

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

Working Paper

Multi-product pricing: Theory and evidence from large retailers in Israel

Bank of Canada Staff Working Paper, No. 2020-12

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Bonomo, Marco Antonio; Carlos Viana de Carvalho; Kryvtsov, Oleksiy; Ribon, Sigal; Rigato, Rodolfo (2020) : Multi-product pricing: Theory and evidence from large retailers in Israel, Bank of Canada Staff Working Paper, No. 2020-12, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2020-12>

This Version is available at:

<https://hdl.handle.net/10419/241178>

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.

Multi-Product Pricing: Theory and Evidence from Large Retailers in Israel

by Marco Bonomo,¹ Carlos Carvalho,² Oleksiy Kryvtsov,³ Sigal Ribon⁴ and Rodolfo Rigato⁵

¹Inspere
bonomo@insper.edu.br

²PUC-Rio
cvianac@decon.ouc-rio.br

³Economic and Financial Research Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9
okryvtsov@bankofcanada.ca

⁴Bank of Israel
sigal.ribon@boi.org.il

⁵Harvard
dinisrigato@g.harvard.edu



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This work may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this note are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

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, the 2019 Luso-Brazilian Macroeconomics Meeting, the ECB conference on "Inflation in a Changing Economic Environment" and Université de Montréal for their comments and discussions. Minnie Cui provided superb research assistance. The views expressed herein are those of the authors and not necessarily those of the Bank of Canada or the Bank of Israel.

Abstract

Standard theories of price adjustment are based on the problem of a single-product firm, and therefore they may not be well suited to analyze price dynamics in the economy with multi-product firms. To guide new theory, we study a unique dataset with comprehensive coverage of daily prices in large multi-product food retailers in Israel. We find that a typical retail store synchronizes its regular price changes around occasional "peak" days when, once or twice a month, it reprices around 10% of its products. To assess the implications of partial price synchronization for inflation dynamics, we develop a new price-setting model in which a firm sells a continuum of products and faces economies of scope in price adjustment. The model generates the partial synchronization pattern with peaks of re-pricing activity observed in the data. We show analytically and numerically that synchronization of price changes attenuates the average price response to a monetary shock; however, only high degrees of synchronization can materially strengthen monetary non-neutrality. Hence, the synchronization of price changes observed in the data is consistent with considerable aggregate price flexibility.

Bank topics: Inflation and prices; Market structure and pricing; Monetary policy

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

Résumé

Les théories habituelles sur l'ajustement des prix sont fondées sur le problème des entreprises à produit unique. Elles pourraient donc ne pas convenir à l'analyse de la dynamique des prix au sein d'une économie composée d'entreprises multiproduits. Pour guider l'élaboration d'une nouvelle théorie, nous étudions un ensemble unique de données comprenant un large éventail de prix affichés quotidiennement par de grands détaillants alimentaires multiproduits en Israël. Nous constatons qu'un détaillant typique synchronise les changements de ses prix réguliers pour les apporter périodiquement lors de jours « de pointe », soit une ou deux fois par mois. Il change alors le prix d'environ 10 % de ses produits. Pour évaluer les effets d'une synchronisation partielle des prix sur la dynamique de l'inflation, nous élaborons un nouveau modèle de fixation des prix dans lequel une entreprise vend un continuum de produits et obtient des économies de gamme grâce aux changements de prix. Le modèle produit un profil de synchronisation partielle à partir des pics de changements de prix observés dans les données. Nous montrons, de façon analytique et numérique, que la synchronisation des changements de prix atténue la réaction moyenne des prix à un choc monétaire. Toutefois, seuls des degrés élevés de synchronisation peuvent renforcer sensiblement la non-neutralité monétaire. Ainsi, la synchronisation des changements de prix observée dans les données concorde avec une flexibilité substantielle des prix globaux.

Sujets : Inflation et prix; Structure de marché et fixation des prix; Politique monétaire

Codes JEL : D21, D22, E31, E52, L11

Non-technical summary

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 these trade-offs jointly using a novel model in which each firm sells a continuum of products. The main question is whether multi-product pricing materially changes a firm's ability to respond to demand or monetary policy disturbances.

To guide new theory, we study a unique dataset with comprehensive coverage of daily prices in large, multi-product food retailers in Israel from May 20, 2015, until October 4, 2019. Our 5% sample contains prices for almost all products sold in the 75 largest stores across 25 retail chains, averaging 7,217 products per store a day.

We exploit comprehensive product and time coverage in this dataset to study the synchronization of price adjustments across products in a store. We find that a typical retail store synchronizes its regular price changes around occasional “peak” days when, once or twice a month, it reprices around 10% of its products. This behavior is not driven by weekly or monthly seasonal effects, holidays, differences between online and offline prices, or retail discounting.

To assess the implications of partial price synchronization for inflation dynamics, we develop a general equilibrium model in which a firm sells a continuum of products and faces economies of scope in price adjustment. Our general framework gives rise to an endogenous degree of synchronization of price changes, and it nests previously studied models in which price changes are not synchronized (no economies of scope) or perfectly synchronized (with maximal economies of scope). The multi-product firm in our model adjusts prices only infrequently. 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 partial synchronization pattern we observe in the data.

We calibrate the model to match the key moments in the micro data and use it to study the responses of the price level and real output to a one-time, unanticipated impulse to money supply. After such a shock, the firm must wait until the next peak adjustment date to change most of its prices, which means that synchronization of price changes mutes the response of the average price and boosts the response of real output. We show analytically and numerically that only high degrees of synchronization can materially weaken monetary non-neutrality. Hence, synchronization of price changes observed in the data is consistent with considerable aggregate price flexibility.

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 these trade-offs jointly using a model in which each firm sells a continuum of products. The main question is whether multi-product pricing materially changes a firm’s ability to respond 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 a unique price dataset from large multi-product food retailers in Israel. A disclosure law enacted in 2014 requires large retailers in Israel to post on their internet sites daily price information for all products sold in their stores. To manage computational constraints, the data used in this paper contain information for stores in the top 5%, by the number of observations, for each retail chain in Israel from May 20, 2015, until October 4, 2019.¹ For all 71 brick-and-mortar stores and 25 retail chains in this sample, we have information about the “base,” or regular, prices for all products sold in those stores on a daily basis. For four retailers, we also have information about their online prices, which we will treat as four additional online stores. In addition to regular prices, we have constructed the final (discounted) prices for the stores owned by Shufersal, the largest food retailer 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. Observations are recorded daily, spanning 1,598 days over the sample period, and prices

¹We excluded store-specific products that do not have a general 13-digit barcode (e.g., fruits, vegetables, bakery goods). We also excluded pharmacy retailers. The remaining dataset contains 506.1 million daily observations.

for almost all products are included, averaging 7,217 per store on a given day.

We exploit comprehensive product and time coverage in our data to study the synchronization of price adjustments across products in a store.² 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 in single-product menu cost economies (Golosov and Lucas, 2007) or the perfect synchronization pattern in Alvarez and Lippi (2014).



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.

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,

²To validate the price-adjustment statistics we obtain in the Israeli data—and, in particular, to gauge the implications of time aggregation for empirical analysis—we also use Information Resources Inc scanner data for weekly price information for 31 products across 50 U.S. cities. We report comparisons of two datasets along two dimensions: the moments of the distribution of price changes and spectral densities of store-level price changes. We find that time aggregation has influence on these moments, which underscores the benefits of using daily price observations to fully explore the implications of synchronization of price changes.

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 frictionless optimal 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, given the large number of products, the individual menu cost c implies that not all prices 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.

We characterize the price-setting policy of the multi-product firm by deterministic dates $\{T_k\}_{k=1}^{\infty}$ when adjustments are made, and thresholds for each adjustment date $\{\bar{x}_k\}_{k=1}^{\infty}$. At each time T_k , the firm adjusts the prices of all products for which the price differs from its frictionless optimal price by a magnitude greater than \bar{x}_k in absolute value. In steady state, the adjustment rule can be characterized by a constant threshold, \bar{x}^* , and the length of the time interval between consecutive adjustment dates, $\tau^* = T_{k+1} - T_k$. We show that the steady-state solution for the distribution of price gaps can be characterized analytically.

We then study the responses of the price level and real output to a one-time unanticipated impulse to money supply. In our model, these responses are determined by the influence of synchronization of price adjustments on the selection effect characteristic in price-setting models (Golosov and Lucas, 2007). In the Golosov-Lucas model, in response to the shock, firms with prices that are farther away from the optimal adjust first, triggering adjustments that are relatively larger and amplifying the response of the aggregate price.³ By contrast, 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, synchronization of price changes, and monetary non-neutrality. We show that the Alvarez-Lippi-Midrigan model corresponds to the extreme case when the selection component is minimal, and that the Golosov-Lucas model is the limiting case in which selection goes to infinity. Consequently, the real effect of the monetary shock in the Golosov-Lucas model is significantly smaller than in the Alvarez-Lippi-Midrigan model.

How much quantitative effect does the synchronization of price changes we document in the data have on the degree of monetary non-neutrality? We address this question numerically by calibrating our model to match three price-setting statistics generated from the Israeli database: the daily fraction of regular price changes, the average absolute size of regular price changes, and the Fisher-Konieczny (FK) synchronization index based on the variation of the fraction of price changes over time (Fisher and Konieczny, 2000). Our numerical results indicate that even small deviations from full synchronization can significantly reduce the persistence of real effects of demand shocks.

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.

For the degree of synchronization observed in the data, our model generates responses very close to the [Goloso 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, extending the result of [Alvarez, Le Bihan, and Lippi \(2016\)](#) to the partial synchronization setting. We show that varying the degree of synchronization in the model endogenously changes the kurtosis of price changes. 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. This empirical challenge is generally pertinent to menu-cost models with Gaussian idiosyncratic shocks ([Midrigan, 2011](#)). To address the discrepancy, we extend the model along two dimensions frequently employed in the literature. We introduce a fraction of free adjustments à la Calvo, as in [Alvarez, Le Bihan, and Lippi \(2016\)](#) and [Nakamura and Steinsson \(2010\)](#), and replace Gaussian shocks to the firm's desired price with fat-tailed Poisson shocks as in [Midrigan \(2011\)](#). We find that, when shocks are non-Gaussian, the determinants of monetary non-neutrality become more complex. 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. When we re-calibrate the Goloso-Lucas (GL), Alvarez-Lippi-Midrigan (ALM), and partial synchronization models to match the observed kurtosis of price changes (around 3.5), they predict quantitatively similar cumulative effects of monetary policy. An important implication of these results is that with fat-tailed shocks, the kurtosis of price changes is no longer a sufficient statistic for monetary non-neutrality, which depends also on synchronization of price changes.

Contribution to related literature

This paper contributes to the extensive literature in monetary economics that aims to reconcile stylized facts about price setting in the data on micro prices with the macroeconomic evidence on inflation and output effects of monetary policy. The micro evidence of relatively high frequency of price adjustments and high volatility of transient idiosyncratic shocks is difficult to reconcile with an inflation process that is stable, persistent, and has low sensitivity to monetary shocks, leading to persistent monetary non-neutrality. The literature has looked for solutions in theory by introducing features such as heterogeneity in nominal rigidity ([Carvalho, 2006](#), [Nakamura and](#)

Steinsson, 2010), imperfect information (Mankiw and Reis, 2002, Reis, 2006, Woodford, 2009), and strategic complementarities (Basu, 1995, Kimball, 1995). A related strand of this literature, recently spurred by Midrigan (2011) and Alvarez and Lippi (2014), asks whether the “multi-product hypothesis” can contribute to 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 changing economic conditions. 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. The synchronization of within-firm prices is corroborated by other studies.⁴ While our evidence broadly corroborates these findings, we provide a more detailed and comprehensive breakdown of multi-product pricing along both time and product dimensions. In particular, we establish a significant degree of partial synchronization of products in a store. We show that this synchronization pattern is only weakly associated with holidays and week/month fixed effects suggesting that synchronization is due to store-specific factors. We also show that time-averaging can lower measured degree of synchronization, emphasizing the importance of high-frequency price data.

To draw implications of price synchronization for 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). 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 show 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. We 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. 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 Golosov and Lucas (2007).⁵ Hence, economies of scope

⁴Studies 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); Bhattarai and Schoenle (2014); Letterie and Øivind Anti Nilsen (2016); Cavallo (2017); Gorodnichenko, Sheremirov, and Talavera (2018); Dedola, Kristoffersen, and Züllig (2019).

⁵Two recent papers highlight other mechanisms that relate multi-product pricing with weak monetary non-neutrality. Pasten and Schoenle (2016) show that under rational inattention, the effect of economies of scope in price adjustment may be reversed as attention to monetary or firm-specific shocks tends to increase with the number of products because of economies of scope in information processing. Kryvtsov and Vincent (2019) provide evidence from the U.K. and U.S. CPI micro data that the frequency of sales is strongly countercyclical. Using a business cycle model with multi-product firms and price discounts, they show that after a monetary contraction, the heightened use of discounts by firms amplifies the fall in the aggregate price level, attenuating by almost 40% the one-year response of real consumption.

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. Using two metrics of synchronization, we document the degree to which price changes are synchronized within a store, across products in a store, across stores in a given chain, and across chains. We then show that the facts are not materially affected by seasonal effects or holidays. Finally, we examine synchronization of prices changes in weekly scanner data from retail stores in the United States and show that time averaging washes out some of the important high-frequency variation in the daily frequency 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 and ensure data quality, we restrict the sample to stores with the number of observations ranked in the top 5% for each retail chain. We exclude store-specific products that do not have a general 13-digit barcode (e.g., fruits, vegetables, bakery goods). We also exclude pharmacy retailers. The remaining dataset contains 506.1 million daily observations for 75 stores and 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

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

⁷A retailer with annual sales exceeding NIS 250 million (about USD 70 million).

