. Two mass points for each major brand and a single mass point for the private label.

[Shilony \(1977\)](#) and [Varian \(1980\)](#) present static models in which sellers use mixed strategies, which results in smooth distributions of prices, Hypothesis 3A. The oligopolistic firms selling homogeneous goods set low (“sale”) prices some of the time to attract customers who have low shopping costs. If the game is replicated independently over many periods, the mixed strategies

produce price variation over time according to an explicit, continuous probability distribution. Firms cut prices solely to compete with rivals rather than to price discriminate.

A smooth distribution with a mass point at the highest price, Hypothesis 3B, can emerge from several theories, all of which require that the good be storable. [Sobel \(1984\)](#) extends the monopoly model in [Conlisk, Grestner, and Sobel \(1984\)](#) to a fixed number of sellers that produce a homogeneous good. Consumers with different preferences for a homogeneous good enter the market in each period and leave when they make a purchase. Again, sellers vary their prices over time, charging relatively high prices most of the time, but occasionally cutting their prices to sell to a large group of customers with relatively low reservation prices. His model depends crucially on consumers having different rates of time preference that are correlated with the intensity of preferences. In both the Sobel and Conlisk, Grestner, and Sobel models, a smooth pattern of price adjustment is observed. In Sobel, the firms’ prices are initially high and each firm sells to high-storage costs (“loyal”) consumers. As time passes, when potentially a large number of low storage consumers (“shoppers”) are in the market, it becomes profitable to decrease prices and to compete for those consumers. Then firm’s prices rise again beginning a new cycle.

In [Pesendorfer \(2002\)](#), some customers consume one unit of the good every period and do not store the product, while others maintain an inventory that they consume if the prices are temporarily high. Consumers who store the good only purchase the product if price drops below a certain threshold. [Pesendorfer](#) predicts a mass point at a high price.⁴

A distribution with two mass points, Hypothesis 3C, is predicted by [Salop and Stiglitz \(1982\)](#) and [Lal \(1990\)](#), both of which require that the good be storable. [Salop and Stiglitz \(1982\)](#) use a two-period model with storage to show that stores may use unannounced, random sales to induce apparently homogeneous consumers to purchase for future consumption. Customers who arrive at a low-price store buy extra units and store them for later consumption, while customers who buy at a high-price store only buy for their immediate needs. Thus, firms can successfully price discriminate by using unannounced sales. With identical firms and consumers,

⁴ In a dynamic oligopoly setting with alternating pricing decisions, [Maskin and Tirole \(1988\)](#) predict a long-run equilibrium (focal) price or Edgeworth price cycles that are similar to the price patterns in [Sobel \(1984\)](#) and [Pesendorfer \(2002\)](#) without the mass points.

we may observe a two-price equilibrium where the low-price store makes more sales than a high-price store, but both stores have the same profit.

In the infinite-horizon model of Lal (1990), national brands alternate promotions to keep private labels from stealing those consumers who are willing to switch brands. In the perfect Nash equilibrium, two national firms implicitly collude to keep out a potential entrant.⁵

Hypothesis 4. Price distributions vary with the durability of the good.

As we discussed with respect to Hypotheses 3, many of these theories about price distributions—particularly those with mass points—require that the good be durable. Thus we expect different price distributions for durable and nondurable goods.

Hypothesis 5. Timing of sales across firms:

- A. Sales patterns are random across firms,
- B. Firms lower prices simultaneously,
- C. National-brand firms alternate sales, but private labels do not have sales,
- D. The probability of a sale rises with the time since the last sale of the relevant brands,
- E. The largest brand is a pricing leader,
- F. Prices fall slowly and then jump to a high level.

In Shilony (1977) and Varian (1980), firms use mixed strategies so that sales are held randomly, Hypothesis 5A. Consequently, firms are unlikely to have sales at the same times, and prices should be neither correlated over time nor predictable (Villas-Boas, 1995). A similar prediction is made by Salop (1977), who shows that a manufacturer may set different prices for a relatively undifferentiated product under several different brand-names. Informed consumers who know the products are identical purchase the least expensive brand, while less-informed consumers may pay higher prices for other virtually-identical goods. Though his model is static, Salop notes that randomly varying the location of the low prices over time might be a feasible dynamic strategy.

In Sobel (1984), all stores may lower their price at the same time and to the same level, Hypothesis 5B. Similarly in Green and Porter's (1984) dynamic oligopoly game with imperfect information, temporary price reductions by all firms are triggered by a price signal caused by fluctuations in demand or other unknown causes. After the trigger is pulled, all firms raise their prices simultaneously.

⁵ Agrawal (1996) adds a retailer pricing decision to the manufacturing pricing decision. Lal and Villas-Boas (1998) derive a model of price promotions in the presence of multiple retailers.

In contrast, in Lal (1990), national brands implicitly collude to alternate promotions to keep private labels from stealing those consumers who are willing to switch brands, Hypothesis 5C. Consequently private labels do not have sales. An alternative explanation for a difference between national brands and private labels is that, because stores receive a larger markup for their private label products, they do not have private label sales or they coordinate them with brand-name sales. Presumably, if manufacturers are determining when sales occur, private label sales are unlikely to affect sales of brand-name products; however, sales of brand-name products may have an effect on other brand-name products.

Pesendorfer (2002) presents Hypothesis 5D. Hypothesis 5E does not emerge from any of the standard theories about sales. However, claims that large firms are price leaders are common in the popular press and in some older industrial organization literature concerning dominant firms. Various theories (Sobel, 1984; Maskin and Tirole, 1988) predict a falling price followed by a sudden jump to a higher price, Hypothesis 5F.

2. Data

We test these hypotheses using a three-year panel of weekly price observations for refrigerated and frozen orange juice from 174 grocery stores in 24 cities from Information Resources Incorporated (IRI).⁶ We focus on the last two years of this data set, 104 observations for 1998 and 1999, reserving observations from the first year of data, 1997, for lags.

We use as many stores from the entire sample as possible. For some stores, we do not have complete price series for all products in each week either because the store did not carry that product that week or due to data errors.⁷ Consequently, we restrict our sample to only those stores that have complete price series for all key orange juice products. The resulting refrigerated orange juice data sample has prices from 22 stores in 13 cities.⁸ The frozen juice prices data set includes 10 stores in Cedar Rapids IA, Eau Claire WI,

⁶ IRI reports that it selects all major chains and a random sample of other stores within their markets.

⁷ Data apparently are missing randomly, as we cannot detect any obvious pattern in the stores that are missing data, and IRI is unaware of any systematic factor.

⁸ Atlanta, GA; Boston, MA; Chicago, IL; Eau Claire, WI; Grand Junction, CO; Los Angeles, CA; Midland, TX; Minneapolis, MN; Pittsfield, MA; San Francisco, CA; Seattle, WA; St. Louis, MO; and Tampa, FL.

and Minneapolis MN. In the refrigerated orange juice sample, some, but not all of the stores, are owned by one of seven national chains. In the frozen sample, two chains each own multiple stores.

