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

Marco Bonomo†
Carlos Carvalho‡
Oleksiy Kryvtsov§
Sigal Ribon¶
Rodolfo Rigato‖

September 2019

Abstract

Standard theories of price adjustment are based on a problem of single-product firm pricing, and therefore they may not be well-suited for studying a more realistic case of firms setting prices for thousands of products. To guide new theory, we study evidence for large multi-product food retailers in Israel. We find that retail stores undertake a majority of their regular price changes during occasional "peak" days, roughly once or twice a month; and on a peak day, stores reprice around 10% of their products. We develop a general equilibrium model of price-setting firms with a continuum of products to assess implications of this evidence for inflation dynamics. In the model, the economies of scope in price adjustment give rise to an endogenous trade-off between adjustment of many prices at a time ("synchronization") and adjustment of misaligned prices ("selection"). By limiting the scope for selection, synchronization of price changes can reduce inflation response to monetary disturbances. The calibrated model, despite matching partial synchronization of price changes in the data, generates only a weak monetary non-neutrality, similar in magnitude to non-neutrality in standard menu cost models. Hence, partial synchronization of price adjustments does not materially deter multi-product firms from responding to monetary disturbances.

JEL classification codes: E31, E51.

Keywords: Inflation, Prices, Multi-product pricing, Monetary Non-neutrality.

*The views expressed herein are those of the authors and not necessarily those of the Central Bank of Brazil, the Bank of Canada, or the Bank of Israel. We thank the IT department at the Bank of Israel for their assistance with the construction of the data. We thank seminar participants at the Bank of Canada, 2019 CEBRA Annual Meeting for their comments and discussions. Minnie Cui provided superb research assistance.

†Insper.Email: *bonomo@insper.edu.br.*

‡Central Bank of Brazil, and Pontifícia Universidade Católica do Rio de Janeiro. Email: *cvianac@econ.puc-rio.br.*

§Bank of Canada, 234 Wellington Street, Ottawa, Ontario K1A 0G9, Canada. Email: *okryvtsov@bankofcanada.ca.*

¶Bank of Israel. Email: *sigal.ribon@boi.org.il* .

‖Harvard.Email: *dinisrigato@g.harvard.edu.*

1 Introduction

The literature about the relation between micro price behaviour and macroeconomic effects has evolved recently due to both the availability of new data sources (e.g., micro data underlying CPIs, scanner data, scrapped online prices) and new analytical models. This literature, however, is short on both empirical studies and macroeconomic theories of behaviours of multiple product firms.[1] Available data do not provide comprehensive information on both the frequency and scope of store-level price changes in national retail industry. Furthermore, standard theories of price adjustment are based on a problem of single-product firm pricing, and therefore they may be ill-suited for studying a more realistic case of firms setting prices for thousands of products. In this paper, we address both of these challenges.

We study new evidence on prices set by large retailers in Israel, and propose a general equilibrium model of multi-product price-setting to account for this evidence. The pattern that emerges in the data is that prices of multi-product firms are partially synchronized due to occasional peaks in firms' repricing activity. To generate this pattern, we propose a theory of multi-product monopolistic firms who adjust prices subject to a price-adjustment technology with endogenous degree of economies of scope in price-setting. In the model, the firm incurs the fixed cost c per price change, and a common cost K for any number of price changes. This technology nests two extreme cases, closely related to models studied in the literature. 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). 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. This case is similar to models with maximal economies of scope in Midrigan (2011) and Alvarez and Lippi (2014), although in those models firms must adjust all prices upon paying the cost K. We calibrate the model to match the key price-setting statistics from the micro data and derive implications for monetary non-neutrality, comparing it with the two extreme nested models.

In our model, synchronization of price adjustments is closely related to the selection effect explored in Golosov and Lucas (2007).[2] If there is a large degree of synchronization, i.e., many prices are adjusted at the same time, firms may bunch adjustments of prices that are away from the adjustment margin, weakening the selection effect. In our data, firms show a considerable degree of price synchronization. The magnitude of synchronization, however, is not sufficient to generate persistent effects of monetary shocks. As a result, aggregate real responses to nominal disturbances in our setting are similar to those in single product models.

Our empirical findings are based on price data from large multi-product food retailers in Israel.

[1] Bhattarai and Schoenle (2014) study pricing behaviour of U.S. producers using US PPI data. They document that that 98.55% of price observations are from multi-product firms. Midrigan (2011), Alvarez and Lippi (2014) develop macroeconomic models of multi-product price setting.

[2] Golosov and Lucas (2007) study selection effects 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.

Thanks to a law enacted in 2014, large retailers are required to publish on their Internet sites daily information for all products sold. To manage computational constraints, the data used in this paper contains information for top 5% of stores (by the number of observations) for each retailer from May 22, 2015 until September 1, 2019. For 71 stores in this sample, we have information about the "base" or regular prices for all individual products on a daily basis. 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 discounts, a third-product-free discount, or two products for 10 NIS (Israeli New Shekel). Based on the available regular price and the discount code, the Bank of Israel constructed the final price for 10 stores owned by Shufersal, the largest food retailer in Israel, with 350 stores servicing about 35 percent of the food retail-chain market.

During the sample period, Israeli headline inflation and food price inflation are approximately zero, and price increases are as likely as price decreases. Regular price changes are large in magnitude, approximately 20%, and the implied duration of price spells is 3.5 months, when calculated from daily observations. Each store in our sample sells a large number of products each day, 7,217 on average. We doubt that the model of a single product firm (or even of firms with a small number of products) is useful for studying price setting of large retailers. Theory of a price-setting firm with an infinite number of products seems more promising.

We exploit comprehensive product coverage in our data to study the synchronization of price adjustments across products in a store. Figure 1 depicts the daily fraction of price adjustments across all stores in the sample, and for four selected stores from different chains. It is apparent that there are recurrent peaks in the fraction of price changes, with values lower than one. These peaks in price adjustments are not due to calendar events. Thus, the data shows a pattern that does not coincide with either the more staggering pattern that one would expect from a single product menu cost economy or with the perfect synchronization generated by Alvarez and Lippi (2014) model, where the payment of a single menu cost entitles the firm to adjust all of its products.

We characterize the degree of synchronization in the data through two indices: the one proposed by Fisher and Konieczny (2000) and another akin to Gini inequality index. The Fisher-Konieczny index is based on the variation of the fraction of adjustment around its average over time. It ranges from 0, corresponding to the absence of synchronization, to 1, representing full synchronization. On average, the Fisher-Konieczny index is 0.236 for daily data at a store level. The second synchronization measure, which we will refer to as Gini synchronization index, is based on the Lorenz curve relating the percentile of days to the fraction of price changes accounted by those days, arranged according to the fraction in ascending order. Again, this index should be zero, in the case of no synchronization and 1 in the case of perfect synchronization. The average Gini synchronization index among stores is 0.747.

