Staggering and Synchronization in Price-Setting: Evidence from Multiproduct Firms

By SAUL LACH AND DANIEL TSIDDON*

Theoretical work on price-setting behavior focuses on the single-product case while, in reality, a single price-setter sells many products. We use retail store-level multiproduct pricing data to learn about price dynamics. We find that, while the timing of a product's price changes is staggered across stores selling the same product, the timing of the price changes of different products sold within the same store is highly synchronized. This finding validates the usual assumption that decisions are staggered across price-setters and suggests that price rigidity is due mostly to "mechanical" reasons and not to informational asymmetries. (*JEL* E10, E31, E52)

One of the most important lessons learned from the Fischer-Taylor analysis of staggered contracts is that the mechanism responsible for the long lag in the response of the aggregate price level to shocks in the money supply relies crucially on the assumption of staggered contracts. If agents fully synchronize their actions, the maximum lag of the aggregate response to shocks in the money supply is the length of the contract.

The logic of this argument applies to the price-setting context as well. Under full information, a necessary condition for changes in the aggregate price level to lag behind shocks in the money supply is that the response of price-setters to the monetary shock is staggered over time. Because not all price-setters change their prices simultaneously, each price-setter takes into account that some of his competitors have not yet changed their prices, which prevents him from changing his own products' prices to accommodate "fully" the change in the money supply. Hence, changes in the aggregate price level lag behind changes in the money supply. But, as shown by Andrew S. Caplin and Daniel F. Spulber (1987), staggering may not be sufficient to generate lags in the response of the aggregate price level, even when price-setters change prices discontinuously. It is, however, always a necessary condition (see also Laurence Ball and Stephen G. Cecchetti, 1988; Ricardo J. Caballero and Eduardo M. R. A. Engel, 1991, 1993; Andrew Caplin and John Leahy, 1991; Tsiddon, 1993).

In a previous paper (Lach and Tsiddon, 1992; hereafter LT), we analyzed store-level monthly price data of 26 food products sold in Israel during high-inflation periods. Figure 1, reproduced from LT, shows that price changes do indeed seem to be staggered: in any single month the proportion of price changes is never close to either 0 or 1, and it is fairly constant over the 18 months analyzed. It hovers around 30 percent, which is consistent with an average duration of a nominal price quotation of 2.5–3 months.

Note that the staggering referred to above, and in the macroeconomic literature in general, is across decision-makers (price-setters), not

* Lach: Department of Economics, The Hebrew University of Jerusalem, Jerusalem 91905, Israel, and NBER; Tsiddon: The Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 69978, Israel, and CEPR. We thank Joshua Angrist, Roland Benabou, Giora Hanoch, and Shlomo Yitzhaki for helpful comments on an earlier version of this paper. Special thanks to Benjamin Daniel for excellent research assistance, and to Miriam Zadik of the Central Bureau of Statistics for providing the data. The comments of two anonymous referees greatly improved the presentation of the paper. We gratefully acknowledge the financial support from the Israel Foundation Trustees, the Israeli Academy of Sciences, and the Maurice Falk Institute for Economic Research in Israel. The paper was revised while Lach was visiting the Industrial Output section at the Federal Reserve Board.

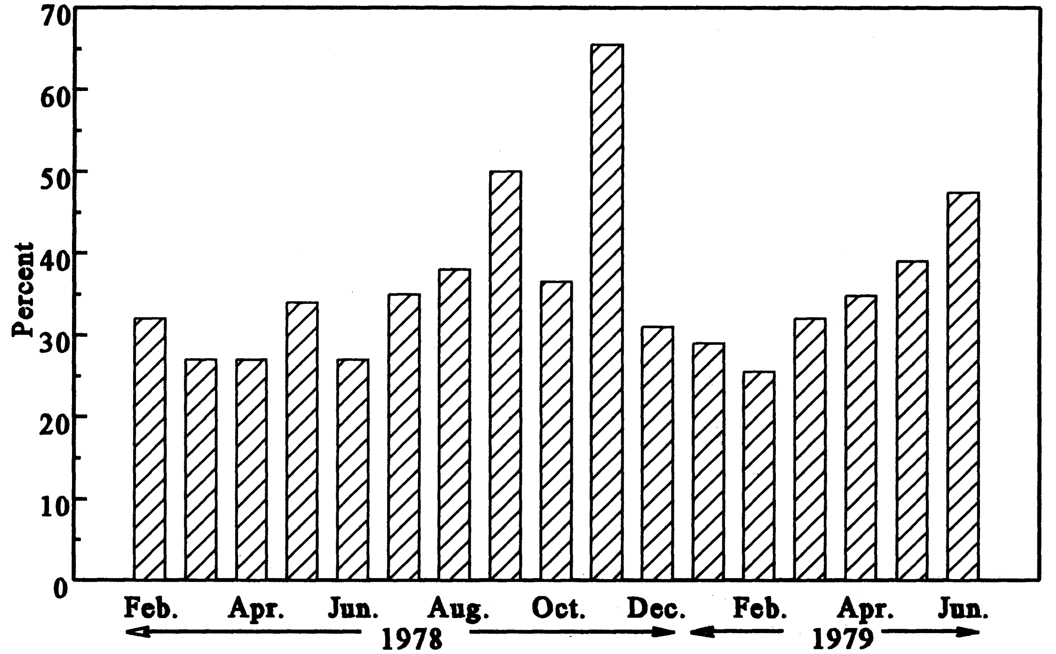

FIGURE 1. PROPORTION OF PRICE CHANGES, FEBRUARY 1978–JUNE 1979

across products. Most of the theoretical and empirical literature on price-setting behavior focuses on the single-product case, thus avoiding any ambiguity in the concept of staggering. Nevertheless, the presence of multiproduct firms raises the possibility that the staggering of price changes occurs across products and not across price-setters.[1] For example, suppose that all price-setters sell the same nine products and all change the prices of the first three products in month 1, of the second three products in month 2, and of the last three in month 3. In month 4 they all start the cycle again. We will then observe a third of all prices changing each month. The data-generating process is characterized by staggering across (groups of) products and perfect synchronization of all price-setters. Of course, the same observed number of price changes is obtained when a third of the stores change all nine prices in month 1, another third in month 2, and the remaining third of the stores in month 3. Now, however, the data-generating process reflects staggering across price-setters (stores) in the timing of their price changes accompanied by perfect synchronization in the timing of price changes within each store.

This extreme example shows that the same observed data can result from diametrically opposed causes. The problem with Figure 1 is that it does not distinguish between changes in products' prices within a store or across stores. The present paper sheds light on the driving forces behind the observed pattern of staggered price changes. A proper account of the distinction between staggering across price-setters vis-à-vis staggering across products is important for macroeconomic analysis since different staggering mechanisms yield different price dynamics.

Our analysis leads us to conclude that Figure 1 is the result of staggering across price-setters, while price changes of different products are synchronized (nonstaggered) within the store. That is, the data exhibit *across-stores staggering* and *within-store syn-*

[1] Mariano Tommasi (1993) seems to be the first to address this issue.

chronization in the timing of price changes. This finding validates the assumption of staggered decisions across price-setters made in most of the "sticky-prices" literature.

The paper is organized as follows. The next section sketches the implications of the main models of price adjustment that guide us in the empirical analysis of the price data described in Section II. In Section III, across-stores staggering in the timing of price changes is analyzed, while the evidence on within-store synchronization is presented in Section IV. Finally, Section V investigates the timing of negative price changes. Conclusions close the paper.

I. Price-Adjustment Models

In this section, we sketch the implications of the main models of price adjustment. Despite the richness of this literature, not much is known about price dynamics in a multiproduct environment.[2] We focus on the restrictions placed by these theories on the behavior of multiproduct firms and draw implications regarding the timing of price adjustments in this setting. These implications will guide us in our analysis of the data.

A. *Signal-Extraction Models*

In the signal-extraction literature (Robert E. Lucas, 1973), a shock hitting all stores at the same time prompts a synchronized response leading to across-stores synchronization. The lack of synchronization observed in Figure 1 does not support this implication. Furthermore, it is difficult to suggest a convincing argument whereby macro shocks lead to within-store synchronization but not to across-stores synchronization. Clearly, this explanation is not consistent with our findings.

One way of reconciling this theory with our empirical findings is to let the effects of macro shocks be unevenly distributed geographically. Even though moving away from a pure macro shock can potentially generate across-stores staggering and within-store synchronization, we show in Section V that the data appear to reject the geographic hypothesis as well.

B. *Search Theory*

In LT we showed that the relative price variability in each (homogeneous) product market is very large. Consequently, consumers have incentives to search for the lowest price. There is, in fact, a rich literature connecting search theory to inflation, but most of its implications cannot be addressed with our data (Roland Benabou, 1988; Tommasi, 1994).

Search, however, is not confined to consumers only. In an inflationary and uncertain environment, sellers also are not fully aware of nominal pressures. Every new price quotation, therefore, brings new information on market conditions to consumers and sellers alike. To the best of our knowledge a model in which consumers and sellers search in the context of multiproduct firms does not yet exist. Hence, we can only offer conjectures about the constraints such a model would impose on the data. In broad terms, and mainly from the firm's own information-gathering perspective, staggering price changes within the firm amounts to following a sequential search procedure, whereas synchronization of price changes is analogous to a fixed-sample search approach. It is well known that, under fairly general conditions, sequential search is a better strategy. We find, nevertheless, that stores synchronize the timing of their products' price changes; that is, they choose the fixed-sample approach, disregarding the strong signal they send to shoppers.

C. *Sticker-Price Model*

Peter A. Diamond (1993) proposes yet another mechanism to justify the sluggishness of the aggregate price level: identical products may have different prices since prices are set at the time of delivery to the store and remain unchanged unless a crucial change in the environment occurs. Our data do not support this hypothesis; for it to be consistent with our findings one needs to assume that all products are delivered simultaneously to each store so as to generate within-store

[2] Notable exceptions are Agnes Sulem (1986) and Eytan Sheshinski and Yoram Weiss (1992).

synchronization, and that there is a sufficiently spread-out distribution of delivery dates, on a monthly basis, across stores. These are strong assumptions which are unlikely to hold for the type of products to be analyzed in this paper.

