

Bonomo, Marco Antonio; Viana de Carvalho, Carlos; Kryvtsov, Oleksiy; Ribon, Sigal; Rigato, Rodolfo

Working Paper

Multi-product pricing: Theory and evidence from large retailers

Texto para discussão, No. 687

Provided in Cooperation with:

Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro

Suggested Citation: Bonomo, Marco Antonio; Viana de Carvalho, Carlos; Kryvtsov, Oleksiy; Ribon, Sigal; Rigato, Rodolfo (2021) : Multi-product pricing: Theory and evidence from large retailers, Texto para discussão, No. 687, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Departamento de Economia, Rio de Janeiro

This Version is available at:

<https://hdl.handle.net/10419/249735>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TEXTO PARA DISCUSSÃO

No. 687

Multi-Product Pricing: Theory and
Evidence From Large Retailers

Marco Bonomo
Carlos Carvalho
Oleksiy Kryvtsov
Sigal Ribon
Rodolfo Rigato



Multi-Product Pricing: Theory and Evidence From Large Retailers*

Marco Bonomo, Carlos Carvalho, Oleksiy Kryvtsov, Sigal Ribon, Rodolfo Rigato[†]

October 2020

Abstract

We study a unique dataset with comprehensive coverage of daily prices in large multi-product retailers in Israel. Retail stores synchronize price changes around occasional “peak” days when they reprice around 10% of their products. To assess aggregate implications of partial price synchronization, we develop a new model in which multi-product firms face economies of scope in price adjustment, and synchronization is endogenous. Synchronization of price changes attenuates the average price response to monetary shocks, but only high degrees of synchronization can substantially strengthen monetary non-neutrality. Our calibrated model generates as little monetary non-neutrality as in Golosov and Lucas (2007).

JEL classification codes: D21, D22, E31, E52, L11.

Keywords: Inflation, Prices, Multi-product pricing, Synchronization of prices, Menu cost, Monetary non-neutrality.

*The views expressed herein are those of the authors and not necessarily those of the Bank of Canada, or the Bank of Israel. We thank the IT department at the Bank of Israel for its assistance with the construction of the data. We would like to thank Information Resources Inc. and Mike Kruger for making their data available. We thank Fernando Alvarez, Cezar Santos, and seminar participants at the Bank of Canada, 2019 CEBRA Annual Meeting, 2019 Luso-Brazilian Macroeconomics Meeting, the ECB conference on “Inflation in a changing economic environment”, Richmond Fed, and Université de Montréal for their comments and discussions. Minnie Cui provided superb research assistance.

[†]Bonomo (Insper): bonomo@insper.edu.br, Carvalho (PUC-Rio): cvianac@econ.puc-rio.br, Kryvtsov (Bank of Canada): okryvtsov@bankofcanada.ca, Ribon (Bank of Israel): sigal.ribon@boi.org.il, Rigato (Harvard): dinirigato@g.harvard.edu.

1 Introduction

The literature on the effects of monetary policy on the economy has relied extensively on the role of sticky prices in the transmission mechanism. In standard models, firms sell only one product, which means that they face an *inter-temporal* pricing decision. For example, a firm that faces an unexpectedly high demand can raise the price right away and incur the cost of price adjustment, or it can delay raising the price and face the rising cost of accommodating extra demand. Most retail firms, however, sell hundreds or thousands of products and, therefore, have an additional *intra-temporal* option to change those prices that are worth adjusting. In this paper, we study joint inter- and intra-temporal pricing decisions using a model in which each firm sells a continuum of products. From a macroeconomic perspective, the main question is whether multi-product pricing materially changes the aggregate price response to economic disturbances.

We propose a theory of multi-product monopolistic firms that use a price-adjustment technology with economies of scope in price setting (Sheshinski and Weiss, 1992). In the model, a firm incurs a common cost K for any number of concurrent price changes, and a fixed cost c for each of those price changes. This technology nests two polar cases studied in the literature. In the case with no economies of scope ($K = 0$), the firm sets the price of each product independently, paying a menu cost c for each price change. This case is equivalent to a continuum of single-product firms, each one subject to a menu cost, as in Golosov and Lucas (2007). In the case with maximal economies of scope in price adjustment ($c = 0$), the firm pays the fixed cost K , which allows it to adjust the price of any number of products at no additional cost. This case is similar to models with maximal economies of scope in Midrigan (2011) and Alvarez and Lippi (2014). In both cases, the degree to which the firm synchronizes price changes across products is fixed by construction—it is either zero (with no economies of scope) or full (with maximal economies of scope). Our general framework gives rise to an endogenous degree of synchronization of price changes that depends on the combination of technology parameters c and K .

To guide this new theory, we examine price data from large multi-product food retailers in Israel. Unlike most available data on micro prices, the dataset provides comprehensive detail along both dimensions of a firm’s price adjustment: time and products. A disclosure law enacted in 2014 requires large retailers in Israel—those with annual sales exceeding NIS 250 million (about \$70 million)—to post on their internet sites daily price information for all products sold in their stores. The data used in this paper contain information for stores representing all large retail chains, 25 in total, from May 20, 2015, until October 4, 2019. To manage computational constraints, we only analyze top 5% of stores for each retailer (by

the number of observations), 71 brick-and-mortar stores in total. For four retailers, we also have information about their online prices, which we treat as four additional online stores. We observe the “base,” or regular, prices for almost all products sold in each store, averaging 7,217 per store on a given day.¹ In addition to regular prices, we have constructed the final (discounted) prices for the stores owned by Shufersal, the largest food retailer in Israel. In all, the dataset contains 506.1 million daily observations.

We exploit comprehensive product and time coverage in our data to study the synchronization of price adjustments across products in a store. Figure 1 depicts the average daily fraction of price adjustments for four selected stores from different chains. It is apparent that stores do a majority of their regular price changes during occasional “peak” days. For example, when we define peaks as the subset of the most active days that jointly account for one-half of all price changes in a store over the entire sample, we find that they occur once every three weeks, on average. On a peak day, a store reprices about 9.9% of its products, 20 times the number of price adjustments on an average off-peak day. This behavior is not driven by weekly or monthly seasonal effects, Jewish holidays, differences between online and offline prices, or retail discounting, suggesting that price synchronization is driven by store-specific fundamentals. In all, the pattern of partial synchronization of price changes that we document is unlike the staggered pattern of price changes for independently produced products, as in [Golosov and Lucas \(2007\)](#), or the perfect synchronization pattern in [Alvarez and Lippi \(2014\)](#).²

To account for partial synchronization of price changes in the data, we develop a continuous-time model of multi-product firms facing economies of scope in price adjustment. In the model, each firm sells a continuum of differentiated goods and faces two types of fixed costs when changing price: a fixed cost K is incurred when at least one price adjustment is made, and an additional cost c is paid for each individual price change. The profit-maximizing price of each product follows a Brownian motion with no drift, independent across goods. Due to the presence of the common cost K , there are intervals of time with no price changes. Additionally, the individual menu cost c implies that only prices far enough from their opti-

¹Since the disclosure law applies to only large retailers, the dataset does not include small food retailers, pharmacy retailers, and mom-and-pop stores. We also exclude store-specific products that do not have a general 13-digit barcode (e.g., fruits, vegetables, bakery goods).

²Extensive empirical literature has studied within-store synchronization of price changes. It typically finds that synchronization of price changes within stores is higher than across stores: [Lach and Tsiddon \(1996\)](#), [Midrigan \(2011\)](#), [Bhattarai and Schoenle \(2014\)](#), [Dedola, Kristoffersenz, and Züllig \(2019\)](#). [Levy et al. \(1997\)](#), [Dutta et al. \(1999\)](#) found that retail prices for similar products tend to be more synchronized. [Goldberg and Hellerstein \(2009\)](#), [Bhattarai and Schoenle \(2014\)](#) use producer price data to document that firms selling more products adjust prices more frequently. [Cavallo \(2017\)](#), [Gorodnichenko and Talavera \(2017\)](#), [Gorodnichenko, Sheremirov, and Talavera \(2018\)](#) find little difference in synchronization of price changes in online and offline stores. [DellaVigna and Gentzkow \(2019\)](#), [Hitsch, Hortaçsu, and Lin \(2019\)](#) provide evidence of uniform pricing by U.S. retail chains.

mal levels will be adjusted simultaneously. So, the multi-product firm in our model adjusts prices only infrequently, and when it does, it adjusts the prices of a substantial number of products at the same time, but never all of them. It thus generates the key features of partial synchronization pattern observed in the data.

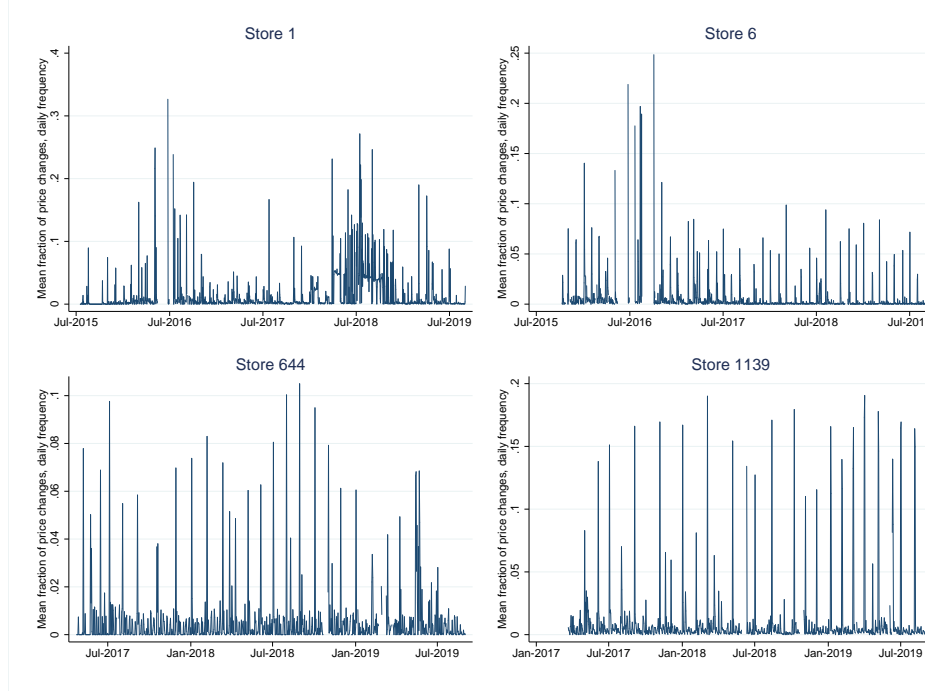


Figure 1: Daily fraction of prices changes, selected stores.

Note: Each plot provides the daily fraction of price changes for one of the four selected (brick-and-mortar) stores in our dataset.

The relevant state variable in the model is the distribution of gaps between the actual and optimal prices across products. Since each firm sells a continuum of products, in steady state this distribution evolves deterministically and we can derive it analytically, even though each individual price gap is stochastic. The firm follows a price-setting policy given by deterministic adjustment dates $\{T_k\}_{k=1}^{\infty}$ and thresholds $\{\bar{x}_k\}_{k=1}^{\infty}$. At each time T_k , the firm adjusts the prices of all products for which the current price differs from its optimal price by a magnitude greater than \bar{x}_k in absolute value. The price-setting policy is characterized by two constants: the time interval between consecutive adjustment dates, $\tau^* = T_{k+1} - T_k$, and threshold, \bar{x}^* .

We then study the responses of the price level and real output to a one-time unanticipated impulse to money supply. These responses are determined by the the selection effect characteristic in price-setting models (Golosov and Lucas, 2007). In a model where single-product firms change prices independently from each other, the shock initially triggers adjustment of

prices that are farther away from the optimal, amplifying the response of the aggregate price.³ In our model, the selection effect is influenced by synchronization of price adjustments. If there is a large degree of synchronization (i.e., many prices are adjusted at the same time) firms may bunch adjustments of prices with smaller price discrepancies, weakening the selection effect. We derive analytically the relationship between the selection effect and monetary non-neutrality, extending Caballero and Engel (2007)’s result to a multi-product setting with endogenous price synchronization. Our model encompasses two existing models as limiting cases. The Alvarez-Lippi-Midrigan model corresponds to the extreme case with no selection effect, and the Golosov-Lucas model is the limiting case in which selection is extremely high and the real effect of the monetary shock is significantly smaller than in the Alvarez-Lippi-Midrigan model.