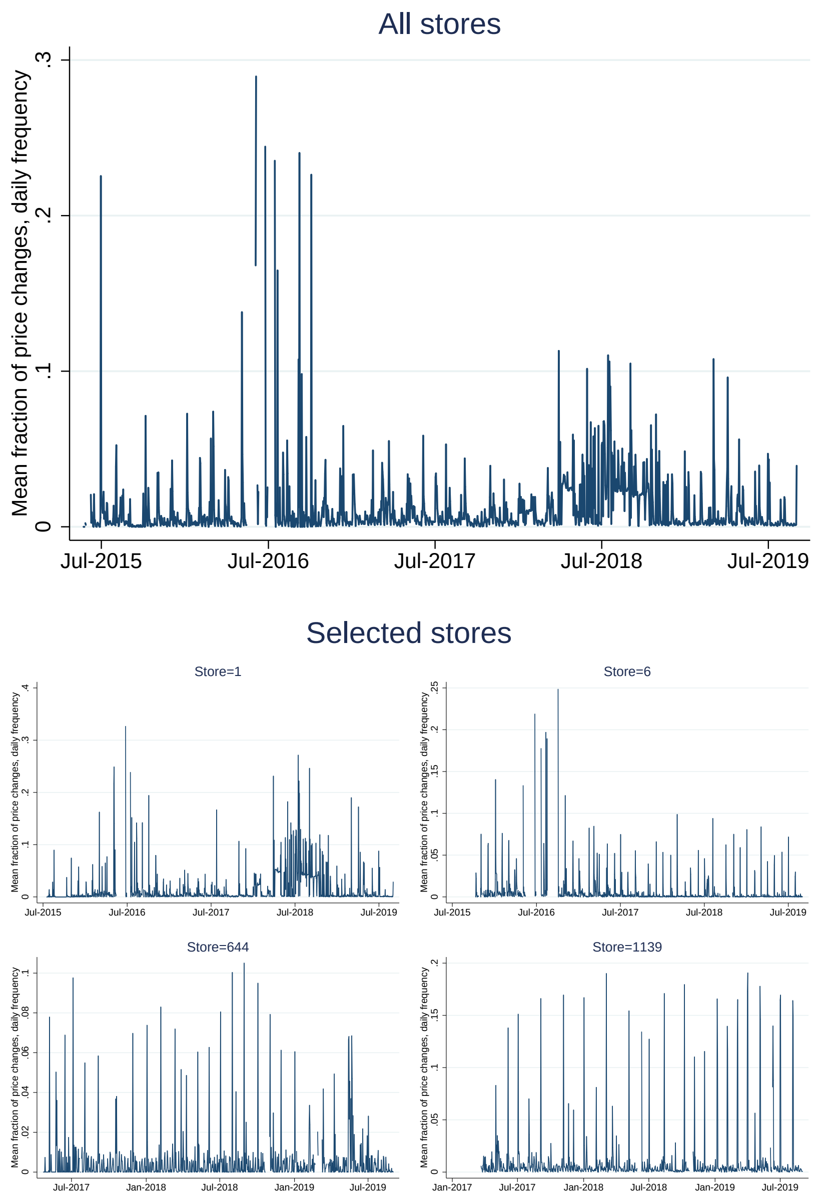

Figure 1: Daily fraction of prices changes.

Note: We compute the fraction of price changes for each store and day. The table provides the fraction for four selected stores (bottom figure) and the weighted mean across stores in the unbalanced panel (top figure). Weights are the average number of products in a store per day.

We develop a continuous-time price-setting model of multi-product firms capable of generating the partial synchronization aspect found in the data. In the model, each firm sells a continuum of different goods and faces two types of costs when changing price: a fixed cost K incurred when at least one price adjustment is made, and an additional cost c paid for each individual price adjustment. For tractability, we assume that the process for the frictionless optimal price of each product has no drift and is subject to idiosyncratic shocks, which are independent across goods. Due to the presence of the common cost K, there are intervals of time where no product has its price adjusted. 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 partial synchronization pattern we observe 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 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 calibrate 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 Fisher-Konieczny synchronization index (FK). Our model is able to reproduce them perfectly with the appropriate choices of the variance of the idiosyncratic shock, the common cost K and individual adjustment cost c. We also calibrate the restricted versions corresponding to no economies of scope ($K = 0$) and maximal economies of scope ($c = 0$) in price adjustment to match the frequency and the size of price adjustments. Those models do not have the flexibility to match the degree of synchronization in the data (FK=0.236), since the simple menu cost model displays no synchronization (FK=0) and the Alvarez-Lippi-Midrigan model generates full synchronization (FK=1). We compute the distribution of price changes for all the three models. Our model features a two-mode price-change distribution, with larger variance than that generated by the simple menu cost model. However, it does not generate the small price adjustments present in the Alvarez-Lippi-Midrigan model.

We evaluate the effect of a monetary policy shock in an economy populated by multi-product firms with price-setting technology calibrated according to the price-setting statistics of the Israeli data. We compare the result to those generated by the single-product menu cost model and the full synchronization model. The magnitude of the real effect of the monetary shock of the Golosov-Lucas model is significantly smaller than that in the Alvarez-Lippi-Midrigan model due to the presence of price selection effect. In response to the shock, firms in the Golosov-Lucas model choose to change first the prices of products that are further away from the optimal, triggering adjustments that are relatively larger and amplifying the response of the aggregate price. By contrast, in the Alvarez-Lippi-Midrigan model there is no price selection since each firm changes all its prices at the same time.

We develop analytical and numerical results that help us answer two questions about how the aggregate price level and real output respond to aggregate demand shocks in a world with multi-product firms. First, how does selection interact with within-store synchronization of price changes? Our analytical results consist of a first order characterization of the response of real output and price level to a demand shock. They allow us to decompose the initial responses of these variables to a shock into an extensive margin component, associated with the frequency of price changes, and an intensive margin component, which we attribute to price selection. We then show how the Alvarez-Lippi-Midrigan model corresponds to the extreme case when the selection component goes to zero, while the Golosov-Lucas model can be seen as the limiting situation in which selection goes to infinity.

Second, how far do we need to move away from the Alvarez-Lippi-Midrigan world in order to have substantially smaller monetary non-neutrality? Our numerical results indicate that small deviations from this case can significantly reduce the persistence of real effects of demand shocks. For the degree of synchronization observed in the data, our model generates responses very close to Golosov-Lucas. Although firms in the partial synchronization model do not change all prices at the same time, they choose to change those prices that are further from the optimal, triggering larger adjustments shortly after the shock. Therefore, the selection effect plays a key role in the partial synchronization model, engineering a faster increase in the aggregate price level and attenuating the monetary policy effect.

Our results also hold for the case in which the fixed cost K is a friction of informational nature. In this alternative specification, the profit-maximizing price of a given product is unobservable, and economies of scope in price adjustments come from the information acquisition technology. There is a fixed cost K whose payment is required in order to observe the profit-maximizing price of all products simultaneously. This informational friction model does not necessarily generate the same responses of output to a monetary shock as the baseline case in which $K$is interpreted as a common menu cost. The reason for this is that in the menu cost case, firms know when the shock happens, and may therefore change their policies instantaneously in response it. On the other hand, in the informational friction world firms would only learn about the aggregate shock after payment of the fixed cost. We prove, however, that changes in policies that arise from an aggregate shock that hits an economy in steady state do not affect aggregates up to the first order. Consequently, both models have the same implications for small shocks. Our analysis therefore can also be seen as a multi-product generalization of Alvarez, Lippi, and Paciello (2011) and Malta et al. (2015), in the limiting case as the number of products go to infinity.

This paper is more broadly related to the extensive recent literature in monetary economics that tries to reconcile stylized facts about price-setting in the micro price data with the macroeconomic evidence about inflation and output effects. 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 (Carvalho, 2006, Nakamura and Steinsson, 2010), imperfect information (Mankiw and Reis, 2002, Reis, 2006, Woodford, 2009), and strategic complementarities (Basu, 1995, Nakamura and Steinsson, 2010). As shown by Midrigan (2011), Alvarez and Lippi (2014), multi-product pricing with economies of scope in price adjustment, attenuates the selection effect and reduces the response of the aggregate price level to shocks.[3]

Midrigan (2011) extends Golosov and Lucas (2007) setting by having a two-product firm with maximal economies of scope in price adjustment subject to fat-tailed distribution of cost shocks. As a result, the model not only fits the micro data[4] well but is also capable of generating real effects of monetary policy that are much greater than in Golosov and Lucas (2007). Alvarez and Lippi (2014) extend Midrigan (2011) maximal economies of scope in price-adjustment model by allowing an arbitrary number of products. They derive analytical expressions for the frequency of adjustment, the hazard rate of price adjustments, and the size distribution of price changes in terms of the structural parameters of the model. They also show analytically that the size of the output response and its duration both increase with the number of products, converging to the response of Taylor's staggered price model when the number of products gets large.

The issue of synchronization versus staggering of price-setting have been discussed in the context of multi-product firms since Lach and Tsiddon (1996). They use a sub-sample of multi-product stores selling wine and meat during the high inflation period in Israel. They found that price adjustments across stores tend to be predominantly staggered, but those within firms tend to be highly synchronized. The synchronization of within firm prices is corroborated by a study of Canadian newspapers by Fisher and Konieczny (2000).

More recently, Bhattarai and Schoenle (2014) use micro data for US producer prices to compute price-setting statistics of multi-product firms. Almost all firms in the sample (98.6%) were multi-product, with a median of 4 products per firm. They found that as the number of products per firm rises the frequency of price adjustments increases, the average size of price adjustments decreases and the dispersion of price changes gets larger. They also document that price changes within firms tend to have the same sign, an effect that increases with the number of products.

Stella (2014) uses weekly data on prices, costs and units sold by a supermarket chain to estimate the costs of changing prices of a multi-product firm. Stella allows for both types of cost included in our model. The total cost from changing prices is estimated to be between 0.22% and 0.59% of revenues, with the common part of the cost accounting for up to 85% of the total menu costs expense. Finally, it is important to emphasize that we model regular price changes only, despite the recent development of theories that account for both regular and temporary price changes and its effects on monetary non-neutrality, such as Alvarez and Lippi (2019).