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 2.5 we review the results for discounted prices.

2.2 Price-adjustment behavior

Two features are unique to this dataset and are 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. We highlight these features in our empirical analysis.

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) 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. 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 (see Appendix E.1).

Monthly frequency of observations, however, filters out high-frequency movement in product prices. Mean duration of price spells implied by the frequencies of price changes in Table 1 would be 8.4 months for monthly observations, 5 months (23 weeks) for weekly observations, and 3.5 months (104 days) for daily observations. Measuring duration directly from observed spells leads to a similar conclusion.

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 (Panel A) 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

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

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	Frequency of observations	Num. of products	Mean fraction of	Mean absolute size	Synchronization of
		per period	price changes	of price changes	price changes
		(1)	(2)	(3)	(4)
<u>A. By store</u>	daily	7,217	0.87%	20.8%	0.236
	weekly	8,170	4.30%	20.3%	0.225
	monthly	9,605	11.46%	20.1%	0.187
	Price increases	7,217	0.45%	19.9%	0.218
	Price decreases	7,217	0.42%	21.4%	0.201
	Same product category*	6,989	0.89%	20.8%	0.289
	Flexible price goods	2,092	2.25%	21.2%	0.376
	Larger stores	10,267	0.89%	21.0%	0.253
	Offline stores**	7897	1.01%	20.69%	0.280
	Online stores**	10,003	0.99%	19.44%	0.343
<u>B. By chain</u>	daily	23,664	0.89%	20.8%	0.220
<u>C. All observations</u>	daily	301,496	0.99%	20.6%	0.142
<u>D. IRI data</u>	weekly	1,854	15.96%	26.2%	0.098
	monthly	2,376	28.56%	23.9%	0.133

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 Konteczny (2000) index defined in Section 2.3. Same product category: daily statistics computed for subsets of products belonging to the same product category (* Shufersal stores only). “Flexible price goods”: statistics computed for subsets of products with the daily fraction of price changes in the top quartile in the store. “Larger stores”: daily statistics computed for subsets of products in stores larger than the median store (by the number of products per day). ** We select only brick-and-mortar stores belonging to the same chains as the four online stores in our dataset.*

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. Panel B in Table 2 shows 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.

Sample	Number of days	Frequency of price changes	
		weighted	unweighted
A. By store			
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%
B. By chain			
Peak days (50%)	105	9.08%	7.57%
Off-peak days (50%)	1393	0.48%	0.39%
Peak days (25%)	36	13.21%	11.50%
Off-peak days (75%)	1462	0.68%	0.54%
All days	1498	0.89%	0.71%

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 (Panel A) and chain (Panel B). 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. For computations by chain (Panel B), we pool observations across stores within the same retail chain.

We also compared product-level frequencies of price changes, separately, for peak and off-peak days (Appendix E.2). If a store sampled prices completely randomly across its products in each day, price-change frequencies across products would be the same for the subsets of peak and off-peak days. We find instead a wide dispersion of product-level price changes between peak and off-peak days. For example, the share of products whose adjustments are twice as likely during peak days than during off-peak days is 0.65 for Shufersal and 0.52 for all stores. Conversely, the share of products whose adjustments are twice as likely during off-peaks days than during peak days is 0.12 for Shufersal and 0.07 for all stores. Hence, on peak re-pricing days, stores increase (decrease) the likelihood of changing prices for some of their products relative to other days.

2.3 Measures of synchronization

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 Table 1 (column 4) provides the weighted mean of the index for the stores in the dataset. At the daily frequency, the index is 0.236 on average, and it varies considerably across stores (see Appendix E.2).

To ensure that our results are not driven by a particular way of measuring synchronization, we analyze an alternative synchronization measure. The second synchronization index is akin to the Gini inequality index. It is based on the Lorenz curve that depicts the distribution of repricing activity by plotting the percentile of days by the number of price changes on the horizontal axis and cumulative fraction of price changes on the vertical axis. Unlike the FK index, the Gini index does not directly depend on the average frequency. Instead, it is a better measure of the spread of re-pricing activity across days in a given store. Nonetheless, we demonstrate in Appendix E.3 that two synchronization measures are largely in line with each other across different dimensions of the data.

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

Synchronization of price changes is a distinct characteristic of firm-level pricing. We observe, in particular, that synchronization within stores is higher than across stores. When we pool observations across stores, the corresponding FK synchronization index values are substantially lower than their store-level averages: 0.142 vs 0.236 (row C in Table 1). 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.

To highlight the variation of price-change synchronization across stores, we study how it depends on similarities across goods, price flexibility, store size and type, and chain effects.

First, we ask whether prices are more synchronized for products within the same broad category 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). The Bank of Israel classified 50 broad product categories for Shufersal stores. Panel A in Figure 2 compares the distribution of the FK synchronization index across stores computed for all products with the distribution conditional on products in the same category. Indeed, price changes are more synchronized within, rather than across, categories. On average, the FK index is 0.289 for products in the same category, versus 0.236 for all products (Table 1).

Similarly, Panel B in Figure 2 compares the distribution for all products with the one for only flexible price products, which we define as products in the top quartile of the frequency of price changes in the store. Flexible price changes are more synchronized, with an FK index of 0.356. 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

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

products adjust prices more frequently.

In our data, larger stores tend to have more synchronized price changes, as evidenced in Panel C in Figure 2. The FK index for the subset of stores above the median is 0.253, versus 0.236 for all stores. 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.

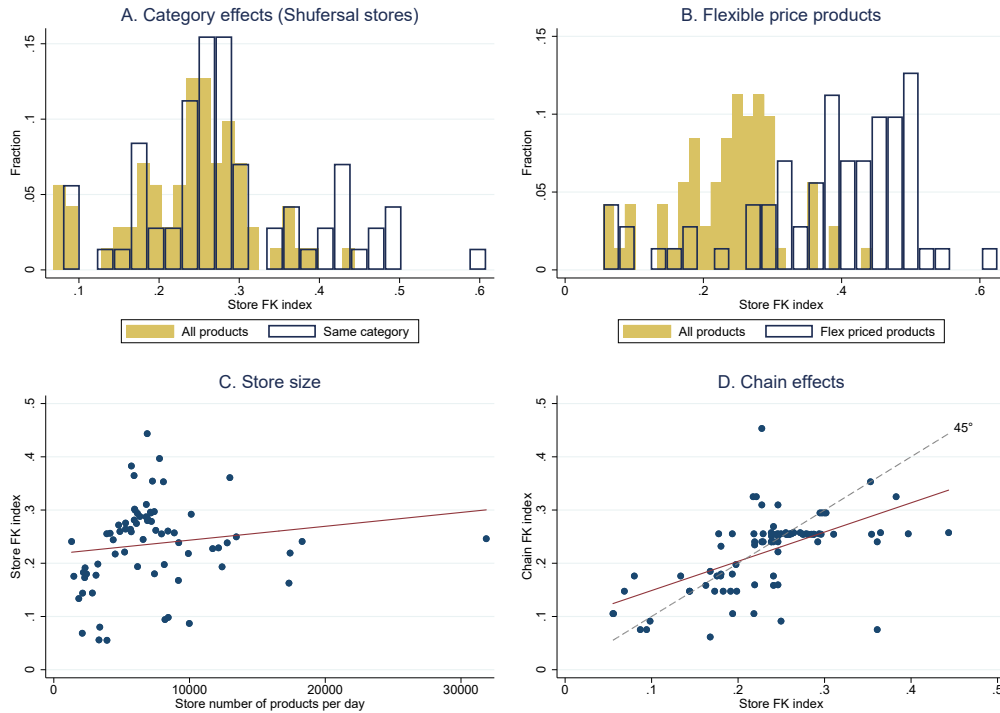


Figure 2: Fisher-Konieczny synchronization across stores.

Note: Panels A and B provide histograms for the Fisher-Konieczny (FK) index across stores. Panel A: all products and those within the same product category (Shufersal stores only). Panel B: all products and those for the products with mean frequency of price changes in the top quartile for that store. Panel C: scatter plot of store size (number of products per day) versus corresponding store FK index values, with the (unweighted) fitted OLS line. Panel D: scatter plot of FK index values for the store (x-axis) and its chain (y-axis), with 45° line and fitted OLS line.

We also compared synchronization of prices in online and brick-and-mortar stores. For cleaner comparison, we only looked at brick-and-mortar stores belonging to the same chains as the four online stores in our dataset. The frequency and size of price changes are very similar for online and offline prices, but online prices are more synchronized, with an FK index of 0.343 versus 0.280 for prices in the brick-and-mortar stores in the same chains. 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 compare FK index values for each store with corresponding FK index values for the chain to which that store belongs. The scatter plot in Panel D of Figure 2 shows that synchronization of price changes across products of the entire chain is somewhat lower than at a store level. The weighted mean FK index goes down to 0.220. 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, Hortagsu, and Lin \(2019\)](#) for U.S. retail chains.

2.5 Price discounts

Do these patterns in adjustment of regular prices also apply to final prices, which incorporate various types of price discounts? To compare the results for final and regular prices, we use the dataset for 10 stores of the largest retail chain in Israel, Shufersal, from January 2016 until mid-2019. The Bank of Israel 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 E.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, 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. The Gini index for the fraction of discounts is 0.174, indicating a much more even distribution of discounts across days than the distribution of the fraction of regular or final price changes, with corresponding Gini index values of 0.708 and 0.797.

2.6 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 E.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 E.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.

2.7 Comparisons with weekly scanner data for the United States

We highlighted that comprehensive coverage of products in a store is a distinguishing feature of this dataset which reveals striking patterns of price adjustments by large retailers. The second distinctive feature of the dataset is its detailed coverage over time—providing daily price quotes for each of the store’s very many products. Most micro data on prices provide less frequent price quotes,

typically at weekly or monthly frequency. To underscore the implications of time aggregation of price observations, 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.¹⁰ The data provide weekly revenues and quantities for products sold in grocery stores across 50 U.S. metropolitan areas. In a given week, unit prices for each product are constructed by dividing weekly revenue by the quantity sold. The share of nominal revenues over the sample period in total revenues over the sample period is used as a store weight.¹¹

For better comparisons with the Israeli data, we focus on stores that belong to large retail chains, and so we exclude independent stores from the IRI sample. The dataset that we use contains around 448 million weekly observations in 518 stores and 75 retail chains, and covers the span of 132 months, from January 2001 to December 2011. Unlike the data for Israel, the IRI dataset includes small stores. Nonetheless, the average number of products sold in a given week is 1,854, reflecting a sufficiently large number of large stores. Another reason for a smaller number of products per week is that the IRI dataset does not provide prices for products that did not register a sale within a week. This implies a higher frequency of weekly price changes in the IRI dataset, 16% (Panel D, Table 1), than the frequency of weekly posted prices in Israel, 4.3%. Conditional on a price change, however, the absolute magnitudes of price changes are not so far apart: 26.2% in the IRI data versus 20.3% in the Israeli data.

Measured synchronization of weekly price changes in IRI data is much lower than synchronization of weekly price changes in the Israeli data, 0.098 versus 0.225. When we further aggregate Israeli price changes to a monthly frequency of observations, synchronization decreases even more, to 0.187. We do not obtain the same aggregation result when we aggregate weekly IRI price changes to a monthly frequency, rather, the FK index increases to 0.133.

Higher frequency of price changes in IRI data than in the Israeli data 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.¹² The differences in the results may also reflect time aggregation

¹⁰More details are provided in [Bronnenberg, Kruger, and Mela \(2008\)](#).

¹¹We define a unique product identifier by matching UPC codes for that product with the product description (e.g., Budweiser lager 355ml). We exclude products that belong to a store’s private label (their coding was changed by IRI in 2007 and 2008), and products that have fewer than two observations per week, and we exclude observations with a unit price less than \$0.10.

¹²The price-setting model we introduce in Section 3 predicts that raising item-specific menu cost *ceteris paribus* increases the average fraction of adjusting prices and increases the synchronization of price changes. Synchronization of price changes in IRI data is *lower* than in the Israeli data (Table 1). Hence, to account for the differences in the average fraction and synchronization of price changes in the U.S. and Israel, the model requires lower store-specific

of the IRI dataset, which only provides average price changes on a weekly basis. Cavallo (2018) argues that even datasets that provide weekly average prices can produce biased results for price change statistics, such as the frequency and size of price changes.

To gauge the extent to which our quantitative results depend on time aggregation, we compare IRI and Israeli price changes along two additional sets of statistics that indirectly depend on the degree of price change synchronization.

Distribution of standardized price changes. First, the literature has emphasized that quantitative success of menu-cost models is associated with their ability to match the distribution of price changes in the data. In particular, 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. Cavallo (2018) shows that time aggregation can materially affect the distribution of price changes.

Following the literature, we filter out heterogeneity across stores by constructing “standardized” log price changes for product i , store s , and time t , z_{ist} , 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 then compute four statistics characterizing the shape of the distribution of standardized price changes that are used in Table II in Alvarez and Lippi (2014). To control for heterogeneity of price changes within a store, we also use three alternative definitions of standardized changes, where for prices within a store we control for differences in price levels, product’s frequency of price changes, or product category. For comparison with IRI data, we also report these statistics at weekly and monthly observation frequencies.