We restrict our attention to the major national brands and the stores' private label products. In the last year of our sample, 1999, the three leading national brands of refrigerated orange juice were Tropicana, with 22.2% market share by sales, Minute Maid, with 16.3% market share, and Florida Natural, with 5.1% market share (www.beverage-digest.com/editorial/991105.html). We restrict our analysis to the leading product for each brand-name: Tropicana Premium Orange Juice, No Pulp; Minute Maid Premium Original Calcium Orange Juice, Low Pulp; Florida Natural Premium Orange Juice; and each store's best-selling private label product. Among frozen concentrates, Minute Maid is the dominant brand, accounting for 40% of the market share. We track Minute Maid, Old Orchard, and each store's best-selling frozen private label product in our analysis.

For each observation, the IRI data set contains price, store identifier, and market (city) identifier. We augmented this data set by adding the corresponding monthly packing-house-door price for juice orange in dollars per box from the USDA-NASS Agricultural Prices Monthly (usda.mannlib.cornell.edu/reports/nassr/price/pap-bb) and the Consumer Price Index (CPI).

Table 1
Orange juice price summary statistics

	Mean	SD	Min	Max
<i>Refrigerated orange juice</i>				
(22 store; 13 cities; 7 chains)				
Price (\$/64 fl oz. carton)				
Tropicana	3.04	0.49	1.68	3.99
Florida Natural	2.97	0.52	1.06	3.99
Minute Maid	2.90	0.54	1.50	3.99
Private label	2.55	1.07	0.50	4.99
Minimum price across stores in a given week				
Tropicana	2.15	0.29	1.68	2.69
Florida Natural	2.02	0.57	1.06	3.49
Minute Maid	2.08	0.44	1.50	2.99
Private label	0.85	0.10	0.50	0.99
<i>Frozen concentrated orange juice</i>				
(10 stores; 3 cities; 2 chains)				
Price (\$/12 fl oz. can)				
Minute Maid	1.56	0.22	0.88	1.77
Old Orchard	1.17	0.20	0.50	1.54
Private label	1.02	0.17	0.50	1.31
Minimum price across stores in a given week				
Minute Maid	1.19	0.17	0.88	1.51
Old Orchard	0.85	0.14	0.50	1.10
Private label	0.73	0.10	0.50	0.92

3. Testing hypotheses 1–4: frequencies, correlations, and distributions

Prices for both refrigerated and frozen orange juice vary substantially over time, even though they exhibit little trend over time. Table 1 shows the mean, standard deviation, minimum price, and maximum price for various types of orange juice. The refrigerated orange juice price is expressed as dollars per 64 ounce carton, while the frozen product price corresponds to 12 fluid ounce cans of frozen concentrate. For refrigerated juice, these summary statistics were calculated over all 22 stores and all 104 weeks for a total of 2288 observations. The “minimum price” rows report the minimum price *by week* as calculated *across* 22 stores for refrigerated products and across 10 stores for frozen products. The standard deviation is for these minimum prices across the 104 weeks.

The highest price by brand was often two or three times the lowest price.⁹ The coefficients of variation of the minimum prices are comparable across the refrigerated and frozen samples.

3.1. Hypothesis 1: frequencies of sales

Most of the variation in price over time for a given product within a store is due to temporary reductions (or, in a few cases, increases) rather than due to a trend.¹⁰ Table 2 reports sale-related summary statistics for several different definitions of a sale price. In the first set of columns, a store has a product “on sale” if its weekly price is at least 25% below the mode price (across all weeks in the two-year period) for that product in that store. When “on sale” is defined as a price that is at least 35% below the mode prices, as in the second set of columns in Table 2, we see that the products are on sale much less frequently. When a sale price was defined as any price more than 50% below the mode, very few of

⁹ Variations in prices *within* single weeks (not shown) were also substantial. On average, within-week standard deviations ranged from 45¢ for refrigerated Tropicana prices to \$1.07 for private label products prices. For frozen products, average within-week standard deviations in prices ranged from 15¢ for private label juice to 19¢ for Minute Maid.

¹⁰ For many, but not all of our grocery stores, orange juice pricing is consistent with the predictions of Hosken and Reiffen (2004) and Nakamura and Steinsson (2007), who argue that the regular price remains fixed over time and the sales price is variable and occurs for short periods of time. These results are also consistent with Dutta, Bergen, Levy's (2002) study of orange juice pricing that there is relative price rigidity and “apparent sluggish and incomplete response of prices to nominal shocks.

Table 2
Orange juice sales summary statistics (percent of weeks on sale by store)

	25% sale			35% sale			50% sale	
	Mean	SD	Min	Max	Mean	SD	Mean	SD
<i>Refrigerated orange juice</i>								
Tropicana	5.85	7.85	0	39.4	0.97	1.81	0.06	0.32
Florida Natural	5.87	7.82	0	29.8	1.83	3.02	0.55	1.30
Minute Maid	3.46	6.31	0	28.9	1.45	3.83	0.06	0.24
Private Label	5.66	7.87	0	40.4	3.92	7.05	1.38	3.49
<i>Frozen orange juice</i>								
Minute Maid	4.8	3.77	0	13.5	3.18	3.07	0	0
Old Orchard	5.39	9.9	0	49	3.78	8.21	0.10	0.29
Private label	10.55	8.14	0	30.8	4.53	5.85	0.20	0.59

A “25% sale” occurs if a store’s price drops by at least 75% below the store’s mode price.

A “35% sale” occurs if a store’s price drops by at least 65% below the store’s mode price.

A “50% sale” occurs if a store’s price drops by at least 50% below the store’s mode price.

the products are “on sale” and the price of frozen Minute Maid is never more than 50% below the mode store price. By any of these definitions, there are a few stores that never have sales and many that have sales that last no more than two weeks per year.

Contrary to the common belief in Hypothesis 1, national brands do not go on sale more frequently than do private label products. Using the 25% criterion for a sale, refrigerated Florida Natural is on sale in approximately 5.87% of the weeks in the sample. Tropicana and the private label products are on sale only slightly less frequently: 5.85% and 5.66% of the weeks, respectively. Minute Maid is on sale only 3.46% of the time.

Results are similar for the frozen products. Minute Maid is on sale relatively infrequently (4.8% of weeks), while private label frozen orange juice is on sale more than 10% of the time. Thus, we do not find a distinction between more and less durable products as one might expect given theories such as those of Sobel (1984) and Conlisk, Gerstner, Sobel (1984).

To summarize:

Result 1. Private labels go on sale roughly as frequently as national brands for refrigerated orange juice, but more frequently for frozen orange juice.¹¹

¹¹ This result is consistent with the prediction of Raju et al. (1990), whose empirical results support their theory for 31 product pairs out of 59 product categories in one store.

3.2. Hypothesis 2: correlations of sales across stores

Our second hypothesis holds that manufacturers rather than chains or individual stores determine the timing of sales. Presumably, if sales are determined independently by each store, sales will not be correlated across stores in different cities, although they might be correlated within a city if stores react to each others’ sales. If sales are determined by chains, then we would be more likely to see correlations among stores in the chain than across other stores. Finally, if manufacturers determine sales, we would be more likely to see the timing of sales correlated across all stores.