The remaining sections of the paper are organized as follows. Section 2 provides new evidence on prices set by large retailers in Israel. We present our model in section 3. The following section discusses the calibration of the three models and presents their predicted monetary effects. The

[3] However, it also reduces the response to idiosyncratic shocks. Thus, it requires a very large variance of idiosyncratic shocks to match the frequency and size of individual price adjustments in the data (Bils, Klenow, and Malin, 2012). 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 and firm-specific shocks tend to increase with the number of products due to economies of scope in information processing.

[4] The assumption that paying once the cost of price adjustment allows the firm to adjust all prices makes the model capable of generating small price adjustments. Lach and Tsiddon (2007) had showed evidence that small price adjustments are more likely when the average size of all simultaneous price change are large—a feature consistent with this price adjustment technology.

last section concludes.

2 Evidence from the retail stores in Israel

2.1 The Israeli retail data

Following the "Social Protest" in Israel[5] 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 retailers[6] operating in Israel are required to publish on their Internet sites daily price information for all products sold, in all points of sale, both brick-and-mortar stores and their Internet site.

The law requires retailers to keep data only for the past 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 1700 stores, which account for most of the volume of food retail sales in the local market. To manage computational constraints, the data used in this paper contains information for top 5% of stores (by the number of observations) for each retailer from May 22, 2015 until September 1, 2019. Some stores have missing data early in the sample, and smaller chains tend to have scarce or poor quality data. Therefore, for sample selection, for each retail chain we sorted stores by the number of days for which data were available, then we chose the top 5% of stores for each retailer, conditional on acceptable data quality. We excluded store-specific products that do not have a general 13-digit barcode (e.g., fruits, vegetables, bakery goods). We also excluded pharmacy retailers and 6 stores that uploaded only from the dataset. The remaining dataset contains 451.7 million daily observations for 71 stores and 21 retail chains.

For each store, have information about the "base" or regular prices for individual products on a daily basis. 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 discounts, a third-product-free discount, or two products for 10 NIS (Israeli New Shekel). Based on the available regular price and the discount code, the Bank of Israel constructed the final price for 10 stores owned by Shufersal, the largest food retailer in Israel, with 350 of its stores servicing about 35 percent of the food retail-chain market.

We report most empirical results for regular prices, and in Section 2.5 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.

[5] See also Chapter 1 in Bank of Israel (2012).

[6] A retailer with annual sales exceeding NIS 250 million (about USD 70 million).

Panel A in Figure 2 shows the distribution of the average number of products per day in a store. 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, a large proportion of products, around 40%, do not register a regular price change in our sample. Nonetheless, there is substantial variation in the average frequency of price changes across stores (Panel B).

Table 1 provides statistics for the frequency and size of price adjustments across retailers. 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.

Monthly frequency of observations, however, filters out high frequency movement in product prices. The duration of price spells implied by the frequencies of price changes in Table 1 would be 8.3 months for monthly observations, 5 months (22 weeks) for weekly observations, and 3.5 months (105 days) for daily observations. Measuring duration directly from observed spells leads to a similar conclusions.

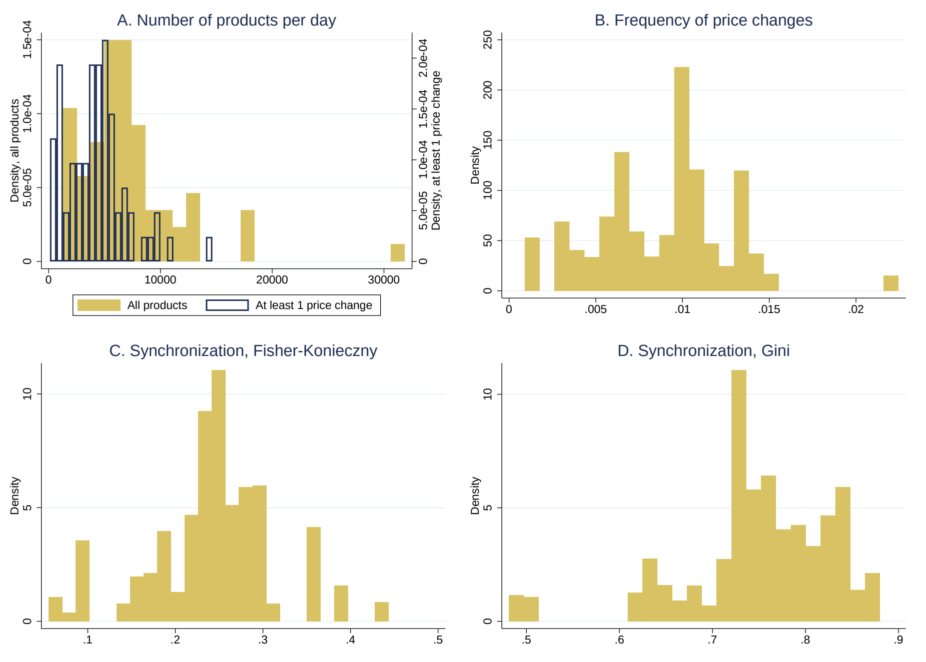

Figure 2: Store-level number of products, frequency and synchronization of price changes.

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 frequencies of price changes across stores. Panel C gives the histogram of Fisher-Konieczny synchronization index across stores, and Panel D provides the histogram of Gini index values across stores.

Sample	Frequency of observations	Number of products per day	Mean fraction of price changes	Mean abs size of price changes	Synchronization statistic	
					Fisher-Konieczny	Gini
		(1)	(2)	(3)	(4)	(5)
A. By store	daily	7,217	0.87%	20.8%	0.236	0.747
	weekly	8,170	4.30%	20.3%	0.225	0.534
	monthly	9,605	11.46%	20.1%	0.187	0.315
Same product category*	daily	6,989	0.89%	20.8%	0.289	0.770
Flexible price goods	daily	2,092	2.25%	21.2%	0.376	0.747
Larger stores	daily	10,267	0.89%	21.0%	0.253	0.752
B. By chain	daily	23,664	0.89%	20.8%	0.220	0.728
C. All observations	daily	301,496	0.99%	20.6%	0.142	0.573

Table 1: Summary statistics for daily price adjustments.

Note: We compute each statistic (in columns) for each store. The table provides weighted means across stores (Panel A), chains (Panel B), or unweighted means (Panel C). Weights are the average number of products in a store per day. "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 frequency 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).*

The most striking pattern of daily price changes is evident in Figure 1 which plots daily fraction of price changes in the dataset, weighted mean across all stores (top panel), and for four selected from different chains (bottom panel). 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 in every 19 days on average. On a peak day a store reprices about 9.8% of all products, *twenty* 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 1.6% of all days are such peaks (one peak in two months), but on such a day a store reprices about 15.3% of all products, *twenty three* times the average number of price adjustments on other days.

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, than across chains. Panel B in Table 2 shows that once we pool observations across stores, synchronization of price changes reduces only marginally.

Peaks in price adjustments are not due to calendar events. We document the fraction of days when a peak overlaps with a holiday in Israel, see Appendix A. Only 6.5% of peak days are holidays, and in turn, only 4.1% of holidays are peaks. Hence, re-pricing peaks are not related to holidays. We also looked at the prevalence of peaks by day of the month and by day of the week. Across all stores and days, price changes are more frequent at the turn of the month and early in the week.

Peaks exhibit a similar but smoother pattern, suggesting that there are no seasonal effects that are specific to peaks.

Sample	Number of days	Frequency of price changes	
		weighted	unweighted
A. By store			
Peak days (50%)	80	9.83%	9.49%
Off-peak days (50%)	1418	0.48%	0.44%
Peak days (25%)	24	15.27%	14.50%
Off-peak days (75%)	1474	0.67%	0.62%
All days	1498	0.87%	0.81%
B. By chain			
Peak days (50%)	91	9.04%	7.57%
Off-peak days (50%)	1407	0.48%	0.39%
Peak days (25%)	30	13.39%	11.50%
Off-peak days (75%)	1468	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 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.

2.3 Synchronization of price changes

To quantify the degree of synchronization of price changes in the data, we are going to use two alternative statistics: Fisher-Konieczny index and a Gini index.

The first index is constructed in the spirit of Fisher and Konieczny (2000) as follows:

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

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, $N_s$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 when they are perfectly synchronized.[7] 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 as shown in Panel C in Figure 2.

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

The second synchronization index is akin to 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, see Figure 3 for the distribution of price changes pooled across all stores. The Figure 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 1 (column 5) shows high degree of synchronization (i.e., high inequality of repricing across days), 0.747 on average, and varying substantially across stores (Panel D, Figure 2).