We summarize the results as follows (see Appendix E.7 for details). First, all four statistics are similar for four alternative ways of controlling for heterogeneity across products. Second, the statistics for IRI data are similar to those produced by Midrigan (2011) using AC Nielsen and Dominick’s scanner data. Third, comparisons of daily, weekly and monthly Israeli statistics indicate that time aggregation from daily to weekly or monthly frequency roughly doubles the mass of small price changes. We also observe that kurtosis and the mass of small price changes are slightly higher than in Israeli data for either weekly or monthly frequency, which could also be in part due to price averaging in the weekly scanner data.

In all, these results underscore the benefits of using daily price observations for fully exploring the implications of synchronization of price changes.

Spectral analysis of store-level price changes. 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.¹³ 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

cost of price adjustment in the U.S. relative to Israel.

¹³We thank Fernando Alvarez for this insight.

density of its daily (weekly, monthly) 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 (see details in Appendix E.8).

There are two main findings. First, for the Israeli stores, spectral density is high for short frequencies of less than 30 days. This remains the case as we aggregate daily price changes into weekly and monthly, and then construct corresponding weekly and monthly 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. Second, including price discounts does not alter the shape of the spectra, which is consistent with our conclusion that peaks are not driven by sales.

Finally, the weekly IRI spectrum is also heavier on the short end, suggesting high-frequency movements in price changes similar to those we find in the Israeli data. This is no longer the case when we aggregate IRI data from weekly to monthly frequency. The monthly IRI spectrum is now tilted upward, unlike the Israeli data, which are still sloping downward. This suggests that time averaging in IRI data washes out some of the important high-frequency variation in the daily frequency of price changes.

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 . Consequently, absent frictions of any nature, a firm always charges the frictionless optimal price. We employ a commonly used second-order approximation for the profit from selling good i around the frictionless optimal price. Therefore, by denoting good i 's price at instant t by $p_{i,t}$, the firm minimizes a loss term L_t of the following form:

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

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. It is simpler, however, 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.¹⁴ 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 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 building block we need to fully characterize the firms' problem is the pricing friction firms face. 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 $K + cm$.

The presence of these different sorts of menu costs gives rise to optimal policies that have three important features. First, a positive fixed cost K generates inaction: a firm's price adjustments occur in dates that are separated by potentially long time intervals. Second, when a firm decides to adjust prices, it does not adjust all of them. Intuitively, it is never optimal for a firm to adjust all prices because there will always be products with arbitrarily small price discrepancies (in absolute values). If a price discrepancy is small enough, it is not optimal to pay the unit menu cost c to set the discrepancy to zero.

Lastly, even though the frictionless optimal price of a single product is stochastic, the relevant state object is the entire distribution of price discrepancies, which evolves in a deterministic way given by the KFE (3). Therefore, 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

¹⁴The 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.

$t = T_k$ firms adjusts 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.¹⁵ 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, imagine the analogous problem for a firm that sells a finite number of products. If, for example, the firm just paid the fixed cost and decided to reset the price of a single product, this product must optimally be the one that has the largest price gap. This is the case because 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, when deciding which goods will have their prices reset, the firm will rank its products according to the size of price discrepancies and, starting from the good with the largest gap, will adjust prices moving down this list until the marginal benefit of adjusting the next price is smaller than the unit cost c . Hence, it is never optimal to adjust a price until all prices with larger discrepancies have been adjusted.

3.2 Interpretation

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), similar to an early study by [Sheshinski and Weiss \(1992\)](#). Empirical literature documents multi-faceted 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 C](#) 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. 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

¹⁵In 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.

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.

3.3 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 A. □

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 A. □

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. 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^* , together with an intermediate distribution g_0^* and optimal policies $\tau^* = \tau(g^*)$ and $\bar{x}^* = \bar{x}(g^*)$, which satisfies the fixed-point problem:

$$g_0^*(x) = g^*(x)1(|x| < \bar{x}^*) + m(\bar{x}, g^*)\delta_0(x), \quad (7)$$

$$g^*(x) = \int_{-\infty}^{+\infty} \frac{1}{\sqrt{\sigma^2\tau^*}} \phi\left(\frac{x-y}{\sqrt{\sigma^2\tau^*}}\right) g_0^*(y) dy. \quad (8)$$

Above, g_0^* is the distribution that arises after prices are reset according to the threshold \bar{x}^* . The relationship between g^* and g_0^* expressed in (8) is a direct application of Lemma (4). The system is easily solvable by creating a discrete grid for possible values of x , since (7) expresses g_0^* as a linear function of g^* , while (8) writes g^* as a linear transformation of g_0^* . If we represent both distributions by the vectors of the values they attain in the x grid, the problem boils down to finding an eigenvector of a large matrix. Figure 3 shows the steady-state distribution of price gaps, before and after price adjustments, 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^* .

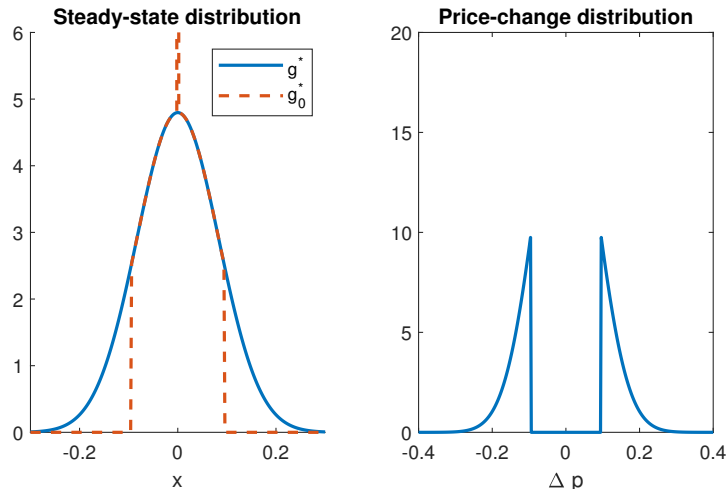


Figure 3: Steady-state and price-change distributions for parameter values $\rho = 0.04$, $\sigma = 0.25$, $K = 0.0001$, $c = 0.001$.

Finally, Figure 4 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 fixed and unit costs alters this pattern.

A high fixed cost K , combined with a low unit 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 4: Daily fraction of prices changes for different combinations of menu costs. Other parameter values: $\rho = 0.04$, $\sigma = 0.25$.

3.4 Calibration

To compare predictions of different models, we calibrate not only the partial synchronization model, but also the Golosov-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. More specifically, our measure of the frequency of adjustments is the average daily share of prices that a firms adjusts. The last moment we pick is the Fisher-Konieczny (FK) index of synchronization.

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 natural since the GL model cannot generate any value for the FK statistic other than zero, while the ALM model can only generate one. Table 3 shows moments for models and data, and Table 4 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 extend our model to match these two latter moments, and discuss their importance.

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	0.236	0.000	1.000	0.236
Additional moments				
Standard deviation of Δp	0.248	0.208	0.260	0.209
Fraction of small adjustments	0.06	0	0.16	0
Kurtosis of Δp	3.53	1	3	1.08

Table 3: 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 E.7.

Parameter	GL	ALM	Partial sync.
σ	0.370	0.464	0.373
K	-	0.011	3.02e-05
c	0.0023	-	0.0020

Table 4: Calibrated parameter values.

Figure 5 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 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 17.8% every 7.5 days.

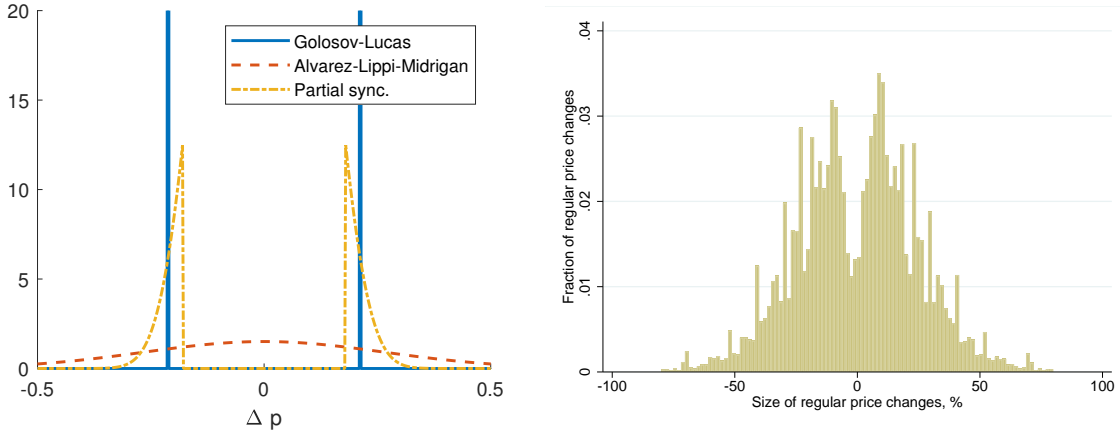


Figure 5: Price-change distributions for all models and data.

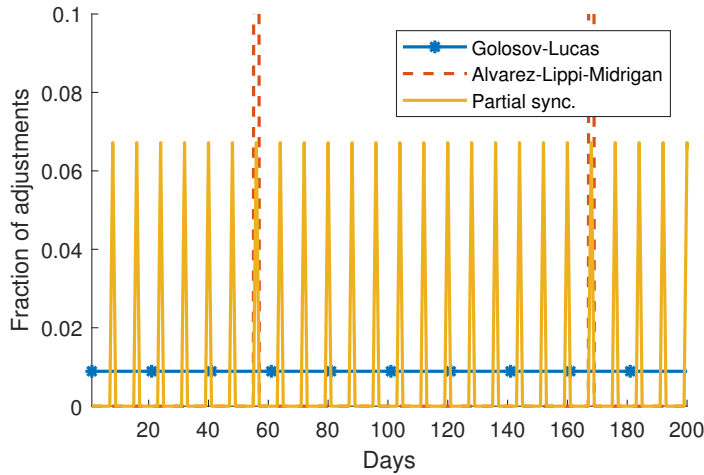


Figure 6: Daily share of adjusted prices for the three models. Spikes in the ALM model have a height of one.

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. Since 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). 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 $t = 0$. To understand how aggregates respond to such a shock, we must first understand what optimal policy following the shock looks like.

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

$$T_1(0) = \tau^* - s,$$

$$T_{k+1}(0) = T_k(0) + \tau^*,$$

$$\bar{x}_k(0) = \bar{x}^*.$$

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 7.

Since we assume that firms are uniformly distributed according to s —the time elapsed since the last adjustment before the shock—it follows that at each instant there is a 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 rate Δ_k/τ^* . More precisely, define $P_\varepsilon(t)$ to be the




Figure 7: Approximate response of the average price of a single firm following a small aggregate shock of size ε .

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 A. □

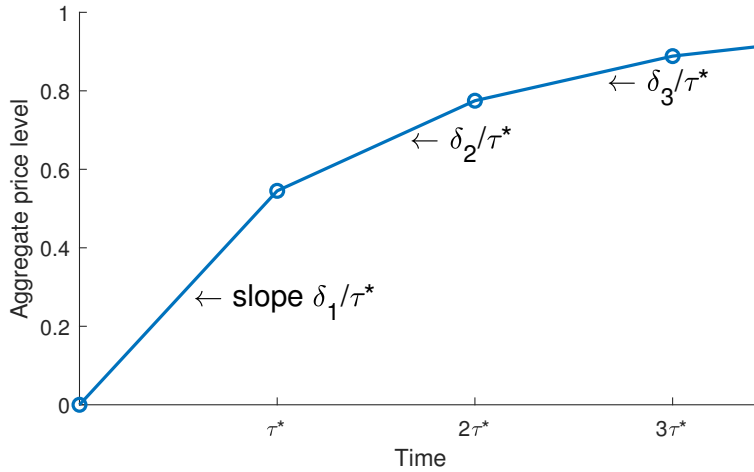


Figure 8: Price-level response to a small aggregate shock.

Figure 8 shows the limiting response of the aggregate price level to a small aggregate shock. Our next result characterizes the slope of the first line segment of the impulse response function. This slope is quantitatively important since, as we shall see, almost 40% of the rise of the price level following an aggregate shock happens in this first segment for our calibration. But first, a few more definitions. 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 F 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 5 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 A.

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*

$$\frac{\partial \Delta_k}{\partial T_j} = \frac{\partial \Delta_k}{\partial \bar{x}_j} = 0. \quad (9)$$

Proof. See Appendix A. □

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*

$$\frac{\delta_1}{\tau^*} = F \times [1 + \bar{x}^* f(\bar{x}^*)]. \quad (10)$$

Proof. See Appendix A. □

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

Building on the analytical results we have presented, we now turn to comparing how different models predict output to respond to demand shocks. Figure 9 shows the impulse response functions for the

calibrated parameter values shown in Table 4. 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. We divide our explanation of this result in two parts. First, we analyze why more synchronization is associated with more persistent real effects of demand shocks. Second, we quantitatively explore how fast the introduction of partial—as opposed to perfect—synchronization reduces monetary non-neutrality. In other words, how far do we need to move away from the ALM perfect synchronization case to have significantly less persistent output responses? The answer is: not very far.




Figure 9: Output response to a small aggregate shock.

Why is partial synchronization associated with smaller 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. We can gain some intuition by looking at expression (10). Since all models are calibrated to match the same frequency of adjustment, the difference must come from the selection component.¹⁶ We can see the ALM model as the limit when $\bar{x}^* \rightarrow 0$. Therefore, 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 emerges thus as a form of selection, as explored by Golosov and Lucas (2007). If a firm adjusts a large fraction of its prices on an adjustment date,

¹⁶Even 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 $F = 365 \times$ daily frequency in steady state, as long as the interval between adjustment dates is longer than one day.

there is not much room for selecting those prices, and monetary shocks have more persistent real effects as a consequence.¹⁷

Now we aim to understand how fast this selection effect operates. 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. As expected from the previous discussion, output decays faster as synchronization fades. Looking only at the selection term $1 + \bar{x}^* f(\bar{x}^*)$ may not be sufficient to understand this relationship because this 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 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.