What is the degree of monetary non-neutrality implied by the synchronization of price changes we document in the data? We address this question numerically by calibrating our model to match three price-setting statistics generated from the Israeli database: the daily fraction and average absolute size of regular price changes, and synchronization of price changes using the Fisher and Konieczny (2000) index. Our numerical results indicate that even small deviations from full synchronization can significantly reduce the persistence of real effects of demand shocks. For the degree of synchronization observed in the data, our model generates responses very close to the Golosov and Lucas (2007) model. Although firms in the partial synchronization model do not change all prices at the same time, they change those prices that are farther from the optimal, triggering larger adjustments shortly after the shock. Hence, the selection effect plays a key role in the partial synchronization model, engineering a faster response of the aggregate price level and thus attenuating the monetary policy effect.

Our numerical simulations imply that monetary non-neutrality is proportional to the kurtosis of price changes for a given frequency of price changes, indicating that the sufficient statistic result of Alvarez, Le Bihan, and Lippi (2016) also holds in a the partial synchronization setting.⁴ In our benchmark setting, the kurtosis of price changes increases monotonically

³Golosov and Lucas (2007) study the selection effect in the recent generation of applied general equilibrium menu-cost models. Earlier theoretical contributions include Caplin and Spulber (1987), Danziger (1999), Caballero and Engel (2007), among others. Carvalho and Kryvtsov (2018) propose a simple, model-free way to measure price selection and its impact on inflation.

⁴The class of models studied in Alvarez, Le Bihan, and Lippi (2016) and Alvarez, Lippi, and Oskolkov (2020) includes the case with a continuum of products but only full price synchronization. Karadi and Reiff (2019), Dotsey and Wolman (2020) provide examples of empirically plausible models in which the conditions for the Alvarez, Le Bihan, and Lippi (2016) result are not satisfied and the sufficient statistic does not hold. In Karadi and Reiff (2019) idiosyncratic shocks are drawn from a mixture of two normal distributions. In Dotsey and Wolman (2020) firms face a multi-state persistent idiosyncratic productivity process and i.i.d. random menu costs.

with the degree of synchronization. Lower synchronization implies a stronger selection effect that, in turn, reduces the kurtosis of price changes relative to the kurtosis of fully flexible price changes. Thus, the degree of synchronization is an independent mechanism that affects monetary non-neutrality in a multi-product setting, which is a generalization of the results in [Midrigan \(2011\)](#) and [Alvarez and Lippi \(2014\)](#).

Our calibrated benchmark model with Gaussian shocks and relatively small synchronization predicts that the kurtosis of price changes is around 1.1, which is smaller than 3.5 in the data. The model also does not generate small price adjustments. Those are empirical hurdles generally pertinent to menu-cost models with Gaussian idiosyncratic shocks ([Midrigan, 2011](#)). To address the discrepancies with the data, we extend the model along two dimensions frequently employed in the literature. First, we replace Gaussian shocks to the firm’s desired price with fat-tailed Poisson shocks as in [Midrigan \(2011\)](#). In this case, the degree to which kurtosis of price changes translates into monetary non-neutrality depends on the degree of price-change synchronization (for the same frequency of price changes). For kurtosis around 3, models with high synchronization generate larger non-neutrality than models with lower synchronization; and the order flips as kurtosis becomes larger. These results imply that with fat-tailed shocks the kurtosis of price changes is no longer a sufficient statistic for monetary non-neutrality, because non-neutrality depends also on synchronization of price changes. When we re-calibrate the Golosov-Lucas, Alvarez-Lippi-Midrigan, and partial synchronization models to match the observed kurtosis of price changes (around 3.5), they predict quantitatively similar cumulative effects of monetary policy, although the shapes of impulse responses are different. Hence, when shocks are non-Gaussian, the determinants of monetary non-neutrality become more complex, but the main takeaway remains intact, i.e., a realistic degree of synchronization implies considerable aggregate price flexibility.

In another extension we introduce a fraction of free adjustments à la Calvo, as in [Alvarez, Le Bihan, and Lippi \(2016\)](#) and [Nakamura and Steinsson \(2010\)](#), to match the fraction of small price adjustments, in addition to other targeted moments. The aggregate results, however, are only slightly changed, because the fraction of small adjustments in our daily data is relatively small. We show that time aggregation in weekly and monthly price data can substantially inflate the share of small price changes, making it a more consequential calibration target.

Contribution to the literature

This paper contributes to recent literature, spurred by [Midrigan \(2011\)](#) and [Alvarez and Lippi \(2014\)](#), which asks whether the “multi-product hypothesis” is material for the effects of monetary shocks. Under this hypothesis, a firm’s ability to adjust prices for its many products may influence how quickly it responds to economic disturbances. The main takeaway from

Midrigan (2011) and Alvarez and Lippi (2014) is that the degree of economies of scope faced by multi-product firms and the resulting synchronization of within-firm price synchronization are *fundamental* for real effects of monetary policy shocks.

The empirical literature has emphasized that firms, especially retailers, synchronize price changes in their stores as a matter of course. Lach and Tsiddon (1996) use a sub-sample of multi-product stores selling wine and meat during the high inflation period in Israel. They find that price adjustments across stores tend to be predominantly staggered, but those within stores tend to be highly synchronized. Bhattarai and Schoenle (2014) document substantial within-firm price synchronization using micro data on U.S. producer prices. Other studies documented that synchronization tends to increase with the number of products, price flexibility, and within categories of similar products.⁵ While our evidence is in line with these findings, we provide a more detailed and systematic breakdown of multi-product pricing along both time and product dimensions. This paper’s empirical contribution is to establish a significant degree of partial price synchronization in a typical retail store which stems from occasional but regular peaks in store’s repricing activity. Since holidays and calendar fixed effects are weak, this synchronization pattern is likely due to store-specific factors, such as fixed cost of changing prices. We show that time-averaging can significantly lower measured degree of synchronization, which underscores the importance of using high-frequency price data.

Theoretical literature has emphasized the fundamental link between within-firm price synchronization and macroeconomic effects of monetary shocks. Midrigan (2011) develops a model with two-product firms facing maximal economies of scope in price adjustment and fat-tailed cost shocks. Midrigan (2011) shows that when the model is calibrated to key moments in micro data, it generates real effects of monetary policy that are five times greater than those effects in the single-firm model in Golosov and Lucas (2007). Bhattarai and Schoenle (2014) demonstrate that models with economies of scope in price adjustment can account for within-firm price synchronization observed for U.S. producers. Alvarez and Lippi (2014) expand Midrigan (2011)’s maximal economies of scope to an arbitrary number of products for the case with Gaussian shocks. They establish analytically that both the size of the output response and its duration increase with the number of products, converging to the response of Taylor (1980) staggered price model when the number of products gets large. Hence, both Midrigan (2011) and Alvarez and Lippi (2014) find that multi-product pricing greatly amplifies the real effect of monetary shocks. This paper’s theoretical contribution is

⁵Studies of within- and across-firm price synchronization include Levy et al. (1997); Levy et al. (1998); Dutta et al. (1999); Fisher and Konieczny (2000); Midrigan (2011); Stella (2014); Letterie and Øivind Anti Nilsen (2016); Cavallo (2017); Gorodnichenko, Sheremirov, and Talavera (2018); Dedola, Kristoffersenz, and Züllig (2019).

to demonstrate that full synchronization of price changes in their models is crucial for this conclusion.

Our framework endogenizes the degree of synchronization by allowing firms to balance the costs of changing their prices across products and over time. Hence, the degree of synchronization in the model is an independent mechanism that affects monetary non-neutrality, with [Goloso and Lucas \(2007\)](#) and [Alvarez and Lippi \(2014\)](#) mechanisms nested as special cases. When we match the general model to the data, we find that the degree of synchronization observed in the data is consistent with weak monetary non-neutrality, similar to non-neutrality in the single-product setup in [Goloso and Lucas \(2007\)](#). Hence, economies of scope in price adjustments that lead to a realistic degree of partial synchronization of price changes do not materially deter multi-product firms from responding to monetary disturbances.

The remaining sections of the paper are organized as follows. Section 2 provides new evidence from price data for large food retailers in Israel. In Section 3, we present a multi-product pricing model with endogenous synchronization of price changes, characterize the firm’s decision functions, and calibrate the model’s parameters. The succeeding section studies analytically and numerically impulse responses to a monetary shock in the model with partial synchronization of price changes and its special cases. The last section concludes.

2 Evidence from the retail stores in Israel

In this section, we describe the novel dataset from large retail stores in Israel. Building on the dataset’s detailed coverage of daily price changes for thousands of products in each store, we analyze the degree to which each store synchronizes price changes across different products in a given day. We document a significant degree of partial price synchronization in a typical retail store which stems from recurrent spikes of the fraction of price adjustments. We show that this synchronization pattern is not due to price discounts, holidays or calendar fixed effects. Finally, we demonstrate that our findings are not exclusive to Israel. We show that synchronization of price changes across stores and across products in a store are in line with previous findings that relied on weekly or monthly price data. We then use weekly scanner data from retail stores in the United States to show that time averaging can significantly wash out high-frequency variation in the daily fraction of price changes.

2.1 The Israeli retail data

Following the “Social Protest” in Israel⁶ and the recommendations of the public committee that was formed in response to this social movement, a “Promotion of Competition in the Food Industry Law” was passed by the Israeli Parliament in 2014. In accordance with this law, large food retailers⁷ operating in Israel are required to publish online daily price information for all products sold in their stores. The law requires retailers to keep data only for the previous three months.

The Bank of Israel scrapes, cleans, and consolidates this information on a daily basis. The historical data include information for all products sold by large retailers—a total of 25 retail chains and around 1,700 stores, which account for most of the volume of food retail sales in the local market from May 20, 2015, until October 4, 2019. To manage computational constraints, for each retailer, we only analyze stores in the top 5%, by the number of observations. Since the disclosure law applies to only large retailers—with annual sales exceeding NIS 250 million (about \$70 million)—the dataset does not include small food retailers, pharmacy retailers, and mom-and-pop stores. We also excluded store-specific products that do not have a general 13-digit barcode (e.g., fruits, vegetables, bakery goods). The remaining dataset contains 506.1 million daily observations for 71 brick-and-mortar stores from 25 retail chains.

For each store, we have information about the “base,” or regular, prices for individual products on a daily basis. For four retailers we also have information about their online prices, which we will treat as four additional online stores. Final prices are based on price discounts (“sales”), which are defined independently and differently by each retailer. Information about price discounts is entered by retailers as a code indicating, for example, a buy-one-get-one-free discount, a third-product-free discount, or two products for 10 NIS (Israeli New Shekel). Based on the available regular price and the discount code, we constructed the final price for the stores owned by Shufersal, the largest food retailer in Israel.⁸

We report most empirical results for regular prices, and in Section 2.3 we review the results for discounted prices.

A combination of two features of this dataset makes it particularly useful for documenting pricing behavior of large retailers: extensive coverage of products in each store and high frequency of product-specific price observations over time. The number of products sold in a given store is large: on average 7,217 products are sold on a given day, with 1,311 (31,847)

⁶See also Chapter 1 in Bank of Israel (2012).

⁷A retailer with annual sales exceeding NIS 250 million (about USD 70 million). According to 2017 Household Income and Expenditure Survey (Table 38 in Israeli Central Bureau of Statistics, 2019), about 60% of food are purchased in chain stores, 15% in groceries, 4% in open market, and the rest in specialty stores.

⁸Shufersal owns 350 stores nationwide, servicing about 35 percent of the food retail-chain market. Eighteen of these stores are in our sample.

products in the smallest (largest) store in the dataset. Due to the relatively short time span of the data and stable near-zero inflation during the sample period, a large proportion of products, around 38%, do not register a regular price change in our sample. Table 1 provides statistics for the fraction and size of price adjustments across retailers.

Sample	Frequency of observations	Num. of products per period	Mean fraction of price changes	Mean absolute size of price changes	Synchronization of price changes
		(1)	(2)	(3)	(4)
<u>A. By store</u>	daily	7,217	0.87%	20.8%	0.259
	weekly	8,170	4.28%	20.2%	0.241
	monthly	9,605	11.41%	20.0%	0.201
Price increases	daily	7,217	0.45%	19.9%	0.218
Price decreases	daily	7,217	0.42%	21.4%	0.201
Offline stores**	daily	7897	1.01%	20.7%	0.280
Online stores**	daily	10,003	0.99%	19.4%	0.343
<u>B. By chain</u>	daily	23,664	0.89%	20.8%	0.235
<u>C. All observations</u>	daily	301,496	0.99%	20.6%	0.206

Table 1: Summary statistics for price adjustments.