2.4 Synchronization across stores

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

Figure 3: Lorenz curve for inequality of re-pricing activity across days, selected stores.
Note: Figure shows the Lorenz curve (cumulative distribution of price changes across days) for four selection stores.

First, we ask whether prices are more synchronized for products within the same broad category than across categories.The Bank of Israel classified 50 broad product categories for Shufersal stores. Panel A in Figure 4 compares the distribution of FK synchronization index across stores computed for all products to the distribution conditional on products belonging to the same category. Price changes are more synchronized within, rather than across categories. On average, FK index is 0.289 for products in the same category, versus 0.236 for all products (Table 1).

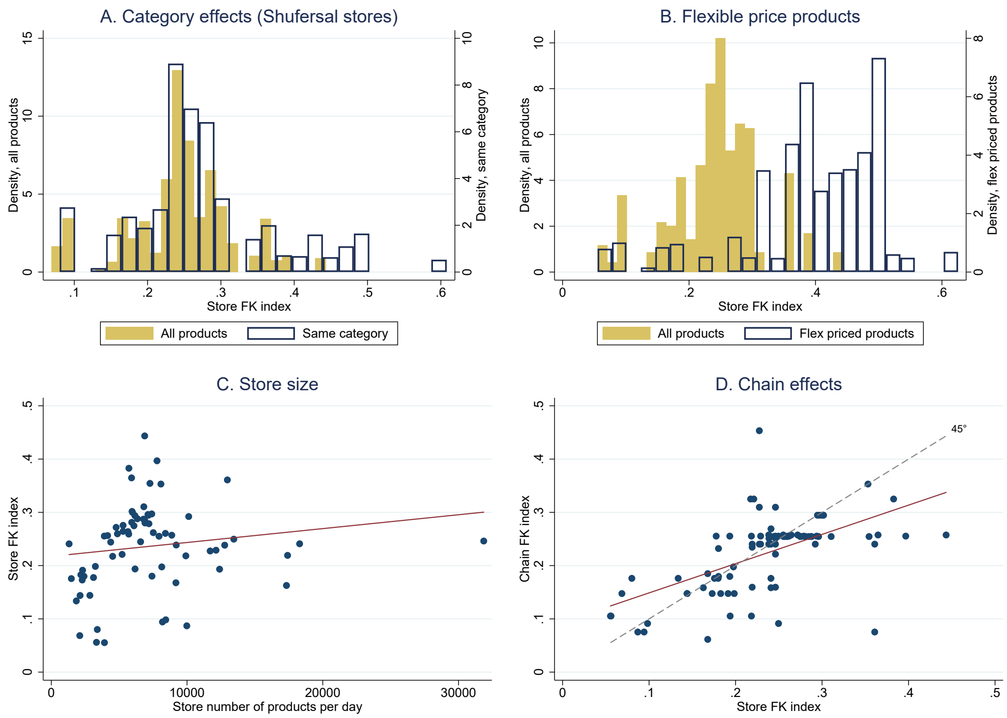

Figure 4: Fisher-Konieczny synchronization across stores.

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 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.

Similarly, Panel B compares distributions for all products and only flexible price products, defined as those in the top quartile of the frequency of price changes in the store. Flexible price changes are more synchronized, with FK index of 0.356. Interestingly, the Gini index is unchanged between all and flexible price goods. This highlights the difference between FK and Gini synchronization: the former increases with the average frequency of price changes, the latter does not.

Panel C shows that larger stores tend to have more synchronized price changes, the (unweighted) fitted line is sloped upward. The FK index for the subset of stores above the median is 0.253, versus 0.236 for all stores. Finally, we compare FK synchronization value computed for each store with values for the chain to which that store belongs. The scatter plot in Panel D shows that synchronization 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 reminiscent of the recent findings by DellaVigna and Gentzkow (2019) for the U.S. retailers.

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 discounts, a third-product-free discount, or two products for 10 NIS (Israeli New Shekel). Corresponding tables and figures are in Appendix A.

Price discounts, "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 and Gini index values for all stores are also similar: 0.355 and 0.708 for final prices, versus 0.262 and 0.797 for regular prices. These results suggest that retailer's decisions to post price discounts are largely independent from 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.

3 A model with multi-product firms and partial synchronization

3.1 An 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 1. 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:

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

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

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

In the above, $W_{i,t}$ 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, which 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 above law of motion 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[8]. We can express the loss term (1) as

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

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 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):

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

The last building block we need to fully characterize the firms' problem are the pricing frictions 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 in order to make any number of price adjustments. Second, there is a unit cost c that must be paid for each individual 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 two important features. First, a positive fixed cost K generates inaction. This means that price adjustments occur in dates that are separated by potentially long time intervals, i.e. firms do not

[8]The 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.

adjust prices continuously over time. Second, when a firm decides to adjust prices, it is never optimal to adjust prices of all its products. Intuitively, it is never optimal 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 in order to set it to zero.

The last important observation before describing the optimal policy is that, even though the frictionless optimal price of a single product is stochastic, the relevant 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 $t = T_k$ 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.[9] 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 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 has just paid the fixed cost and has 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 would rank its products according to the size of price discrepancies and, starting from the good with the largest gap, move down the list adjusting prices 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.

Our model nests two other cases previously studied in the price setting literature. 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 in order to reset all prices at once. In this case, economies of scope of adjusting prices is maximal and, at any given instant, the share of prices that a certain firm resets is either zero or one. Therefore, if we set $c = 0$ our model becomes an infinite product limit of the Alvarez-Lippi-Midrigan framework.

The competing extreme 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 one responsible for adjusting the price of a single good, as in Golosov 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

[9] In the presence of a nonzero drift, the optimal sequence of thresholds has to be split into sequences of upper thresholds $\{\overline{x}_k\}_{k=1}^{\infty}$, 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 $x \geq \overline{x}_k$ 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 price gaps to be expected to 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.

day, the share of products of a given firm that had theirs prices adjusted is constant. Our model therefore flexibly captures pricing behaviors ranging from within-firm perfect synchronization in price adjustments to variety-specific price adjustments.

3.2 Recursive formulation

The main difficulty in solving the partial synchronization model is that the relevant state variable in the dynamic optimization problem is the entire distribution of price discrepancies. Alvarez and Lippi (2014) show that, in the perfect synchronization case ($c = 0$), there is no need to keep track of the whole distribution. In their model, all the 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 our 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:*

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

Proof. See appendix. □

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

$$L_t = L_0 + \sigma^2 t \tag{5}$$

Proof. See appendix. □

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 very precise, approximation for this distribution, as explained in Appendix B. The function V then satisfies:

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

In the above, 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:

$$m(\bar{x}, g) = \int\limits_{|x| \geq \bar{x}} g(x)\, dx$$

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:

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

In the expression above, $1(\cdot)$ 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 ($|x| > \bar{x}$), which amount to a mass $m(\bar{x}, g)$ of products, and consequently 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 costs 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 τ^* time periods 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:

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

$$g^*(x) = \int\limits_{-\infty}^{+\infty} \frac{1}{\sqrt{\sigma^2\tau^*}}\phi\left(\frac{x-y}{\sqrt{\sigma^2\tau^*}}\right) g_0^*(y)\, dy \tag{8}$$

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). It is easier to see why we call the set of equations above a fixed-point problem by substituting the first equation into the second. Nevertheless, the system is easily solvable by creating a discrete grid for the 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 vectors of the values they attain in the x grid, the problem boils down to finding an eigenvector of a large matrix. Figure 5 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^*.

Finally, Figure 6 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 K cost 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.

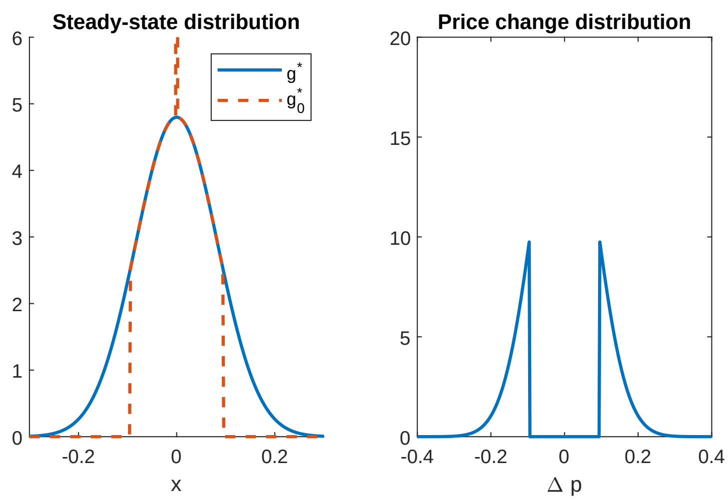

Figure 5: Steady-state and price change distributions for parameter values $\rho = 0.04$, $\sigma = 0.25$, $K = 0.0001$, $c = 0.001$.