D. *Adjustment-Costs Models*

Models with convex costs of price adjustments yield discontinuities in nominal price changes. It is reasonable to expect that, in a multiproduct context, these adjustment costs have a store-specific component; that is, these costs also depend on the characteristics of the price-setter. The term "menu cost" comes alive: the cost of printing a new menu is shared by all products.

Clearly, store-specific menu costs will induce the price-setter to synchronize its price changes, an implication consistent with our findings. Note, however, that store-specific costs should only induce synchronization in the *timing* of price changes, and not in the *size* of the price changes for individual products.[3] This justifies our focus on the synchronization in the timing of such changes.[4]

II. Description of the Data

The data used in this work is a subsample of the data used in LT, where it is described in detail. The original data set consists of nominal price quotations for 26 food products collected monthly from a sample of stores[5] by the Central Bureau of Statistics (CBS) for the purpose of computing the consumer price index (CPI). That is, for each product we have a panel of prices extending across stores and over time. Alternatively, for each store we have a panel of prices extending over products and over time.

The products in the sample are all homogeneous and did not change substantially either in quality or in market structure. Furthermore, their prices were not controlled by the government during the period of investigation.

Since part of the focus of the current study is on issues related to the comovement of prices within stores, we selected 21 products that could be grouped into two broad classes: wines and meat products.[6] Note that each store in our data sells either wine or meat products, and none of them sells both wines and meat.

The periods for which most of the data are available are 1978–1979, 1981–1982, and the first nine months of 1984, before across-the-board price controls were first put into effect. The data for 1980 and 1983 have disappeared from the CBS archives. The analysis in this paper is restricted to the 1978–1979:6 subperiod, corresponding to a single inflationary step as defined by Nissan Liviatan and Silvia Piterman (1986), because the price-duration data are less affected by the one-month truncation introduced by the sampling interval.[7]

In the latter part of this paper we analyze the within-store dimension of the data. For this to be meaningful we focus on stores selling three or more products. The upper graph in Figure 2 plots the number of stores meeting this requirement by product class. There are twice as many stores selling meat products as stores selling wines, and the number of stores

[3] In Lach and Tsiddon (1996) we analyze the size distribution of the within-store average price change.

[4] Another explanation that fits the within-store synchronization of price changes is based on informational externalities. Ball and Cecchetti (1988) show that a mechanism in which each price-setter derives information on inflationary pressures from observing the decisions made by other price-setters may generate an equilibrium with staggered price-setting. This explanation works best in conjunction with store-specific menu costs. In this case, it amplifies the within-store synchronization phenomenon, and yields an intuitive and plausible mechanism that explains staggering across stores.

[5] These are grocery or liquor stores; supermarkets and chain stores are not included in the sample.

[6] Wine products consist of nine wines and liquors: arrack (anise), white vermouth, liquor, champagne, vodka, red wine, rosé wine, hock wine, and sweet red wine. The 12 meat products, including three types of fish, are fresh beef, frozen goulash, frozen beef liver, fresh beef liver, chicken breast, chicken liver, turkey breast, beefsteak, drumsticks, fish fillet, buri fish, and codfish.

[7] During this period the mean inflation rate was 3.9 percent per month, with a standard deviation of 1.9 percent. The median rate was 3.5 percent per month. The object of study is the occurrence of a price change. In order to take account of round-off errors this event is defined to occur whenever an observed price change exceeds 0.5 percent.

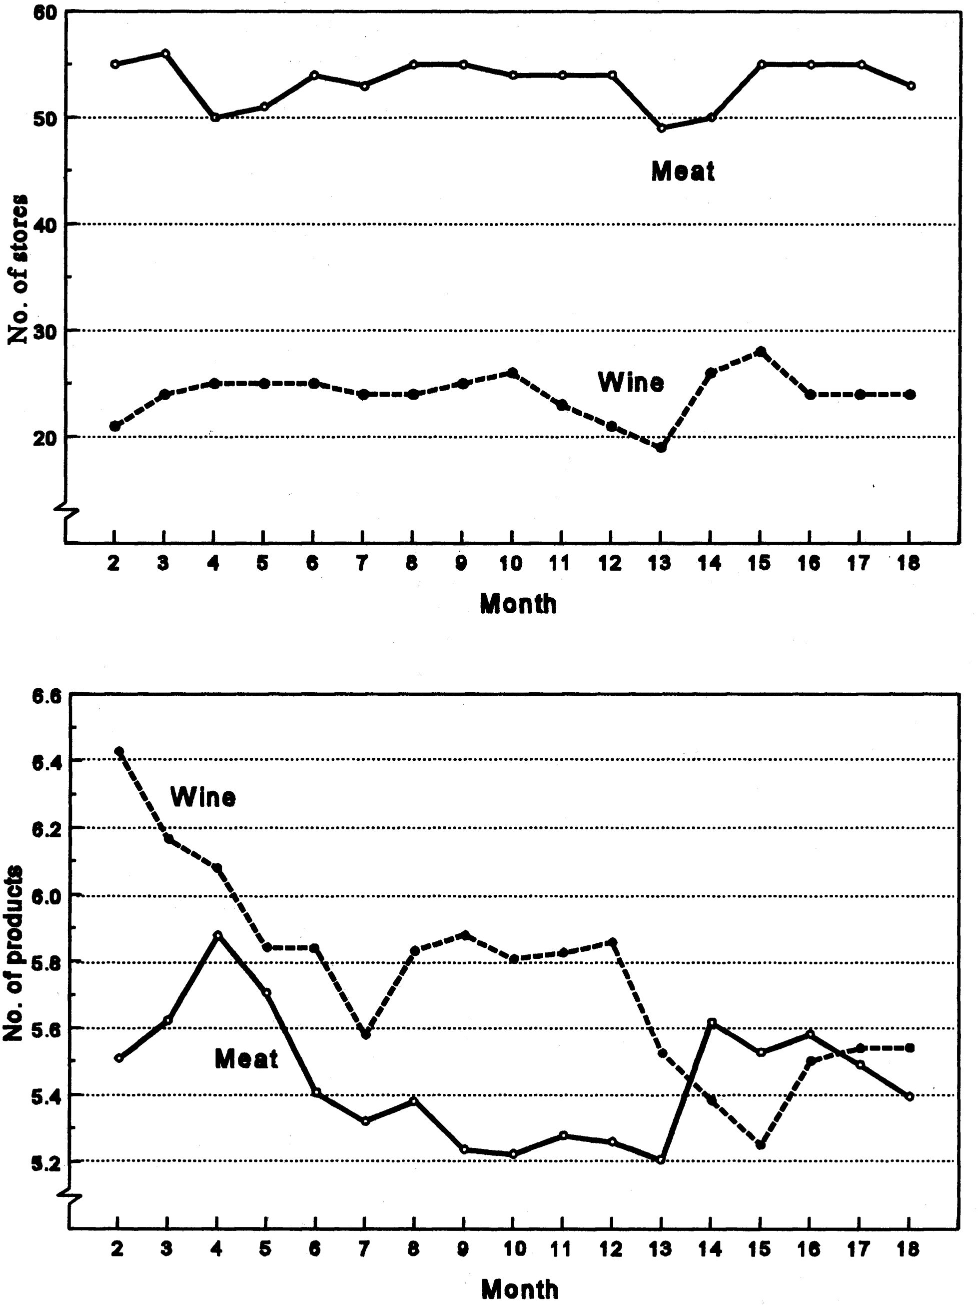

FIGURE 2. NUMBER OF STORES AND PRODUCTS, FEBRUARY 1978–JUNE 1979

Note: The upper graph is for the number of stores with three or more products; the lower graph is for the number of products per store.

is stable over time.[8] The number of products, averaged over stores, is plotted in the lower graph. The average number of products per store lies in the 5.2–6 range, with some variation over time in these averages. The standard deviation of the number of products per store is 2–2.5 products.

III. Across-Stores Staggering

A necessary condition for an effective monetary policy in a rational-expectations environment is that all price-setters do not change their prices simultaneously in response to a monetary shock. This lack of simultaneity is termed *across-stores staggering*. The term refers to staggering in the *timing* of price changes across different stores for a *given* product. In most macroeconomic models (e.g., Stanley Fischer, 1977), across-stores staggering implies more than mere lack of simultaneity; it also embodies the notion of "regular cyclicity" in the response of price-setters to a shock. One group of price-setters is first to change prices, followed by another group of different price-setters; at some point in time, however, before the second group acts again, the first group of price setters changes its prices a second time. Our data are uniquely suited to check the extent to which these phenomena prevail. This is the purpose of this section, which is divided into three parts, each one presenting empirical evidence on different features of across-stores staggering in the timing of price changes.[9]

A. *Proportion of Price Changes*

The first step is to examine, for each product, the time series of the proportion of stores that changed prices. Lack of staggering or simultaneity implies that stores either change their prices together or do not, that is, that the observed proportions are close to 1 or to 0.

Figures 3 and 4 present such a time series for the 17 months between February 1978 and June 1979, for each product. At first glance, the proportion of stores changing prices is well below 1 in all months, with the exception of November 1978. The requirement that these proportions be above zero is also satisfied to a lesser extent. Omitting the November 1978 observation, meat products do not show much variability over time compared to wines. Wines, on the other hand, exhibit a lower proportion of stores changing prices during the first half of the year.

In a stationary-inflation environment a store following an $(S - s)$ pricing policy is expected to change its prices by the same amount every δ months (δ being determined by the parameters of the inflation process and profit function) (Sheshinski and Weiss, 1992). What does this imply for the observed proportions of stores changing prices? If stores are expected to change prices every δ months, and there is sufficient heterogeneity in the initial conditions, then after a long enough number of months, the proportion of stores changing prices every month should stabilize around $1/\delta$. The horizontal dashed line in each panel of Figures 3 and 4 is $1/\delta$, where δ is the average duration of a price quotation taken from table 4 in LT. The "fit" seems to be much better for meat products than for wines.

B. *Simultaneous Price Changes*

The issue of staggering can be tackled from another angle by asking: how many stores change prices simultaneously? Let M_{ijt} be the number of stores changing the price of product j during month t simultaneously with store i. Table 1 presents the values of M averaged over months, products, and stores.