*Note: The table provides weighted means for statistic in columns (1)–(4) across stores (Panels A, D), chains (Panel B), or unweighted means (Panel C). Weights are the average number of products in a store per day. Synchronization of price changes in column (4) is measured by Fisher and Konieczny (2000) index defined in Section 2.2. ** We select only brick-and-mortar stores belonging to the same chains as the four online stores in our dataset.*

In addition to daily price changes, we construct statistics for weekly and monthly price changes defined on the basis of the last available price quotes within each week or month. At a monthly frequency, price behavior of Israeli retailers resembles the behavior previously documented in other surveys (Klenow and Malin, 2010). During the sample period, inflation in Israel fluctuated roughly around zero. In a given month, around one-tenth of prices would change across stores, with about an even split between price increases and decreases. Each change is quite large, around 20% in absolute magnitude.

2.2 Synchronization of price changes

The most striking pattern of daily price changes is evident in Figure 1, which plots daily fraction of price changes in the dataset for four selected brick-and-mortar stores from different chains. It shows that occasionally stores reprice a bulk of their products. To be concrete, we define the peaks in price-adjustment activity for a given store as the set of days with the highest number of price changes that together account for half of price changes in all days for that store. Table 2 shows the breakdown of frequencies of price changes for peaks

and the remaining days (off-peaks). Only 5.3% of all days are peaks, i.e., one peak every 19 days on average. On a peak day a store reprices about 9.9% of all products, 20 times the number of price adjustments on an average off-peak day. To emphasize unequal distribution of re-pricing activity, the table also shows the results for the subset of peaks that together account for 25% of all price changes in a store. Only 1.9% of all days are such peaks (one peak every two months), but on such a day a store reprices about 15.1% of all products, 22 times the average number of price adjustments on other days.

Sample	Number of days	Frequency of price changes	
		weighted	unweighted
Peak days (50%)	79	9.89%	9.49%
Off-peak days (50%)	1419	0.48%	0.44%
Peak days (25%)	28	15.09%	14.50%
Off-peak days (75%)	1470	0.67%	0.62%
All days	1498	0.87%	0.81%

Table 2: Frequency of price changes for peak and off-peak days.

Note: Table 2 shows the frequency of price changes for peaks and off-peaks by store. For each store, compute the number of price changes in each day. Order days by this number in ascending order. Divide days into two groups, where “peaks” are the days with the highest number of price changes that together account for 50% (or 25%) of all price changes in the store, and “off-peaks” are the remaining days. For each group, compute the weighted and unweighted mean fraction of price changes across stores. Weights are the average number of products in a store per day.

To quantify the degree of synchronization of price changes in the data, we use the index constructed in the spirit of [Fisher and Konieczny \(2000\)](#) as follows:

$$FK_s \equiv \sqrt{\frac{\frac{1}{N_s} \sum_t (Fr_{s,t} - \overline{Fr_s})^2}{\overline{Fr_s} \cdot (1 - \overline{Fr_s})}},$$

where FK_s is the index for store s , $Fr_{s,t}$ is the fraction of price changes in store s period t , $\overline{Fr_s}$ is its mean over t , and N_s is the number of price changes in store s . By construction, $FK_s = 0$ when price changes are perfectly staggered, and $FK_s = 1$ when they are perfectly synchronized.⁹ Figure 2 provides the distribution of FK index values for daily frequency across stores. The weighted mean of the index is 0.259 on average, and it varies considerably across stores, with interquartile range of 0.098. Intuitively, the increase in synchronization of price changes between stores in the 25th and 75th percentiles of the FK index (0.218 versus 0.295) corresponds to an increase in the average distance between the

⁹See also [Dias et al. \(2005\)](#) for the properties of this index and a statistical test for staggering of price-setting based on it.

repricing peaks of 19 days, i.e., roughly doubling that distance.¹⁰




Figure 2: Distribution of the Fisher-Konieczny synchronization index across stores.

To measure price synchronization, the literature also relied on statistics derived from probit or logit regressions (Midrigan, 2011; Bhattarai and Schoenle, 2014; Dedola, Kristoffersenz, and Züllig, 2019). However, these statistics are no longer reliable when the distribution of the fraction of price changes places a lot of weight on extremes, as is the case in our analysis. We demonstrate this result in Appendix A.3, where we also provide the estimated probit regression coefficients.¹¹

Peaks are present in all stores in the sample, although there is substantial variation in the timing of peaks across stores. There are apparent chain effects in price adjustment: peak days are highly synchronized across stores belonging to the same chain, rather than across chains. Table 1 reports that once we pool observations across stores in the same chain, synchronization of price changes reduces only marginally. We do not find significant differences between synchronization of regular price increases and decreases.

Finally, re-pricing peaks do not appear to reflect price adjustments for a specific subset of products. When we compared product-level frequencies of price changes, separately, for

¹⁰Using the model in Section 3 we show analytically that the FK index is proportional to $\sqrt{\tau^*}$, where τ^* is the distance between repricing peaks (assuming fixed frequency of price changes and no off-peak price adjustments). This mapping fits the empirical data well; for example, it predicts that the distance between peaks for stores in the 25th and 75th percentiles of the FK index differs by 83%, in line with what we observe directly.

¹¹For robustness, we analyze yet another synchronization measure that is well-behaved when the fraction of adjustments has peaks. This measure is akin to the Gini inequality index. It is based on the Lorenz curve that maps cumulative fraction of price changes against the percentile of days according to the number of price changes in a day, ranked in ascending order. We demonstrate in Appendix A.3 that FK and Gini synchronization measures are qualitatively in line with each other across different cuts of the data.

peak and off-peak days, we found that for most products prices tend to adjust both on peak and off-peak days (Appendix A.2).

2.3 Price discounts

Do these patterns in adjustment of regular prices also apply to final prices, which incorporate various types of price discounts? Although our main dataset contains only regular prices, we collected information on price discounts for 10 stores of the largest retail chain in Israel, Shufersal, from January 2016 until mid-2019. We constructed the final price based on the available regular price and the discount code, indicating, for example, a buy-one-get-one-free discount, a third-product-free discount, or two products for 10 NIS (Israeli New Shekel). Corresponding tables and figures are in Appendix A.4.

Price discounts, or “sales,” are common in the data, accounting for 26% of all price observations. A typical price discount is a large and temporary reduction in price (Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008). A sale is associated with a discounted price that is on average 24% lower than the corresponding regular price, and it lasts around 49 days. Since final prices incorporate discounts, they change more frequently and by a larger magnitude than regular prices. The mean fraction of final price changes is 25.7% per month (10.0% for regular price changes), and the mean absolute size of those changes is 23.2% (19.2% for regular price changes).

We find that our measures of synchronization yield similar results for final and regular prices. Only 4.8% (1.9%) of days in the Shufersal sample account for 50% (25%) of all final price changes, which is close to 3.3% and 1.1% of days for regular price changes. The mean Fisher-Konieczny index values for final prices, 0.355, are only a bit higher than 0.262 for regular prices, possibly reflecting higher average fraction of adjusting final prices. These results suggest that retailer’s decisions to post price discounts are largely independent of decisions to change regular prices. In particular, there are no clear peak days for changing price discounts in the store like we observe for regular prices.

2.4 Calendar fixed effects

Calendar events—holidays and week/month fixed effects—are not important factors for peaks in price adjustments.

Holidays. We study the overlap of repricing peaks with holidays or holiday eves in Israel. Out of 1598 days in the sample, 138 are national holidays. On some of the holidays (e.g., religious holidays) stores are closed. The overlap between the peaks in store and the holidays is small: on average, only 6.5% of peak days are holidays and, in turn, only 4.1% of

holidays are peaks. These frequencies are similar if we aggregate store observations by chain or the entire sample (see Appendix A.5).

During the week before a holiday, the likelihood of a peak day falls by about one-quarter. Using the subset of the Israeli data for the largest food retailer, Shufersal, we find that during weeks before the holidays, retailers increase the frequency and size of price discounts, especially before two major holidays—Rosh Hashana (New Year) and Pessach (similar to Easter). The effect of weeks prior to holidays on the final price (the sum of the regular price and the discount) is not significantly different from zero. This evidence suggests that prior to holidays, retailers shift their pricing activity toward price discounts. This is consistent with evidence in [Warner and Barsky \(1995\)](#) that retailers time their markdowns to periods of high demand and fierce competition between chains.

Week/month fixed effects. We also looked at the prevalence of peaks by day of the week and by day of the month. Appendix A.6 provides distributions of the daily price changes and their frequency by day of the week (month). Peaks tend to fall more frequently early in the week (Mondays and Tuesdays) and are the least frequent during the weekend (Fridays and Saturdays). This pattern is somewhat offset by the higher fraction of price adjustments during peaks that fall on Fridays (11%) and Saturdays (14%) relative to working days (around 8%). Peaks are twice as likely to occur during the first week of the month than on any other week. The number of adjustments during a peak day does not depend on the day of the month.

This evidence is in line with earlier work by [Levy et al. \(1997\)](#), [Levy et al. \(1998\)](#), [Dutta et al. \(1999\)](#) who found some evidence of time-dependent pricing using weekly price data for U.S. supermarket and drugstore chains. For example, [Levy et al. \(1997\)](#) reports that most prices in a store are changed on Sundays and Mondays. Since our data is daily, we can provide a more accurate account of how much calendar events overlap with re-pricing peaks. We find that re-pricing peaks are specific to each store, and price changes show little synchronization across retailers. In all, only some of the partial synchronization of price changes by stores can be associated with holidays and week/month fixed effects.

2.5 How representative are these findings?

We argue that overall these findings are not specific to Israel. We first find that synchronization of price changes across stores and across products in a store is in line with what was documented in the previous literature. Second, we point out that comparisons with other micro data are obscured by time aggregation of daily price observations. Using weekly transactions level data for the U.S. as the case study, we show that time aggregation can double the share of small price changes, but does not change their dispersion or kurtosis. Finally,

we demonstrate that spectral densities of the fraction of regular price changes display a high mass at high frequencies for both Israeli and U.S. data. In this section, we summarize the results, relegating details to Appendices A.2 and A.7.

Synchronization of price changes across stores and across products in a store

We observe, first, that synchronization within stores is higher than across stores. This finding is in line with findings from previous studies of various datasets: [Lach and Tsiddon \(1996\)](#) using selected food prices in Israel, [Midrigan \(2011\)](#) using U.S. Dominick’s scanner data, [Bhattarai and Schoenle \(2014\)](#) using U.S. PPI data, and [Dedola, Kristoffersenz, and Züllig \(2019\)](#) using Danish PPI data.

Next, we show that price changes are more synchronized within, rather than across, categories. Previous empirical studies of retail micro data found that prices for similar products tend to be more synchronized ([Levy et al., 1997](#); [Dutta et al., 1999](#)). Furthermore, we find that price changes for products with more flexible prices are more synchronized.

Third, price changes in larger stores tend to be more synchronized than in smaller stores. This finding is consistent with [Goldberg and Hellerstein \(2009\)](#) and [Bhattarai and Schoenle \(2014\)](#), who use producer price data from the Bureau of Labor Statistics to document that firms selling more products adjust prices more frequently. Using Danish PPI micro data, [Dedola, Kristoffersenz, and Züllig \(2019\)](#) find that synchronization of price changes is stronger within firms than across firms in the same industry, and that within-firm synchronization increases with the number of goods per firm.

Next, the frequency and size of price changes are very similar for online and offline prices, but online prices are more synchronized. These findings are consistent with [Cavallo \(2017\)](#), who studies price adjustments of large retailers that sell both online and offline. He finds that the frequency and size of online and offline price changes are similar, and that there is little synchronization between online and offline price changes. [Gorodnichenko and Talavera \(2017\)](#) and [Gorodnichenko, Sheremirov, and Talavera \(2018\)](#) use online price data from price-comparison websites and an online shopping platform. They find that relative to prices in brick-and-mortar stores, online prices change more frequently, but otherwise they exhibit similar behavior; in particular, prices posted by the same seller are similarly synchronized. Online stores in our dataset have a much larger scale, offering 10,003 products per day, which may explain the higher synchronization of price changes relative to prices in conventional stores. We also compare only stores of the same chain, which helps us control for chain effects.

Finally, we show that synchronization of price changes across products of the entire chain is only slightly lower than at a store level. This suggests that retailers actively synchronize

price changes across their stores. This is consistent with recent findings of uniform pricing by DellaVigna and Gentzkow (2019) and Hitsch, Hortacısu, and Lin (2019) for U.S. retail chains.

Moments of the distribution of price changes

We compare our results for Israel with those from the Information Resources, Inc (IRI) data for the United States. The IRI marketing and market research dataset contains weekly transactions (“scanner”) data for 31 grocery products, such as milk, beer, coffee, razors, laundry detergent, and frozen pizza, across 50 U.S. metropolitan areas.¹²

Measured synchronization of weekly regular price changes in IRI data is much lower than synchronization of weekly regular price changes in the Israeli data, 0.098 versus 0.241. In addition to differences between datasets,¹³ some of the difference could be due to time aggregation of daily price changes to a weekly frequency. To better account for time aggregation in dataset comparisons, we produce two additional sets of statistics: the spectrum of store-level price changes and the fraction of small price changes.