We calculate the correlations between prices at the national, market, and chain levels and report summary statistics of those correlations in Table 3. First, for every pair of stores in our sample, we found the correlation coefficient between prices for the same brand and type of juice in the store pairs. Then we computed the mean national-level correlation coefficient as the average (for each brand) of the store pair coefficients. The standard deviation, minimum, and maximum were computed in the same way. For the market level statistics, we computed the correlation coefficients for store pairs in the same market. For the chain statistics, we computed the correlation for the store pairs in the same chain. The national level provides the broadest or largest sample, capturing the correlations between prices of a given product across all stores. The market level statistics groups stores by city, and the chain statistics summarizes correlations across stores belonging to the same chain.¹²

There is surprisingly little price correlation at the national level, with the correlation coefficients averaging between 0.30 and 0.45 for the various brands of refrigerated orange juice. The correlation coefficients for frozen orange juice are even lower. Indeed, the correlation across all the refrigerated private label products at the national level is not much lower than the correlation coefficient for the national brands, even though the private labels prices presumably are set independently across chains or retailers. For frozen products, the private label and Minute Maid

¹² Suppose there are four stores A, B, C, and D. Stores A, B, and D are located in the same market, and B and D belong to the same chain. The national mean correlation for Tropicana would be the average of correlation coefficients for Tropicana prices in A and B, A and C, A and D, B and C, B and D, and C and D. The market statistics would be calculated from the correlation coefficients for Tropicana prices in A and B, A and D, and B and D. The chain correlation average would be the correlation coefficient between Tropicana prices in B and D.

Table 3
Refrigerated orange juice price correlations at the national, market, and chain levels

	Mean	SD	Min	Max
<i>Refrigerated orange juice</i>				
National level				
Tropicana	0.33	0.22	0.00	1.00
Florida Natural	0.34	0.23	0.00	0.98
Minute Maid	0.42	0.22	0.00	0.92
Private Label	0.30	0.23	0.00	0.86
Market level				
Tropicana	0.63	0.34	0.11	1.00
Florida Natural	0.61	0.34	0.05	1.00
Minute Maid	0.64	0.33	0.05	1.00
Private Label	0.52	0.36	0.01	1.00
Chain level				
Tropicana	0.82	0.17	0.62	1.00
Florida Natural	0.77	0.18	0.58	1.00
Minute Maid	0.78	0.24	0.38	1.00
Private Label	0.66	0.31	0.17	1.00
<i>Frozen orange juice</i>				
National level				
Minute Maid	0.28	0.26	0.00	0.87
Old Orchard	0.49	0.18	0.09	0.97
Private Label	0.33	0.23	0.00	0.82
Market level				
Minute Maid	0.38	0.25	0.13	0.73
Old Orchard	0.58	0.13	0.39	0.83
Private Label	0.34	0.29	0.06	0.73
Chain level				
Minute Maid	0.37	0.15	0.16	0.56
Old Orchard	0.60	0.14	0.45	0.83
Private Label	0.37	0.10	0.22	0.51

Average, standard deviation, minimum and maximum of the correlation coefficients. Between stores that are at the national level, in the same market, or in the same chain.

have similar correlations, while Old Orchard has a higher coefficient. Thus, it seems unlikely that manufacturers are determining when sales occur at a national level.

The average market level correlation for both refrigerated and frozen orange juice is substantially greater than the national-level correlation. The correlations are significantly higher at the chain level than at the city level for refrigerated juice, whereas the correlations are roughly equal for frozen juice. For refrigerated juice, two of several chains had prices that were perfectly correlated across the stores of the chain. For the remaining five chains, the correlations across the stores of the chains were above 0.5 for national brands. Thus, it seems that refrigerated prices are set by the chains with some local variation while frozen juice prices are set at either the city or chain level.

Result 2. Manufacturers do not determine the timing of sales for refrigerated or frozen orange juice.

3.3. Hypotheses 3 and 4: price distributions

To investigate our third set of hypotheses concerning the distribution of prices of a given product over time, we examined each store's histogram of 1998 prices. To illustrate the various patterns, we display histograms from two typical stores. The top half of Fig. 1 shows the price distribution for each refrigerated brand in a typical store in Grand Junction, Colorado.


Tropicana's price distribution may be (loosely) described as uniform, as several prices occur with roughly equal frequency in the dataset. The distributions are virtually unimodal for Florida Natural and Minute Maid, while almost all of the weight is in two adjacent prices (essentially a single mass point rather than a bimodal distribution) for the private label. The bottom half of Fig. 1 shows three distributions for frozen orange juice in another typical store in Cedar Rapids, Iowa. Minute Maid has a nearly unimodal distribution, the private label is basically bimodal, and Old Orchard has two adjacent dominant prices.

For refrigerated orange juice, there are 4 brands in 22 stores, or 88 price distributions. Of these 88 price distributions, 26 had a bimodal pattern, 31 had one price virtually every week, 20 had a strong central tendency (roughly normal distribution), and 4 appear to be uniformly distributed (several prices occurred with virtually-identical frequencies).


There are 30 frozen orange juice price distributions: 3 brands in 10 stores. Of these, 5 had a clear bimodal pattern, 19 had essentially one price all the time, 3 exhibited a strong central tendency, and 2 had a relatively uniform distribution of prices.

Fewer than one-third of the refrigerated and one-sixth of the frozen price distributions are possibly consistent with stores pursuing mixed strategies in pricing, which implies that prices should not remain constant over long periods of time and should lack mass points, Hypothesis 3A. Similarly, Villas-Boas (1995) tests and rejects Varian's (1980) closed-form distribution hypothesis. Slade (1994) and Pesendorfer (2002) present strong evidence against sales being random over time (Hypothesis 3A and 5B).

There is virtually no evidence supporting Hypothesis 3B that the distribution is smooth except for a mass point at the highest price. None show exactly that pattern, fewer than half of the refrigerated product histograms show a single mass point, while two-thirds of the frozen product price histograms were unimodal.



Note: Values are 1998 prices for 64 fl. oz. cartons of refrigerated orange juice.



Note: Values are 1998 prices for 12 fl. oz. containers of frozen orange juice concentrate.

Fig. 1. Histograms of orange juice prices in representative stores.

Likewise, we find only a little evidence supporting the alternating sales, Hypothesis 3C, which should lead to a distribution with two mass points. Fewer than a third of the stores for refrigerated orange juice and a sixth for frozen orange juice exhibit the bimodal pattern

(two, nonadjacent peaks) that one would expect if sales were alternated with a regular price. Thus, the price histograms provide descriptive evidence that no single model of sales describes pricing patterns across all stores.

Result 3. Price distributions differ in no obvious pattern by brand and by store. Little support was found for any one of the three clearly predicted price distributions.

We also find little support for Hypothesis 4 that pricing patterns should differ by the durability of the good. A comparison of the histograms for refrigerated and frozen orange juice indicates that sale price patterns do not vary systematically by product storability.