When a store selling wines changes one of its prices, it usually does so together with 6.9 other stores (42 percent of its competitors) on

[8] In LT we showed that, even though there were some changes over time in the identity of the stores, there is a sizable core of stores that remained in the sample for long periods of time.

[9] Obviously, the timing of price changes is correlated across stores. This correlation is a result of responses to factors common to all stores (e.g., to an increase in the aggregate rate of inflation), and not because of strategic behavior. Given that our data are composed of small grocery stores located all over the country (we do not sample supermarkets), this is not a bad assumption. Letting Z_t denote all the common factors alluded to above, we assume that *conditional* on Z_t the probability of store i changing the price of some product is not affected by what another store does or did. Hereafter, independence across stores refers to this form of conditional independence.

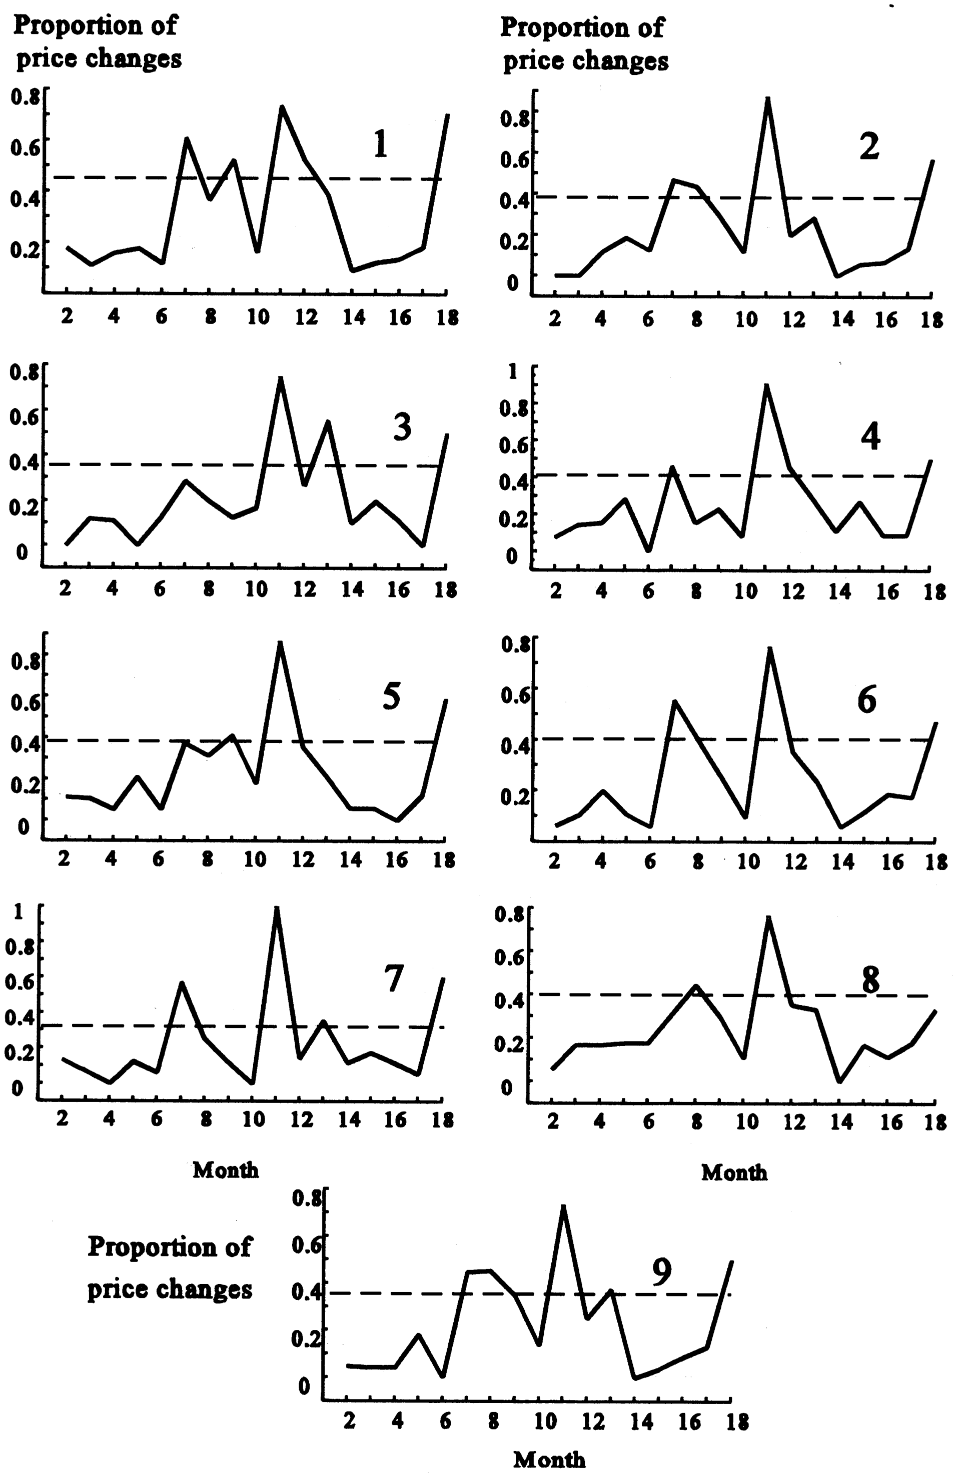

FIGURE 3. PROPORTION OF PRICE CHANGES FOR WINES (NINE DIFFERENT PRODUCTS), FEBRUARY 1978–JUNE 1979

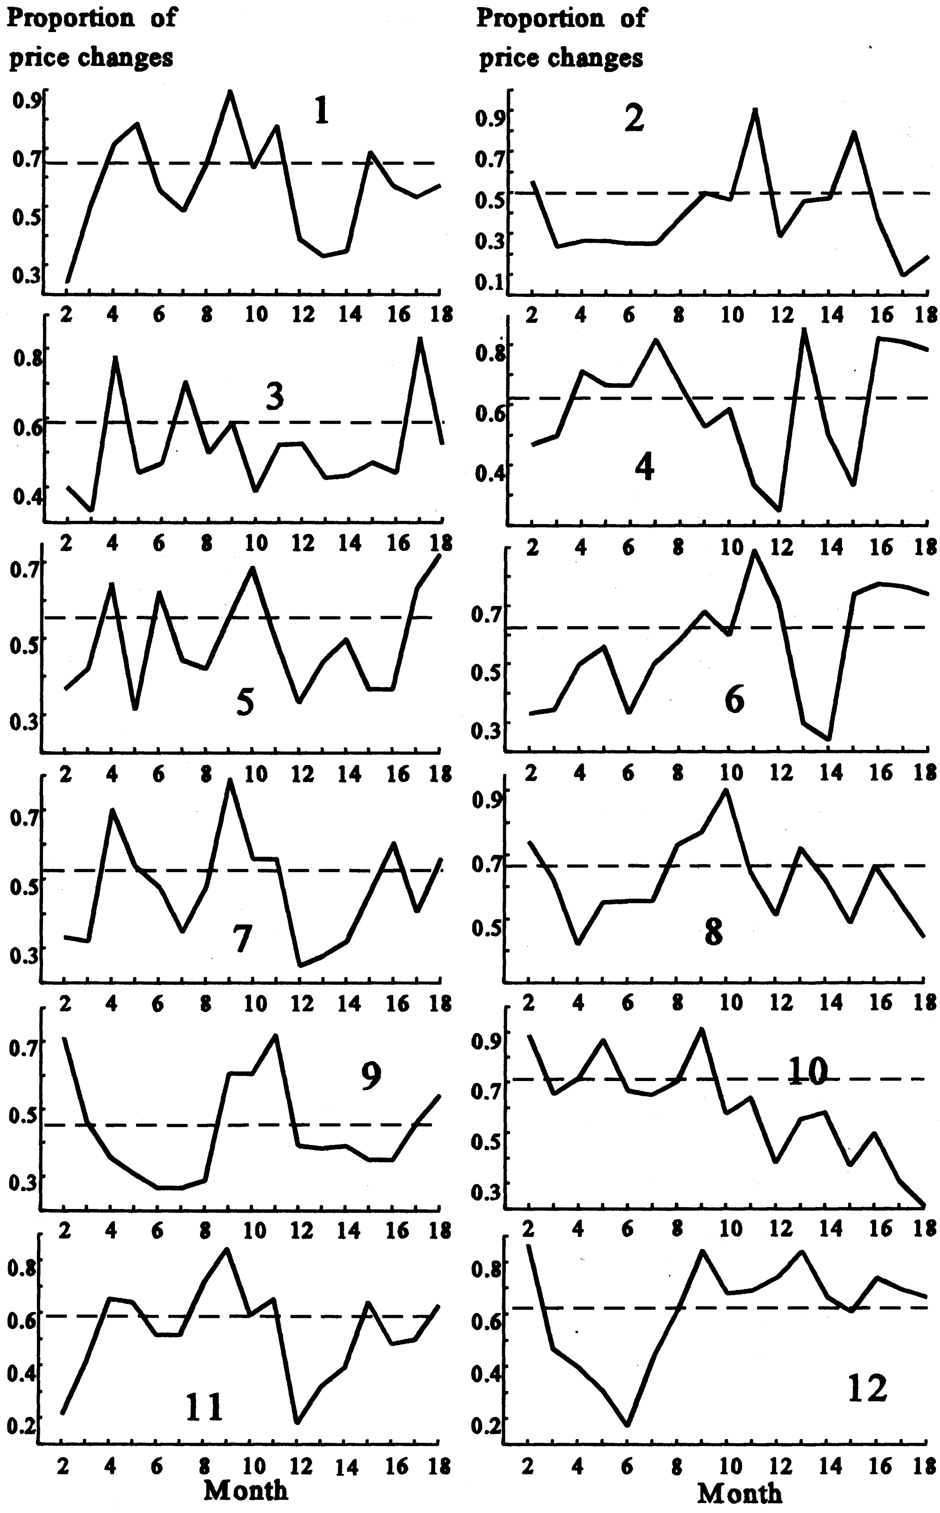

FIGURE 4. PROPORTION OF PRICE CHANGES FOR MEAT PRODUCTS (12 DIFFERENT PRODUCTS), FEBRUARY 1978–JUNE 1979

TABLE 1—SIMULTANEOUS PRICE CHANGES: SUMMARY STATISTICS

Product type	Variable	Mean	Median	Standard deviation	Minimum	Maximum
Wines	$M_{i\cdot\cdot}$	6.90	6.32	2.50	2.83	14.00
	$\text{Share}_{i\cdot\cdot}$	0.42	0.40	0.14	0.20	0.87
Meat	$M_{i\cdot\cdot}$	15.24	17.31	4.66	6.80	23.40
	$\text{Share}_{i\cdot\cdot}$	0.56	0.57	0.06	0.37	0.71

Notes: $M_{i\cdot\cdot}$ equals the number of stores changing price of product j in month t simultaneously with store i (M_{ijt}) averaged over the number of products j sold by store i and over the number of months in which these products were sold. $\text{Share}_{i\cdot\cdot}$ equals M_{ijt} divided by the number of stores selling product j during month $t-1$, averaged over products and months.

average. From the store-level data underlying the results in Table 1, we know that 62.5 percent of the stores change price simultaneously with 5–9 other stores, or that 50 percent of the stores change prices at the same time that 32–52 percent of their competitors do.[10] Most stores selling meat products usually change prices simultaneously with 56 percent of their competitors, or with 15.2 other stores on average.