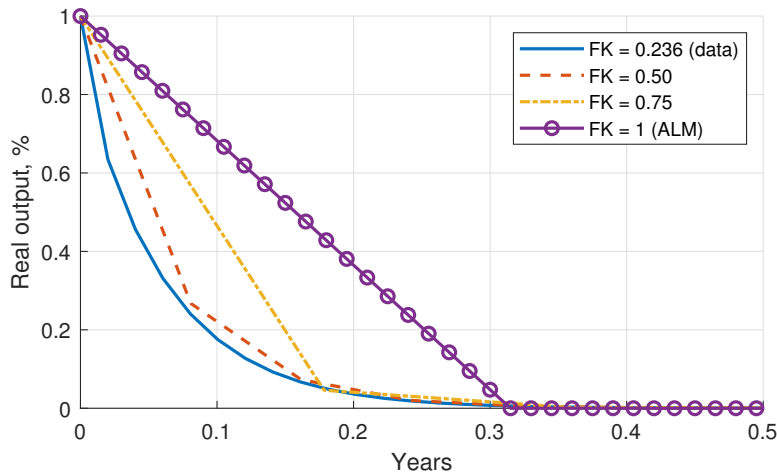


Figure 10: Real output response and synchronization.

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 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). Our partial synchronization model is not included in that class of models. Nevertheless, the ALL result appears to hold for our partial synchronization model as well. Figure 12 shows that as kurtosis of price changes varies across models with different degrees of synchronization of price changes (solid blue line), the ratio

¹⁷Since we study the limiting case in which the shock size goes to zero, the aggregate response does not come from anticipation of adjustment episodes.

of kurtosis to the area under the impulse response function is roughly constant (red dotted line). 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. In particular, with infinitely many products but imperfect synchronization, kurtosis is always less than 3, except in the case of the perfect synchronization ALM model, where all price changes are equal to (Gaussian) increments in frictionless optimal prices, and hence, are distributed with a kurtosis of exactly 3.

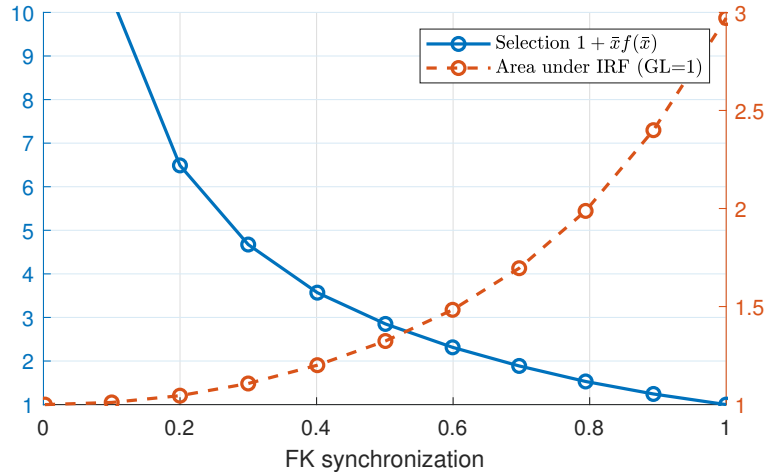


Figure 11: Selection component $1 + \bar{x}^* f(\bar{x}^*)$ and cumulative response for various values of the FK synchronization measure. Cumulative response of the GL model is normalized to 1.

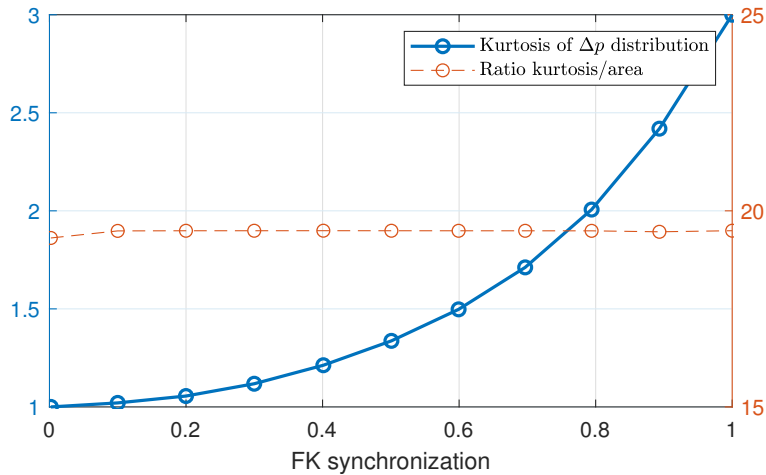


Figure 12: Kurtosis of price-change distribution for various values of the FK synchronization measure.

Large kurtosis of the price-change distribution of our sample poses the following challenge to

existing theories. As can be seen in Table 3, 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 12 shows, is associated with low kurtosis. There are two possible modifications to our model that have been frequently employed in the literature and could possibly 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 3 and Figure 5, 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 D. The other alternative, which we explore next, is the introduction of fat-tailed shocks 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 ,¹⁸ 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 13 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.

In this case, to simplify the numerical procedure, we do not solve for the optimal policy of the firm.¹⁹ Instead, we calibrate the model by choosing policies \bar{x} and τ and parameters of the stochastic process η and β .²⁰ 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

¹⁸The 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 $1/\eta$.

¹⁹When 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.

²⁰We solve the model in the following way. Given parameters η and β and a policy \bar{x} and τ , we use equations (7) and (8) 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.

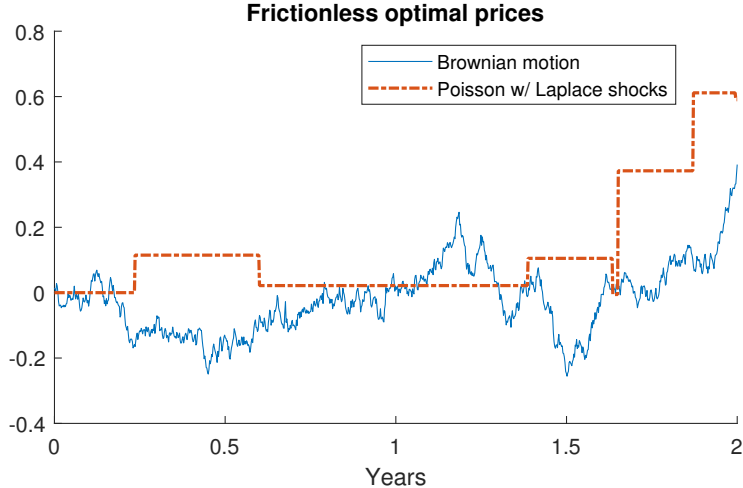


Figure 13: Sample paths of Brownian motion and Poisson process with Laplace shocks.

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 4.

As previously, we calibrate GL, ALM, and partial synchronization models with this new stochastic process.²¹ 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 5 and 6. 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 6). 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.

²¹In 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 .

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	0.236	0.000	1.000	0.236
Kurtosis of Δp	3.53	3.52	3.53	3.53
Additional moments				
Standard deviation of Δp	0.248	0.253	0.266	0.253
Fraction of small adjustments	0.06	0	0.17	0

Table 5: 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	0.059
τ	-	0.315	0.020
η	4.70	35.95	4.90
β	0.150	0.056	0.147

Table 6: Calibrated parameter values for models with fat-tailed shocks.

Turning to monetary non-neutrality, Figure 14 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 synchronization cases, compensating for the selection effect embedded in those models.

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 15 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 nonneutrality when kurtosis is 3, and the degree of non-neutrality across models is reversed for higher values of kurtosis.²²

²²By coincidence, when kurtosis is exactly equal to that of the distribution of price changes in our sample, our partial synchronization model generates almost exactly the same area under IRF as the ALM model.

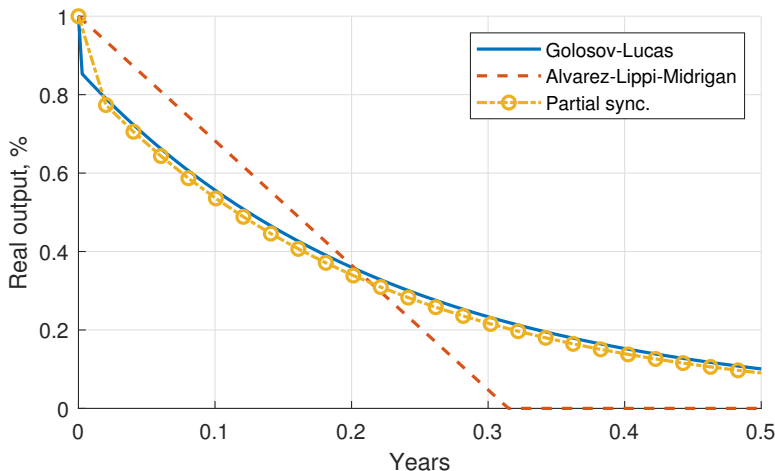


Figure 14: Impulse response functions for all models with fat-tailed shocks.

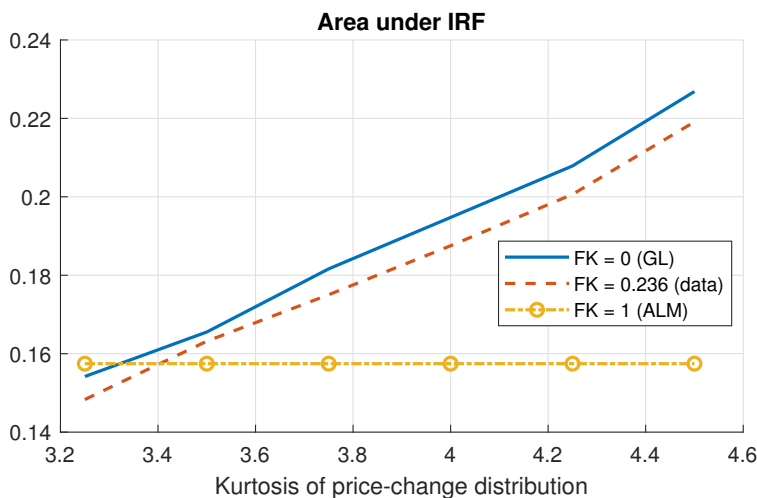


Figure 15: Area under impulse response functions as a function of the kurtosis of price-change distribution.

The results displayed in Figure 15 show that the determinants of monetary non-neutrality become more complex when shocks cease to be Gaussian, and that synchronization of price changes is one of these determinants. In particular, for ALM, monetary non-neutrality does not vary with kurtosis. On an adjustment date, a firm in the ALM model resets all its prices, so the increase in its average price after such an episode depends on the mean of its price-discrepancy distribution g but is independent of its shape. This is visible in equation 10 for the first slope of the IRF (since ALM has only one slope). Since in the ALM model $\bar{x}^* = 0$, the slope of the impulse response function depends only on the frequency of adjustments. In general, the degree of monetary non-neutrality displayed in the graph is the result of the interaction between exogenous features of the process for $p_{i,t}^*$ and the strength of the selection effect in the model that varies with the degree of synchronization. Although the kurtosis of price changes depends also on those two factors, it ceases to be a

sufficient statistic for non-neutrality.

Figure 16 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 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, a given level of price-change kurtosis may be generated by different combinations of exogenous kurtosis and degrees of synchronization. To determine monetary non-neutrality, however, it is necessary to understand how much of the kurtosis is due to the shock process, and how much of it emerges endogenously as a function of price selection, which in turn is a function of synchronization.

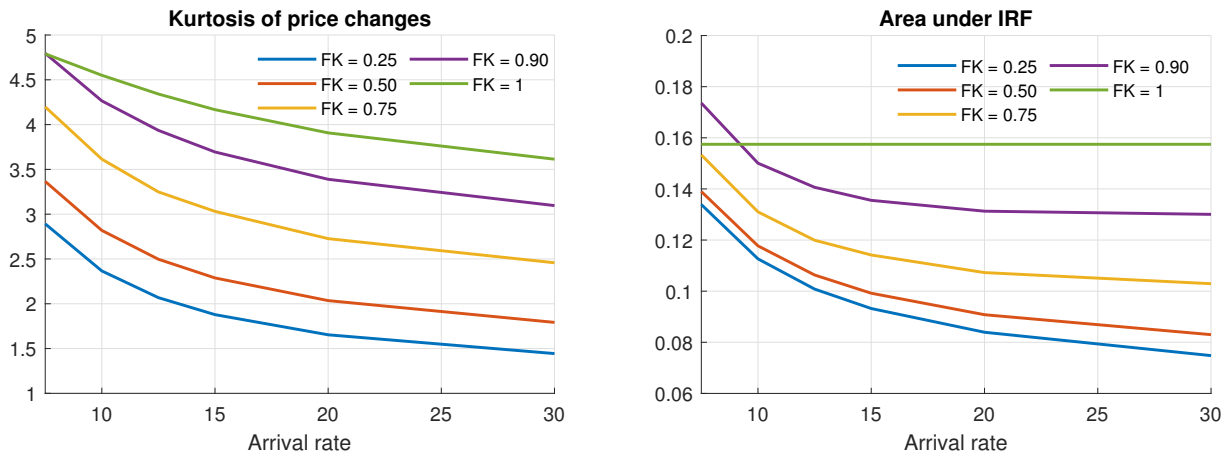


Figure 16: Area under impulse response functions as a function of the kurtosis of price-change distribution.