3.3 Calibration

In order 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 price-setting literature, we use the frequency of price adjustments and the average size of price changes. More specifically, our measure of 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 lacks one kind of menu costs. 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

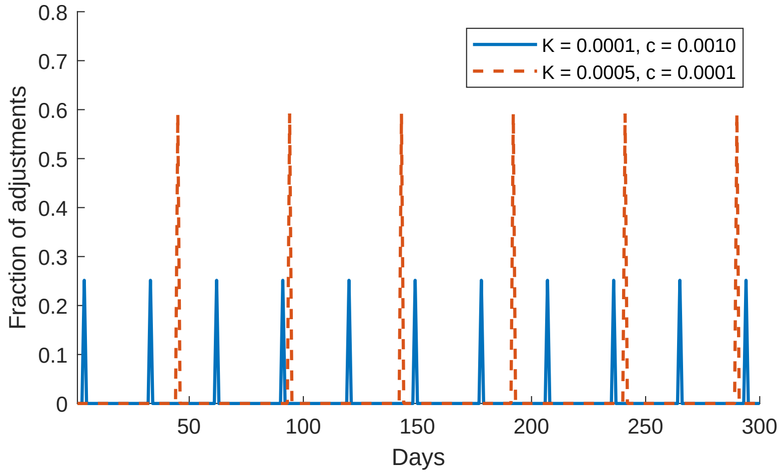

Figure 6: Daily fraction of prices changes for different combinations of menu costs. Other parameter values: $\rho = 0.04$, $\sigma = 0.25$.

can only generate 1. Table 3 shows moments for models and data, and Table 4 shows calibrated parameter values. The optimal policy for the partial synchronization model in our calibration consists of adjusting every 7.5 days the prices with corresponding gaps larger than 17.8%.

Moment	Data	GL	ALM	Partial sync.
Daily fraction of price changes	0.0089	0.0089	0.0089	0.0084
Avg. abs size of price changes	0.209	0.209	0.209	0.216
Fisher-Konieczny index	0.236	0.000	1.000	0.236

Table 3: Moments from data and calibrated models.

Note: Values in the Data are weighted means across stores in the data, provided n Table 1, first row, column (2)–(4).

Parameter	GL	ALM	Partial sync.
σ_i	0.3774	0.4730	0.3834
K	-	0.0106	3.24e-05
c	0.0022	-	0.0020

Table 4: Calibrated parameter values.

Figure 7 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 need only pay the fixed cost K to adjust all prices, even those 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 8 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 closer to replicating the frequent and short peaks seen in the data (Figure 1).

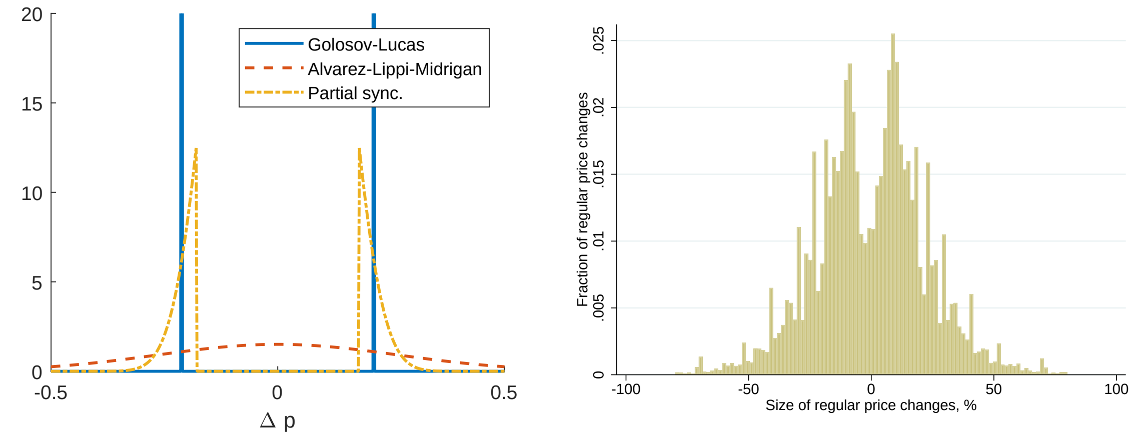

Figure 7: Price change distributions for all models and data.

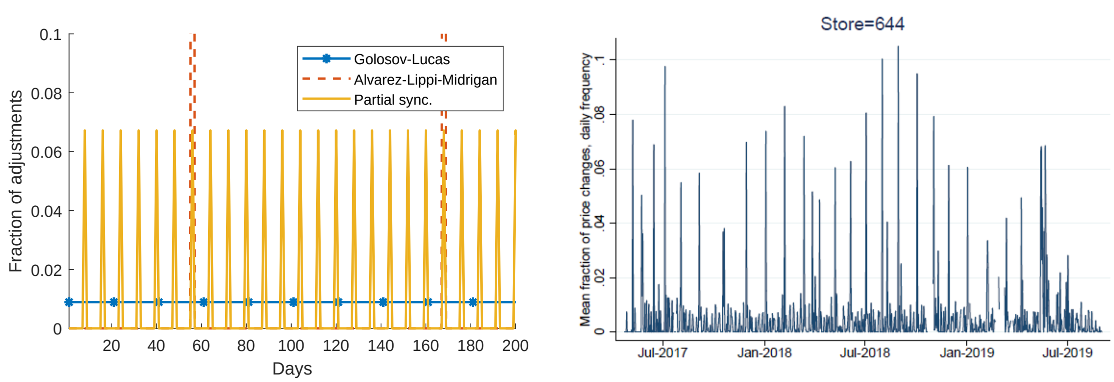

Figure 8: Daily share of adjusted prices for the three models and data. Spikes in the ALM model have height 1.

4 Real effects of demand shocks

4.1 Some analytics for the partial synchronization model

In our model, a frictionless optimal price is the sum of a product-specific shock that follow Brownian motions and an aggregate demand component M_t that has so far been held constant. Now we consider 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 different firms and products in the model economy. Real output Y_t, also in logs, is then given by:

$$Y_t = M_t - P_t$$

We study an economy with a continuum of identical firms, which is hit by an unanticipated aggregate shock of size ε at $t = 0$. This shock shifts aggregate demand M_t to $M_t + \varepsilon$ and, as a consequence, the price gap distribution of all firms is also shifted: $g_t(x)$ becomes $g_t(x + \varepsilon)$. Prior to the occurrence of the aggregate shock, we naturally consider the situation in which all firms are in steady-state, adjusting prices every τ^* periods. Moreover, we start with a situation in which firms are uniformly distributed according to the time elapsed since the last adjustment date, that is, a constant flow of firms adjusts prices over time before time $t = 0$. In order 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, which can be obtained by setting $\varepsilon = 0$, the firm would resume its steady-state policy, characterized by:

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

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

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

$$\Delta_k = 0$$

To obtain analytical results about the responses of aggregate price level and output to demand shocks, we focus on the limit as $\varepsilon \to 0$. 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 average price will take discrete steps at dates of the form $k\tau^* - s$. Figure 9 illustrates how this response would look like.

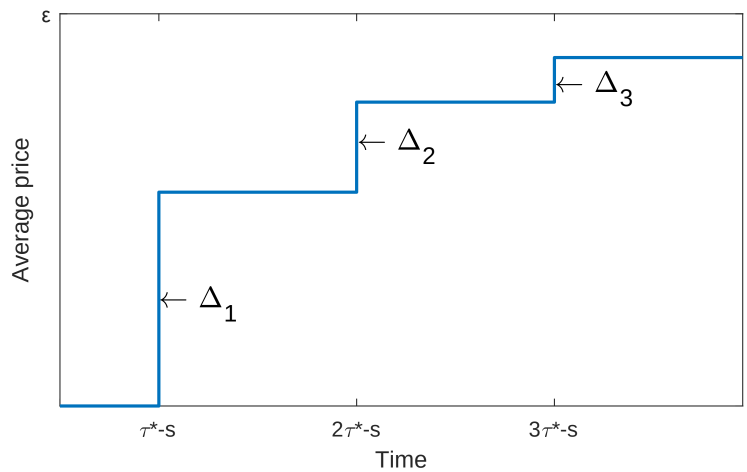

Figure 9: Approximate response of the average price of a single firm following a small aggregate shock of size ε.