The preceding analysis indicates that, in general, the proportion of stores is away from the zero-one boundaries. Furthermore, from the point of view of the individual store, a change in prices does not indicate that all of its competitors follow suit, even though a sizable share of them do.

C. *Regular Cyclicity*

Another characteristic of across-stores staggering is not captured either by the observed proportions of price changes or by the number of simultaneous moves. As the opening paragraph of this section suggested, having the *same* group of firms change prices every month during the first, say, six months, while another group does so during the second part of the year, is not the kind of staggering economists have in mind when analyzing price dynamics; it does not conform with the notion of regular cyclicity. Staggering embodies a notion of regular cyclicity.

The "perfect" across-stores staggering is one in which, in response to a monetary shock, a different $1/\delta$ of all stores change prices every month. After δ months all stores have responded to the shock, and the cycle starts again. Letting $X_{ijt} = 1$ indicate that store i changed the price of product j during month t (otherwise, $X_{ijt} = 0$), the "perfect" time series of X_{ijt} is composed of 1's every δ months and 0's everywhere else.

We examine the X_{ijt} time series for each store i and product j for evidence of regular cyclicity.[11] We count the number of times prices were observed to change consecutively at least twice, at least three times, and so on, and divide this count by the potential number of consecutive price changes.[12] These ratios can be interpreted as unconditional probabilities of observing, say, at least K consecutive price changes. Table 2 presents summary statistics.[13]

For stores selling wines, the estimated probability of spreading out a price change over two or three consecutive months is quite low. The mean probability, averaged across stores, is 8.3 percent, while the median probability is only 5 percent, reflecting the large number of 0 values. Even without their standard errors, these estimates suggest

[10] The values of M_{ijt} for each store i, averaged over months and products, appear in tables A1 and A2 of the working-paper version of this paper (Lach and Tsiddon, 1994).

[11] There are potentially 360 and 1,080 such series for wines and meat products, respectively. For example, there are 40 different stores selling some of the nine wine products. However, a nonnegligible number of the series are missing, since most stores do not sell all nine products.

[12] Appendix B explains how the counting is done.

[13] See tables A3 and A4 in Lach and Tsiddon (1994) for the store-level data on the C's.

TABLE 2—PROBABILITY OF CONSECUTIVE PRICE CHANGES: SUMMARY STATISTICS

Statistic	Product type	C2	C3	C4	C5	C6	C7	C8	C9
Mean	W	0.083	0.026	0.003	0.002	0	0	0	0
	M	0.189	0.105	0.103	0.048	0.021	0.016	0.020	0.030
Median	W	0.050	0	0	0	0	0	0	0
	M	0.176	0.097	0.062	0	0	0	0	0
Percentage of zeros	W	41.0	66.7	94.9	94.9	100	100	100	100
	M	10.1	35.2	45.5	67.1	84.5	89.0	92.5	91.1
Standard deviation	W	0.119	0.045	0.013	0.010	0	0	0	0
	M	0.137	0.109	0.150	0.091	0.058	0.053	0.115	0.131
Maximum	W	0.571	0.200	0.074	0.05	0	0	0	0
	M	1.000	0.500	1.000	0.500	0.333	0.333	1	1

Notes: The entries in a column labeled CK, $K = 2, \ldots, 9$ are the averages across stores of the unconditional probabilities of observing a spell of K consecutive price changes in wine prices (W) and meat-product prices (M). The average probability for C10–C17 is less than 0.01. See Appendix B for details.

that consecutive price changes are not a prevalent phenomenon in liquor stores. The store-level probabilities reveal that close to 41 percent of the liquor stores never spread out a change in the price of any products over two or more months (C2–C17 have zero entries in 15 stores).

Most stores selling meat products experienced two, three, and even four consecutive price changes at least once. These events are quite infrequent: the mean probability is 19 percent for two consecutive changes. Note, however, that a nonnegligible number of stores do have five or more consecutive price changes. Unlike liquor stores, the notion of across-stores staggering in the timing of price changes finds less support in the meat-products market.

A different perspective on the issue of regular cyclicity is also instructive. A modest requirement for staggering to occur is that stores alternate their decisions to change prices, that is, that stores changing a price in the current month did not change that price during the previous month, and conversely, stores that changed a price last month do not change it again this month.[14]

A simple 2 × 2 contingency table with two rows for the values of X_{ijt-1} and two columns for the values of X_{ijt} summarizes this information for each store-product-month observation. Assuming that X_{ijt} is (conditionally) independent and identically distributed across stores allows us to aggregate the tables over stores.[15] This still leaves us with 17 × 21 tables for each product-month combination. Assuming that the distribution of X_{ijt} is time-invariant during the 17 months reduces the information to 21 2 × 2 contingency tables (9 for wines and 12 for meat products) shown in Tables 3 and 4.

The frequency counts over stores and months of the event represented by each cell are used to compute probabilities. The top entry in each cell is the row percentage which is the maximum-likelihood estimator of the probability of X_{ijt} given X_{ijt-1}. The bottom entry is the column percentage, which is the maximum-likelihood estimator of the probability of X_{ijt-1} given X_{ijt}. Letting the first coordinate be the value of X at $t - 1$ and the second its value at t, then 0.26 in the

[14] An empirical check of this assertion is meaningful only if stores are sampled more often than the frequency of price changes, which is the case in the 1978–1976:6 period. Note that this behavior is not sufficient to generate staggering. If stores do behave this way, and if there is sufficient heterogeneity in the stores' initial conditions, across-stores staggering will occur. Otherwise, the result may be that all stores do indeed alternate their price changes but do so in a synchronized fashion.

[15] Note that this assumption allows for heterogeneity in the size of the price change.

TABLE 3—2 × 2 CONTINGENCY TABLES OF PRICE CHANGES FOR WINE PRODUCTS

	Product W1			Product W2			Product W3		
	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total
$X_{t-1} = 0$	0.74	0.26	228	0.76	0.24	193	0.79	0.21	163
	0.78	0.72		0.79	0.81		0.81	0.87	
$X_{t-1} = 1$	0.68	0.32	71	0.78	0.22	50	0.86	0.14	35
	0.22	0.28		0.21	0.19		0.19	0.13	
Total	217	82	299	186	57	243	159	39	198

	Product W4			Product W5			Product W6		
	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total
$X_{t-1} = 0$	0.78	0.22	138	0.77	0.23	209	0.76	0.24	214
	0.80	0.73		0.80	0.79		0.78	0.75	
$X_{t-1} = 1$	0.71	0.29	38	0.75	0.25	53	0.73	0.27	63
	0.20	0.27		0.20	0.21		0.22	0.25	
Total	135	41	176	200	62	262	208	69	277

	Product W7			Product W8			Product W9		
	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total
$X_{t-1} = 0$	0.76	0.24	190	0.76	0.24	202	0.75	0.25	264
	0.77	0.88		0.77	0.75		0.79	0.88	
$X_{t-1} = 1$	0.88	0.13	48	0.74	0.26	62	0.85	0.15	62
	0.23	0.12		0.23	0.25		0.21	0.12	
Total	186	52	238	199	65	264	252	74	326

Notes: $X_t = 1$ if the price changed in month t, and $X_t = 0$ if the price did not change. Each entry is the relative frequency of the cell-specific event counted over stores and over 17 months. The top entry in each cell is the row percentage: the maximum-likelihood estimator of the probability of X_{ijt} given X_{ijt-1}. The bottom entry in each cell is the column percentage: the maximum-likelihood estimator of the probability of X_{ijt-1} given X_{ijt}.

(0, 1) cell of the first contingency table is the probability of a liquor store changing the price of product 1 at t given that he did not change that price in the previous month, while 0.72 is the probability that a store did not change the price in the last period, given that the price is changed at t.

The only implications of across-stores staggering are that the bottom (0, 1) and top (1, 0) entries are large relative to the corresponding (1, 1) entry. In probability terms, the probability of *no* price change at $t - 1$ conditional on a change at t is larger than the probability of a price change at $t - 1$ conditional on a change at t (column percentage); and the probability of *no* price change at t conditional on a change at $t - 1$ is larger than the probability of a price change at t conditional on a change at $t - 1$ (row percentage). Liquor stores easily satisfy these implications, a finding consistent with the relatively small number of price changes in consecutive months; meat stores do not satisfy these implications.

In this section we analyzed three features of the data: the proportion of price changes, the number of simultaneous moves, and the phenomenon of regular cyclicity. The behavior of liquor stores matches the predictions of a model in which stores are staggered in the timing of each product's price changes. These new findings reinforce the conclusion reached in LT that prices of wine products are slow to adjust, with the proviso that the timing of the price changes is staggered across liquor stores.