Cavallo (2018) who shows that time aggregation can materially affect the distribution of price changes. Furthermore, the mass of small price changes (Klenow and Kryvtsov, 2008) and the thickness of the tails of the distribution (Midrigan, 2011; Alvarez, Le Bihan, and Lippi, 2016) have been related to menu-cost models’ predictions for the degree of monetary non-neutrality.

Following the literature, we filter out heterogeneity across stores by constructing “standardized” log price changes for each product. Our results show important influence of time averaging on the distribution of price changes. Time aggregation from daily to weekly frequency roughly doubles the mass of small price changes in Israel (Table 3). At weekly or monthly frequency, kurtosis and the mass of small price changes is more in line with the scanner data, although slightly lower in Israeli data. The results are not driven by heterogeneity across products.

¹²More details are provided in Bronnenberg, Kruger, and Mela (2008).

¹³Higher frequency of weekly regular price changes in IRI data, 16.0%, than in the Israeli data, 4.3%, can be attributed to mandatory item price tags in Israel. Levy et al. (1997) examine the process for changing prices in five U.S. supermarket chains. They show that supermarkets situated in states that require a separate price tag on each item face two-and-a-half times higher menu cost and adjust their prices twice less frequently than supermarkets not facing such requirements. At the same time, lower synchronization of price changes in the U.S. data suggests that store-specific cost of price adjustments—costs of collecting information, making and implementing pricing decisions—are also lower in the U.S. relative to Israel.

	Regular prices, by store			Bhattarai and Schoenle	Midrigan AC Nielsen	Midrigan Dominick's
	Daily	Weekly	Monthly	Monthly	Weekly	Weekly
A. Israel						
$\text{std}(z_{ist}) / \text{mean}(z_{ist})$	0.68	0.67	0.69			
$\text{kurtosis}(z_{ist})$	3.5	3.3	3.4			
$\text{fraction } z_{ist} < 0.5 * \text{mean}(z_{ist})$	0.14	0.27	0.28			
$\text{fraction } z_{ist} < 0.25 * \text{mean}(z_{ist})$	0.06	0.10	0.11			
B. United States						
$\text{std}(z_{ist}) / \text{mean}(z_{ist})$		0.85	0.91	1.55	0.72	0.81
$\text{kurtosis}(z_{ist})$		3.4	4.7	17	3.6	4.5
$\text{fraction } z_{ist} < 0.5 * \text{mean}(z_{ist})$		0.37	0.38	0.50	0.25	0.31
$\text{fraction } z_{ist} < 0.25 * \text{mean}(z_{ist})$		0.17	0.20	0.38	0.10	0.14

Table 3: Standardized regular price changes, Israel and the United States.

Note: Standardized log price changes z_{ist} for product i , store s and time t are defined by subtracting from log price p_{ist} the mean μ_s of non-zero log price and dividing by their standard deviation σ_s changes at a store level: $z_{ist} \equiv \frac{p_{ist} - \mu_s}{\sigma_s} \Big|_{p_{ist} \neq 0}$. We also report these statistics at weekly and monthly observation frequencies. Entries for [Bhattarai and Schoenle \(2014\)](#) (10 goods) and [Midrigan \(2011\)](#) (regular prices) are taken from Table II in [Alvarez and Lippi \(2014\)](#).

Spectral densities of the fraction of price changes in a store

A useful way of circumventing the time aggregation problem is to examine the spectral densities of store-level time series of the frequency of price changes.¹⁴ We apply spectral analysis of three store samples: regular price changes in all retail stores in the Israeli data, regular and discounted price changes in Shufersal stores in the Israeli data, and regular price changes in IRI data. For each store in the sample, we compute spectral density of its daily and weekly fraction of price changes. We then compute the weighted mean of the store-level spectra, where weights are the average number of products in a store per day.

For the Israeli stores, spectral density is high for short frequencies of less than 30 days (Figure 3). This remains the case as we aggregate daily price changes to weekly frequency and construct corresponding weekly spectra. Heavy short-frequency spectra are consistent with the time series patterns we document for daily price changes—a bulk of them occur during peak days, which occur more frequently than once a month. For the most part, peaks do not occur at regular time intervals, which elevates the whole short-frequency part of the spectrum. Including price discounts does not alter the shape of the spectra, which is consistent with our conclusion that peaks are not driven by sales. The weekly IRI spectrum is also heavier on the short end, suggesting high-frequency movements in price changes similar

¹⁴We thank Fernando Alvarez for this insight.

to those we find in the Israeli data.

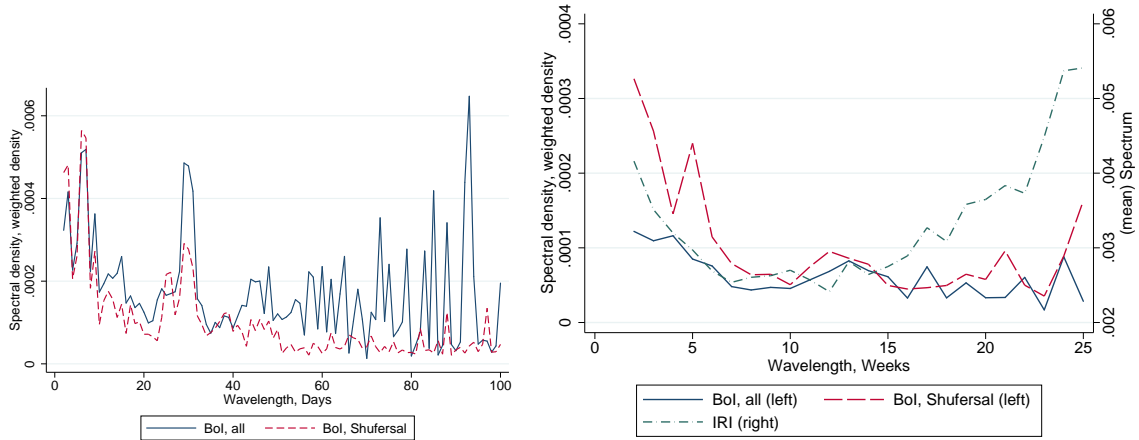


Figure 3: Spectral density of regular prices changes.

Note: Figures provide average spectral density of store-level regular price changes for all retail stores in the Israeli data (“BoI”); regular and discounted price changes in Shufersal stores in daily Israeli data (“BoI, Shufersal”); and regular price changes in weekly IRI data. For each store in the sample, we compute spectral density of its daily and weekly fraction of price changes. We then compute the weighted mean of the store-level spectra, where weights are the average number of products in a store per day. Spectrum frequencies are daily (left) and weekly (right).

3 A model with multi-product firms and partial synchronization

3.1 Overview

We develop a continuous time model of price setting in which each firm sells a continuum of differentiated goods, with total mass normalized to one. We also refer to these goods as varieties or products. Given the large number of products that stores in our data sample sell, the assumption that firms in our model sell a continuum of products is suitable for our purposes. Each variety is indexed by $i \in [0, 1]$ and has a frictionless optimal price $p_{i,t}^*$, where t indexes time. All prices are in log units. The frictionless optimal price $p_{i,t}^*$ is the profit-maximizing price for good i . Absent frictions of any nature, a firm always charges the frictionless optimal price. In the presence of adjustment costs, good i ’s price at instant t , $p_{i,t}$, may differ from the frictionless price, leading to a profit loss.¹⁵ For a firm with a continuum of products, we can write the profit loss at time t , using a second-order approximation, as

¹⁵Appendix B.1 provides the microeconomic foundations for our general equilibrium model, and is based on Alvarez and Lippi (2014) and Bonomo et al. (2019).

follows:

$$L_t = \int_0^1 (p_{i,t} - p_{i,t}^*)^2 di. \quad (1)$$

Intuitively, this expression is the sum of profit losses associated with sub-optimal prices across all products. Firms discount future costs at a rate ρ . We assume that each product i 's frictionless optimal price follows a Brownian motion:

$$dp_{i,t}^* = -\sigma dW_{i,t},$$

where $W_{i,t}$ is a variety-specific standard Brownian motion assumed to be independent across goods, and σ is a parameter that captures the volatility of this process that is common across varieties. In order to retain tractability, we assume that there are no strategic complementarities, i.e. the optimal price of a given product does not depend on the prices of other varieties. This greatly simplifies solving for the aggregate price level following a common shock to all price gaps, as we do in Section 4, because it eliminates feedback effects from firms' pricing decisions to desired prices. Although feasible, the introduction of such complementarities would make the relationship between synchronization and monetary non-neutrality less transparent. We therefore leave it to future work.

It is simpler to express the firm's problem in terms of *price discrepancies*, or *price gaps*, which are defined as $x_{i,t} = p_{i,t} - p_{i,t}^*$. The law of motion above for $p_{i,t}^*$ implies that price discrepancies, in the absence of price adjustments, are also Brownian motions of the form

$$dx_{i,t} = \sigma dW_{i,t}.$$

It is convenient to state the loss function in terms of the distribution of these price discrepancies. Let $g_t(x)$ be the probability density function (p.d.f) that describes the distribution of price gaps.¹⁶ We can express the loss term (1) as

$$L_t = \int_{-\infty}^{+\infty} x^2 g_t(x) dx. \quad (2)$$

This is essentially a change of variables in equation (1). Instead of summing the losses associated with each product, we now sum the loss x^2 associated with each price gap $x = p - p^*$, multiplied by the number of times (or density, more specifically) that such a gap occurs $g_t(x)$. The evolution of $g_t(x)$, given an initial distribution $g_0(x)$, is given by a Kolmogorov

¹⁶The distribution of price discrepancies may have atoms following adjustment dates, and would thus not be expressible as a probability density function. We omit this for simplicity.

forward equation (KFE):

$$\frac{\partial g_t}{\partial t}(x) = \frac{\sigma^2}{2} \frac{\partial^2 g_t}{\partial x^2}(x). \quad (3)$$

The last component of the firms' problem is the specification of the pricing constraints. We assume that firms face menu costs of two different kinds. First, there is a fixed cost K that firms are required to pay to make any number of price adjustments. Second, there is a unit cost c that must be paid for each price adjustment. More precisely, since firms sell a continuum of varieties, c is a cost per measure of adjusted prices. Therefore, a firm that adjusts the prices of a measure m of its products in a single date must pay $K + cm$.

These pricing constraints give rise to optimal policies that have two important features. First, a positive common cost K generates common inaction: there will be time intervals in which the firm does not adjust any price. Second, individual menu cost c implies that when a firm adjusts its prices, it does not adjust all of them. Intuitively, there will always be products with arbitrarily small price discrepancies (in absolute value) for which it is not optimal to pay the unit menu cost c and set the discrepancy to zero.

The relevant state variable in the model is the distribution of price gaps. Since each firm sells a continuum of products, in steady state this distribution evolves deterministically, according to (3). Given an initial distribution $g_0(x)$, the optimal policy consists of sequences of deterministic adjustment dates $\{T_k\}_{k=1}^\infty$ and thresholds $\{\bar{x}_k\}_{k=1}^\infty$ such that at instant $t = T_k$ firms adjust all prices that have price gaps x larger, in absolute terms, than \bar{x}_k , that is $|x| \geq \bar{x}_k$. Since there is no drift in the discrepancies' Brownian motions, all reset prices have discrepancies optimally set to zero.¹⁷ Consequently, the distribution of discrepancies will feature a Dirac mass at $x = 0$ at adjustment dates. These Dirac masses are, however, instantly dissolved by the diffusive nature of the Brownian motion.

To see the intuition about why the optimal policy takes the form of thresholds $\{\bar{x}_k\}_{k=1}^\infty$ just described, notice that given that the common cost K is paid, adjusting the price of a given good does not affect other goods' price gaps and the expected flow of future costs that arise from them. Therefore, after paying the common cost, the firm will decide independently to adjust the price of a product whenever its price gap is high enough to make the benefit of adjustment higher than c . Since the benefit of adjusting a product's price is increasing in its price gap, there is a minimum price gap level, \bar{x}_k , such that it is optimal to adjust all

¹⁷In the presence of a non-zero drift, the optimal sequence of thresholds has to be split into sequences of upper thresholds $\{\bar{x}_k\}_{k=1}^\infty$, lower thresholds $\{\underline{x}_k\}_{k=1}^\infty$, and targets $\{x_k^*\}_{k=1}^\infty$ such that a price is adjusted at date T_k only if the corresponding discrepancy x satisfies either $x \geq \bar{x}_k$ or $x \leq \underline{x}_k$. The discrepancy is set to x_k^* , which is not necessarily zero. In other words, a positive (negative) inflation rate causes the expectations that price gaps will fall (rise). In this case, adjusting firms will optimally reset prices to a level above (below) their frictionless optima, even though costs are not being instantaneously minimized by this decision.

products with price gaps larger than it.