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

In this paper, we bring new evidence on prices set by large retailers and propose a general equilibrium model of multi-product price setting to account for this evidence. The data provide unique detail on day-to-day price changes across thousands of products sold in large food retailers. The pattern that emerges in the data is that prices of multi-product firms are partially synchronized due to occasional peaks in these firms' repricing activity, where a large share of the prices are simultaneously adjusted. To generate this pattern, we propose a theory of multi-product pricing that incorporates a price-adjustment technology with an endogenous degree of the economies of scope. We calibrate the model based on a synchronization index and other key price-setting statistics from the micro data and derive implications for monetary non-neutrality, comparing it with two extreme nested models. In the case with no economies of scope, the firm sets the price of each product independently,

paying a separate menu cost for each price change. This case is equivalent to a single-product pricing model. In the case with maximal economies of scope, the firm pays the fixed cost to adjust all of its prices. We show that the calibrated model, despite displaying considerable synchronization of price changes, generates relatively small monetary non-neutrality, close to that generated by an economy populated by single-product firms. The selection effect in the partial synchronization model—with the monetary shock triggering initial adjustments from the group of products with infra-marginal prices—is similar in magnitude to the selection effect in standard menu-cost models.

References

- Bank of Israel. “The Bank of Israel Annual Report for 2012.” URL <http://www.boi.org.il/en/NewsAndPublications/PressReleases/Pages/02042013.aspx>.
- 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 Luigi Paciello. 2011. “Optimal Price Setting With Observation and Menu Costs.” *The Quarterly Journal of Economics* 126 (4):1909–1960.
- Basu, Susanto. 1995. “Intermediate Goods and Business Cycles: Implications for Productivity and Welfare.” *The American Economic Review* 85 (3):512–531.
- 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. 2006. “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks.” *The B.E. Journal of Macroeconomics* 2.
- 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.
- 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.
- 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.
- Kimball, Miles S. 1995. “The Quantitative Analytics of the Basic Neomonetarist Model.” *Journal of Money, Credit and Banking* 27 (4):1241–1277.
- 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 Nicolas Vincent. 2019. “The Cyclicalities of Sales and Aggregate Price Flexibility.” Mimeo, October.
- 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.
- Mankiw, Gregory N. and Ricardo Reis. 2002. “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve.” *Quarterly Journal of Economics* 117 (4):1295–1328.
- Midrigan, Virgiliu. 2011. “Menu Costs, Multi-Product Firms and Aggregate Fluctuations.” *Econometrica* 79:1139–1180.
- Nakamura, Emi and Joñ 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.
- Pasten, Ernesto and Raphael Schoenle. 2016. “Rational inattention, multi-product firms and the neutrality of money.” *Journal of Monetary Economics* 80 (C):1–16.
- Reis, Ricardo. 2006. “Inattentive Producers.” *Review of Economic Studies* 73 (3):793–821.

- 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.
- Woodford, Michael. 2009. “Information-constrained state-dependent pricing.” *Journal of Monetary Economics* 56 (S):100–124.
- 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.

A Proofs

A.1 Lemma 1

Substitute (4) into (3) and use $\phi''(x) = -x\phi'(x) - \phi(x)$.

A.2 Lemma 2

From (4), it follows that

$$\begin{aligned} \int_{-\infty}^{+\infty} x^2 g_t(x) dx &= \int_{-\infty}^{+\infty} \left\{ \int_{-\infty}^{+\infty} x^2 \frac{1}{\sqrt{\sigma t}} \phi\left(\frac{x-y}{\sqrt{\sigma t}}\right) dx \right\} g_0(y) dy \\ &= \int_{-\infty}^{+\infty} \{y^2 + \sigma^2 t\} g_0(y) dy \\ &= \int_{-\infty}^{+\infty} y^2 g_0(y) dy + \sigma^2 t. \end{aligned}$$

The second line follows because we have the p.d.f. of a Normal distribution with mean y and variance $\sigma^2 t$, and the third line follows from the fact that $g_0(x)$ integrates to one.

A.3 Proposition 1

Let $\bar{p}_\varepsilon^s(t)$, for $t \geq 0$, be the average price of a firm whose last price adjustment before the shock was at date $-s$ (i.e., s periods before the shock hits) for $0 \leq s < \tau^*$. Since firms are uniformly distributed according to the time elapsed between the last adjustment date and $t = 0$, it follows that

$$P_\varepsilon(t) = \int_0^{\tau^*} \frac{1}{\tau^*} \bar{p}_\varepsilon^s(t) ds.$$

Therefore,

$$\lim_{\varepsilon \rightarrow 0} \frac{P_\varepsilon(t)}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \int_0^{\tau^*} \frac{1}{\tau^*} \frac{\bar{p}_\varepsilon^s(t)}{\varepsilon} ds. \quad (\text{A.1})$$

Let $\{T_k^s(\varepsilon)\}_{k=1}^\infty$ be the optimal adjustment dates for firm s , and Δ_k^s be the change in the firm's average price after the k -th adjustment episode following the shock. Of course, Δ_k^s depends on ε and on the optimal policy, but we shall omit this dependence for now. Δ_k^s will be studied in more detail in the next proof. We have:

$$\bar{p}_\varepsilon^s(t) = \begin{cases} 0 & t \in [0, T_1^s(\varepsilon)) \\ \sum_{j=1}^{k-1} \Delta_j^s & t \in [T_{k-1}^s(\varepsilon), T_k^s(\varepsilon)) \text{ and } k \geq 2. \end{cases}$$

If adjustment dates are continuous in ε at point $\varepsilon = 0$, we $T_k^s(\varepsilon) \rightarrow T_k^s(0) = k\tau^* - s$ as $\varepsilon \rightarrow 0$. As a consequence,

$$\frac{\bar{p}_\varepsilon^s(t)}{\varepsilon} \rightarrow \begin{cases} 0 & t \in (0, \tau^* - s) \\ \sum_{j=1}^{k-1} \delta_j & t \in ((k-1)\tau^* - s, k\tau^* - s) \text{ and } k \geq 2. \end{cases}$$

Two observations are important here. First, the limit in points of the form $k\tau^* - s$ may be undefined. It could be either $\sum_{j=1}^{k-1} \delta_j$ if $T_k^s(\varepsilon)$ converges to $k\tau^* - s$ from below, $\sum_{j=1}^k \delta_j$ if it converges from above, or otherwise undefined. Nevertheless, this is irrelevant since it is a set of measure zero. Second, δ_k does not depend on s . This happens because firms with higher s do exactly the same as firms with lower s , only with a delay. This could fail for large shocks if, for example, two firms characterized by different s respond immediately to a large shock, but this is irrelevant as the shock size goes to zero and $T_k^s(\varepsilon) \rightarrow k\tau^* - s$.

Moreover, we have $\frac{\bar{p}_\varepsilon^s(t)}{\varepsilon} \in [0, 1]$, i.e., the firm's average price level does not overshoot the increase in demand from the aggregate shock.²³ So the Dominated Convergence Theorem allows us to exchange the order of integration and limit operators in (A.1), which gives us

$$\lim_{\varepsilon \rightarrow 0} \frac{P_\varepsilon(t)}{\varepsilon} = \begin{cases} \frac{\delta_1}{\tau^*} \frac{t}{\tau^*} & t \in [0, \tau^*) \\ \sum_{j=1}^{k-1} \frac{\delta_j}{\tau^*} + \frac{\delta_k}{\tau^*} \left[\frac{t}{\tau^*} - (k-1) \right] & t \in [(k-1)\tau^*, k\tau^*) \text{ and } k \geq 2. \end{cases}$$

A.4 Lemma 3

For any given sequence of adjustment dates $\{T_j\}_{j=1}^\infty$ and thresholds $\{\bar{x}_j\}_{j=1}^\infty$, not necessarily optimal, let $g_k(x; \varepsilon, \{T_j\}_{j=1}^k, \{\bar{x}_j\}_{j=1}^{k-1})$ be the distribution of price gaps that emerges immediately before the k -th adjustment episode. Note that g_k is determined by the first k values of the $\{T_j\}_{j=1}^\infty$, but only by the first $k-1$ values of $\{\bar{x}_j\}_{j=1}^\infty$, since it is the distribution at date T_k *before* adjustments take place. We can then express Δ_k as

$$\Delta_k(\varepsilon, \{T_j\}_{j=1}^k, \{\bar{x}_j\}_{j=1}^k) = - \int_{|x| > \bar{x}_k} x g_k(x; \varepsilon, \{T_j\}_{j=1}^k, \{\bar{x}_j\}_{j=1}^{k-1}) dx. \quad (\text{A.2})$$

For a small aggregate shock, we can use the chain rule to obtain the following approximation:

$$\Delta_k(\varepsilon, \{T_j(\varepsilon)\}_{j=1}^k, \{\bar{x}_j(\varepsilon)\}_{j=1}^k) \approx \Delta_k(0, \{T_j(0)\}_{j=1}^k, \{\bar{x}_j(0)\}_{j=1}^k) + \left[\frac{\partial \Delta_k}{\partial \varepsilon} + \sum_{j=1}^k \frac{\partial \Delta_k}{\partial T_j} T_j'(0) + \sum_{j=1}^k \frac{\partial \Delta_k}{\partial \bar{x}_j} \bar{x}_j'(0) \right] \varepsilon.$$

In the above, all the partial derivatives are evaluated at point $(0, \{T_j(0)\}_{j=1}^k, \{\bar{x}_j(0)\}_{j=1}^k)$, but this argument is omitted for conciseness. Note, moreover, that in the absence of any innovation ($\varepsilon = 0$), we are in a steady state with a constant price level, so $\Delta_k(0, \{T_j(0)\}_{j=1}^k, \{\bar{x}_j(0)\}_{j=1}^k) = 0$, and the

²³This can be shown from the fact that adjustments happen symmetrically around zero (i.e., a price with gap x is adjusted if and only if a price with gap $-x$ is adjusted) and that reset prices are set so as to have zero gap.

approximation becomes

$$\Delta_k(\varepsilon, \{T_j(\varepsilon)\}_{j=1}^k, \{\bar{x}_j(\varepsilon)\}_{j=1}^k) = \left[\frac{\partial \Delta_k}{\partial \varepsilon} + \sum_{j=1}^k \frac{\partial \Delta_k}{\partial T_j} T_j'(0) + \sum_{j=1}^k \frac{\partial \Delta_k}{\partial \bar{x}_j} \bar{x}_j'(0) \right] \varepsilon. \quad (\text{A.3})$$

Now observe that, if we set $\varepsilon = 0$, the distribution g_k will be symmetric for all k regardless of the sequences $\{T_j\}_{j=1}^\infty$ and $\{\bar{x}_j\}_{j=1}^\infty$. Therefore, the integral (A.2) will always be zero and is consequently independent of T_j and \bar{x}_j for any j , so the partial derivatives with respect to policy variables are zero and our first order approximation becomes

$$\Delta_k(\varepsilon, \{T_j(\varepsilon)\}_{j=1}^k, \{\bar{x}_j(\varepsilon)\}_{j=1}^k) \approx \frac{\partial \Delta_k}{\partial \varepsilon} \varepsilon.$$

A.5 Proposition 2

In steady state, we have:

$$\begin{aligned} \frac{\partial \Delta_1}{\partial \varepsilon} &= \frac{\partial}{\partial \varepsilon} \left(- \int_{|x| > \bar{x}^*} x g^*(x + \varepsilon) dx \right) \\ &= - \int_{|x| > \bar{x}^*} x g^{*l}(x) dx. \end{aligned}$$

Integrating by parts, we have:

$$\begin{aligned} \frac{\partial \Delta_1}{\partial \varepsilon} &= 2\bar{x}^* g(\bar{x}^*) + m \\ m &= \int_{|x| > \bar{x}^*} g^*(x) dx. \end{aligned} \quad (\text{A.4})$$

The result follows from rearranging (A.4) and using $F = m/\tau^*$ and

$$f(x) = \frac{2g(x)1(x \geq \bar{x}^*)}{m}.$$

B Numerical method for solving the price-setting problem

The strategy we adopt to solve the recursive problem 6 is to approximate the distribution g by a member of some parametric family of probability distributions. What motivates our choice is the following. Given an initial condition g_0 , equation (4) tells us that the solution to the KFE g_t can be written as the p.d.f. of a sum of two independent random variables: one with p.d.f. g_0 and the other normally distributed as $N(0, \sigma^2 t)$. As a consequence, it is possible to show that after properly scaling and shifting g_t so that it becomes the distribution of a zero-mean, unit-variance random variable, the resulting function converges to the p.d.f. of a standard normal distribution as $t \rightarrow \infty$.

Interestingly, for our numerical purposes the convergence happens fast enough that g_τ in (4) can be well approximated by a normal distribution. The fit is better for higher values of T and σ , but even for T as low as 0.02, as the approximation is good in our calibration. We choose our state variable in the recursive formulation of the problem to be the distribution of prices gaps immediately after the payment of the fixed cost because, at that instant, the time elapsed since the last adjustment episode is maximal and so the distribution of discrepancies is as close as possible to a Gaussian curve.