The results for stores selling meat products are mixed. The proportion of price changes are bounded away from 0 and 1 and, on average, stores change prices at the same time as 56 percent of their competitors. Therefore, these

TABLE 4—2 × 2 CONTINGENCY TABLES OF PRICE CHANGES FOR MEAT PRODUCTS

	Product M1			Product M2			Product M3		
	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total
$X_{t-1} = 0$	0.47	0.53	284	0.66	0.34	417	0.48	0.52	130
	0.49	0.39		0.64	0.53		0.47	0.49	
$X_{t-1} = 1$	0.37	0.63	368	0.55	0.45	280	0.50	0.50	139
	0.51	0.60		0.36	0.47		0.53	0.51	
Total	269	383	652	431	266	697	131	138	269
	Product M4			Product M5			Product M6		
	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total
$X_{t-1} = 0$	0.39	0.61	83	0.52	0.48	141	0.44	0.56	178
	0.39	0.43		0.56	0.53		0.46	0.43	
$X_{t-1} = 1$	0.43	0.57	117	0.49	0.49	117	0.41	0.59	227
	0.61	0.57		0.44	0.44		0.54	0.57	
Total	82	118	200	131	127	258	171	234	405
	Product M7			Product M8			Product M9		
	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total
$X_{t-1} = 0$	0.57	0.43	214	0.41	0.59	217	0.63	0.37	265
	0.58	0.49		0.41	0.37		0.61	0.50	
$X_{t-1} = 1$	0.48	0.52	186	0.38	0.62	350	0.53	0.47	203
	0.42	0.51		0.59	0.63		0.39	0.50	
Total	211	189	400	222	345	567	275	193	468
	Product M10			Product M11			Product M12		
	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total	$X_t = 0$	$X_t = 1$	Total
$X_{t-1} = 0$	0.53	0.47	143	0.52	0.48	210	0.48	0.52	159
	0.47	0.30		0.56	0.43		0.46	0.36	
$X_{t-1} = 1$	0.35	0.65	241	0.40	0.60	219	0.38	0.62	235
	0.53	0.70		0.44	0.57		0.54	0.64	
Total	160	224	384	197	232	429	165	229	394

Notes: $X_t = 1$ if the price changed in month t, and $X_t = 0$ if the price did not change. Each entry is the relative frequency of the cell-specific event counted over stores and over 17 months. The top entry in each cell is the row percentage: the maximum-likelihood estimator of the probability of X_{ijt} given X_{ijt-1}. The bottom entry in each cell is the column percentage: the maximum-likelihood estimator of the probability of X_{ijt-1} given X_{ijt}.

stores' behavior exhibits characteristics of across-stores staggering. In many cases, however, their behavior does not accord with the concept of regular cyclicity.[16]

Across-stores staggering occurs in the market for each individual product, that is, in the across-stores dimension. In foodstuffs, as in most products, stores (price-setters) sell many different products. Do stores make use of their multiproduct characteristics to learn about the

[16] The difference in behavior between the two markets, meat products and wines, is certainly interesting, but before we speculate on the reasons for this difference, we should consider the possibility that it is an artifact of the data. For across-stores staggering to be *observed* we require, in addition to staggering in the timing of price changes across stores, that the duration of a price quotation be greater than one month (the sampling interval). Since the average duration of a price within this group of stores is less than two months—the unweighted mean is 1.72 with a standard deviation of 0.22 (table 4 in LT)—stores selling meat products may fail to show staggering simply because we cannot detect this behavior given our sampling frequency, not because the timing of prices changes are not staggered.

inflationary process? Or, are adjustment costs lumpy enough to prevent such search activity?

Answers to these questions are crucial when trying to discriminate among the different models of price-setting behavior. If stores do change different products' prices on different dates one could interpret regular cyclicity as a costly search process where each change in a specific price is an investment in discovering the aggregate shock. Extensions of the signal-extraction model to multiproduct firms should imply the presence of "within-store staggering" in addition to the observed across-stores staggering: when a firm sells many products it should tend to change the prices of some products at each date rather than bunching all price changes together. This implication motivates the analysis in the next section.

IV. Within-Store Synchronization

Do stores tend to change the prices of different products simultaneously? Is the change in the price of a particular product in a particular store usually accompanied by changes in other products' prices in the same store? If such simultaneity exists we call it *within-store synchronization.* Note that we investigate synchronization in the *timing* of changes in the prices of different products sold in a single store. Other related issues, such as the cross correlation in the size of price changes, are not explored here.[17]

A. *Proportion of Price Changes*

A natural measure of the degree of within-store synchronization is the proportion of products whose prices changed during a month. In our notation, this proportion is

$$\varphi_{it} = \frac{1}{|\mathcal{G}_{it}|} \sum_{j \in \mathcal{G}_{it}} X_{ijt} \tag{1}$$

where $\mathcal{G}_{it}$ is the set of products whose prices were recorded in store i during months $t-1$ and t (that is, the number of products sold at $t-1$ and t, equal to the number of nonmissing values of X_{ijt}) and $|\mathcal{G}_{it}|$ is the number of products in $\mathcal{G}_{it}$, that is, the cardinality of the set. We actually define φ_{it} for the subsample of stores that sell at least three products in each class, $|\mathcal{G}_{it}| \geq 3$. Recall that the stores in our sample sell either meat products or wines, but not both. Hence, synchronization between classes of products cannot be addressed with these data.[18]

We start by asking what values of φ_{it} should be expected when there is within-store synchronization. Clearly, we cannot provide a definite answer to this question without a structural model, but we can be fairly confident that when inflation is as high as it was during the period—3.9 percent per month—the probability of a store not changing *any* of its products' prices during a month is very low when the decision to change price is independent across products. Hence, observing many φ_{it}'s equal to zero should be indicative of within-store synchronization.[19,20] Table 5

[17] See Lach and Tsiddon (1996) for an analysis of the size distribution of the within-store average price change.

[18] A problem with our data is that we do not know whether more than one change in price occurred within the month, so that an observed $\varphi = 1$ is consistent with lack of within-store synchronization on, say, a weekly basis. Our definition of synchronization implies that two products are synchronized even if one's price is changed on the first day of the month and the other's on the last day of the same month. Another issue is that we sample a small fraction of the products sold by the store so that the true φ may be very different from the observed φ. Our results are, of course, conditional on the sample. To the extent that the sample of products is random, our conclusions can be carried over to the entire population.

[19] It is important to recall that there was no slowdown in the rate of inflation during this period.

[20] The same conclusion could be reached if all prices were changed during the month, $\varphi_{it} = 1$. A problem with this conclusion is that, given a positive rate of inflation, and with a long enough interval of time between samplings, a store will eventually change all its prices and we will observe $\varphi_{it} = 1$. To deduce that there is synchronization across products is, of course, misleading. In this case, $\varphi_{it} = 1$ is evidence of nothing but the fact that the frequency of sampling is too low relative to the rate of inflation. Hence we should be cautious in the interpretation of φ's equal to 1. We do not, however, believe this is an issue in our data. Recall that, in this period, when the average monthly rate of inflation was 3.9 percent, the average duration of a price quotation was 2.2 months. Had we used quarterly data, our definition of synchronization guarantees that we would have found perfect within-store synchronization in the data. But since we use monthly data, the severeness of this problem is reduced. In

TABLE 5—THE DISTRIBUTION OF φ_{it}

Product type		Quintile 0	0.0–0.2	0.2–0.4	0.4–0.6	0.6–0.8	0.8–1	1	Total
Wines	Observations:	236	27	51	18	18	12	46	408
	Percentage:	57.8	6.6	12.5	4.4	4.4	2.9	11.3	100.0
Meat	Observations:	139	28	152	169	243	66	111	908
	Percentage:	15.3	3.1	16.7	18.6	26.8	7.3	12.2	100.0

Note: We define φ as the proportion of products whose price changed during month t in store i. See equation (1).

presents the frequency distribution of φ_{it} for wines and meat products.

The difference between wine and meat products is quite striking. While most of the observations on wines correspond to $\varphi_{it} = 0$, the distribution of φ_{it} for meat products is much more balanced. If $\varphi_{it} = 0$ is the only credible evidence for synchronization in the timing of price changes across products we must conclude that most wine stores synchronize the timing of the price changes of their products, while stores selling meat products do so to a lesser extent.[21]

In the remainder of this section we present additional evidence favoring the within-store-synchronization hypothesis. We provide formal tests of the hypothesis, but doing so requires a series of compromising assumptions. It is therefore comforting to note that the direct evidence and conclusions from Table 5 are consistent with the more formal analysis.

The expected value of φ_{it} is

$$(2) \qquad E(\varphi_{it}) = \frac{1}{|\mathcal{G}_{it}|} \sum_{j \in \mathcal{G}_{it}} P_{ijt}$$

where $P_{ijt} = \Pr\{X_{ijt} = 1\}$ is the unconditional probability of observing a price change in product j at store i during month t.

The null hypothesis to be tested is the lack of within-store synchronization. This is interpreted as stating that the sequence $\{X_{ijt}\}$ is pairwise independent over products j. Under this hypothesis, the variance of φ_{it} is

$$(3) \qquad V(\varphi_{it}) = \frac{1}{|\mathcal{G}_{it}|^2} \sum_{j \in \mathcal{G}_{it}} P_{ijt}(1 - P_{ijt})$$

and for large enough $|\mathcal{G}|$, $[\varphi_{it} - E(\varphi_{it})]/V(\varphi_{it})^{1/2}$ is approximately distributed as a standard normal variable. If $\{X_{ijt}\}$ is a sequence of independent random variables over stores i, then

$$(4) \qquad T_t = \sum_{i \in \mathcal{N}_t} \frac{[\varphi_{it} - E(\varphi_{it})]^2}{V(\varphi_{it})} \to \chi^2_{|\mathcal{N}_t|}$$

where $\mathcal{N}_t$ is the set of stores with nonmissing φ in month t.

We focus on T_t since $|\mathcal{G}|$ is relatively smaller than $|\mathcal{N}|$ in our data. T_t is not a statistic because it depends on unknown parameters. Note that neither $E(\varphi)$ nor $V(\varphi)$ is observed, nor does the null hypothesis specify their values. $E(\varphi)$ and $V(\varphi)$ have to be estimated, and for this we need estimators of the probabilities of a price change in all the products.

P_{ijt} cannot be estimated for every store-product-month observation.[22] Hence, we have

particular, note that for wine products the average duration of a price quotation is four months.