3.2 Recursive formulation

The main difficulty in solving the partial synchronization model is that the relevant state variable in the dynamic optimization problem is the entire distribution of price discrepancies. [Alvarez and Lippi \(2014\)](#) show that, in the perfect synchronization case ($c = 0$), there is no need to keep track of the whole distribution. In their model, all relevant information for the firm can be summarized by a one-dimensional object, namely the loss term (1). This does not apply in our framework, and we must state the Bellman equation for a value function that takes as input an infinite dimensional object. Before proceeding to the recursive formulation, however, it is convenient to go through two simple mathematical results.

Lemma 1. *Let $\phi(\cdot)$ denote the p.d.f. of a standard normal distribution. Given an initial condition $g_0(x)$, the solution of the KFE (3) is*

$$g_t(x) = \int_{-\infty}^{+\infty} \frac{1}{\sqrt{\sigma^2 t}} \phi\left(\frac{x-y}{\sqrt{\sigma^2 t}}\right) g_0(y) dy. \quad (4)$$

Proof. See Appendix B.2. □

Lemma 2. *In the absence of price adjustments, the loss term (2) evolves linearly according to*

$$L_t = L_0 + \sigma^2 t. \quad (5)$$

Proof. See Appendix B.2. □

Lemma 1 expresses the distribution of price gaps at time t as the convolution of the initial distribution g_0 and a normal p.d.f. This is intuitive: the distribution at time t corresponds to the “sum” of the initial distribution and an independent Gaussian shock, which reflects the fact that price gaps follow independent stochastic processes with normally distributed increments. Lemma 2 simply reflects the fact that the variance of a Brownian motion increases linearly with time, as it is the sum of independent, equally distributed components.

Now let $V(g)$ denote the value function of a firm, which takes as input the distribution g of price gaps. We shall state the problem recursively for the case in which g is the distribution of price discrepancies immediately after the payment of the fixed cost K , but *before any price adjustments take place*. Such a choice for the state variable is convenient for the numerical procedure we adopt, which involves using a simple, yet precise, approximation for

this distribution, as explained in Appendix B.3. The function V then satisfies

$$V(g) = \min_{\bar{x}, \tau} c m(\bar{x}, g) + \int_0^\tau e^{-\rho t} (L_0 + \sigma^2 t) dt + e^{-\rho \tau} [K + V(g_\tau)], \quad (6)$$

where the choice variable \bar{x} is the threshold such that prices with gaps larger than \bar{x} are reset, and τ is the amount of time the firm decides to wait until the next price-adjustment date. The function $m(\bar{x}, g)$ is the mass of reset prices, defined as

$$m(\bar{x}, g) = \int_{|x| \geq \bar{x}} g(x) dx.$$

L_0 is the instantaneous loss associated with the intermediate distribution $g_0(x)$, which is the distribution of price discrepancies after adjustments are made, given by

$$g_0(x) = g(x)1(|x| < \bar{x}) + m(\bar{x}, g)\delta_0(x),$$

where $1(\cdot)$ is an indicator function and $\delta_0(x)$ is the Dirac function centered at the point $x = 0$. Since prices with discrepancies larger than \bar{x} are reset, the distribution $g_0(x)$ is simply $g(x)$ with the tails removed and their mass sent to the origin, as adjusted prices have zero discrepancies. Finally, $g_\tau(x)$ is the solution of the KFE (3) at the next adjustment date τ , given the initial condition $g_0(x)$ and computed using (4).

The meaning of (6) is the following. After paying the fixed cost K , the firm adjusts prices that correspond to the tails of the distribution of price gaps ($|x| > \bar{x}$) that amount to a mass $m(\bar{x}, g)$ of products, and consequently, the firm pays $c m(\bar{x}, g)$ in unit costs. After resetting prices, the firm is left with a new distribution of price discrepancies $g_0(x)$ that generates instantaneous loss L_0 . Since the evolution of $g_0(x)$ is deterministic, given by (3), the firm then chooses how long to wait (τ units of time) until the next price-adjustment date, when it pays the fixed cost K and obtains continuation value $V(g_\tau)$. In the meantime, the firm incurs losses that grow linearly over time, as given by (5).

Finally, solving the Bellman equation above gives us optimal policies $\bar{x}(g)$ and $\tau(g)$. We then define a *steady-state distribution* g^* as a p.d.f., with corresponding optimal policies $\tau^* = \tau(g^*)$ and $\bar{x}^* = \bar{x}(g^*)$, which remains unchanged after the process of resetting prices according to the discrepancy threshold \bar{x}^* and waiting time periods τ^* until the next adjustment date. Therefore, when the system starts from the distribution g^* , the trajectory of the state distribution repeats itself every τ^* periods. Formally, we have:

Definition 1. A *steady-state distribution* is a p.d.f. g^* and optimal policies $\tau^* = \tau(g^*)$ and

$\bar{x}^* = \bar{x}(g^*)$, which satisfies the fixed-point problem:

$$g^*(x) = \int_{-\infty}^{+\infty} \frac{1}{\sqrt{\sigma^2 \tau^*}} \phi\left(\frac{x-y}{\sqrt{\sigma^2 \tau^*}}\right) [g^*(y)1(|y| < \bar{x}^*) + m(\bar{x}, g^*)\delta_0(y)] dy. \quad (7)$$

The system is easily solvable by creating a discrete grid for possible values of x , since (7) writes g^* as a linear transformation of the term in square brackets. This term, which is itself a linear function of g^* , is the distribution of price gaps immediately after prices are adjusted. If we represent $g^*(.)$ by a vector of the values it attains in the x grid, the problem boils down to finding an eigenvector of a large matrix. Figure 4 shows the steady-state distribution of price gaps g^* and g_0^* , before and after price adjustments respectively, and price changes for illustrative parameter values. The spike at $x = 0$ is the finite grid analog of a Dirac mass, and the price-change distribution corresponds to the tails of the stationary distribution g^* .




Figure 4: Steady-state and price-change distributions for parameter values $\rho = 0.04$, $\sigma = 0.25$, $K = 0.0001$, $c = 0.001$.

Finally, Figure 5 shows the share of reset prices on a daily basis for two different parameterizations. We can see that, similar to the data, our model generates a spiky pattern for this statistic over time. It is also interesting to notice how the combination of common and individual costs alters this pattern. A high common cost K , combined with a low individual cost c is associated with higher but infrequent spikes, as expected. On the other extreme, a low K , high c parameterization generates frequent but small peaks.




Figure 5: Daily fraction of prices changes for different combinations of menu costs. Other parameter values: $\rho = 0.04$, $\sigma = 0.25$.

3.3 Interpretation of price-adjustment costs

Our theory incorporates two-dimensional cost of price adjustment, in which one component is product-specific (cost c) and the other is common for products in the store (cost K), as in the early study by [Sheshinski and Weiss \(1992\)](#). Empirical literature documents multifaceted nature of price adjustments by retailers, including the labor costs of printing and changing item price tags, the cost of gathering relevant information, and the cost of making and implementing managerial decisions ([Levy et al., 1997](#), [Zbaracki et al., 2004](#)). Our model can be interpreted as a multi-product extension of models with both information frictions and menu costs ([Alvarez, Lippi, and Paciello, 2011](#), [Bonomo et al., 2019](#)), in the limiting case in which the number of products goes to infinity. In such a case, profit-maximizing prices are not observed, and the common cost K is interpreted as a fixed cost of obtaining information about all frictionless optimal prices simultaneously, rather than a common menu cost, while c remains a unitary menu cost. Although in such a model the impulse responses to an aggregate shock are in general different from responses in our benchmark model, we show in Appendix B.4 that they are equal up to a first order. Hence, the nature of the common component of price-adjustment cost—physical menu cost or the cost of acquiring and processing information—is not material to the responses of aggregate price and output of an industry dominated by large multi-product retailers, as long as the monetary shock is small. This result is also related to [Alvarez, Lippi, and Passadore \(2017\)](#) who show that for a small monetary shock the responses in state- and time-dependent single-product models are the same up to a first order. Like us, they conclude that for small shocks the underlying nature of the nominal rigidity is irrelevant.

Our model nests two other cases previously studied in the literature on price-setting. [Midrigan \(2011\)](#) and [Alvarez and Lippi \(2014\)](#) study models in which firms sell a finite number of products and are required to pay a single menu cost to reset all prices at once. In those cases, the economies of scope of adjusting prices are maximal and, at any given instant, the share of prices that a certain firm resets is either zero or one. In our model, setting $c = 0$ yields the infinite-product limit of the Alvarez-Lippi-Midrigan framework.

The other extreme case is $K = 0$. In this case, there are no economies of scope of adjusting prices and we can imagine each firm in our model as a continuum of independent firms subject to idiosyncratic shocks, each responsible for adjusting the price of a single good, as in [Goloso and Lucas \(2007\)](#). In this extreme, firms continuously reset prices that reach certain adjustment thresholds and the law of large numbers thus guarantees that in any given time interval, e.g., a day, the share of products of a given firm that had their prices adjusted is constant. Our model therefore flexibly captures pricing behaviors ranging from perfect within-firm synchronization in price adjustments to variety-specific price adjustments.

3.4 Calibration

To compare predictions of different models, we calibrate not only the partial synchronization model, but also the Goloso-Lucas (GL) and Alvarez-Lippi-Midrigan (ALM) cases. Since we fix the time discount rate at $\rho = 0.04$, there are three parameters left to be calibrated in the partial synchronization model: the volatility σ , the fixed cost K and the unit cost c . We need therefore three moments from the data. As usual in the literature on price setting, we use the frequency of price adjustments and the average absolute size of price changes. Our measure of the frequency of adjustments is the average daily share of prices that a firms adjusts. To capture the degree of synchronization of price adjustments by a multi-product firm, we use the FK synchronization index as our third moment.

The GL and ALM models have two parameters each, since each of these settings has only one menu-cost parameter. We therefore drop the FK index when calibrating these models. This is a natural choice since FK statistics are fixed in those models: zero in the GL and one in the ALM.

[Table 4](#) shows moments for models and data, and [Table 5](#) shows calibrated parameter values. All models are able to match the average frequency and magnitude of price changes. However, as discussed, only the partial synchronization model can match the FK index. As for the untargeted moments, all models generate reasonable values for the standard deviation of the distribution of price changes. None of the models employed here succeed in matching the fraction of small adjustments, defined as those smaller (in absolute value) than one-quarter of the average adjustment size. The GL and partial synchronization models are not

able to generate any small adjustment, an issue well known in menu-cost models (Midrigan, 2011), while the ALM model significantly overshoots it. As for the kurtosis of the price-change distribution, only the ALM model is able to generate a value close to the data, as the price-change distribution in this model is always normal, thus implying a kurtosis of 3. In Section 4, we discuss extensions that match these two additional moments.

Moment	Data	GL	ALM	Partial sync.
Targeted moments				
Daily fraction of price changes	0.0087	0.0087	0.0087	0.0087
Average $ \Delta p $	0.208	0.208	0.208	0.208
Fisher-Konieczny index	0.259	0.000	1.000	0.259
Additional moments				
Standard deviation of Δp	0.248	0.208	0.260	0.210
Fraction of small adjustments	0.06	0	0.16	0
Kurtosis of Δp	3.53	1	3	1.09

Table 4: Moments from data and calibrated models.

Note: Values in the data are weighted means across stores in the data, provided in Table 1, first row, column (2)–(4). Small adjustments refers to standardized price adjustments that are smaller (in absolute magnitude) than one-quarter of the average standardized adjustment. Details for the moments of the distribution of standardized price changes in the data are contained in Appendix A.7.

Parameter	GL	ALM	Partial sync.
σ	0.370	0.464	0.374
K	-	0.011	3.70e-05
c	0.0023	-	0.0020

Table 5: Calibrated parameter values.

Importantly, two features of our calibration may generate an artificially lower value of the FK statistic. First, our model lacks aggregate and sectoral shocks. Shocks that affects more than one product simultaneously may cause the fraction of price changes in adjustment dates to vary, increasing our synchronization measure. Second, we work in the limit case as the number of products go to infinity. Statistical variation in price adjustments that comes from a finite sample of products would have similar consequences. In order to understand the effects of these assumptions, we simulate a panel of 7200 prices, in line with the average store in our sample, and introduce Brownian aggregate shocks, common to all products, with

a 5% yearly standard deviation. We then compute the implied FK statistic for a firm that follows the same steady state policy as in our baseline model.¹⁸ In this setting, we have a modest increase in synchronization, from 0.236 to 0.237, which suggests that results are not driven by these assumptions.