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 $t \in (k\tau^*, (k+1)\tau^*)$, adjusting firms' average prices 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 aggregate price level at instant t following a shock of size ε and let

$$\delta_k = \lim_{\varepsilon \to 0} \frac{\Delta_k}{\varepsilon}$$

We have the following result.

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

Proof. See appendix. □

Figure 10 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, more than half of the rise of the price level following an aggregate shock happens in this first segment for our calibration. But first, a couple more definitions. Let F be the steady-state instantaneous frequency of price adjustments, defined as

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

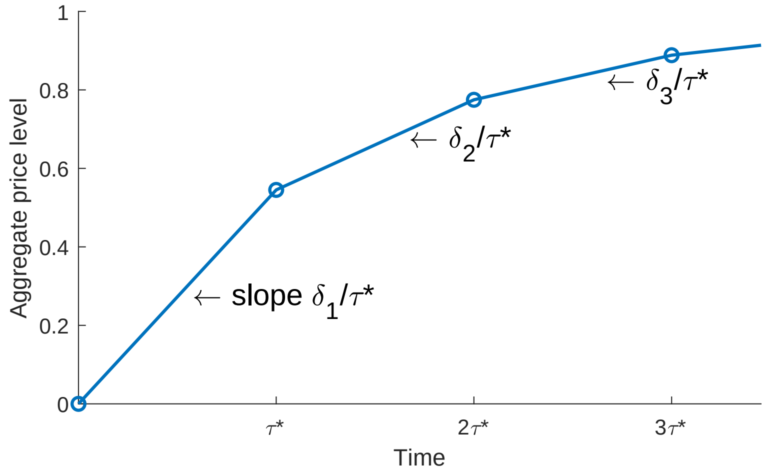

Figure 10: Normalized price level response to a small aggregate shock.

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 (7) associated with $x \geq 0$, and scaled to integrate to 1. In order to characterize the slope of the first line segment, we make use of the following lemma, whose proof and necessary definitions are presented in the appendix.

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:*

$$\frac{\partial \Delta_k}{\partial T_j} = \frac{\partial \Delta_k}{\partial \bar{x}_j} = 0 \tag{9}$$

Proof. See appendix. □

As a consequence of the lemma above, we have the following result.

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

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

Proof. See appendix. □

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 Some comments about a model with informational 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 in order 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). 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 in 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. Impulse responses of aggregate output and price level would be the same in both cases for small enough shocks. This is a consequence of two facts. First, the steady state optimal policy is the same for both economies. Second, changes in optimal policies do not affect aggregate outcomes to a first order, which is the consequence of 9 that allows us to prove 10.

4.3 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 11 shows the impulse response functions for calibrated parameter values shown in 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 phenomenon 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 synchronization reduces monetary non-neutrality. In other words, how far do we need to move away from the ALM perfect synchronization case in order to have significantly less persistent output responses? The answer is: not very far.

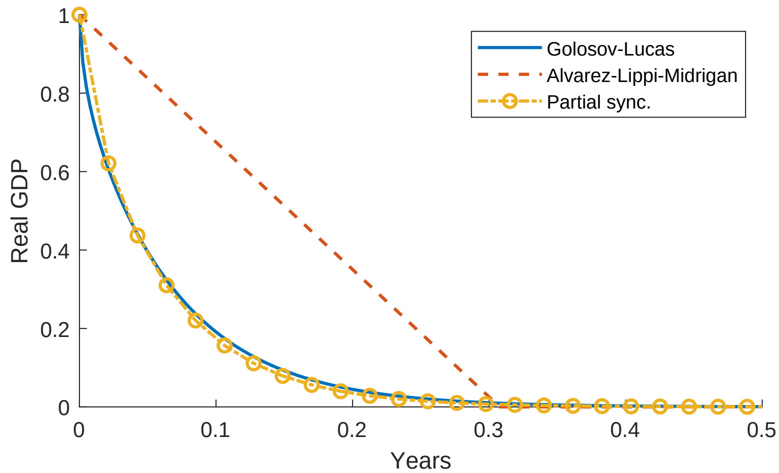

Figure 11: Normalized real output response to a small aggregate shock.

Why is partial synchronization associated with smaller non-neutrality? First, note that more than half of the decay of output following the shock happens in the first round of adjustments, i.e., 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[10]. We can see the ALM model as the limit when $\bar{x}^* \to 0$. 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 limit case in which $\bar{x}^* \to \mathbb{E}|\Delta p|$ and $f(x)$ approaches a Dirac mass at $\bar{x}^*$. Therefore, for 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 relation 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, there is not much room for selecting those prices, and monetary shocks have more persistent real effects as a consequence. It is also important to emphasize that, since we study the limit case in which the shock size goes to zero, the aggregate response does not come from anticipation of adjustment episodes.

Now we aim to understand how fast this selection effect operates. Figure 12 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 as it is

[10]Even 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.

related only to the initial, rather than total, degree of aggregate price flexibility. For this reason, Figure 13 shows both the selection term and the cumulative output response, measured as the area under the impulse response function, for varying FK values. The message here is clear: the marginal effect of synchronization on aggregate price stickiness is increasing in synchronization, in such a way that small departures from the ALM framework can reduce considerably the degree of monetary non-neutrality.

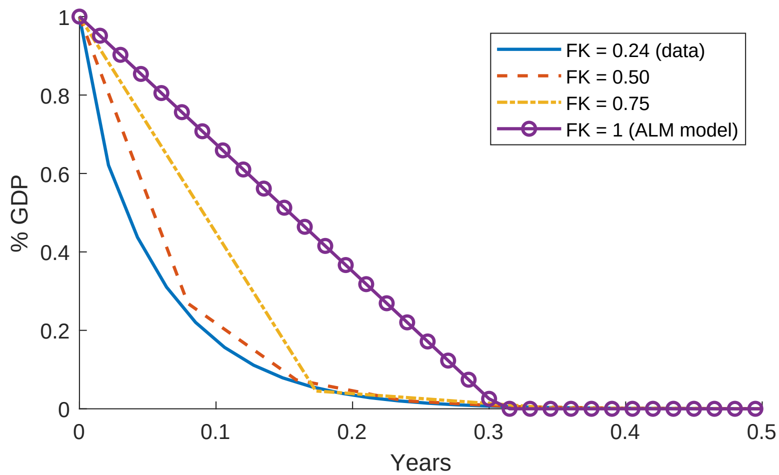

Figure 12: Real output response and synchronization.

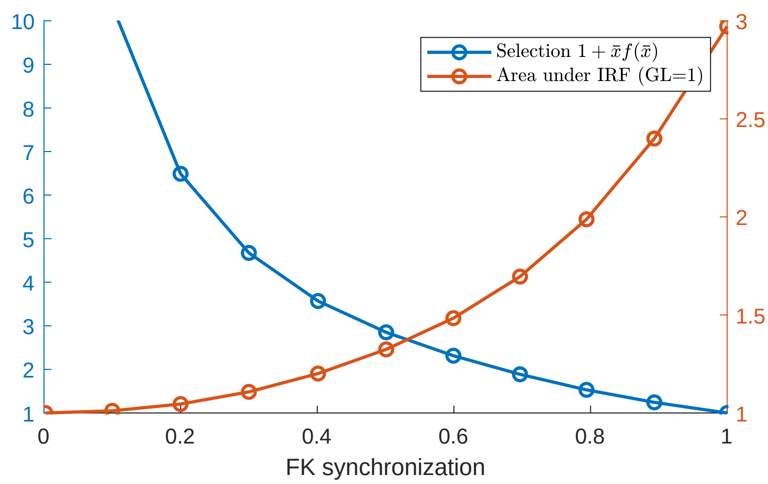

Figure 13: Selection component $1 + \bar{x}^* f(\bar{x}^*)$ and cumulative response for various values of the FK synchronization measure. Cumulative response of the GL model is normalized to 1.

5 Conclusions

In this paper we bring new evidence on prices set by large retailers in Israel, and propose a general equilibrium model of multi-product price-setting to account for this evidence. Each price retailer in our sample prices on average 7000 products. The pattern that emerges in the data is that prices of multi-product firms are partially synchronized due to occasional peaks in 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 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, as in Golosov and Lucas (2007). In the case with maximal economies of scope, the firm pays the fixed cost to adjust a bulk of its prices, as in Midrigan (2011) and Alvarez and Lippi (2014). We show that the calibrated model, despite displaying considerable synchronization of price changes, generates monetary effects that are relatively small, close to those generated by an economy populated by single-product firms. The selection effect in the partial synchronization model—with the monetary shock triggering initially 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 and Francesco Lippi. 2014. "Price Setting with Menu Cost for Multiproduct Firms." *Econometrica* 82 (1):89–135.