If we approximate our infinite dimensional state variable by a normal distribution, we are left with a two-dimensional problem, since normally distributed variables may be characterized by two parameters. In fact, since in steady state the mean of the price gap distribution is zero, our problem becomes unidimensional. Finally, to show the goodness of fit, Figure B.1 compares the steady-state distribution that would arise from following the optimal policy for calibrated parameter values, given by equations (7) and (8), with the corresponding normal approximation. It also shows the difference between the c.d.f. of the approximating normal $\Phi(x)$ and the one obtained numerically $G(x)$. The right-hand panel of Figure B.1 shows that for no interval $[a, b]$, the approximation predicts a mass of price gaps that is more than 0.0234 away from the true value, which happens for the interval $[-0.0890, 0.0915]$.




Figure B.1: Steady-state distribution and normal approximation. $\sigma = 0.3834$, $\tau^* = 0.0207$. The value of σ shown here is the volatility of the Brownian motion, not the standard deviation of the normal approximation.

C Interpretation of price-adjustment frictions

Suppose we have an economy in which profit-maximizing prices are unobservable and economies of scope come not from the price-adjustment technology, but from the information-acquiring process. More precisely, imagine that K now is a cost whose payment is required to observe frictionless optimal prices for all products simultaneously, while c remains a standard, product-specific menu cost. One could ask how this economy would differ from the one we have studied so far.

First, observe that, since the profit-maximizing price of a given product is a martingale (i.e., it satisfies $\mathbb{E}_t p_{i,t+h}^* = p_{i,t}^*$ for $h > 0$), no adjustments will be made without new information. Intuitively, a firm would not adjust the price of a product if the frictionless optimal price does not change in expectation. Adjustments are only made when the fixed cost for acquiring information is paid, as in [Alvarez, Lippi, and Paciello \(2011\)](#) and, in the limiting case of infinitely many products, in [Alvarez and Lippi \(2014\)](#). However, since the firm sells a continuum of products, the distribution of price gaps is perfectly predictable and given by the KFE (3), as long as there are no aggregate shocks. In other words, the law of large numbers makes the distribution of price gaps perfectly predictable, even though each individual price gap is unobserved. Therefore, the steady-state optimal policy and stationary distributions are the same as in our baseline menu-cost framework.

The menu cost and the informational friction economies would differ, however, in response to an aggregate shock. In the menu-cost economy, when the shock hits at $t = 0$, firms are allowed to instantaneously recalculate the optimal policy, whereas in the costly information economy, firms would only learn about the shock after collecting information. Nonetheless, we could follow the same steps for proving 9 and 10 to obtain exactly analogous results to this case. As a consequence, even though the two economies could respond differently to an aggregate shock, there is no difference up to the first order.

This result offers an important insight into the nature of the frictions in product markets and their implications for monetary non-neutrality. When monetary shocks are sufficiently small, the responses of aggregate price and output of an industry dominated by large multi-product retailers are not likely to be sensitive to whether price stickiness is due to physical menu cost or the cost of acquiring and processing information.

D A model with random free adjustments

We consider here an extension of the partial synchronization model (with Brownian shocks) in which free adjustment opportunities arrive independently for each product at rate λ . Since free adjustments arrive for all prices, regardless of the size of the price gap, it generates a mass of small price changes. We consequently choose the fraction of small price adjustments as the extra empirical moment we need to calibrate the additional parameter λ .

Tables D.1 and D.2 show moments and calibrated parameter values. A very small λ is enough to match the fraction of small adjustments; this is because, in our sample, they are not too frequent. The calibrated value $\lambda = 0.47$ corresponds roughly to one free adjustment per product every two years, on average. Consequently, the distribution of price changes, shown in figure D.1 is not too different from the original one, except for the mass of small adjustments. Thus, impulse response functions for monetary shocks, shown in figure D.2, are almost indistinguishable.

Moment	Data	Partial sync.
Targeted moments		
Daily fraction of price changes	0.0087	0.0087
Average $ \Delta p $	0.208	0.208
Fisher-Konieczny index	0.236	0.236
Fraction of small adjustments	0.06	0.06
Additional moments		
Standard deviation of Δp	0.248	0.218
Kurtosis of Δp	3.53	1.22

Table D.1: Moments from data and model with random free adjustments.

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	Partial sync.
σ	0.390
K	5.04e-05
c	0.0028
λ	0.477

Table D.2: Calibrated parameter values for model with random free adjustments.

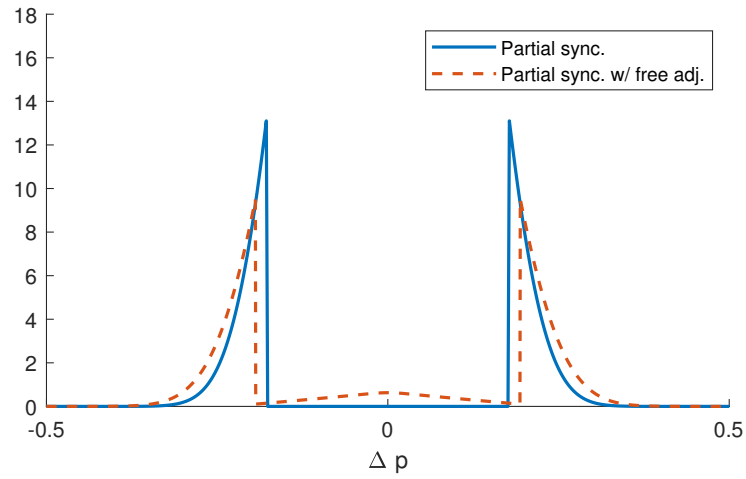


Figure D.1: Price-change distributions.

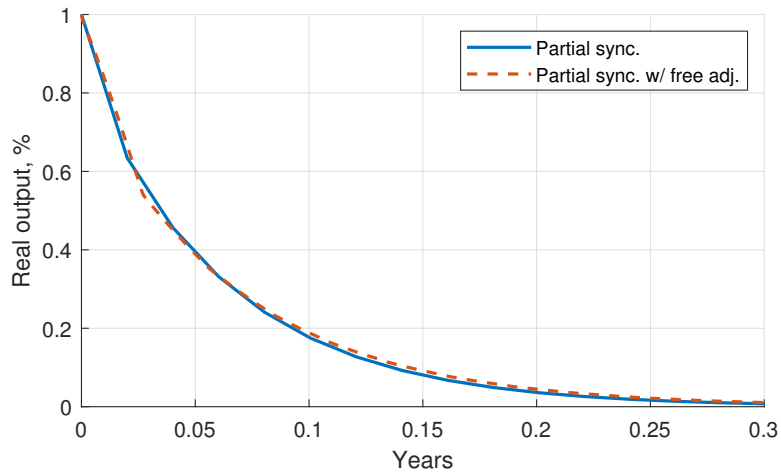


Figure D.2: Impulse response functions.

E Data Appendix

E.1 Inflation and frequency of price changes

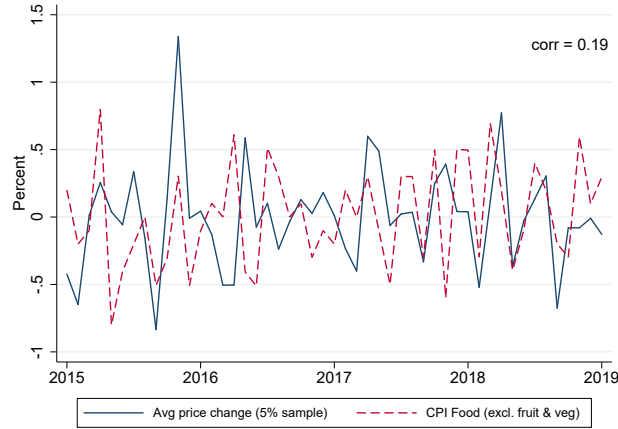


Figure E.1: Food inflation in Israel.

Note: We compute the average price changes for each store at monthly frequency. The figure provides the weighted mean fraction across stores and compares it with inflation for official CPI for food products (excluding fruits and vegetables). Weights are the average number of products in a store per day.

	Inflation, %	Fraction of price changes	Fraction of price increases	Fraction of price decreases	Absolute size of price changes, %	Size of price increases, %	Absolute size of price decreases, %
<u>A. Daily data</u>							
N	1498	1498	1498	1498	1480	1472	1476
mean	-0.005	0.010	0.005	0.005	20.9	20.1	21.5
sd	0.202	0.020	0.011	0.011	3.4	4.0	3.8
<u>B. Weekly data</u>							
N	217	217	217	217	217	217	217
mean	-0.020	0.043	0.022	0.021	20.0	19.1	20.7
sd	0.400	0.038	0.022	0.020	2.2	2.5	2.3
<u>C. Monthly data</u>							
N	53	53	53	53	53	53	53
mean	-0.056	0.112	0.058	0.054	19.9	18.9	20.9
sd	0.541	0.042	0.028	0.023	1.6	1.8	1.6

Table E.1: Summary statistics for regular price changes.

Note: Entries include the number of observations and mean and standard deviation for time series at daily, weekly and monthly frequencies. For each day (week or month), we compute the statistics for each column by pooling price-change observations across stores. We then take weighted means and standard deviations across days (weeks or months). Weights are the average number of products in a store per day.

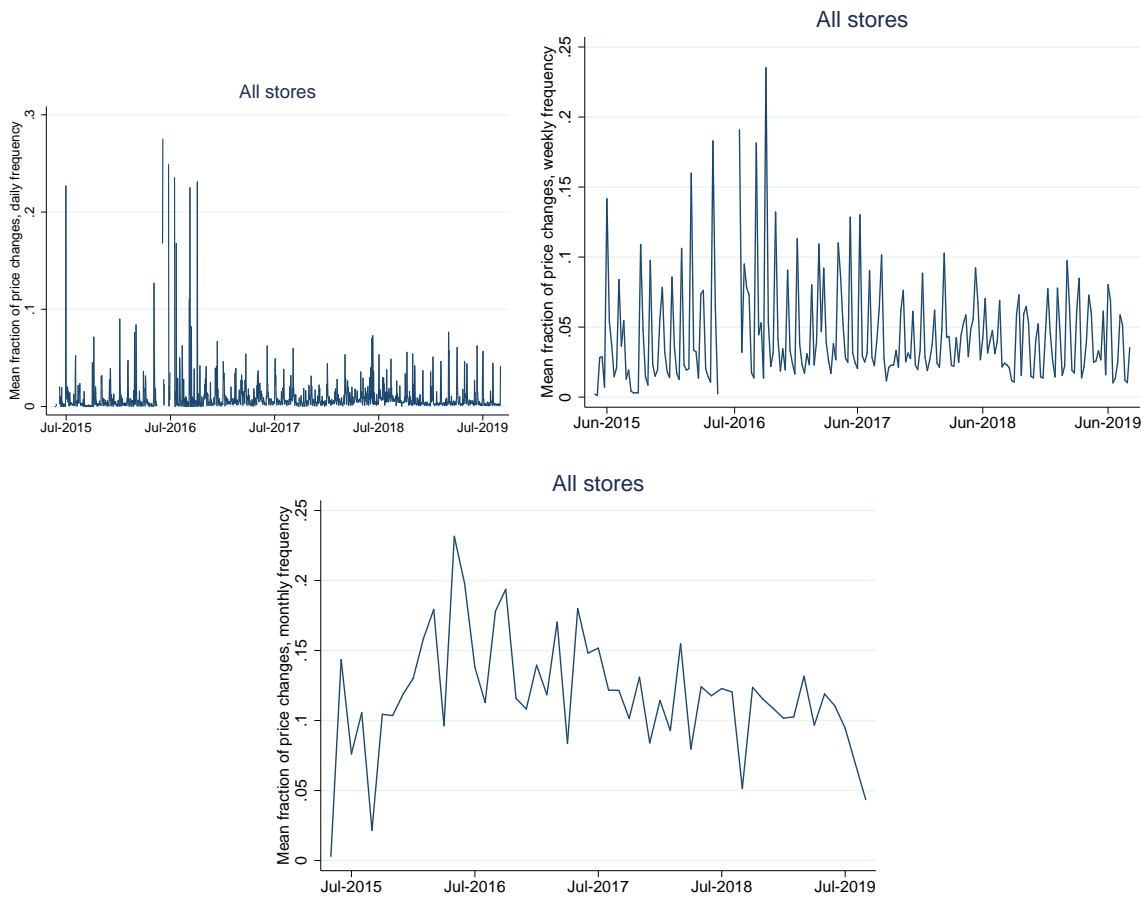


Figure E.2: Fraction of regular prices changes, all stores.

Note: We compute the fraction of price changes for each store at daily, weekly and monthly frequency. The figures provide the weighted mean fraction across stores for daily frequency (top-left panel), weekly frequency (top-right panel) and monthly frequency (bottom panel). Weights are the average number of products in a store per day.

E.2 Synchronization of price changes across products

Figure E.3 provides distribution of the number of products across stores, average fraction of adjusting prices across stores, and two measures of synchronization across stores. The number of products sold in a given store is large: on average 7,217 products are sold on a given day, 1,311 (31,847) 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.