Result 4. We do not find any systematic difference in the price distributions between storable (frozen) and nonstorable (refrigerated) orange juice products.

4. Testing hypothesis 5: timing of sales

Hypothesis 5 contains a variety of predictions about the timing of sales. By direct observation (and the evidence about price distributions from above), we can reject Hypothesis 5F (Sobel, 1984) that prices slowly fall then jump to the highest price—even for the durable product.

We take three approaches to examining the remaining hypotheses. First, we use runs tests and summary statistics to examine the first three hypotheses that sales are random (Hypothesis 5A), occur simultaneously across firms (5B), or the national brands alternate (5C). Second, we estimate a probit model, similar to the one presented by Pesendorfer (2002), that attempts to predict sales based on the time since a previous sale (5D). Then, we turn to a vector autoregressive (VAR) model to look more formally at several of the hypotheses.¹³

4.1. Hypotheses 5A–5C: runs tests and counts

Our previous examination of price distributions casts doubt on Hypothesis 5A that sales are random. Similarly, Villas-Boas (1995), Slade (1994), and Pesendorfer (2002) present strong evidence against sales being random over time.

To determine if products go on sales simultaneously, Hypothesis 5B, we counted them. As Appendix A shows, multiple sales occur; however, they are not very frequent. Those that occur are not always between the same brands. Moreover, there is no obvious difference between the durable and nondurable goods as some theories suggest. Finally, multiple brands on sale fre-

quently involved the private label. For the refrigerated orange juice, there were 2288 store-week observations. In 613, a sale occurred. There were 181 multiple sales (30% of all sales), of which 104 involved the private label (57% of multiple sales). One out of six times that a brand-name product was on sale, the store also had its private label on sale.

The Wald–Wolfowitz runs test provides a nonparametric approach to testing for alternating sales, Hypothesis 5C. Consider two brands, Brand 1 and Brand 2. Suppose Brand 1 has the first sale, which we code as a 1 one in our list of sales. If Brand 2 goes on sale before Brand 1’s sale ends, we do not count that event as a separate sale. After the first brand’s sale ends, we find the next sale. We record a 1 if it is Brand 1, and a 2 if it is Brand 2. We continue coding through all the sales for each store.

The null hypothesis is that sales are random. We can calculate the probability that a brand is on sale at a store based on the frequency that brand is on sale in during the entire sample period. In the Wald–Wolfowitz runs test, we examine whether the number of runs—consecutive sales in our list of sales of any number by a given brand—could be generated randomly.

In general when we compare pairs of brands, we found the evidence quite mixed. For refrigerated orange juice, Tropicana and Minute Maid had 109 runs out of a maximum of 129 runs. The expected number of runs if the process is random is 82. Thus, we reject the null hypothesis of random sales with a $z=4.5$. Similarly, Tropicana and Florida Natural had 162 runs out of 227 possible alternating sales with 115 expected runs, so we reject the null ($z=6.2$). On the other hand, if we compare Tropicana to the combined “brand” of Florida Natural or Minute Maid, we also find excessive runs ($z=7.6$). So Tropicana is more likely to alternate sales with the other two brands than would be predicted by a random process. Of the private label brand pairs, only Tropicana has a slight excess of runs ($z=2.4$) and there are no excess runs between Florida Natural and Minute Maid ($z=0.5$).

For frozen orange juice, the results are less supportive of a theory that brands alternate sales. The national brands Minute Maid and Old Orchard had significantly too few runs ($z=-2.2$), Minute Maid and the private label had too many ($z=4.1$), and Old Orchard and the private label had the expected number.

4.2. Hypothesis 5D: probit analyses

Past sales may “cause” current sales. That is, sale pricing of a certain product may be due to sale of another

¹³ Dutta, Bergen and Levy (2002) similarly examine the orange juice market, though they do not focus on sales. They provide a detailed time-series analysis of the wholesale and retail prices of refrigerated and frozen orange juice. One key difference between their time series analysis and ours is that they abstract from cross-product price effects, which is the focus of our research.

Table 4
 Probit: probability of sale of refrigerated orange juice (price less than 75% of the mode price)

	Sale of product			
	Tropicana	Florida Natural	Minute Maid	Private Label
<i>Weeks elapsed since the last sale for product in a given store</i>				
Tropicana	-0.026* (0.007)	-0.014 (0.008)	-0.033* (0.013)	-0.005 (0.009)
Florida Natural	-0.012 (0.012)	-0.015 (0.012)	-0.029* (0.014)	0.008 (0.015)
Minute Maid	0.011 (0.008)	-0.003 (0.007)	-0.005 (0.008)	0.015 (0.011)
Private Label	0.031* (0.013)	-0.012 (0.012)	0.002 (0.011)	-0.041* (0.015)
<i>Weeks elapsed since the last sale for product in a given market</i>				
Tropicana	0.028* (0.011)	0.018 (0.013)	0.045* (0.017)	-0.003 (0.023)
Florida Natural	0.010 (0.012)	-0.001 (0.013)	0.033* (0.014)	-0.013 (0.016)
Minute Maid	-0.022* (0.009)	0.004 (0.009)	-0.002 (0.009)	-0.017 (0.015)
Private Label	-0.031* (0.013)	0.003 (0.014)	0.001 (0.011)	0.020 (0.017)

Reported values are marginal effects. Number of observations=2288. Marginal effects for the store and weekly dummies are not reported. Asymptotic standard errors are in parentheses. An “***” indicates that we can reject the null hypothesis at the 5% level.

product in the same store or chain, or across stores, chains and cities. To determine the probability that a sale of one brand “causes” a sale in another brand within a store, we estimate a probit model for each product where the dependent variable equals one if the product in a store in a given week is on sale. Again, we define a sale as occurring when a store sets a price that is at least 25% below the mode price in that store.¹⁴

We use probit regressions to examine the occurrence of sales for the four refrigerated and three frozen orange juice products. Thus, for each specification, we run seven probits. For each product, such as Tropicana refrigerated orange juice, an observation in our probit regressions is a particular store in a particular week. The dependent variable is one if there is a sale in that store that week for that orange juice product. In each probit, we assume that the error terms are uncorrelated across time. Our chief explanatory variables are the number of weeks since the last sale of this product and the last rival product sales in this store and the minimum number of weeks since a sale in this market (city) for each product.

¹⁴ We obtain virtually identical results when we use the 35% threshold. The 50% definition of a sale results in too few sales for this analysis to be enlightening.

The latter measure is included to capture the effect of incentives provided by the manufacturer that affect all stores within a city. Our results are virtually the same if we replace these city-level measures with the minimum number of weeks since the last sale in the entire sample for each product.

Thus, each refrigerated probit equation eight such independent variables: the number of weeks since the last sale in that store for each of the four products and the number of weeks since the last sale in that market for each of the four products. There are only six such variables in the frozen regressions, because there are only three frozen products.