[21] In an attempt to see whether the heterogeneity in the number of products sold by the store, $|\mathcal{G}_{it}|$, has an effect on the conclusions because of the possible effect of $|\mathcal{G}|$ on φ, we divided the observations into those corresponding to stores having $3 \le |\mathcal{G}_{it}| \le 5$ and those with $|\mathcal{G}_{it}| \ge 6$. Our conclusions were left unchanged.

[22] Under the null hypothesis we do not need to estimate the joint probability of $X_{i1t}, \ldots, X_{iJt}$ and then integrate out the marginal probabilities. Thus, under the null hypothesis, estimation of P_{ijt} is greatly simplified, since it allows us to ignore the information embodied in the behavior of the other products.

TABLE 6—CHI-SQUARE TESTS OF WITHIN-STORE SYNCHRONIZATION

Significance level	Number of rejections	
	Wines	Meat products
5 percent	11	16
10 percent	13	17

Notes: Entries are the number of times the hypothesis that T_t has a χ^2 distribution with $|\mathcal{N}_t|$ degrees of freedom [see equation (4)] is rejected, where $\mathcal{N}_t$ is the set of stores with no missing data in month t. The test was conducted 17 times, once for each month between February 1978 and June 1979. A rejection means that the "no synchronization" hypothesis is rejected.

to make some assumptions. The first assumption is that $\{X_{ijt}\}$ is independently and identically distributed over i. This assumption restricts the probability of a price change to be the same across stores but allows for heterogeneity in the size of the price change. For all stores i, $P_{ijt} = P_{jt}$. Note that $E(\varphi_{it})$ and $V(\varphi_{it})$ may vary across stores because of differences in the number and composition of products sold at time t. Next, we assume that $\{X_{ijt}\}$ is independent, not identically distributed, over t.[23]

Under these assumptions, a consistent estimator of P_{jt} under the null hypothesis is the sample mean of X_{ijt} over stores,

$$(5) \qquad \bar{X}_{jt} = \frac{1}{|\mathcal{N}_{jt}|} \sum_{i \in \mathcal{N}_{jt}} X_{ijt}$$

where $\mathcal{N}_{jt}$ is the set of stores selling product j in months t and $t - 1$.

It should be noted that the absence of within-store synchronization does not rule out the possibility that a large proportion of products behave in the same way. This is, in fact, expected due to the high level of inflation during the period. Lack of synchronization merely says that the joint probability of $X_{i1t}, \ldots, X_{iJt}$ is the product of the marginal probabilities for each product; it can be anything between 0 and 1.

T_t was computed 17 times, from February 1978 until June 1979, using $\bar{X}_{jt}$ in place of P_{ijt}. Table 6 presents the number of rejections of the null hypothesis embedded in (4) at the 5- and 10-percent significance levels. Using $\bar{X}_{jt}$ as the estimator of P_{jt} has a simple interpretation: it compares a measure of within-store synchronization (φ_{it}) with an average measure of across-stores synchronization.[24]

If within-store synchronization is the result of a matching between the products sold by the store and an inflationary shock, then φ_{it} and $\bar{X}_{jt}$ should follow similar patterns. In addition, heterogeneity in the inflation process, across products or over time, should not cause much of a difference between φ_{it} and $\bar{X}_{jt}$. Put differently, the results of the test mean that the observed within-store synchronization is unrelated to the actual path of inflation. Since within-store synchronization does not mirror the inflationary process, the reasons for its existence lie somewhere else.

The above arguments can be depicted graphically once we restrict the process $\{X_{ijt}\}$ to be independently and identically distributed over stores, products, and time. Then φ_{it} and $\bar{X}_{jt}$ are identically distributed for any i, j, and t. Figure 5 shows the histograms of φ_{it} and $\bar{X}_{jt}$ in the sample. Note that for both wines and meat products, the distribution of φ has thicker tails than the distribution of $\bar{X}$. In particular, the mass at 0 and at 1, is significantly higher for φ than for $\bar{X}$. As mentioned above, if the observed pattern of φ merely reflected the inflationary process, the same should be true of $\bar{X}$. Figure 5 strongly rejects this possibility.

If the unconditional probability of a price change is the same across products, then, under the null hypothesis, the number of price changes in each store in any month should be distributed as a binomial random variable with parameters $|\mathcal{G}_{it}|$ and the common probability

[23] See Appendix A for the case where X_{ijt} follows a first-order stationary Markov process. The conclusions are not affected by the change in assumptions.

[24] The across-stores synchronization measure is based on $\bar{X}_{jt}$, the proportion of stores changing the price of product j during a month. This measure was, in fact, used to characterize across-stores staggering in Subsection III-A. Table A1 in Appendix A provides these estimates: for wines (meat products) the lowest value of $\bar{X}_{j\cdot}$, the temporal mean of $\bar{X}_{jt}$, is 0.21 (0.40), and the highest is 0.31 (0.62); the mean value is 0.24 (0.53).

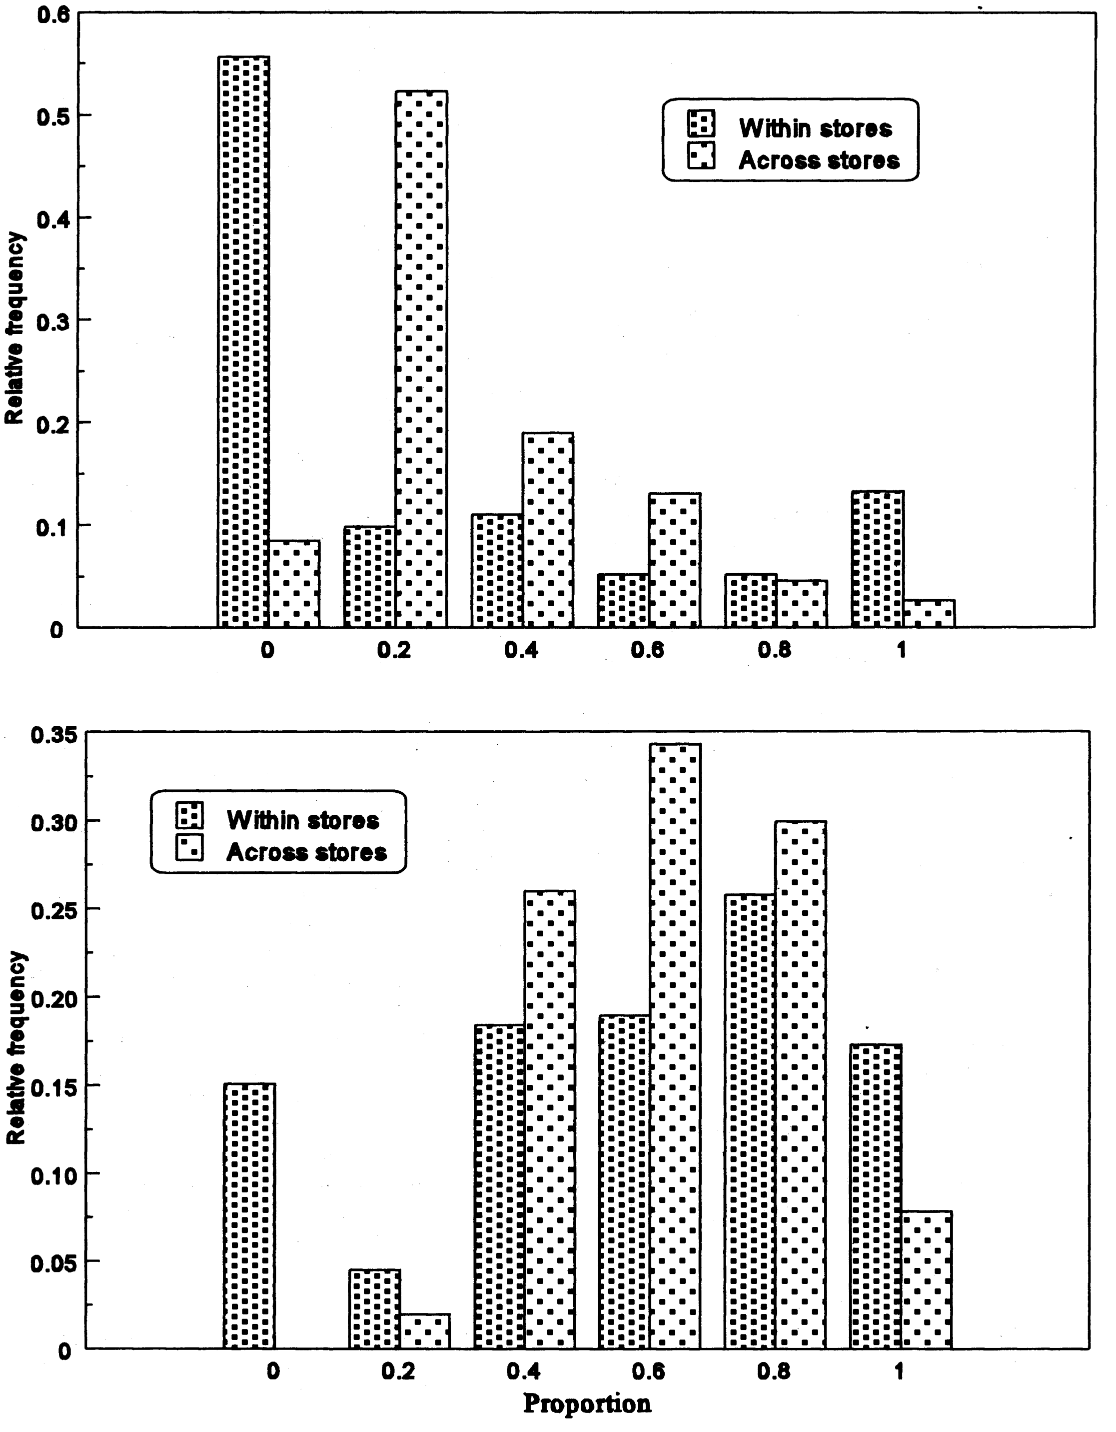

FIGURE 5. PROPORTION OF PRICE CHANGES WITHIN AND ACROSS STORES

Note: The top graph gives the proportion of price changes for wine. The bottom graph gives the proportion of prices changes for meat products.