Figure 6 compares price-change distributions for all three models and data. Since in the GL setting firms continuously adjust prices that reach certain thresholds, the price-change distribution consists simply of two mass points placed on these limits. On the other hand, in the ALM model, firms only pay the fixed cost K to adjust all prices, even those prices that are close to their profit-maximizing levels. Thus, the price-change distribution features many small price adjustments. The partial synchronization model features more variability in the size of price changes than the GL case and, contrary to the ALM setting, no small price adjustments. Only the partial synchronization model can match the two modes placed approximately in a symmetrical manner around the origin in the empirical price-change distribution.




Figure 6: Price-change distributions for all models and data.

Figure 7 shows the daily fraction of adjustments for all models over time. As expected, it is constant for the GL case and assumes only the values zero and one for the ALM case. The partial synchronization model is therefore the one that comes closest to replicating the frequent and short peaks seen in the data (Figure 1). The optimal policy for the partial synchronization model in our calibration consists of adjusting the prices with corresponding gaps larger than 17.4% every 8.6 days.

¹⁸This must be interpreted as an approximation. The optimal policy for a model with a finite number of products and aggregate shocks is state-dependent and cannot be represented as deterministic adjustment dates and thresholds.

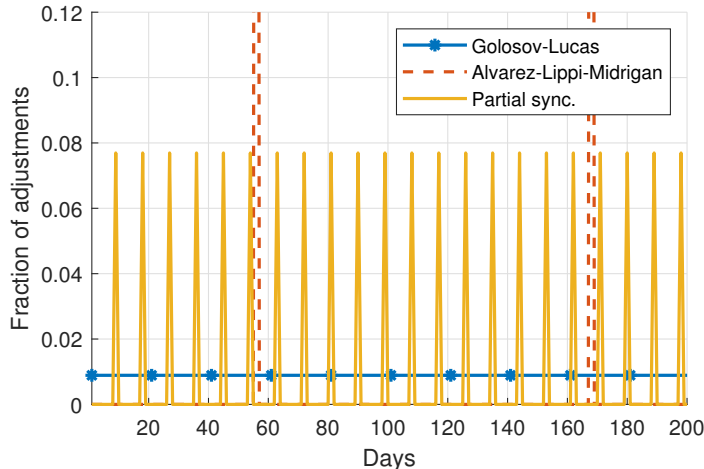


Figure 7: Daily share of adjusted prices for the three models. Spikes in the ALM model have a height of one.

4 Real effects of demand shocks

4.1 Analytical results for the partial synchronization model

In our model, a frictionless optimal price is the sum of a product-specific shock that follows Brownian motion and an aggregate demand component M_t that has so far been held constant. Now we consider the responses of the price level and real output to a one-time, unpredictable shock to M_t . Let P_t denote the aggregate price level, in logs, which is simply the average price across all firms and products in the model economy. Real output Y_t , also in logs, is then given by

$$Y_t = M_t - P_t.$$

We study an economy with a continuum of identical firms, which is hit by an unanticipated aggregate shock of size ε at $t = 0$. This shock shifts aggregate demand M_t to $M_t + \varepsilon$ and, as a consequence, the price gap distribution of all firms is also shifted: $g_t(x)$ becomes $g_t(x + \varepsilon)$. Prior to the occurrence of the aggregate shock, we naturally consider the situation in which all firms are in steady state, adjusting prices every τ^* periods. Moreover, we start with a situation in which firms are uniformly distributed according to the time elapsed since the last adjustment date, that is, a constant flow of firms adjusts prices over time before time $t = 0$. In the absence of aggregate shocks, this distribution will be invariant over time. To understand how aggregates respond to such a shock, we must first understand what optimal policy following the shock looks like.¹⁹

¹⁹Since we employ the second-order approximation to the objective function, there are no general equilibrium feedback effects of an aggregate shock to a firm's optimal policy (Alvarez and Lippi, 2014).

Consider a firm that had its last adjustment date at instant $t = -s$, for a given $0 < s < \tau^*$. Recall that the firm's problem is deterministic, since the evolution of the relevant state variable, namely the distribution of price gaps, is perfectly predictable and can be computed using the KFE (3). Therefore, after the unanticipated shock of size ε is realized at $t = 0$, the optimal policy can be represented by a new deterministic sequence of adjustment dates $\{T_k(\varepsilon)\}_{k=1}^\infty$ and thresholds $\{\bar{x}_k(\varepsilon)\}_{k=1}^\infty$ that depend only on ε and (implicitly) on s . Also, define Δ_k as the change in the firm's average price in the k -th adjustment episode following the shock, which is a function of the shock size ε and the optimal policy, although we omit this dependence for simplicity. In the absence of any changes to aggregate demand ($\varepsilon = 0$), the firm would follow its steady-state policy characterized by zero changes in its average price ($\Delta_k = 0$), and

$$\begin{aligned} T_1(0) &= \tau^* - s, \\ T_{k+1}(0) &= T_k(0) + \tau^*, \\ \bar{x}_k(0) &= \bar{x}^*. \end{aligned}$$

To obtain analytical results for the responses of aggregate price level and output to demand shocks, we focus on the limit as $\varepsilon \rightarrow 0$. As ε decreases, the adjustment dates converge to the ones that would arise in steady state, as long as these dates vary continuously with ε . Consequently, the firm's average price will take discrete steps at dates of the form $k\tau^* - s$, as illustrated in Figure 8.




Figure 8: Approximate response of the average price of a single firm following a small aggregate shock of size ε .

Since we assume that firms' adjusting times are uniformly distributed, and that adjusting intervals are τ^* , it follows that at each instant there is a density flow $1/\tau^*$ of adjusting firms.

Moreover, for $t \in (k\tau^*, (k+1)\tau^*)$, the average prices of adjusting firms change by an amount Δ_k , so the aggregate price level changes at a constant rate Δ_k/τ^* . More precisely, define $P_\varepsilon(t)$ to be the aggregate price level at instant t following a shock of size ε and let

$$\delta_k = \lim_{\varepsilon \rightarrow 0} \frac{\Delta_k}{\varepsilon}.$$

We have the following result.

Proposition 1. *The normalized aggregate price response $P_\varepsilon(t)/\varepsilon$ converges to a piece-wise linear function with kinks at positive multiples of τ^* as $\varepsilon \rightarrow 0$. Moreover, the slope of the k -th line segment is δ_k/τ^* .*

Proof. See Appendix B.2. □




Figure 9: Price-level response to a small aggregate shock.

Figure 9 shows the limiting response of the aggregate price level to a small positive aggregate shock, as established in Proposition 1. The slope of the first line segment of the impulse response function is quantitatively important since, as we show below, it accounts for almost 40% of the rise of the price level following an aggregate shock. To characterize the slope analytically, we first establish another important property of the impulse response.

Let F be the steady-state instantaneous frequency of price adjustments, defined as

$$F = \lim_{\Delta t \rightarrow 0} \frac{\text{Fraction of prices that change in } (t, t + \Delta t)}{\Delta t}.$$

Note that since F is a time density, it can take any positive value, including values greater than one. Also, define $f(x)$ as the density of the distribution of absolute size of price changes,

which is simply the part of the curves shown in Figure 6 associated with $x \geq 0$, and scaled to integrate to 1. To characterize the slope of the first line segment, we make use of the following lemma, whose proof and necessary definitions are presented in Appendix B.2.

Lemma 3. *Changes in policies in response to an aggregate shock do not have first-order effects on Δ_k around steady state. More precisely, for any positive integers j and k we have*

$$\frac{\partial \Delta_k}{\partial T_j} = \frac{\partial \Delta_k}{\partial \bar{x}_j} = 0. \quad (8)$$

Proof. See Appendix B.2. □

Lemma 3 is reminiscent of Proposition 1 in Alvarez and Lippi (2014), and earlier insights from Caballero and Engel (1993, 2007). Alvarez and Lippi show that a firm’s optimal policy after a monetary shock differs from the steady-state policy only up to second-order terms, because aggregate variables do not interact with price gaps up to second order. In turn, Lemma 3 establishes that in partial equilibrium, the effect of the shock on the optimal policy is second order.

As a consequence of Lemma 3, we have the following result.

Proposition 2. *The slope of the first line segment in the impulse response function is*

$$\frac{\delta_1}{\tau^*} = F \times [1 + \bar{x}^* f(\bar{x}^*)]. \quad (9)$$

Proof. See Appendix B.2. □

The intuition for this result is the following. The immediate response of the price level is the product of an *extensive margin* component F and an *intensive margin*, or *selection component*, $1 + \bar{x}^* f(\bar{x}^*)$, which depends on the size of the marginal adjustment \bar{x}^* multiplied by its density $f(\bar{x}^*)$. Caballero and Engel (2007) prove a similar result for a single-product model.

4.2 Synchronization and real effects of demand shocks

We now turn to the analysis of the relationship between synchronization of price changes and the magnitude of output response to demand shocks. Figure 10 shows impulse response functions of several models calibrated to match the same frequency of price adjustments and average size of price changes, but different values of the FK synchronization measure. Output decays faster as synchronization fades. We can gain some intuition for this result by looking

at expression (9). Since all models are calibrated to match the same frequency of adjustment, the differences in responses must come from the selection component.²⁰




Figure 10: Real output response and synchronization.

We can see the ALM model as the limit when $\bar{x}^* \rightarrow 0$. For a fixed frequency, the ALM model delivers the slowest initial response. In fact, since the effect of the monetary shock is fully reverted in the ALM model after the first round of adjustments, the first segment corresponds to the whole positive part of the impulse response function. Since \bar{x}^* is the smallest possible adjustment size, we have $\bar{x}^* \leq \mathbb{E}|\Delta p|$. We can thus see the GL model as the limiting case in which $\bar{x}^* \rightarrow \mathbb{E}|\Delta p|$ and $f(x)$ approaches a Dirac mass at \bar{x}^* . Therefore, for a fixed frequency and mean adjustment size, the GL model can be seen as the limit in which the speed of the initial price-level response blows up. The relationship between synchronization and persistence of real effects of nominal shocks is thus directly associated with selection effects explored by Golosov and Lucas (2007). If a firm adjusts a large fraction of its prices on an adjustment date, there is not much room for selecting those prices, and monetary shocks have more persistent real effects as a consequence.²¹

Looking only at the selection term $1 + \bar{x}^* f(\bar{x}^*)$, however, may not be sufficient to understand the relationship between synchronization and non-neutrality because the selection term is related only to the initial, rather than total, degree of aggregate price flexibility. For this reason, Figure 11 shows both the selection term and the cumulative output response, measured as the area under the impulse response function, for varying FK values. We can

²⁰Even though the instantaneous frequency F is not the same as the daily adjustment frequency we use in our calibration, there is a one-to-one mapping between both, given by $F = 365 \times \text{daily frequency}$ in steady state, as long as the interval between adjustment dates is longer than one day.

²¹Since we study the limiting case in which the shock size goes to zero, the aggregate response does not come from anticipation of adjustment episodes.

see that the marginal effect of synchronization on aggregate price stickiness is such that small departures from the ALM framework can considerably reduce the degree of monetary non-neutrality.

These results bear important implications for the degree of monetary non-neutrality associated with the synchronization of price changes we document in the data. Figure 12 shows the impulse response functions for the calibrated parameter values shown in Table 5. The most salient feature of this figure is that the partial synchronization model, when calibrated to match our data, is very close to GL, generating much less monetary non-neutrality than ALM. Why is the calibrated partial synchronization model associated with so little non-neutrality? First, note that almost 40% of the decay of output following the shock happens in the first round of adjustments, underlying the first line segment of the impulse response function. Less synchronization is associated with less persistent real effects, but as shown in Figure 11, no synchronization or moderate levels of synchronization produce close results in terms of non-neutrality. It turns out the degree of synchronization found in the data is moderate, resulting in a partial synchronization model with enough selection to be quantitatively very close to GL.

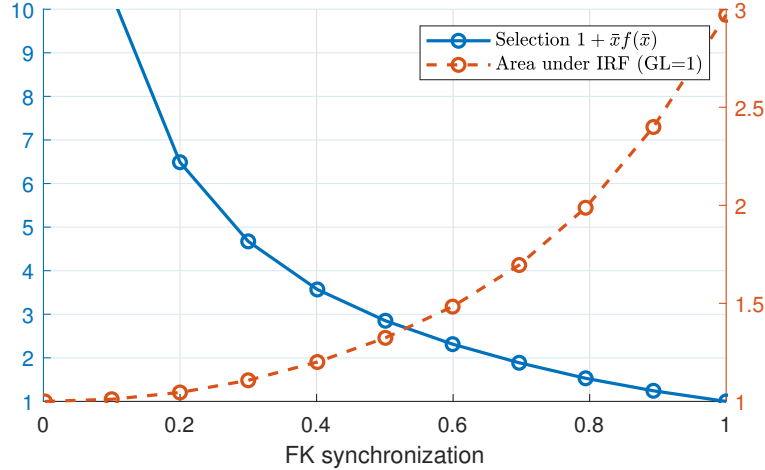


Figure 11: Selection component $1 + \bar{x}^* f(\bar{x}^*)$ and cumulative response for various values of the FK synchronization measure. Cumulative response of the GL model is normalized to 1.