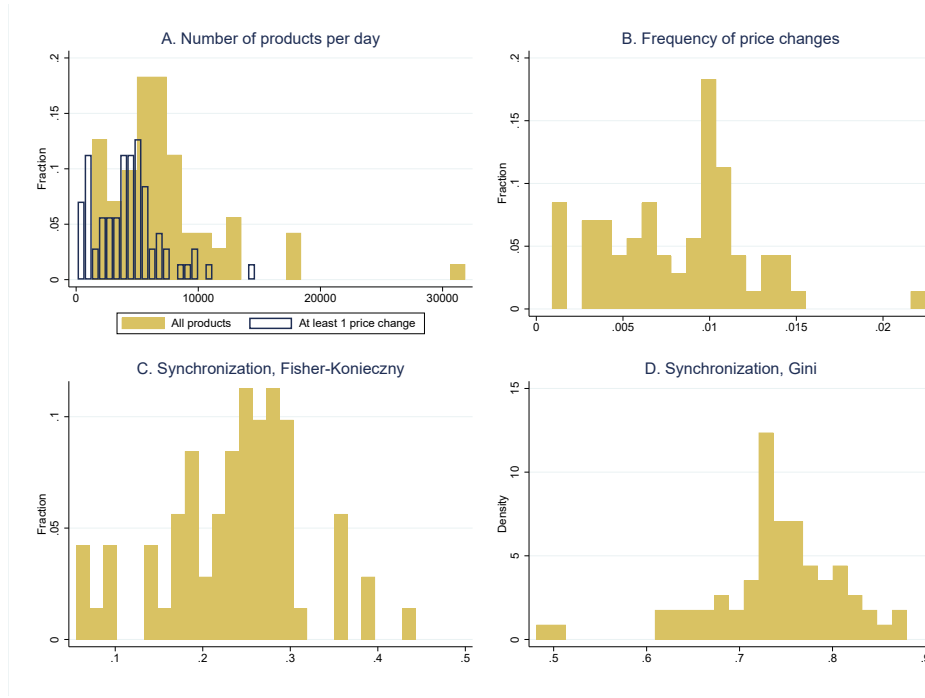


Figure E.3: Store-level number of products, frequency and synchronization of price changes.

Note: Panel A provides histograms for the average number of price observations per store per day (all products and products with at least one price change over the sample period). Panel B shows the histogram of the fraction of price changes across stores. Panel C gives the histogram of the Fisher-Konieczny synchronization index across stores, and Panel D provides the histogram of the Gini index values across stores (Gini index is defined in Appendix E.3).

Figure E.4 contains two scatter plots. Each point on a scatter plot represents a pair of two frequencies for a product (or product category) in a store: frequency of adjustment on off-peak days (x -axis) and a frequency of adjustments on peak days (y -axis). These scatter plots are convenient for gauging the balance between adjustment frequencies on peak versus off-peak days at a product (or product category) level. Products for which these two frequencies are equal would show as a dot on the 45° line. Products for which adjustments fall predominantly on peak days are above the 45° line. If a store sampled prices completely randomly across its products each day, product-level price-change frequencies would be the same for the subsets of peak and off-peak days.

The left plot is for Shufersal stores only, where each point of the scatter plot is for products in the same product category (50 categories). The right plot is for all individual products in the

dataset, with at least 10 price-change observations. Both plots show that the mass of dots is spread above and below the 45° line: there are products that adjust prices mostly during peak days, but there are also products that adjust mostly during off-peak days. The share of products whose adjustments are more likely during peak days than during off-peak days is 0.82 for Shufersal and 0.76 for all stores. Similarly, the share of products whose adjustments are *twice* as likely during peak days than during off-peak days is 0.65 for Shufersal and 0.52 for all stores. Conversely, the share of products whose adjustments are twice as likely during off-peaks days than during peak days is 0.12 for Shufersal and 0.07 for all stores. Hence, on peak re-pricing days, stores increase (decrease) the likelihood of changing prices for some of their products relative to other days.

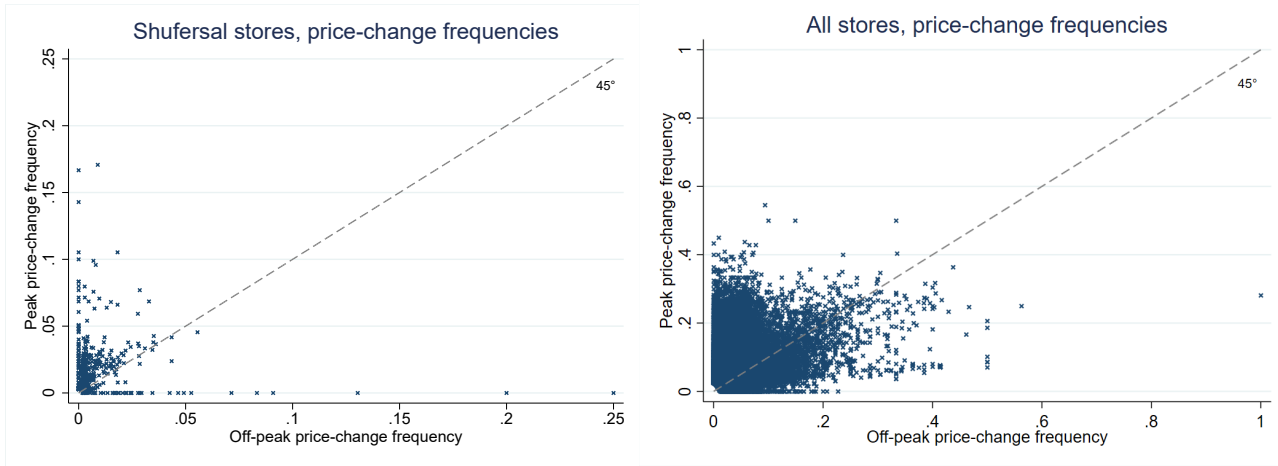


Figure E.4: Fraction of price changes during peak and off-peak days, by product.

E.3 Gini index of synchronization of price changes

To ensure that our results are not driven by any particular way of measuring synchronization, we analyze an alternative synchronization measure. The second synchronization index is akin to the Gini inequality index. It is based on the Lorenz curve that depicts the distribution of repricing activity by plotting the percentile of days by the number of price changes on the horizontal axis and cumulative fraction of price changes on the vertical axis. Figure E.5 shows that re-pricing activity is very unequal across days, with around 6% of days accounting for half of all price changes. The Gini statistic is the size of the area between the Lorenz curve and the 45° line divided by a half. Table E.2 shows a high degree of synchronization (i.e., high inequality of repricing across days), 0.747 on average, and which varies substantially across stores (Panel D, Figure E.3).

The two measures capture different synchronization patterns. For example, the FK index is sensitive to the average frequency of price changes: more frequent adjustments imply a higher FK index value. The Gini index does not directly depend on the average frequency. Instead, it is a better measure of the spread of re-pricing activity across days in a given store. Nonetheless, Table E.2 shows that two synchronization measures are broadly in line with each other across different dimensions of the data.

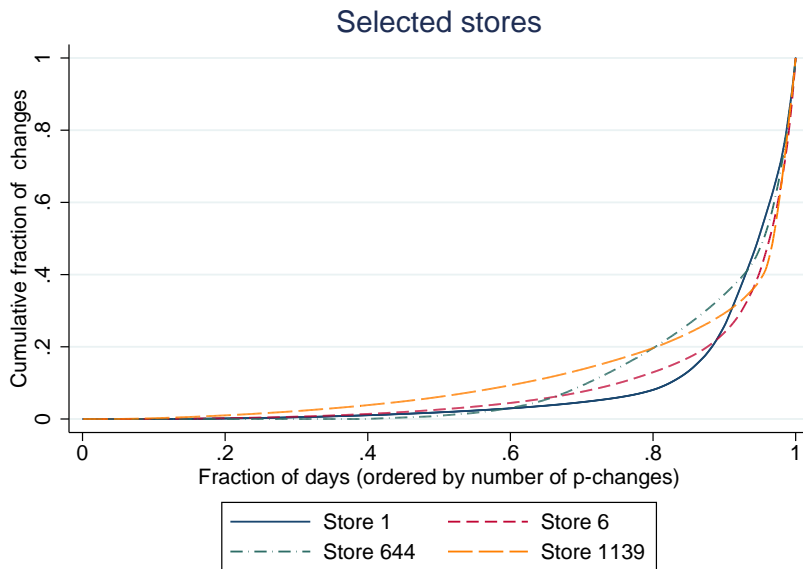


Figure E.5: Lorenz curve for inequality of re-pricing activity across days, selected stores.

Note: Figure E.5 shows the Lorenz curve (cumulative distribution of price changes across days) for four selected stores.

Sample	Frequency of observations	Synchronization statistic	
		Fisher-Konieczny	Gini
<u>A. By store</u>	daily	0.236	0.747
	weekly	0.225	0.534
	monthly	0.187	0.315
Price increases	daily	0.218	0.784
Price decreases	daily	0.201	0.767
Same product category*	daily	0.289	0.770
Flexible price goods	daily	0.376	0.747
Larger stores	daily	0.253	0.752
Offline stores**	daily	0.280	0.751
Online stores**	daily	0.343	0.835
<u>B. By chain</u>	daily	0.220	0.728
<u>C. All observations</u>	daily	0.142	0.573
<u>D. IRI data</u>	weekly	0.098	0.139
	monthly	0.133	0.138

Table E.2: Comparison of Fisher-Konieczny and Gini synchronization indices.

Note: The table provides weighted means for synchronization indices 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 is measured by Fisher and Konieczny (2000) index defined in Section 2.3 and Gini index defined in this Appendix. Same product category: daily statistics computed for subsets of products belonging to the same product category (* Shufersal stores only). Flexible price goods*: statistics computed for subsets of products with the daily fraction of price changes in the top quartile in the store. Larger stores*: daily statistics computed for subsets of products in stores larger than the median store (by the number of products per day). ** We select only brick-and-mortar stores belonging to the same chains as the four online stores in our dataset.*

E.4 Price discounts, final prices and price synchronization

	Inflation, %	Fraction of price changes	Fraction of price increases	Fraction of price decreases	Absolute size of price changes, %	Size of price increases, %	Absolute size of price decreases, %
A. Regular prices							
N	978	978	978	978	920	896	890
mean	0.001	0.007	0.004	0.003	18.9	18.2	19.8
p50	0.000	0.002	0.001	0.001	19.3	18.4	20.3
sd	0.158	0.023	0.012	0.012	6.5	7.5	6.0
iqr	0.014	0.003	0.001	0.002	7.5	8.9	8.0
B. Final prices							
N	978	978	978	978	976	975	964
mean	-0.003	0.024	0.012	0.012	23.9	22.8	24.7
p50	-0.022	0.009	0.003	0.005	23.8	22.9	24.1
sd	0.638	0.052	0.030	0.027	4.4	5.3	5.7
iqr	0.123	0.011	0.004	0.007	4.2	5.8	5.1

Table E.3: Final versus regular price changes for Shufersal.

Note: Entries include the number of observations and mean and standard deviation for time series at daily frequencies. The data cover 10 Shufersal retail stores from January 2016 until mid-2019. For each day, we compute the statistic for each column by pooling price-change observations across stores. We then take means and standard deviations across days. Panel A (B) provides statistics for regular (final) price changes.

Store	Sub-chain	Fraction of discounts	Mean abs size of discounts, %	Duration of discounts, days	
				Complete spells	All spells
6	Premium stores	0.288	26.4	56	59
23	Premium stores	0.289	26.3	50	60
39	Premium stores	0.243	26.3	51	60
168	Premium stores	0.295	26.4	53	62
55	Discount stores	0.302	22.4	41	47
113	Discount stores	0.305	22.2	46	55
188	Discount stores	0.316	22.4	46	52
217	Express (convenience)	0.166	26.4	49	48
296	Express (convenience)	0.158	26.2	49	49
606	Ultra-discount stores	0.218	22.0	56	61
Mean (stores)		0.258	24.7	50	55
Mean (pooled)		0.259	24.4	49	56

Table E.4: Price discounts at Shufersal.

Note: The data cover 10 Shufersal retail stores from January 2016 until mid-2019. Shaded areas outline different store types.

Store	Sub-chain	Fisher-Konieczny index		Gini index	
		Regular prices	Final prices	Regular prices	Final prices
6	Premium stores	0.298	0.380	0.810	0.717
23	Premium stores	0.282	0.365	0.798	0.687
39	Premium stores	0.304	0.364	0.838	0.757
168	Premium stores	0.319	0.380	0.826	0.720
55	Discount stores	0.319	0.380	0.827	0.706
113	Discount stores	0.276	0.369	0.814	0.698
188	Discount stores	0.305	0.368	0.827	0.688
217	Express (convenience)	0.177	0.290	0.694	0.630
296	Express (convenience)	0.212	0.298	0.790	0.698
606	Ultra-discount stores	0.124	0.353	0.751	0.778
Mean (stores)		0.262	0.355	0.797	0.708
Mean (pooled)		0.274	0.363	0.798	0.709

Table E.5: Synchronization of regular and final price changes for Shufersal.

Note: The data cover 10 Shufersal retail stores from January 2016 until mid-2019. Column entries compare synchronization statistics (Fisher-Konieczny and Gini) for daily final and regular price changes. Shaded areas outline different store types.

Store	Sub-chain	Peaks (Regular)		Off-peaks (Regular)		Peaks (Final)		Off-peaks (Final)		All	
		# days	Freq	# days	Freq	# days	Freq	# days	Freq	# days	Freq
6	Premium stores	22	0.156	898	0.004	42	0.265	878	0.013	920	0.024
23	Premium stores	25	0.135	880	0.004	44	0.250	861	0.013	905	0.024
39	Premium stores	18	0.155	888	0.003	38	0.226	868	0.010	906	0.019
168	Premium stores	19	0.177	897	0.004	42	0.257	874	0.012	916	0.024
55	Discount stores	19	0.179	890	0.004	51	0.271	858	0.016	909	0.030
113	Discount stores	30	0.109	886	0.004	43	0.267	873	0.013	916	0.025
188	Discount stores	22	0.155	900	0.004	50	0.265	872	0.015	922	0.029
217	Express (convenience)	51	0.033	747	0.002	50	0.137	748	0.009	798	0.017
296	Express (convenience)	31	0.070	903	0.002	46	0.168	888	0.009	934	0.016
606	Ultra-discount stores	67	0.012	894	0.001	31	0.217	930	0.007	961	0.014
Mean (stores)		30	0.118	878	0.003	44	0.232	865	0.012	909	0.022
Mean (pooled)			0.091		0.003		0.232		0.012		0.022

Table E.6: Re-pricing peaks for final and regular prices for Shufersal.