———. 2019. "Temporary Price Changes, Inflation Regimes and the Propagation of Monetary Shocks." *forthcoming in American Economic Journal: Macroeconomics* .

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.

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2012. "Reset Price Inflation and the Impact of Monetary Policy Shocks." *American Economic Review* 102 (6):2798–2825.

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.

Danziger, Leif. 1999. "A Dynamic Economy with Costly Price Adjustments." *American Economic Review* 89:878–901.

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.

Fisher, Timothy C. G. and Jerzy D. Konieczny. 2000. "Synchronization of price changes by multi-product firms: evidence from Canadian newspaper prices." *Economics Letters* 68 (3):271–277.

Golosov, Mikhail and Robert E. Lucas. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115:171–199.

Klenow, Peter J. and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Econoimcs* 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.

Lach, Saul and Daniel Tsiddon. 1996. "Staggering and Synchronization in Price-Setting: Evidence from Multiproduct Firms." *American Economic Review* 86 (5):1175–1196.

———. 2007. "Small price changes and menu costs." *Managerial and Decision Economics* 28 (7):649–656.

Malta, Vivian, Rene Garcia, Carlos Carvalho, and Marco Bonomo. 2015. "Persistent Monetary Non-neutrality in an Estimated Model with Menu Costs and Partially Costly Information." 2015 Meeting Papers 1339, Society for Economic Dynamics.

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 Joń 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.

Stella, Andrea. 2014. "The Magnitude of Menu Costs: A Structural Estimation." 2014 Meeting Papers 436, Society for Economic Dynamics.

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

A Data Appendix

Figure 14: 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 to inflation for official CPI for food products (excluding fruits and vegetables). Weights are the average number of products in a store per day.

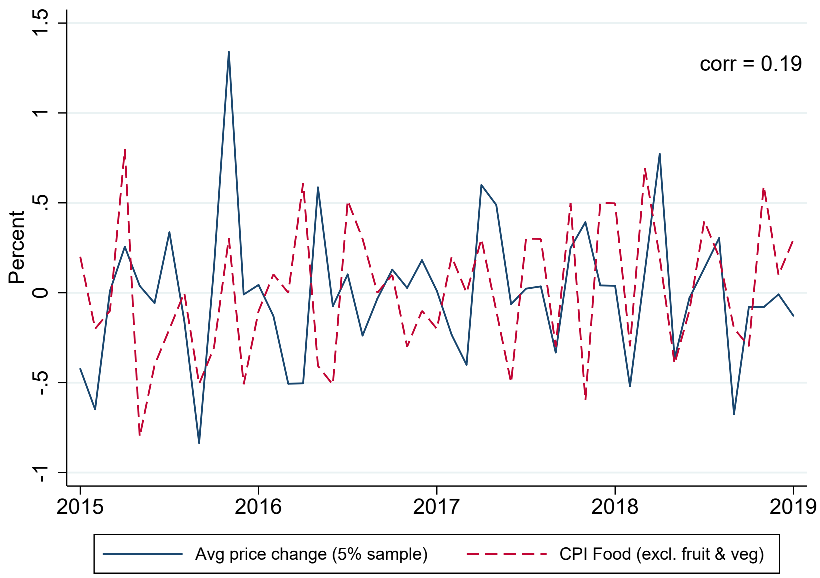

Figure 15: Weekly and monthly fraction of regular prices changes, pooled across stores.

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

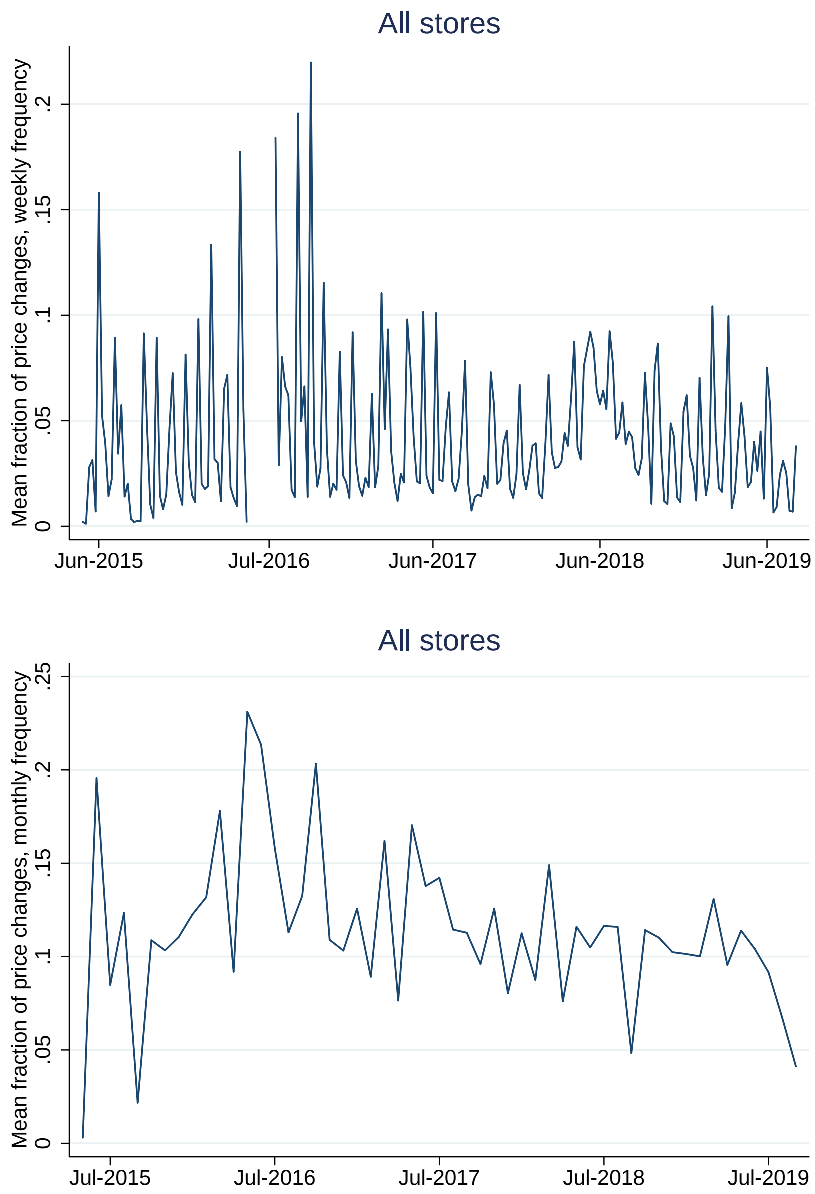

Table 5: Peaks and holidays.

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.

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%

Figure 16: Distribution of daily price changes and their frequency, by day of the month.

Note: (Weighted) Distribution of the daily price changes and their frequency by day of the month, 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).

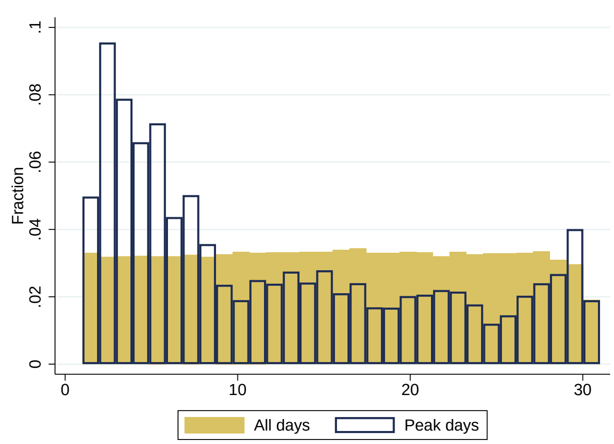

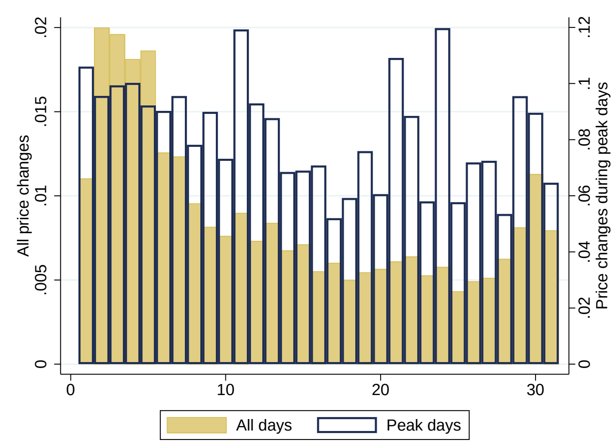

Figure 17: Fraction of daily price changes and their frequency, by day of the week.