TABLE 7—OBSERVED AND EXPECTED COUNTS OF EXTREME EVENTS ($\varphi_{it} = 0$ AND $\varphi_{it} = 1$)

Product type		Zero changes $\varphi_{it} = 0$	All changes $\varphi_{it} = 1$
Wines	Observed:	236	46
	Expected:	150.5	10.7
Meat	Observed:	139	111
	Expected:	41.6	69.4

Notes: The observed counts of $\varphi_{it} = 0$ and $\varphi_{it} = 1$ come from Table 5. The expected count of each event under the hypothesis that no within-store synchronization is obtained from the binomial distribution of φ_{it} under this hypothesis. See Subsection IV-A for details.

P_{wt} or P_{mt} for wines and meat products, respectively, estimated by averaging $\bar{X}_{jt}$ over products.

The most compelling evidence in favor of within-store synchronization in the timing of price changes is given by the frequent occurrence of the events $\varphi_{it} = 0$ and $\varphi_{it} = 1$. We can now compare the observed (absolute) frequency of these events (Table 5) with the expected (absolute) frequency under the binomial assumption.[25] Within-store synchronization predicts that the observed frequencies will be higher than the expected ones. Table 7 corroborates this prediction.

B. *Pairwise Correlation in the Timing of Price Changes*

So far our approach to the measurement of within-store synchronization captures behavior within a period. Another—perhaps more dynamic—approach is the coevolution of two different products j and k, X_{ijt} and X_{ikt}, within each store over time. An additional implication of pairwise independence in the timing of price changes is that the covariance over time between any two pairs of products sold in the same store is zero. This issue is analyzed in this subsection, thereby putting together, the concept of regular cyclicity with the static notion of within-store synchronization.

We focus our analysis on the behavior of the cross product $X_{ijt}X_{ikt}$. We define the indicator function $S_{it}(j, k)$ as follows: when both products behave similarly $S_{it}(j, k) = 1$; else $S_{it}(j, k) = 0$. That is, when either $X_{ijt} = X_{ikt} = 1$ or $X_{ijt} = X_{ikt} = 0$, $S_{it}(j, k) = 1$. The mean value of $S_{it}(j, k)$ over time, $S_{i\cdot}(j, k)$, is the proportion of synchronization or matching between two products j and k in store i. Over all stores and pairs of distinct products we obtained 579 and 1,069 $S_{i\cdot}(j, k)$'s for wines and meat products, respectively. Table 8 displays features of the distribution of $S_{i\cdot}(j, k)$. Recall that within-store synchronization implies "high" values of $S_{i\cdot}$.

A rough benchmark figure for the expected proportion of matchings, $S_{i\cdot}(j, k)$, under the assumption of no within-store synchronization can be obtained from the mean and maximum values of $\bar{X}_{jt}$ reported in Table A1 in Appendix A. For wines we are led to expect an $S_{i\cdot}(j, k)$ around 0.0576 ($=0.24^2$) and no larger than 0.0961 ($=0.31^2$), while for meat products $S_{i\cdot}(j, k)$ should hover around 0.281 ($=0.53^2$) and be no more than 0.384 (0.62^2). It is clear that the observed proportions of matchings are larger than the expected ones.[26]

[25] For each value of $|\mathcal{G}_{it}| \geq 3$ and for every month we compute the binomial probabilities of observing zero and $|\mathcal{G}_{it}|$ price changes using the estimated P_{wt} and P_{mt}. These probabilities are multiplied by the number of stores selling $|\mathcal{G}_{it}|$ products to obtain the expected absolute frequency, or count, of zero or $|\mathcal{G}_{it}|$ price changes in each month. This calculation holds for any finite value of the sample size $|\mathcal{G}_{it}|$.

[26] This, of course, does not constitute a formal testing procedure. Lack of within-store synchronization means that the joint probability of observing a price change in both products j and k equals the product of the marginal probabilities of a price change in goods j and k. This means that the covariance over time between X_{ijt} and X_{ikt} is zero. Testing for zero covariances is not pursued here because (i) it is difficult to assign a reliable standard error to the

TABLE 8—CUMULATIVE DISTRIBUTION OF $S_{i\cdot}(j, k)$

Product type	N	Mean	Minimum	Quantile				Maximum
				5	10	25	50	
Wines	579	0.87	0.33	0.59	0.67	0.80	0.89	1.00
Meat	1,069	0.58	0.00	0.27	0.36	0.47	0.59	1.00

Notes: $S_{it}(j, k) = 1$ when products j and k sold by store i in month t behave similarly with respect to the timing of their price change; else $S_{it}(j, k) = 0$. That is, when either $X_{ijt} = X_{ikt} = 1$ or $X_{ijt} = X_{ikt} = 0$, $S_{it}(j, k) = 1$. The mean value of $S_{it}(j, k)$ over time, $S_{i\cdot}(j, k)$, is the proportion of synchronization or matching between two products j and k in store i. See Subsection IV-B for details.

Clearly, meat products and wines do not behave in the same way. Recall that we are analyzing the same time period for each product so that the stores selling these products operate in the same macroeconomic environment. It may be that aggregate variables, such as those related to monetary expansion, or even the average rate of inflation, transmit into meat products with much more noise. In other words, meat products are subject to more idiosyncratic shocks. This may be responsible, at least in part, for the fact that synchronization within the store is not as complete in stores selling meat products as it is in liquor stores.

V. Geographical Shocks and the Presence of Negative and Positive Price Changes

This section examines the coexistence of positive and negative nominal price changes within the store. The phenomenon of negative nominal changes during a period of high inflation is interesting. With an inflation rate of about 3.9 percent per month, one is tempted to think that very few nominal prices, if any, are likely to adjust downward. Our data show that this is not so. About 12 percent of all changes in our sample are downward changes in this period (11.1 percent in meat products and 14.7 percent in wines).

estimator of the covariance, since it depends on the serial correlation pattern of each $\{X_{ijt}\}$ sequence, and (ii) a formal procedure would be based on large-sample theory whose finite-sample properties are unknown. This is a problem since each $S_{i\cdot}(j, k)$ is an average of at most 17 observations and usually considerably fewer than that.

We have argued before that within-store synchronization can result from the existence of store-specific costs of adjusting prices. However, there may be other explanations for this observation. Our competing hypothesis is that monetary shocks are distributed unevenly across geographical regions.[27] The timing of the negative nominal price changes offers a viable way of contrasting the two hypotheses.

Suppose there exist idiosyncratic shocks that are independently distributed across products as well as across stores. Suppose that a store "observes" a negative shock in the market for product j. If there were no store-specific component to the costs of adjusting prices, then the store would adjust the price of product j downward only at the moment the product-specific negative shock arrives. This implies that the timing of negative price changes is *uncorrelated* with the timing of positive price changes. If a negative shock to a particular product in a specific store coincides with a positive regional monetary shock affecting the store—the geographically unevenly distributed shock—then there are weaker incentives to accommodate the negative idiosyncratic shock, since it is partly or fully compensated for by the positive regional shock. In this case, the timing of negative and positive changes in prices within a store ought to be *negatively correlated.*

All the above implications hold under the assumption of no store-specific adjustment costs. If there are store-specific costs to chang-

[27] Note, however, that Israel's area is just under 22,000 square kilometers.

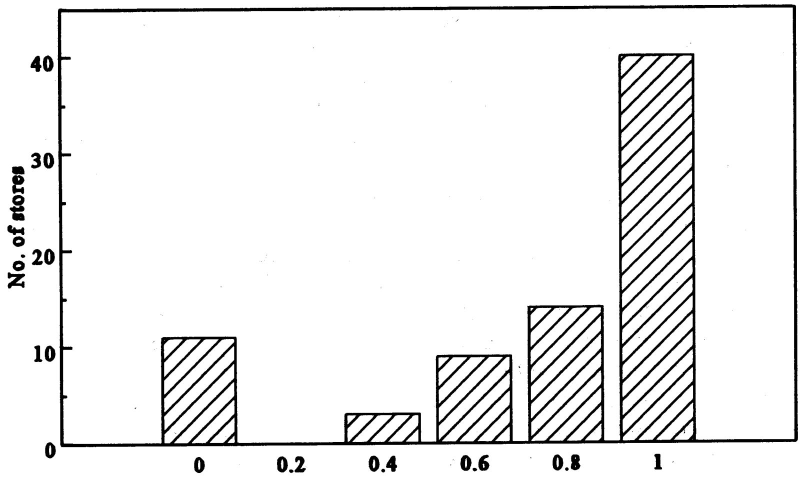

FIGURE 6. COEXISTENCE OF NEGATIVE AND POSITIVE PRICE CHANGES

ing prices, the store should try to bunch together negative and positive changes in prices, implying a *positive* correlation between the timing of positive and negative price changes.

Figure 6 presents the degree to which the timing of positive and negative price changes coincide. The horizontal axis shows the proportion of negative price changes that occur simultaneously (in the same month) with positive price changes within the same store. The vertical axis indicates the frequency counts. In 40 stores all negative price changes coincided with positive price changes, while in 11 stores negative price changes were not accompanied by positive ones.[28] We interpret the left-skewness of Figure 6 as favoring the menu-cost explanation of the existence of within-store synchronization over the explanation of a geographically uneven macroeconomic shock.

VI. Conclusions

A price-setter usually sets prices for many different products. This obvious fact is an aspect of price-setting behavior that has been neglected in most of the theoretical and empirical work on the subject. The purpose of our paper is to draw attention to this issue. We do this by empirically investigating a rich body of data on prices of meat products and wines collected at the store level in Israel.

The data show that, when stores (price-setters) change prices, they change the prices of most of the products they sell. That is, there exists within-store synchronization in the timing of price changes. In addition, stores are staggered in the timing of their price changes. These findings justify the use of staggered price-setting mechanisms in analyses of the inflationary process.

We also contrast the implications of some of the prominent models of price-setting behavior with the data. Among the potential explanations, the one suggested by the menu-cost model seems to be the one most consistent with the data. While we do not formally test the menu-cost model against the other alternative models, we tend to conclude that, at least for foodstuffs, the menu-cost approach describes the data well. The results from LT reinforce this conclusion.