4.3 The relationship between synchronization, the kurtosis of price changes, and monetary non-neutrality

How do our conclusions relate to the sufficient statistic result derived by [Alvarez, Le Bihan, and Lippi \(2016\)](#) (henceforth, ALL)? ALL show that, in a large class of pricing models, under




Figure 12: Output response to a small aggregate shock.

the assumption that frictionless optimal prices follow Brownian motions, the area under the impulse response function generated by a small aggregate shock should be proportional to the kurtosis of the price-change distribution (for a given frequency of price changes).²² Our partial synchronization model is not included in that class of models. Nevertheless, the ALL result holds numerically for our partial synchronization model as well. Figure 13 shows that as kurtosis of price changes varies across models with the same frequency but different synchronization of price changes (solid blue line), the ratio of kurtosis to the area under the impulse response function is roughly constant (red dotted line).²³ For a given number of products, varying the degree of synchronization endogenously changes the degree of kurtosis produced by the model. Less synchronization leads to lower kurtosis because the selection effect embedded in the price-adjustment mechanism of the model reduces the kurtosis with respect to that of the frictionless price process.

Since in our baseline setup synchronization and kurtosis of price changes are directly related, we conjecture that synchronization is an alternative sufficient statistic for monetary non-neutrality. This is no longer the case in a setup where ALL theorem does not apply. In Section 4.4, we introduce fat-tailed shocks, which allows us to disconnect variation in synchronization of price changes from their kurtosis. In that case, variation in price synchronization

²²Alvarez, Lippi, and Oskolkov (2020) extend the sufficient statistic result to menu cost models with a generalized hazard function. Karadi and Reiff (2019), Dotsey and Wolman (2020) provide examples of empirically plausible models in which the conditions for the ALL result are not satisfied and the ALL sufficient statistic does not hold. In Karadi and Reiff (2019) idiosyncratic shocks are drawn from a mixture of two normal distributions. In Dotsey and Wolman (2020) firms face a multi-state persistent idiosyncratic productivity process and i.i.d. random menu costs.

²³We verify numerically that the variation in the standard deviation of price changes does not alter the monotone relationship between kurtosis of price changes and monetary non-neutrality. This is consistent with ALL's theorem for small monetary shocks.




Figure 13: Kurtosis of price-change distribution for various values of the FK synchronization measure.

Note: For all values of FK, we calibrate the models to generate the same frequency and average absolute size of price changes.

has an independent impact on real effects of monetary policy.

Large kurtosis of the price-change distribution documented in the Israeli data poses a challenge to our model. As can be seen in Table 4, with Gaussian shocks it is not possible to simultaneously match the kurtosis of price changes and the degree of synchronization of adjustments. The degree of synchronization in the data is relatively small, which, as Figure 13 shows, is associated with low kurtosis. There are two possible modifications to our model that have been frequently employed in the literature and that could help reconcile kurtosis and synchronization, as we see in the data.

The first is the introduction of free adjustments *à la* Calvo, as in ALL and Nakamura and Steinsson (2010). In this framework, each price is subject to random free adjustment opportunities that arrive at a constant rate and independently across products. It turns out that this modification of the model is not very helpful. As shown both in Table 4 and Figure 6, the share of small price changes is small in our sample, and therefore introducing free adjustments to match this statistic does not make a substantial difference for the kurtosis of the price-change distribution. Not surprisingly, the impulse response function of a monetary shock remains almost unchanged in this case. We therefore leave this extension to Appendix B.5.²⁴ The other alternative, which we explore next, is the introduction of fat-tailed shocks

²⁴As we discussed in Section 2.5 and Appendix A.7, time aggregation in weekly and monthly price data can substantially inflate the share of small price changes, making it a more consequential calibration target. For example, in our daily data, only 6% of adjustments are smaller in absolute magnitude than one-quarter of the mean absolute size of price changes. When we time-aggregate observations to a weekly (monthly) frequency, this share goes up to 10% (11%). In two weekly scanner data sets for the U.S. prices reported in

to frictionless optimal prices.

4.4 Model with fat-tailed shocks

We consider a modification of our model in which frictionless optimal prices no longer follow Brownian motions. We suppose instead that each $p_{i,t}^*$ follows a Poisson process, receiving shocks with Laplace distribution (with zero mean and scale parameter β) and arrival rate η ,²⁵ which is the average number of shocks that arrive in a year for a given product. The simple introduction of infrequent shocks would by itself increase the kurtosis of the price-change distribution even with Gaussian shocks, as in Midrigan (2011). However, to generate the degree of excess kurtosis observed in the data, we add fat-tailed shocks based on the Laplace distribution, which has a kurtosis of 6. Figure 14 illustrates this new stochastic process for $p_{i,t}^*$. Importantly, the Brownian motion can still be obtained as a limit case of this Poisson process. If we take the limit as $\eta \rightarrow \infty$ and $\beta \rightarrow 0$ (at a certain joint velocity), we converge to a Brownian motion. Intuitively, this is the limit in which an infinite number of infinitesimally small shocks arrive in each time interval.




Figure 14: Sample paths of Brownian motion and Poisson process with Laplace shocks.

In this case, to simplify the numerical procedure, we do not solve for the optimal policy of the firm.²⁶ Instead, we calibrate the model by choosing policies \bar{x} and τ and parameters of

Midrigan (2011), this share is 10% and 14%; and this share is 10% in the weekly IRI data.

²⁵The probability density function of the Laplace distribution we employ is therefore given by $f(x) = (2\beta)^{-1} \exp(-|x|/\beta)$, and has mean zero and variance $2\beta^2$. Time intervals between two consecutive shocks have exponentially distributed lengths with mean $1/\eta$.

²⁶When shocks are Gaussian, the distribution just before the lump-sum cost is paid is approximately Gaussian. Then the fact that the Gaussian distribution depends only on two parameters reduces the dimensionality of the state space from infinity to two.

the stochastic process η and β .²⁷ We can do so without loss of generality as long as, for each quadruple $(\bar{x}, \tau, \eta, \beta)$, there exist parameters (K, c, η, β) that generate \bar{x} and τ as optimal policies (note that η and β are the same in both quadruples). In other words, we can directly choose the optimal policy to calibrate our model if, given any policy \bar{x} and τ , there exist parameters under which it is indeed the optimal policy. We choose this approach because it greatly simplifies the calibration exercise. The drawback, however, is that it does not allow us to recover the values of the underlying menu costs K and c . Nevertheless, given that we match the same moments, we expect these values to be very close to the ones we obtain in Table 5.

As previously, we calibrate GL, ALM, and partial synchronization models with this new stochastic process.²⁸ Since we now have four instead of three model parameters, we need an additional calibration target. Naturally, we choose the kurtosis of the price-change distribution. Calibration results are shown in Tables 6 and 7. The partial synchronization model is able to match the targeted moments and generate reasonable values for the standard deviation of the price changes. As for the fraction of small adjustments, results do not change much, with ALM model still generating too many small price adjustments.

One key feature of the calibration of the ALM case, which is going to play an important role in understanding monetary non-neutrality, is that it needs a smaller departure from the Brownian motion case to generate the same kurtosis of price changes as in the data. This can be seen from the high arrival rate of shocks η (around 36 per year, or three per month) with a small standard deviation (Table 7). With Brownian shocks, the ALM model generates normally distributed price changes with a kurtosis of 3, which is close to the value of 3.53 we measure empirically. On the other extreme, the GL model always generates a kurtosis of 1 when shocks are Brownian, and consequently, it needs a substantial departure from the Gaussian case. The partial synchronization case needs slightly less leptokurtic shocks than GL to match the kurtosis of price changes found in the data.

Turning to monetary non-neutrality, Figure 15 shows impulse response functions of real output to monetary shocks for all models. Interestingly, differences in the degree of non-neutrality are substantially smaller than with the Gaussian shocks discussed earlier, because the underlying process for $p_{i,t}^*$ is significantly more fat-tailed for the GL and partial synchro-

²⁷We solve the model in the following way. Given parameters η and β and a policy \bar{x} and τ , we use equation (7) to obtain the stationary distribution of price discrepancies (the latter equation must be slightly modified for non-Brownian shocks). Having the stationary distribution, it is easy to compute any moments. Then we optimize over the parameter space to find values that generate moments as close as possible to the data.

²⁸In the knife-edge case in which there is no inflation, even in the ALM not all prices would be adjusted when the fixed cost is paid. The reason is that, with infrequent shocks, some products will not have received any shocks between two consecutive adjustment dates and will therefore have a zero price gap. We therefore consider the limiting case $\mu \rightarrow 0$ for the ALM model, so that all prices are adjusted after the firm pays K .

nization cases, compensating for the selection effect embedded in those models.

Moment	Data	GL	ALM	Partial sync.
Targeted moments				
Daily fraction of price changes	0.0087	0.0087	0.0087	0.0087
Average $ \Delta p $	0.208	0.208	0.208	0.208
Fisher-Konieczny index	0.259	0.000	1.000	0.259
Kurtosis of Δp	3.53	3.52	3.53	3.55
Additional moments				
Standard deviation of Δp	0.248	0.253	0.266	0.247
Fraction of small adjustments	0.06	0	0.17	0

Table 6: Moments from data and models calibrated with fat-tailed shocks.

Note: Values in the data are weighted means across stores in the data, provided in Table 1, first row, columns (2)–(4). Small adjustments are defined as those smaller than one-quarter of the average size.

Parameter	GL	ALM	Partial sync.
\bar{x}	0.059	0	0.074
τ	-	0.315	0.024
η	4.70	35.95	5.78
β	0.150	0.056	0.129

Table 7: Calibrated parameter values for models with fat-tailed shocks.

To better understand how synchronization and kurtosis are related to monetary non-neutrality, we calibrate each model to match the same frequency and average absolute size of price changes, while varying the kurtosis of price changes. Figure 16 displays the relation between the area under IRFs and the kurtosis of the price-change distribution for different levels of synchronization. Results show that kurtosis of the price-change distribution is no longer a sufficient statistic for the area under the impulse response function, as the curves have different slopes according to FK levels. In addition, the effect of synchronization ceases to be monotonic. The ALM model generates the strongest non-neutrality when kurtosis is 3, and the degree of non-neutrality across models is reversed for higher values of kurtosis.

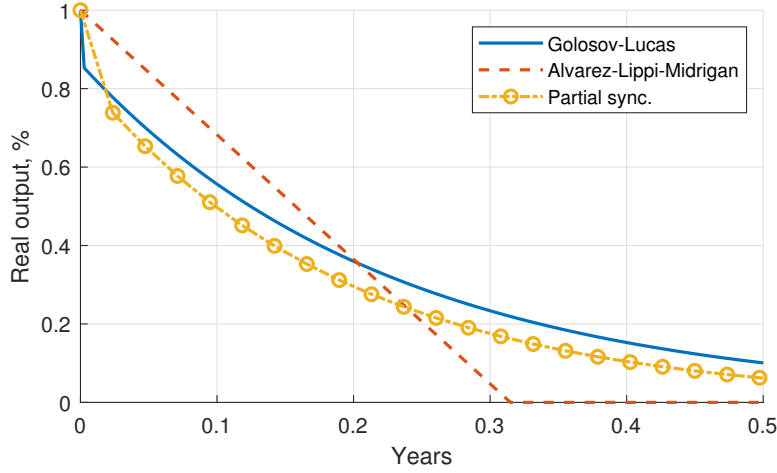


Figure 15: Impulse response functions for all models with fat-tailed shocks.