Each equation also has two types of dummy variables. Weekly dummies control for the possibility of national-level pricing, nationwide promotional campaigns, or seasonality in the cost of oranges or other critical inputs. Store dummies capture systematic differences in the use of sales across stores. We consider three specifications where we include only the weekly dummies, only store dummies, or both. When we examine only refrigerated products and include both weekly and store dummy variables, none of the dummy variables was statistically significantly different from zero. However, in the frozen product regressions, some dummies were statistically significantly different from zero.

In regressions that include only the weekly dummies, none of these dummies is statistically significantly different from zero for refrigerated orange juice. Thus,

Table 5
 Probit: probability of sale of frozen orange juice (price less than 75% of the mode price)

	Sale of product		
	Minute Maid	Old Orchard	Private Label
<i>Weeks elapsed since the last sale for product in a given store</i>			
Minute Maid	0.006 (0.005)	-0.007 (0.003)	-0.006 (0.003)
Old Orchard	-0.031* (0.008)	-0.006 (0.004)	-0.003 (0.004)
Private Label	-0.003 (0.010)	0.009 (0.006)	-0.007 (0.006)
<i>Weeks elapsed since the last sale for product in a given market</i>			
Minute Maid	0.000 (0.008)	-0.012* (0.006)	-0.001 (0.006)
Old Orchard	0.041* (0.009)	-0.002 (0.007)	0.010 (0.006)
Private Label	0.012 (0.016)	-0.022 (0.011)	-0.013 (0.013)

Reported values are marginal effects. Number of observations=1040. Marginal effects for the store and weekly dummies are not reported. Asymptotic standard errors are in parentheses. An “***” indicates that we can reject the null hypothesis at the 5% level.

we conclude that no time dependent, national phenomenon influences refrigerated orange juice prices. For the frozen product, a few weekly dummies were statistically significant for Minute Maid, suggesting that Minute Maid may engage in national promotional activities in a few weeks of the year. When only store indicators were included in the regressions, the dummies were statistically significant only in the refrigerated and frozen Minute Maid regressions.

Tables 4 and 5 report the probit marginal effects (at the means) and asymptotic standard errors for refrigerated and frozen products, respectively, where both sets of dummy variables are included. If the coefficients in the regression other than the constant term were zero, then sales are IID events, and this model would provide no useful prediction of sales. In fact, the likelihood-ratio (LR) test for this hypothesis is rejected in all cases.

This model does not provide the type of results one might expect. In both the refrigerated and frozen orange juice probit regressions, very few of the coefficients for weeks since the last sale in the own store are statistically significantly different from zero: 4 out of 16 for the refrigerated sample (2 were for the same given product) and 1 out of 9 for the frozen (and it was not for the same given product). All of the statistically significant coefficients are negative. Thus, either previous sales have no effect on current sales or the longer it has been since a sale in a store, the less likely it is that a sale will occur this week. Of course, if a store rarely held sales in the past, it may be relatively unlikely to hold sales now.

The probit regressions also include a variable representing the number of weeks since a given product was on sale in any store in the given market. For the minimum number of weeks since the last sale in the market, only 5 out of 16 coefficients are statistically significant for the refrigerated sample, and only 2 out of 9 are significant for the frozen sample. In the refrigerated product regressions, three of the five statistically significant coefficients are positive and the rest are negative. For the frozen sample, one is positive and one is negative. In short, most coefficients are statistically insignificant, and no obvious sign pattern emerges for the statistically significant coefficients.

Pesendorfer (2002) estimates a similar model for ketchup. His results for ketchup are very different from our results for orange juice. Pesendorfer reports positive coefficients for the variables representing time since the last sale for both a store and the market. Thus, he concludes that there is a systematic pattern of sales that is inconsistent with the lack of a pattern predicted by Varian's mixed strategy story.

Our results do not support Pesendorfer's conclusion. We find a very mixed pattern where past sales with a store or a market may encourage, inhibit, or have no effect on sales. We do not believe that our results differ from Pesendorfer's because of differences in our models, as we have replicated his specification as closely as possible.¹⁵ However, our results may differ from his because we study different products.

According to Nakamura and Steinsson's (2007) theoretical model of price rigidity, we might expect asymmetric retail price adjustments in response to supply shocks. Firms may lower price when cost falls, but hesitate to raise price above a certain price cap when costs rise. Thus, it is possible that there are asymmetric price adjustments. To examine this possibility, we divided the data into two periods where the monthly ("near term") orange juice futures price (U.S. Department of Agriculture) were rising or were falling. We then re-estimated the probit models by including a dummy when the future price increased from the previous month or by estimating two separate probits for the two subperiods. The differences in the coefficients between the two periods were not statistically significant. Thus, we fail to confirm Hypothesis 5D that the probability of a sale increases with the time since the last sale.

4.3. Hypotheses 5A–5F: temporal ordering of sales

Our probit analysis investigates the question: Do the prices of one brand "cause" changes in the prices of another brand? That is, does the earlier sale of one brand consistently precede or "cause" the sale of another? Rather than use a probit model to forecast the probability of a sale given the timing of past sales, we recast the problem as a vector autoregressive (VAR) process. By doing so, we can forecast price conditional on past prices of a given brand and those of its rivals. Given estimates of a vector autoregressive model, we test the hypotheses that past prices of a given brand and its rivals "cause" the present price of a brand. We then use these estimates to forecast product prices taking into account of all the feedback effects between brands. For instance, our vector autoregressive model—but not our

¹⁵ Pesendorfer (2002) uses time trends and squared time trends instead of our weekly fixed effects. He also includes the square of the minimum number of weeks since last sale. We replicated this specification with our data, but did not get a significant coefficient on the squared term of weeks since last sale. Unlike Pesendorfer, we use levels rather than logs of the time variables. With logs, none of the probit coefficient estimates were statistically significantly different from zero.

probit model—can detect the pattern that Brand 1 has a sale that “causes” Brand 2 to have a sale which in turn causes Brand 1 to have a sale, Hypothesis 5B. We use our VAR analysis to investigate “Granger tests of temporal ordering.” Following the standard but inaccurate terminology, we refer to these as Granger causality tests.

Before we conducted our analyses, we tested for and found little evidence of trends, and we tested for unit roots (see Appendix B). Because deflating prices by the CPI makes no appreciable difference in this two-year period, we report all statistics using nominal prices. We estimate vector autoregressive (VAR) models for both prices and the log of prices of each brand within a store, separately for refrigerated and for frozen products.¹⁶ Because the results are qualitatively similar, we discuss only the log price results.

We estimate a refrigerated and a frozen products VAR systems for each store. For refrigerated orange juice, we estimate a four-equation VAR, one equation for each product. For frozen, the VAR system has three equations. A typical refrigerated regression has the logged price of one refrigerated (frozen) product as the dependent variable and the lagged logged prices of all the refrigerated (frozen) products as independent variables. We also include quarterly dummy variables, current packing-house-door price for juice oranges, and the current CPI. Based on the Schwarz Information Criterion (SIC), we use eight-week lags.¹⁷

Based on likelihood-ratio tests, we reject the null hypothesis that all slope coefficients are equal to zero in 83 of 88 regressions for refrigerated prices and for 28 of 30 cases for frozen prices. That is, there are only a few stores for which the prices are best explained as random variation about a constant term. Thus, we again reject Hypothesis 5A that sales patterns are random. The average R^2 measures for the refrigerated juice VARs are 0.70 (standard deviation is 0.17, minimum is 0.36, and maximum is 0.96) for Tropicana, 0.71 (0.20, 0.31, 0.99) for Florida Natural, 0.79 (0.15, 0.50, 0.96) for Minute Maid, and 0.72 (0.21, 0.41, 0.98) for the private labels.