We do not interpret the data as suggesting that the effects of partial information on price dynamics are minimal. The data only suggest

[28] Of these, seven are liquor stores, and four are stores selling meat products.

that, at high rates of inflation, the consequences of incomplete information are overshadowed by the economic implications of the presence of frictions in setting new prices. Thus, this is simply another costly aspect of inflation: at high rates of inflation price-setters must pay more attention to frictions than to gathering and processing information. Inflation therefore makes price-setting behavior a more mechanical process.

To sum up, we believe the empirical findings of within-store synchronization and of across-stores staggering are important because they validate the assumption made in much of the "sticky-prices" literature that decisions are staggered across price-setters, and not across products. Second, they provide further empirical support for the conjecture that price rigidity is due to mechanical reasons, that is, to menu costs, and not to informational asymmetries. And last, they indicate that further research on the dynamics of prices should take into account the multiproduct character of the price-setter.

Appendix A

This appendix presents an alternative scenario for estimating the P_{jt}'s in Subsection IV-A which takes explicit account of dynamics in the X_{ijt} process. Assume that the probability of observing X_{ijt} conditional on all the relevant information available to store i at time t, I_{it}, is the same as that probability conditional only on information on what happened to product j during the previous period:

$$P(X_{ijt}/I_{it}) = P(X_{ijt}/X_{ijt-1}).$$

This assumption embeds the restriction imposed by the null hypothesis of lack of within-store synchronization jointly with a Markovian assumption. Note, also, that the conditional probability is time-invariant, which may be a strong restriction even though the macroeconomic environment (the inflation rate) was quite stable during the period. In a sense this assumption is the complement of the one used in the text. It assumes a particular type of time-dependence for the $\{X_{ijt}\}$ process but restricts it to be stationary over time, whereas the assumption used in the text allows for nonstationarity but assumes independence over time.

Under the assumptions, we can dispense with the store and time subscripts and denote the probability of a price change in product j conditional on X_{ijt-1} as $P_j(0)$ or $P_j(1)$, according to whether X_{ijt-1} is 0 or 1. The stochastic process $\{X_{ijt}\}$ is a time-invariant Markov chain over t. There are different chains for different products, but all stores follow identical processes. The one-step transition probabilities matrix is

$$\mathbf{P}_j = \begin{bmatrix} 1 - P_j(0) & P_j(0) \\ 1 - P_j(1) & P_j(1) \end{bmatrix}.$$

The maximum-likelihood estimators of the one-step transition probabilities are the row percentages in Tables 3 and 4 (top entry in each cell). In order to get the unconditional probabilities appearing in equations (2) and (3) we need to know the probability distribution of the initial state X_{ij0}. Given the initial distribution we can obtain the unconditional probability of a price change at any time t in product j:

$$P_{jt} = (1 - P_j^{(0)})P_j(0)^{(t)} + P_j^{(0)}P_j(1)^{(t)}$$

where, say, $P_j^{(0)}$ is the probability of a price change at $t = 0$ and $P_j(0)^{(t)}$ is the probability of a price change at time t conditional on no price change at $t = 0$. More precisely $P_j(0)^{(t)}$, for example, is the (0, 1) element in the t-step transition-probabilities matrix $\mathbf{P}_j^{(t)}$ ($\mathbf{P}_j$ multiplied by itself t times).

It turns out that, for all practical purposes, there is no need to compute P_{jt}. The limiting probabilities π_j and $1 - \pi_j$ of the Markov chains associated with the matrices in Tables 3 and 4 are arrived at very rapidly: irrespective of the values of the initial probabilities, it takes at most two or three periods to get within three decimal places of the limiting probabilities. That is, P_{jt} is very close to π_j for $t \geq 3$. We therefore use estimates of π_j to estimate P_{jt}. These are given by

$$\text{(A1)} \qquad \hat{\pi}_j = \frac{\hat{P}_j(0)}{1 + \hat{P}_j(0) - \hat{P}_j(1)}$$

where the $\hat{P}$'s are read off directly from Tables 3 and 4.

Table A1 summarizes the features of the different estimates of P_{jt}. The entries are sta-

TABLE A1—UNCONDITIONAL PROBABILITIES OF A PRICE CHANGE

Probability	Mean	Standard deviation	Minimum	Maximum
Wines $\hat{\pi}_j$	0.23	0.02	0.20	0.28
Wines $\bar{X}_{j\cdot}$	0.24	0.03	0.21	0.31
Meats $\hat{\pi}_j$	0.53	0.07	0.38	0.61
Meats $\bar{X}_{j\cdot}$	0.53	0.07	0.40	0.62

Notes: $\hat{\pi}$ is the estimated probability of a change in product j's price assuming that: (i) $\{X_{ijt}\}$ is a time-invariant Markov chain over months t; (ii) $\{X_{ijt}\}$ is pairwise independent over products j (lack of within-store synchronization); and (iii) $\{X_{ijt}\}$ is independently and identically distributed over stores i. [See equation (A1).] $\bar{X}_{j\cdot}$ is the estimated probability of a change in product j's price assuming that: (i) $\{X_{ijt}\}$ is independent, not identically distributed over months t; (ii) $\{X_{ijt}\}$ is pairwise independent over products j (lack of within-store synchronization); and (iii) $\{X_{ijt}\}$ is independently and identically distributed over stores i. $\bar{X}_{j\cdot}$ is the average of $\bar{X}_{jt}$ from equation (5) over t.

tistics corresponding to the distribution of the product-specific estimates. The table indicates that the estimated probability of a price change is robust to the assumptions and method of estimation.

APPENDIX B

Here we explain how we count the potential number of K consecutive price changes in the X_{ijt} time series. Counting is done for each store separately over all products sold by the store and over months for which data are available (the maximum being 17 months per product). We count nonoverlapping spells of consecutive price changes. For example, a spell of four consecutive changes is counted only once as a spell of 4, and not as three spells of 2, or two spells of 3. The qualifier "at least" is important. Since our data are censored from both right and left, there are many instances in which a spell of two consecutive price changes is preceded or followed by a missing value. The censoring results either from a store not being included in the sample in a particular month or from the store having run out of the product at the time of sampling.

The potential number of K consecutive price changes is derived as follows: first, we identify the spells of L consecutive observations on X_{ijt}, $2 \leq L \leq 17$. The reason for having spells of varying size is the presence of many missing values in the X's. Next, for each spell we count the number of possible ways K nonoverlapping consecutive price changes, $X_{ijt} = 1$, can occur. We then sum the number of potential K consecutive prices changes over all observed spells.

For example, the first store sells seven wine products. Two consecutive price changes are observed twice in the first product and once each in the fourth and seventh product. The observed number of two consecutive price changes is, therefore, four.

What is the potential number of two consecutive price changes? First, we identified the spells of consecutive observations. We found two spells of seven consecutive nonmissing X's (in products 1 and 9), three spells of six consecutive nonmissing X's (in products 2, 6, and 7), one spell of five (in product 8), one spell of three (in product 3), and one spell of two consecutive nonmissing X's (in product 3). Simple counting shows that the number of ways to accommodate exactly two consecutive price changes in each spell of length seven, six, and five is two, while the potential number of two consecutive price changes in spells of length three and two is one. Adding over all the observed spells results in potentially 14 pairs of consecutive price changes. The C2 entry for the first store is, therefore, $^4/_{14} = 0.286$.

REFERENCES

Ball, Laurence and Cecchetti, Stephen G. "Imperfect Information and Staggered Price Setting." *American Economic Review*, December 1988, *78*(5), pp. 999–1018.

Benabou, Roland. "Search, Price Setting and Inflation." *Review of Economic Studies*, July 1988, *55*(3), pp. 353–76.

Caballero, Ricardo J. and Engel, Eduardo M.R.A. "Dynamic (*S*, *s*) Economies." *Econometrica*, November 1991, *59*(6), pp. 1659–86.

______. "Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy." *Review of Economic Studies*, January 1993, *60*(1), pp. 95–119.

Caplin, Andrew S. and Leahy, John. "State Dependent Pricing and the Dynamics of Money and Output." *Quarterly Journal of Economics,* August 1991, *106*(3), pp. 683–708.

Caplin, Andrew S. and Spulber, Daniel F. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics*, November 1987, *102*(4), pp. 703–25.

Diamond, Peter A. "Search, Sticky Prices and Inflation." *Review of Economic Studies*, January 1993, *60*(1), pp. 53–68.

Fischer, Stanley. "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule." *Journal of Political Economy*, February 1977, *85*(1), pp. 191–205.

Lach, Saul and Tsiddon, Daniel. "The Behavior of Prices and Inflation: An Empirical Analysis of Disaggregated Data." *Journal of Political Economy*, April 1992, *100*(2), pp. 349–89.

______. "Staggering and Synchronization in Price-Setting: Evidence from Multiproduct Firms." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 4759, June 1994.

______. "Small Price Changes and Menu Costs." Maurice Falk Institute for Economic Research in Israel (Jerusàlem, Israel) Discussion Paper No. 96.06, March 1996.

Liviatan, Nissan and Piterman, Silvia. "Accelerating Inflation and Balance of Payments Crises, 1973–1984," in Yoram Ben-Porath, ed., *The Israeli economy: Maturing through crises*. Cambridge, MA: Harvard University Press, 1986, pp. 320–46.

Lucas, Robert E. "Some International Evidence on Output–Inflation Trade-offs." *American Economic Review*, June 1973, *63*(3), pp. 326–34.

Sheshinski, Eytan and Weiss, Yoram. "Staggered and Synchronized Price Policies under Inflation: The Multiproduct Monopoly Case." *Review of Economic Studies*, April 1992, *59*(2), pp. 331–59.

Sulem, Agnes. "Explicit Solution of a Two-Dimensional Deterministic Inventory Problem." *Mathematics of Operations Research*, February 1986, *11*(1), pp. 134–46.

Tommasi, Mariano. "Inflation and Relative Prices: Evidence from Argentina," in Eytan Sheshinski and Yoram Weiss, eds., *Optimal pricing, inflation, and the cost of price adjustment*. Cambridge, MA: MIT Press, 1993, pp. 485–511.

______. "The Consequences of Price Instability on Search Markets: Toward Understanding the Effects of Inflation." *American Economic Review*, December 1994, *84*(5), pp. 1385–96.

Tsiddon, Daniel. "The (Mis) Behavior of the Aggregate Price Level." *Review of Economic Studies*, December 1993, *60*(4), pp. 889–902.