Note: The data cover 10 Shufersal retail stores from January 2016 until mid-2019. Column entries compare mean frequency of daily final price changes with regular price changes for peak and off-peak days. Peaks are the days with the highest fraction of price changes that jointly account for half of all price changes in the store. Shaded areas outline different store types.

E.5 Frequency of price changes and holidays in Israel

Table E.7 tabulates the frequencies of joint occurrences of Jewish holidays and the peaks of repricing activity in a store. On average, only 6.5% of peak days in a store are holidays, and in turn, only 4.1% of holidays are peaks in a store.

Store	Fraction of daily observations			
	Not-Peak / Not-Holiday	Peak / Not-Holiday	Not-Peak / Holiday	Peak / Holiday
By store	85.07%	5.52%	9.02%	0.38%
By chain	85.92%	4.78%	8.98%	0.32%
All observations	85.89%	4.88%	8.94%	0.28%

Table E.7: Peaks and holidays.

Note: Peaks are the days with the highest fraction of regular price changes that jointly account for half of all price changes in the store. For each store we compute the fraction of days that are peaks (not peaks) and fall (do not fall) on a holiday. Entries are weighted means of these fractions across stores (first row), chains (second row), or unweighted means (last row). Weights are the average number of products in a store per day.

E.6 Frequency of price changes and weekly/monthly fixed effects

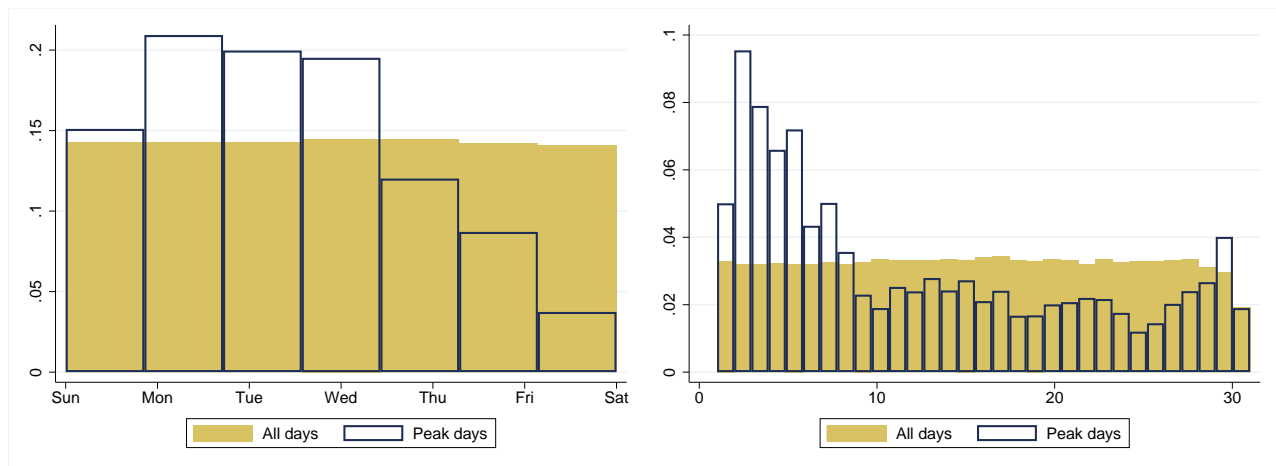


Figure E.6: Fraction of observations, by day of the week and by day of the month.

Note: Figure E.6 shows the (weighted) fraction of price observations, by day of the week (left plot) and by day of the month (right plot), for all days and only peak days. Weights are the average number of products in a store per day. Peaks are the days with the highest fraction of regular price changes that jointly account for half of all price changes in the store. Shaded (empty) bars correspond to changes for all days (for the subset of peak days).




Figure E.7: Fraction of price changes, by day of the week and by day of the month.

Note: Figure E.7 shows the (weighted) fraction of price changes, by day of the week (left plot) and by day of the month (right plot), for all days and only peak days. Weights are the average number of products in a store per day. Peaks are the days with the highest fraction of regular price changes that jointly account for half of all price changes in the store. Shaded (empty) bars correspond to changes for all days (for the subset of peak days).

E.7 Distribution of standardized price changes

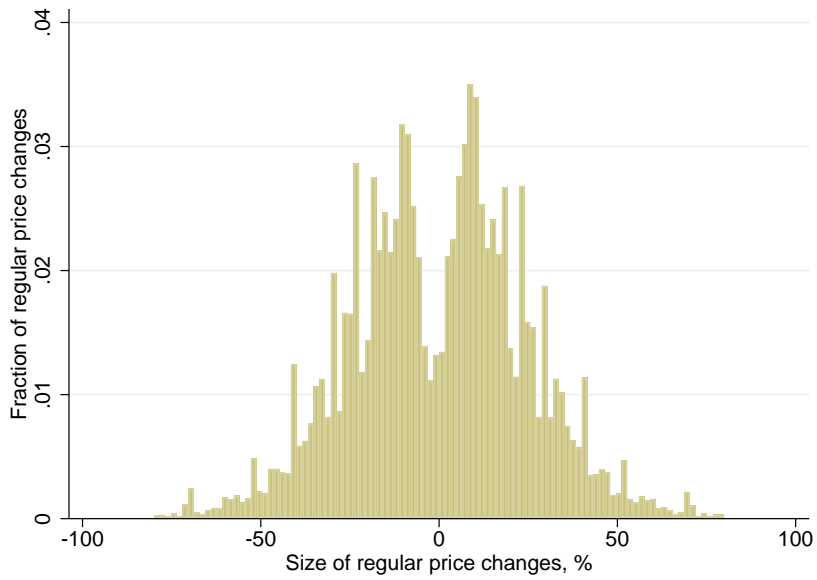


Figure E.8: Distribution of standardized price changes, Israel.

Note: Figure E.8 shows the unweighted distribution of standardized non-zero price changes across stores in Israeli data. Define “standardized” log price changes for product i , store s and time t , z_{ist} , 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: $z_{ist} \equiv \frac{p_{ist} - \mu_s}{\sigma_s} \Big|_{p_{ist} \neq 0}$.

	All prices, by store			All prices, by store and price point deciles			All prices, by store and frequency of price change deciles			By store and product category			Bhattarai and Schoenle	Midrigan AC Nielsen	Midrigan Dominick's	
	Daily	Weekly	Monthly	Daily	Weekly	Monthly	Daily	Weekly	Monthly	Daily	Weekly	Monthly	Monthly	Weekly	Weekly	
A. Israel																
$\text{std}(z_{ist}) / \text{mean}(z_{ist})$	0.68	0.67	0.69	0.69	0.67	0.69	0.68	0.67	0.70	0.67	0.67	0.72				
$\text{kurtosis}(z_{ist})$	3.5	3.3	3.4	3.6	3.4	3.4	3.8	3.6	3.7	4.1	4.1	4.0				
fraction $ z_{ist} < 0.5 * \text{mean}(z_{ist})$	0.14	0.27	0.28	0.15	0.26	0.28	0.14	0.27	0.28	0.15	0.26	0.28				
fraction $ z_{ist} < 0.25 * \text{mean}(z_{ist})$	0.06	0.10	0.11	0.07	0.11	0.11	0.05	0.10	0.11	0.08	0.13	0.15				
B. United States																
$\text{std}(z_{ist}) / \text{mean}(z_{ist})$		0.85	0.91		0.84	0.89		0.85	0.90		0.83	0.89	1.55	0.72	0.81	
$\text{kurtosis}(z_{ist})$		3.4	4.7		3.6	5.1		3.4	4.9		4.0	5.1	17	3.6	4.5	
fraction $ z_{ist} < 0.5 * \text{mean}(z_{ist})$		0.37	0.38		0.34	0.35		0.37	0.37		0.34	0.36	0.50	0.25	0.31	
fraction $ z_{ist} < 0.25 * \text{mean}(z_{ist})$		0.17	0.20		0.18	0.19		0.17	0.20		0.16	0.19	0.38	0.10	0.14	

Table E.8: Standardized price changes, Israel and the United States.

Note: We then compute four statistics characterizing the shape of the distribution of standardized price changes that are used in Table II in [Alvarez and Lippi \(2014\)](#). Define “standardized” log price changes for product i , store s and time t , z_{ist} , by subtracting from log price $p_{i,ist}$ the mean μ_s of non-zero log price and dividing by their standard deviation σ_s changes at a store level: $z_{ist} \equiv \frac{p_{i,ist} - \mu_s}{\sigma_s} \Big|_{p_{i,ist} \neq 0}$. To control for heterogeneity of price changes within a store, we also use three alternative definitions of standardized changes, where for prices within a store we control for (i) differences in price levels, (ii) product’s frequency of price changes, or (iii) product category. In cases (i) and (ii), we sort individual products within each store by their price level (or frequency). We define standardized price changes by store and decile of price level (or frequency). In case (iii), for Israel we use the data for 10 Shufersal retail stores from January 2016 until mid-2019, which includes information on product categories. In IRI data we use 31 broad product categories. In case (iii), we define standardized price changes by store and product category. 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 \(2014\)](#), who note, “For the Bhattarai and Schoenle (2011) data, the number of products n is the mean of the categories considered based on the information in Table 1, the ratio $\text{Std}(|\Delta p_i|) / E(|\Delta p_i|)$ is from Table 2 (firm-based), the fraction of $|\Delta p_i|$ that is small is from Table 14, and the kurtosis is from Figure 7. The data from [Midrigan \(2009\)](#) are taken from the distribution of standardized prices in Table 2b.”

E.8 Spectral analysis of store-level price changes

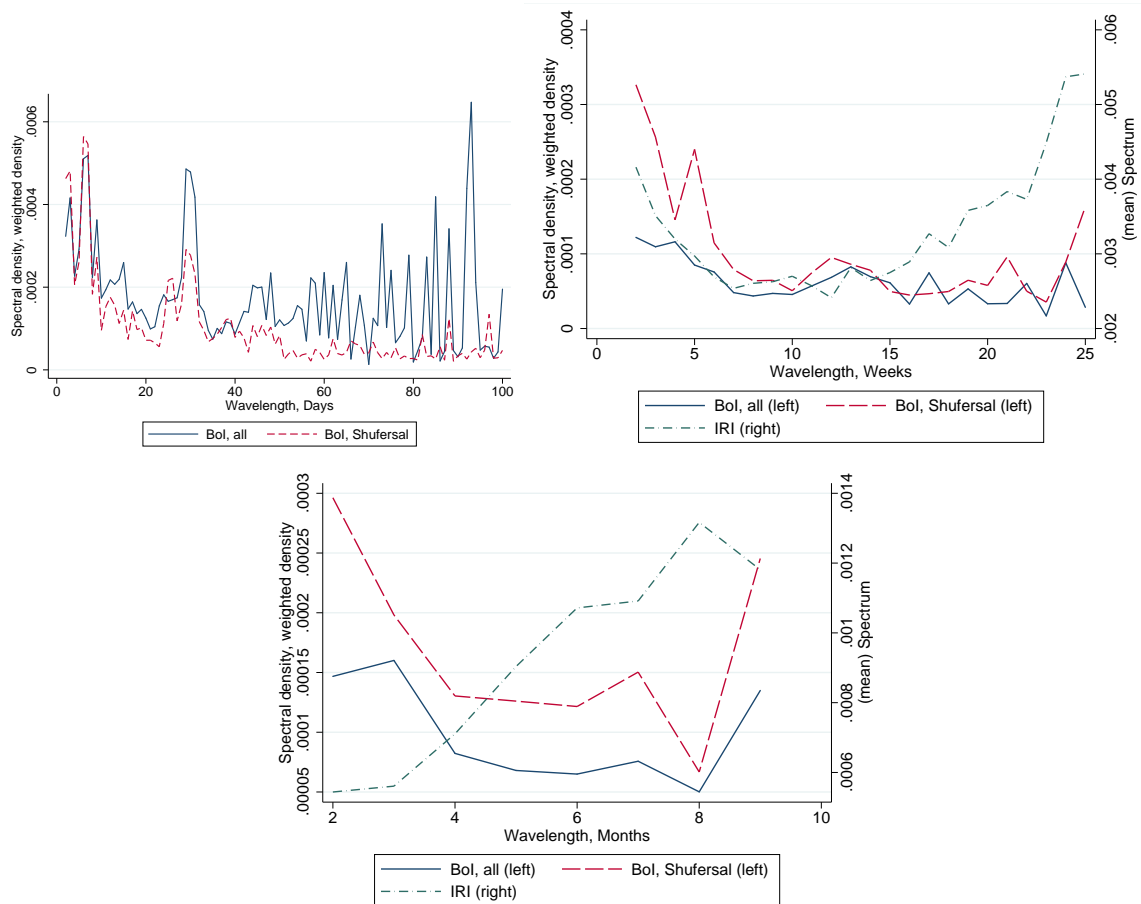


Figure E.9: Spectral density of regular prices changes.

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 the Israeli data (“BoI, Shufersal”); and regular price changes in IRI data. For each store in the sample, we compute spectral density of its daily (weekly, monthly) 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 (top left), weekly (top right), and monthly (bottom).