¹⁶ We do not difference. One might think that if a series has a unit root, it is better to first difference the variable before using it in the VAR. However, taking the differences first does not improve prediction if a series has a unit root, and differencing decreases prediction accuracy in the absence of a unit root.

¹⁷ The SIC procedure suggested eight, eight, seven and six week lag lengths for Tropicana, Florida Natural, Minute Maid, and the private label juice, respectively. Given that our results for frozen products were similar, we used an eight-week lag in all of these regressions. As an experiment, we tried 10- and 12-week lags and found that our results were virtually unchanged.

The corresponding average R^2 measures for frozen products are 0.49 (0.12, 0.23, 0.64) for Minute Maid, 0.65 (0.14, 0.46, 0.80) for Old Orchard, and 0.54 (0.18, 0.29, 0.85) for the private labels.

We also examined the variance–covariance matrix of the residuals of these VARs to determine if firms lower their prices simultaneously, Hypothesis 5B. For a pair of products, we test whether the residuals in one equation (e.g., Tropicana’s) are independent of the residuals in another equation (e.g., Minute Maid’s). If the prices of the two products move together, the residuals should be correlated. We test the null hypothesis that the correlation coefficient of the residuals is zero using a 5% criterion. For refrigerated orange products, the Tropicana and Florida Natural pair and the Tropicana and private label pair have the highest correlations. In both of these cases, we reject the null hypothesis of no correlation of the residuals for 30% of the stores. The other equation pairs for refrigerated products have a lower percentage of stores in which the null hypothesis is rejected. Indeed, in only 10% of the stores can we reject the null hypothesis that the Minute Maid and the private label residuals are uncorrelated. For frozen orange juice, we reject the null in at most 40% of the stores for the Minute Maid and private label pair. The Minute Maid and Old Orchard pair has a 20% rejection rate, while the Old Orchard and private label pair has a 10% rate. Thus, as with our

Table 6
Summary of Granger causality test results ($\alpha=5\%$)

Refrigerated orange juice			
Price of	Granger-causes	Price of	% of stores
Tropicana	→	Florida Natural	63.64
	→	Minute Maid	36.36
	→	Private label	50.00
Florida Natural	→	Tropicana	50.00
	→	Minute Maid	36.36
	→	Private label	45.45
Minute Maid	→	Tropicana	36.36
	→	Florida Natural	40.91
	→	Private label	45.45
Private label	→	Tropicana	31.82
	→	Florida Natural	31.82
	→	Minute Maid	27.27
Frozen orange juice			
Price of	Granger-causes	Price of	% of stores
Minute Maid	→	Old Orchard	60.00
	→	Private label	60.00
Old Orchard	→	Minute Maid	60.00
	→	Private label	30.00
Private label	→	Minute Maid	30.00
	→	Old Orchard	20.00

counting data analysis, the evidence does not strongly support the hypothesis that sales are simultaneous.

4.4. Granger causality tests

Table 6 summarizes our Granger causality tests. We say that Brand 1 “Granger-causes” Brand 2’s price if we reject the hypothesis (at the 5% level) that the VAR (lagged price) coefficients for Brand 1 in the equation for Brand 2’s price are collectively zero using a likelihood-ratio test. The first row of Table 6 shows that Tropicana’s refrigerated orange juice price Granger-causes the Florida Natural price in 64% of the stores. The refrigerated private label price Granger-causes the brand prices 27% to 32% of the time, while the prices of the national brands Granger-cause the prices of the other brands and private label between 36% and 64% of the time. Tropicana—the brand with the largest market share—has the largest effect on the prices of rival products, followed by Florida Natural, and Minute Maid.

For frozen orange juice brands, Minute Maid’s price Granger-causes the price of Old Orchard and the private label 60% of the time, the Old Orchard price Granger-causes Minute Maid’s price 60% of the time but private label’s price only 30%, and the private label price Granger-causes Minute Maid’s price 30% and Old Orchard’s price 20% of the time.

A necessary condition for Hypothesis 5C (firms alternate sales) to be true is that one national brand Granger causes another. There is a high level of mutual Granger causality for Tropicana and Florida Natural, but not for Minute Maid among refrigerated products.

Among frozen products, Old Orchard and Minute Maid frequently cause price changes in each other. While an examination of the raw data shows that private labels have sales (Table 2), their sales are often Granger caused by the major labels and infrequently cause them. Taken together, this evidence does not support 5C that national brands have alternating sales and private firms do not have sales.

All the large brands Granger cause the pricing of other brands in some, but not all stores. Thus, there is no clear evidence of a single price leader (Hypothesis 5E).

4.5. Impulse response functions

While the bivariate Granger causality tests above provide a statistical test of whether one brand’s price is useful in predicting another brand’s price, they do not show the degree of the response. Consequently, we turn to an impulse–response analysis to show how much a brand’s price changes over time as a function of a change in the price of either itself or another brand. We compute impulse–responses for each VAR and construct graphical representations of the impulse and responses to show how a set of prices reacts over time to a price shock. The response takes account of changes throughout the system. For example, if the Tropicana price “changes” the Minute Maid price and the Minute Maid price changes the private label price, but Tropicana does not directly change the private label price, then these indirect price effects will appear in the impulse–response but not in the bivariate Granger tests. Because our VAR estimation uses the log of product prices, impulses and




Fig. 2. Frozen orange juice impulse–response.

Table 7
Number of stores in which impulse response is greater than half of the impulse

Refrigerated orange juice		Stores with Response > 1/2 Impulse	
Impulse brand	Response brand	Number out of 22	Percentage
Tropicana	Tropicana	5	23
Tropicana	Florida Natural	4	18
Tropicana	Minute Maid	2	9
Tropicana	Private label	6	27
Florida Natural	Tropicana	2	9
Florida Natural	Florida Natural	10	45
Florida Natural	Minute Maid	2	9
Florida Natural	Private label	6	27
Minute Maid	Tropicana	6	27
Minute Maid	Florida Natural	7	32
Minute Maid	Minute Maid	13	59
Minute Maid	Private label	9	41
Private label	Tropicana	2	9
Private label	Florida Natural	2	9
Private label	Minute Maid	3	14
Private label	Private label	8	36

Frozen orange juice		Stores with Response > 1/2 Impulse	
Impulse brand	Response brand	Number out of 10	Percentage
Minute Maid	Minute Maid	1	10
Minute Maid	Old Orchard	2	20
Minute Maid	Private label	2	20
Old Orchard	Minute Maid	1	10
Old Orchard	Old Orchard	3	30
Old Orchard	Private label	0	0
Private label	Minute Maid	1	10
Private label	Old Orchard	1	10
Private label	Private label	2	20

responses can be interpreted as percentage changes in price.