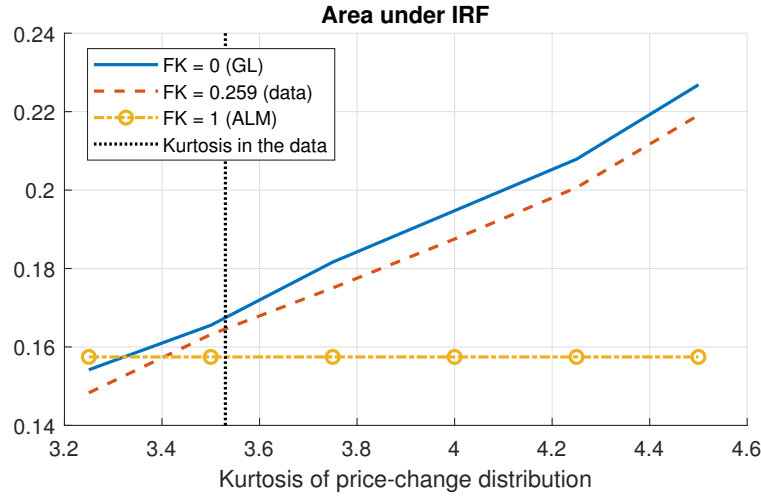


Figure 16: Area under impulse response functions as a function of the kurtosis of price-change distribution.

Note: For all values of FK and kurtosis, we calibrate the models to generate the same frequency and average absolute size of price changes.

Figure 17 displays two graphs, relating the kurtosis of the price-change distribution, in the first graph, and monetary non-neutrality, in the second, to the rate of arrival of shocks η . The latter is inversely related to the kurtosis generated by the shock processes, with higher values corresponding to lower kurtosis processes, closer to Brownian motion. Not surprisingly, higher exogenous kurtosis of shocks (lower η) generates higher kurtosis of price changes, when we keep the degree of synchronization fixed. A higher synchronization (higher FK) shifts the curve upwards. Hence, when shocks are non-Gaussian, the determinants of monetary non-neutrality become more complex. On the one hand, monetary non-neutrality

is determined jointly by the kurtosis of the idiosyncratic shock process and price selection, which is a function of synchronization. On the other hand, kurtosis of price changes is no longer a sufficient statistic for non-neutrality, because it depends on a particular combination of the kurtosis of shocks and the degree of synchronization of price changes.




Figure 17: Area under impulse response functions as a function of the kurtosis of price-change distribution.

Note: For all values of FK and rate of arrival of shocks, we calibrate the models to generate the same frequency and average absolute size of price changes. The rate of arrival of shocks η is inversely related to the kurtosis of the shock process, with higher values corresponding to processes more similar to Brownian motions.

5 Conclusions

The literature on multi-product pricing is still scarce and relatively new; there is still dearth of evidence and theory to help us understand behavior of large retailers and their impact on the economy. This paper contributes on both fronts. On the empirical end, we exploit unique detail on day-to-day price changes across thousands of products sold in large food retailers in Israel. A typical retailer synchronizes its regular price changes around occasional peak days when a large share of store prices are simultaneously adjusted. This behavior suggests that retailers actively exploit substantial economies of scope in their day-to-day price adjustment. Based on available evidence, we conclude that this finding is robust to incorporating price discounts, online prices, chain effects, and calendar effects. Future research can advance empirical evidence along several dimensions. First, data on quantities of sold products complement price data and help gauge the substitution of consumption spending across products within the same store or certain category in the store (Chevalier and Kashyap, 2019), or products supplied by the same vendor. Second, improved datasets can balance representation of the economy by incorporating small- and medium-size stores, and by including durables,

services, and other goods. Third, advancements in identification of monetary policy shocks and longer data can provide corroborating evidence on the role of price synchronization in monetary transmission (Hong et al., 2020). Finally, adding information about retailers’ production costs (Dedola, Kristoffersen, and Züllig, 2019), inventories and stockouts (Kryvtsov and Midrigan, 2013), relationships with suppliers (Anderson et al., 2017), and the costs of price adjustments (Levy et al., 1997) is crucial for further understanding the elements of the price-decision problem faced by retailers.

On the theory side, we provide a generalization of existing models to a more flexible setting with economies of scope in price adjustment. Depending on the parameter configuration, our model can match an arbitrary degree of partial synchronization. We rely on a few key simplifying assumptions to make the model tractable. We assume that shocks to each product’s desired price are independent across products in the same store and across stores. Based on our analysis, we posit that relaxing these assumptions will not qualitatively change the main insight in the paper: the selection effect in partial synchronization models brings their aggregate behavior closer to models with perfect staggering and substantial price flexibility. Future models can refine quantitative implications of incorporating store- or chain-specific disturbances and add strategic complementarities in price decisions for products within and across stores.

References

- Bank of Israel. “The Bank of Israel Annual Report for 2012.” URL <http://www.boi.org.il/en/NewsAndPublications/PressReleases/Pages/02042013.aspx>.
- Israeli Central Bureau of Statistics. “Data From the 2017 Household Expenditure Survey: General Summary.” URL https://www.cbs.gov.il/he/publications/DocLib/2019/households17_1755/e_print.pdf.
- Alvarez, Fernando, Hervé Le Bihan, and Francesco Lippi. 2016. “The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach.” *American Economic Review* 106 (10):2817–2851.
- Alvarez, Fernando and Francesco Lippi. 2014. “Price Setting with Menu Cost for Multiproduct Firms.” *Econometrica* 82 (1):89–135.
- Alvarez, Fernando, Francesco Lippi, and Juan Passadore. 2017. “Are State- and Time-Dependent Models Really Different?” *NBER Macroeconomics Annual* 31 (1):379–457.

- Alvarez, Fernando E., Francesco Lippi, and Aleksei Oskolkov. 2020. “The Macroeconomics of Sticky Prices with Generalized Hazard Functions.” NBER Working Papers 27434, National Bureau of Economic Research, Inc.
- Alvarez, Fernando E., Francesco Lippi, and Luigi Paciello. 2011. “Optimal Price Setting With Observation and Menu Costs.” *The Quarterly Journal of Economics* 126 (4):1909–1960.
- Anderson, Eric, Benjamin A. Malin, Emi Nakamura, Duncan Simester, and Jón Steinsson. 2017. “Informational rigidities and the stickiness of temporary Sales.” *Journal of Monetary Economics* 90 (C):64–83.
- Bhattarai, Saroj and Raphael Schoenle. 2014. “Multiproduct firms and price-setting: Theory and evidence from U.S. producer prices.” *Journal of Monetary Economics* 66:178–192.
- Bonomo, Marco, Carlos Carvalho, Rene Garcia, and Vivian Malta. 2019. “Persistent Monetary Non-neutrality in an Estimated Model with Menu Costs and Partially Costly Information.” mimeo.
- Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela. 2008. “Database Paper—The IRI Marketing Data Set.” *Marketing Science* 27 (4):745–748.
- Caballero, Ricardo J. and Eduardo M. R. A. Engel. 1993. “Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy.” *Review of Economic Studies* 60 (1):95–119.
- Caballero, Ricardo J. and Eduardo M.R.A. Engel. 2007. “Price stickiness in Ss models: New interpretations of old results.” *Journal of Monetary Economics* 54 (Supplement):100–121.
- Caplin, Andrew S. and Daniel F. Spulber. 1987. “Menu Costs and the Neutrality of Money.” *Quarterly Journal of Economics* 102 (4):703–726.
- Carvalho, Carlos and Oleksiy Kryvtsov. 2018. “Price Selection.” Staff Working Papers 18-44, Bank of Canada.
- Cavallo, Alberto. 2017. “Are Online and Offline Prices Similar? Evidence from Large Multi-channel Retailers.” *American Economic Review* 107 (1):283–303.
- . 2018. “Scraped Data and Sticky Prices.” *The Review of Economics and Statistics* 100 (1):105–119.
- Chevalier, Judith A. and Anil K. Kashyap. 2019. “Best Prices: Price Discrimination and Consumer Substitution.” *American Economic Journal: Economic Policy* 11 (1):126–159.

- Danziger, Leif. 1999. “A Dynamic Economy with Costly Price Adjustments.” *American Economic Review* 89:878–901.
- Dedola, Luca, Mark Strøm Kristoffersen, and Gabriel Züllig. 2019. “Price synchronization and cost pass-through in multiproduct firms: Evidence from Danish producer prices.” Tech. rep.
- DellaVigna, Stefano and Matthew Gentzkow. 2019. “Uniform Pricing in U.S. Retail Chains.” *The Quarterly Journal of Economics* .
- Dias, Daniel, Carlos Marques, Pedro Neves, and João Santos Silva. 2005. “On the Fisher-Konieczny index of price changes synchronization.” *Economics Letters* 87 (2):279–283.
- Dotsey, Michael and Alexander L. Wolman. 2020. “Investigating Nonneutrality in a State-Dependent Pricing Model With Firm-Level Productivity Shocks.” *International Economic Review* 61 (1):159–188.
- Dutta, Shantanu, Mark Bergen, Daniel Levy, and Robert Venable. 1999. “Menu Costs, Posted Prices, and Multiproduct Retailers.” *Journal of Money, Credit and Banking* 31 (4):683–703.
- Fisher, Timothy C. G. and Jerzy D. Konieczny. 2000. “Synchronization of price changes by multiproduct firms: evidence from Canadian newspaper prices.” *Economics Letters* 68 (3):271–277.
- Goldberg, Pinelopi K. and Rebecca Hellerstein. 2009. “How rigid are producer prices?” Staff Reports 407, Federal Reserve Bank of New York.
- Golosov, Mikhail and Robert E. Lucas. 2007. “Menu Costs and Phillips Curves.” *Journal of Political Economy* 115:171–199.
- Gorodnichenko, Yuriy, Viacheslav Sheremirov, and Oleksandr Talavera. 2018. “Price Setting in Online Markets: Does IT Click?” *Journal of the European Economic Association* 16 (6):1764–1811.
- Gorodnichenko, Yuriy and Oleksandr Talavera. 2017. “Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration.” *American Economic Review* 107 (1):249–82.
- Hitsch, Günter J., Ali Hortaçsu, and Xiliang Lin. 2019. “Prices and Promotions in U.S. Retail Markets: Evidence from Big Data.” NBER Working Papers 26306, National Bureau of Economic Research, Inc.

- Hong, Gee Hee, Matthew Klepacz, Ernesto Pasten, and Raphael Schoenle. 2020. “The Real Effects of Monetary Shocks: Evidence from Micro Pricing Moments.” Working Papers Central Bank of Chile 875, Central Bank of Chile.
- Karadi, Peter and Adam Reiff. 2019. “Menu Costs, Aggregate Fluctuations, and Large Shocks.” *American Economic Journal: Macroeconomics* 11 (3):111–46.
- Klenow, Peter J. and Oleksiy Kryvtsov. 2008. “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” *Quarterly Journal of Economics* 123 (3):863–904.
- Klenow, Peter J. and Benjamin A. Malin. 2010. “Microeconomic Evidence on Price-Setting.” *Handbook of Monetary Economics* 16 (1):231–284.
- Kryvtsov, Oleksiy and Virgiliu Midrigan. 2013. “Inventories, Markups, and Real Rigidities in Menu Cost Models.” *Review of Economic Studies* 80 (1):249–276.
- Lach, Saul and Daniel Tsiddon. 1996. “Staggering and Synchronization in Price-Setting: Evidence from Multiproduct Firms.” *American Economic Review* 86 (5):1175–1196.
- Letterie, Wilko and Øivind Anti Nilsen. 2016. “Price Changes - Stickiness and Internal Coordination in Multiproduct Firms.” CESifo Working Paper Series 5701, CESifo Group Munich.
- Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable. 1997. “The Magnitude of Menu Costs: Direct Evidence from Large U. S. Supermarket Chains.” *The Quarterly Journal of Economics* 112 (3):791–824.
- Levy, Daniel, Shantanu Dutta, Mark Bergen, and Robert Venable. 1998. “Price adjustment at multiproduct retailers.” *Managerial and Decision Economics* 19 (2):81–120.
- Midrigan, Virgiliu. 2011. “Menu Costs, Multi-Product Firms and Aggregate Fluctuations.” *Econometrica* 79:1139–1180.
- Nakamura, Emi and Jón Steinsson. 2008. “Five Facts About Prices: A Reevaluation of Menu Cost Models.” *Quarterly Journal of Economics* 123 (4):1415–1464.
- . 2010. “Monetary Non-neutrality in a Multi-Sector Menu Cost Model.” *Quarterly Journal of Economics* 125 (3):961–1013.
- Sheshinski, Eytan and Yoram Weiss. 1992. “Staggered and Synchronized Price Policies Under Inflation: The Multiproduct Monopoly Case.” *Review of Economic Studies* 59 (2):331–359.

- Stella, Andrea. 2014. “The Magnitude of Menu Costs: A Structural Estimation.” 2014 Meeting Papers 436, Society for Economic Dynamics.
- Taylor, John. 1980. “Aggregate Dynamics and Staggered Contracts.” *Journal of Political Economy* 88 (1):1–23.
- Warner, Elizabeth J. and Robert B. Barsky. 1995. “The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays.” *The Quarterly Journal of Economics* 110 (2):321–352.
- Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen. 2004. “Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets.” *Review of Economics and Statistics* 86 (2):514–533.