Note: (Weighted) Distribution of the daily price changes and their frequency by day of the week, 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).

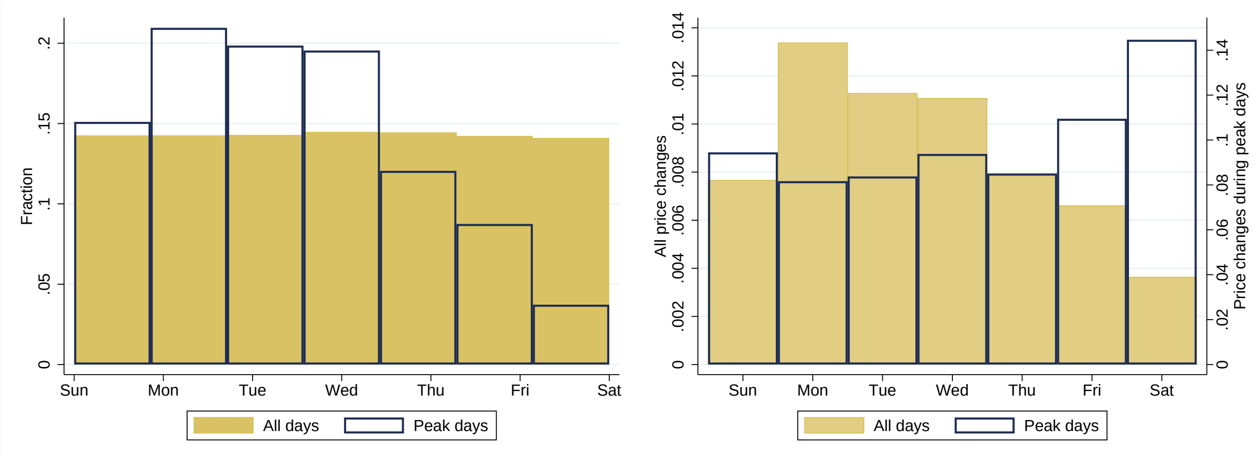

Table 6: Final versus regular price changes for Shufersal.

Note: Entries include the number of observations, mean and standard deviation for time series at daily frequencies. The data covers 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.

	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 7: Price discounts at Shufersal.

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

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 8: Synchronization of regular and final price changes for Shufersal.

Note: The data covers 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	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 9: Re-pricing peaks for final and regular prices for Shufersal.

Note: The data covers 10 Shufersal retail stores from January 2016 until mid–2019. Column entries compare mean frequency of daily final and 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.

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

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 $t \to \infty$.

Interestingly, for our numerical purposes the convergence happens fast enough so that g_T 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 in our calibration, the approximation is good. We choose our state variable in the recursive formulation of the problem to be the distribution of prices gaps immediatelly 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 18 compares the steady-state distribution that would arise from following the optimal policy for calibrated parameter values, given by equations (7) and (8), and 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 18 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 $[-0.0890, 0.0915]$.

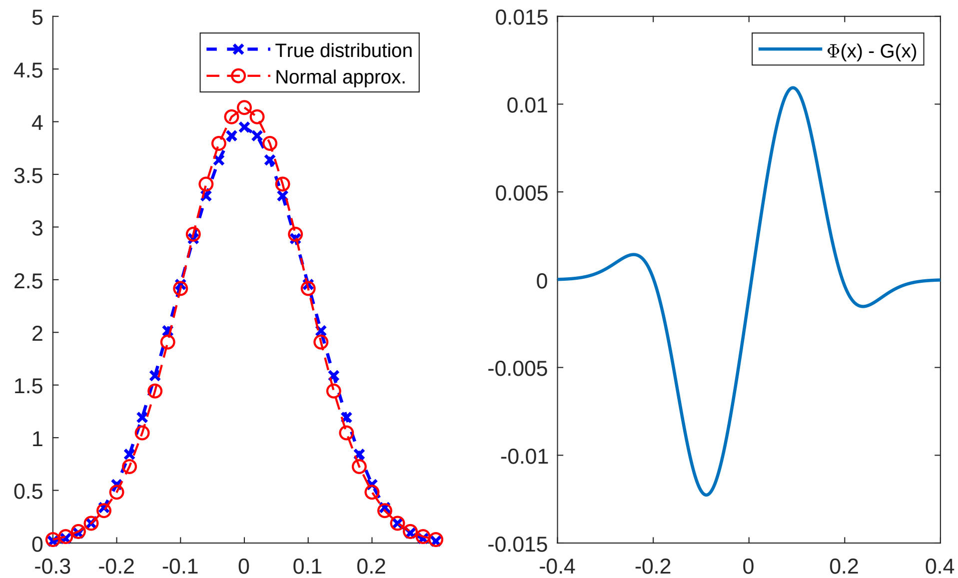

Figure 18: Steady-state distribution and normal approximation. $\sigma = 0.3834$, $T^* = 0.0207$. The value of σ shown here is the volatility of the Brownian motion, not the standard deviation of the normal approximation.

C Proofs

C.1 Lemma 1

Substitute (4) into (3) and use the fact that $\phi''(x) = -x\phi'(x) - \phi(x)$

C.2 Lemma 2

From (4), it follows that

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

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.

C.3 Proposition 1

Let $\bar{p}_\varepsilon^s(t)$, for $t \geq 0$, be the average price changed by a firm whose last price adjustment before the shock was at date $-s$, i.e. s periods before the shock hits, for $0 < s < \tau^*$. 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

$$\lim_{\varepsilon \to 0} \frac{P_\varepsilon(t)}{\varepsilon} = \lim_{\varepsilon \to 0} \int_0^{\tau^*} \frac{1}{\tau^*} \frac{\bar{p}_\varepsilon^s(t)}{\varepsilon}\,ds \tag{11}$$

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 $T_k^s(\varepsilon) \to T_k^s(0) = k\tau^* - s$ as $\varepsilon \to 0$. As a consequence,

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

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 does not happen as the shock size goes to zero and $T_k^s(\varepsilon) \to k\tau^* - s$.

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[11]. So the Dominated Convergence Theorem allows us to exchange the order of integration and limit operators in (11), which gives us:

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

C.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 is it the distribution at date T_k *before* adjustments take place. We can then express Δ_k as

$$\Delta_k(\varepsilon,\{T_j\}_{j=1}^{k},\{\bar{x}_j\}_{j=1}^{k}) = -\int\limits_{|x|>\bar{x}_k} x\, g_k(x;\varepsilon,\{T_j\}_{j=1}^{k},\{\bar{x}_j\}_{j=1}^{k-1})\,dx \tag{12}$$

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 constant price level, so $\Delta_k(0,\{T_j(0)\}_{j=1}^{k},\{\bar{x}_j(0)\}_{j=1}^{k}) = 0$ and the approximation becomes

$$\Delta_k(\varepsilon,\{T_j(\varepsilon)\}_{j=1}^{k},\{\bar{x}_j(\varepsilon)\}_{j=1}^{k}) = \left[\frac{\partial\Delta_k}{\partial\varepsilon} + \sum_{j=1}^{k}\frac{\partial\Delta_k}{\partial T_j}T_j'(0) + \sum_{j=1}^{k}\frac{\partial\Delta_k}{\partial\bar{x}_j}\bar{x}_j'(0)\right]\varepsilon \tag{13}$$

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 (12) will always be zero and is consequently

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

independent of T_j and $\bar{x}_j$ for any j, so the partial derivatives with respect to policy variables is zero and our first order approximation becomes

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

C.5 Proposition 2

In steady-state, we have:

$$\begin{aligned} \frac{\partial \Delta_1}{\partial \varepsilon} &= \frac{\partial}{\partial \varepsilon}\left(- \int\limits_{|x|>\bar{x}^*} x g^*(x+\varepsilon)\, dx \right) \\ &= - \int\limits_{|x|>\bar{x}^*} x g^{*\prime}(x)\, dx \end{aligned}$$

Integrating by parts, we have:

$$\frac{\partial \Delta_1}{\partial \varepsilon} = 2\bar{x}^* g(\bar{x}^*) + m \tag{14}$$

$$m = \int\limits_{|x|>\bar{x}^*} g^*(x)\, dx$$

The result follows from rearranging (14) and using $F = m/\tau^*$ and

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