Fig. 2 provides a sample of the impulse responses for frozen orange juice in a single store. It shows the response of prices of frozen Minute Maid, Old Orchard, and a private label product to a one-standard-deviation (of the two-year log price series for a given product in a given store) upward shock to the Minute Maid price. The figure illustrates that a Minute Maid price increase of 7% leads to a price increase of approximately 3% three weeks later in the private label. The maximum price response from Old Orchard is an increase of less than half of 1%.

The impulse responses for frozen and refrigerated orange juice prices are summarized in Table 7. We test whether the effects are statistically significantly different from zero using asymptotic standard errors.

The table shows the number of stores that experienced at least one statistically significant (at the 5% level) price responses exceeding 50% of the impulse within 10 weeks of the shock. For each product, store, and week, we test the null hypothesis that a response equals zero (i.e. no response using a 5% criterion). When the test led to rejection of the null and the response was greater than 50% of the impulse value, we included the event in our count. For example, if the price of Tropicana in a given store in given week changed by more than 6% after a 12% increase in the price of Minute Maid in that store, then we would count that impulse–response in the table.

In general, the price responses between brands within a single store are weak. The strongest effects are own-brand price effects. By examining the 88 impulse response figures, we find that Florida Natural and Minute Maid, in 10 and 13 stores respectively, have a pattern where a decrease in a product's own-price leads first to an own-price increase and then a decrease.

The strongest effect on another refrigerated brand was the effect of the price of refrigerated Minute Maid on the private label price. A Minute Maid price decrease led to a price decrease of at least half the magnitude in the private label product in 41% of the stores. Price changes for Tropicana, the refrigerated orange juice market leader, have their greatest effect on private label (27% of stores). The private label has little effect (at most, in 14% of stores) on any of the name brands.

For frozen orange juices, again the largest effect is an own effect. An Old Orchard sale results in another sale of at least half the first sales magnitude within 10 weeks in 30% of the stores. The largest cross effect is that of the market leader, Minute Maid, on both Old Orchard and the private label, where 2 stores showed a price response of 50% of the impulse. To summarize, for both frozen and refrigerated orange juice, a sale of one brand of orange juice triggers a price response by other brands of at least half the magnitude of the original sale in fewer than half the cases.

Turning again to Hypothesis 5E, price leadership, and 5C, alternating sales, we get a clearer picture than that given by the causality tests. When we look at how many stores have a response that is half or more of the impulse (rather than the statistically detectable response in the causality tests), we find that for the refrigerated product, Minute Maid, has the largest effect on the other brands and the other brands have little effect on Minute Maid. Thus if any firm is a price leader, Hypothesis 5E, among the refrigerated brands, it is Minute Maid, which is not the largest brands. Nor is there a clear pattern of alternating sales. For frozen, the magnitudes of the

response are small, which is not supportive of either the price leadership or alternating sales hypotheses.

Our most striking result is that the best predictor of a sale for a particular brand is a previous sale of that brand. This result comports with neither price leadership, nor the alternating sales hypothesis.

We summarize our findings on patterns of sales in:

Result 5. A. Evidence from other papers, price distributions, runs tests, the probit, and the VAR analysis all strongly reject the hypothesis that sales are random.

B. The impulse response functions from the VAR predict that a price decrease in one brand will lead to a price decrease in another brand with some lag. The low magnitude of the residual autocorrelation matrix from the VAR does not lend much support to the hypothesis that firms lower prices simultaneously. Moreover, a count of the number of simultaneous sales within a store shows that they are relatively rare.

C. Direct observation of frequency of sales contradicts the hypothesis that private labels do not have sales. The runs test finds that runs (alternating sales) occur more often than if sales were random. The VARs show that a sale by one national brand often causes a sale by another national brand, though the magnitude of this effect is small.

D. The probit equations do not support the hypothesis that the probability of a sale increases as the time from last sale increases. Indeed, we usually find the opposite result.

E. There is weak evidence from the VARs that a change in price by the largest brands cause a change in price, with some time lag, in the other brands.

F. Direct observation shows that prices do not fall slowly and then jump to a high level.

5. Conclusions

Based on a variety of summary statistics, runs test, probit regressions, and time-series analyses, we conclude that no one theory of sales fully describes pricing practices for either refrigerated or frozen orange juice. We tested five sets of hypotheses.

First, contrary to popular belief and two theories, sales are not more frequent for national brands than for house brands. Private labels hold sales as frequently as do national brands for refrigerated orange juice and more frequently than does Minute Maid, the dominant national frozen product. This later result is consistent with the theory of [Raju et al. \(1990\)](#).

Second, we reject the maintained hypothesis in most theories of sales that manufacturers determine the sales. Based on correlations and other evidence, sales of

orange juice brands are determined by stores or chains rather than by manufacturers nationally. This result may explain why the leading theories of sales, which presume that manufacturers determine pricing, do not predict sales patterns for orange juice very well. This result also has implications for modeling and estimating price cost margins along the distribution chain. It highlights the role retailers play in determining price variations—retailers are not passive (cf. [Chevalier et al., 2003](#)).

Third, the various sales theories make conflicting predictions about the distribution of prices, none of which holds universally or even commonly. Rather, the frequency of sales, the depth of the sale, and the resulting price distribution vary substantially across most stores. We see histograms with smooth, bell shapes (as with the mixed strategies in [Varian, 1980](#)), distributions with one or two mass points ([Lal, 1990](#)), and other patterns. We do not observe a smooth histogram with a mass point at the highest price ([Sobel, 1984](#)).

Fourth, we examined whether there was a systematic difference in price distributions for storable, frozen orange juice versus nonstorable, refrigerated juice. Several of these theories apply only for durable goods and lead to very different predictions than the other theories. However, we found no systematic difference between frozen and refrigerated juice.

Fifth, we rejected many of the theories about the timing of sales and found only weak evidence for the others. Sales are not random, sales of various brands occur simultaneously occasionally, prices do not fall slowly and then jump, and there is little evidence of pure price leadership. Our results are the exact opposite of [Pesendorfer's \(2002\)](#) prediction that the probability of a sale increases with the time since the last sale of various brands. For some, but not all pairs of brands, we find evidence of alternating sales. This pattern is not a strong one, however.

We have some strong results. Most of the variation in orange juice prices over time is due to temporary price reductions, rather than to a time trend or patterns of smooth adjustment. Our Granger causality (temporal ordering) tests indicate that a sale of a major national brand is more likely to “cause” later sales of other brands than is a sale of a private label product, but the magnitude of the response is less than in proportion. A sale of a major national brand “causes” sales of other brands in at least half the stores for only one-quarter of the other refrigerated brands and half the frozen brands.

An impulse response analysis shows substantial differences across brands in the size of the response of rival brands to a sale. A past sale by a brand has a much

larger effect than a past sale of a rival brand on its own later sales. Compared to other refrigerated or frozen orange juice brands, Minute Maid sales result in more large-price responses in other brands' prices. Nonetheless, for both frozen and refrigerated orange juice, a sale of one brand of orange juice triggers a price response of other brands of at least half the magnitude of the original sale in fewer than half of the cases, which does not support the notion of alternating sales.

Appendix A. Simultaneous sales

We counted how many times brands have sales in the same week across all stores (2288 observations). In the following table, the diagonal elements are the number of times only the listed refrigerated brand was on sale. The off diagonal elements are the counts of multiple products on sale at the same time (Tropicana and Florida Natural were on sale in the same week 38 times).

Refrigerated	Tropicana	Florida Natural	Minute Maid	PL
Tropicana	115			
Florida Natural	38	107		
Minute Maid	20	16	75	
PL	29	35	21	135

In addition, Tropicana, Florida Natural, and the private label had 10 common sales; Tropicana, Florida Natural, and Minute Maid had 3; Tropicana, Minute, Maid, and the private label had 4; Florida Natural, Minute Maid, and the private label had 4; and all were on sale 1 time.

The following is the table for frozen products. All products were on sale 8 times. There are 319 total sales and 1040 store weeks. There are 51 multiple sales.

Frozen	Minute Maid	Old Orchard	PL
Minute Maid	59		
Old Orchard	3	71	
PL	20	20	138

All = 8 times.

Appendix B. Trends and unit roots of price data

Using price in an augmented Dickey–Fuller test, we reject the unit root hypothesis at the 0.05% level in 68 of 88 brand-store combinations: 19 for Tropicana, 18 for Florida Natural, 15 for Minute Maid, and 16 for the private label. Using the log of price, we get similar results, rejecting the unit root in 66 of 88 brand-store combinations: 19 for Tropicana, 18 for Florida Natural, 15 for Minute Maid, and 14 for the private label.

Using price, we reject the unit root hypothesis at the 0.05% level in 68 of 88 brand-store combinations: 19 for Tropicana, 18 for Florida Natural, 15 for Minute Maid, and 16 for the private label. Using the log of price, we get similar results, rejecting the unit root in 66 of 88 brand-store combinations: 19 for Tropicana, 18 for Florida Natural, 15 for Minute Maid, and 14 for the private label.

Testing for a price trend (where the null hypothesis is a random walk) by regressing price on a constant and a time trend, we fail to reject the null hypothesis of no trend in 34 of 88 cases: 6 for Tropicana, 9 for Florida Natural, 12 for Minute Maid, and 7 for the private label. The summary of results using the log of price is similar.

The estimated trend coefficients are very small (regardless of whether they are or are not statistically significantly different from zero). On average, trend coefficients with price are less than 0.003 (0.3¢ per week): 0.0018 for Tropicana, 0.0022 for Florida Natural, 0.0023 for Minute Maid, and 0.0027 for the private label. On average, the trend coefficients with log prices are 0.0006 for Tropicana and 0.0009 for the other products. We conclude that there is no economically or statistically meaningful trend in these data and that the variation within the price series must be due to something other than a steady change in price.

References

- Agrawal, D., 1996. Effect of brand loyalty on advertising and trade promotions: a game theoretic analysis with empirical evidence. *Marketing Science* 15 (1).
- Aguirregabiria, V., 1999. The dynamics of markups and inventories in retail firms. *Review of Economic Studies* 66 (2), 275–308.
- Blattberg, R.C., Neslin, S.A., 1989. *Sales Promotion: Concepts, Methods and Strategies*. Prentice Hall, Englewood Cliffs, N.J.
- Blattberg, R.C., Briesch, R., Fox, E.J., 1995. How promotions work. *Marketing Science* 14, 122–132.
- Chevalier, J.A., Kashyap, A.K., Rossi, P.E., 2003. Why don't prices rise during periods of peak demand? Evidence from scanner data. *American Economic Review* 93 (1), 15–37.
- Conlisk, J., Gerstner, E., Sobel, J., 1984. Cyclic pricing by a durable goods monopolist. *Quarterly Journal of Economics* 99, 489–505.
- Dutta, S., Bergen, M., Levy, D., 2002. Price flexibility in channels of distribution: evidence from scanner data. *Journal of Economic Dynamics and Control* 26 (11), 1845–1900.
- Green, E., Porter, R., 1984. Non-cooperative collusion under imperfect price information. *Econometrica* 52, 87–100.
- Hendel, I., Nevo, A., 2006a. Sales and consumer inventory. *Rand Journal of Economics* 37 (3), 543–561.
- Hendel, I., Nevo, A., 2006b. Measuring the implications of sales and consumer inventory behavior. *Econometrica* 74 (6), 1637–1673.
- Hosken, D., Reiffen, D., 2004. Patterns of retail price variation. *Rand Journal of Economics* 35 (1), 128–146.
- Lal, R., 1990. Price promotions: limiting competitive encroachment. *Marketing Science* 9, 247–262.

- Lal, R., Villas-Boas, J.M., 1998. Price promotions and trade deals with multiple retailers. *Management Science* 44 (7), 935–949.
- Maskin, E., Tirole, J., 1988. A theory of dynamic oligopoly II: price competition, kinked demand and Edgeworth cycles. *Econometrica* 56 (3), 571–599.
- Nakamura, E., Steinsson, J., 2007. Price setting in forward-looking customer markets. Harvard University Working Paper.
- Perloff, J.M., Wu, X., 2007. Tax incidence varies across the price distribution. *Economic Letters* 96 (1), 116–119.
- Pesendorfer, M., 2002. Retail sales: a study of pricing behavior in supermarkets. *Journal of Business* 75, 33–66.
- Raju, J.S., Srinivasan, V., Lal, R., 1990. The effects of brand loyalty on competitive price promotional strategies. *Management Science* 36 (3), 276–304.
- Salop, S.C., 1977. The noisy monopolist. *Review of Economic Studies* 44, 393–406.
- Salop, S.C., Stiglitz, J.E., 1982. The theory of sales: a simple model of equilibrium price dispersion with identical agents. *American Economic Review* 72, 1121–1130.
- Shilony, Y., 1977. Mixed pricing in oligopoly. *Journal of Economic Theory* 14, 373–388.
- Slade, M., 1994. Optimal Pricing with Costly Adjustment and Persistent Effects: Empirical Evidence, mimeo, University of British Columbia.
- Sobel, J., 1984. The timing of sales. *Review of Economic Studies* 51, 353–368.
- Varian, H., 1980. A model of sales. *American Economic Review* 70, 651–659.
- Villas-Boas, J.M., 1995. Models of competitive price promotions: some empirical evidence from the coffee and saltine crackers markets. *Journal of Economics and Management Strategy* 4, 85–107.
- Villas-Boas, S.B., 2007. Vertical relationships between manufacturers and retailers: inference with limited data. *Review of Economic Studies* 74 (2), 625–652.
- Warren, E., Barsky, R., 1995. The timing and magnitude of retail store markdowns: evidence from weekends and holidays. *Quarterly Journal of Economics* 110, 321–352.