Price Dynamics with Customer Markets*

Luigi Paciello

Einaudi Institute for Economics and Finance and CEPR

Andrea Pozzi

Einaudi Institute for Economics and Finance and CEPR

Nicholas Trachter

Federal Reserve Bank of Richmond

February 26, 2015

Abstract

We provide new evidence of the relationship between firm pricing and customer dynamics. We build and quantify a model where firms choose prices taking into account their effect on the evolution of their customer base, and customers face frictions to reallocate to other firms. The model accounts for salient features of retail price data, such as the high kurtosis of the price distribution and the low pass-through of cost shocks. We also study aggregate dynamics in our framework and show that microfounding customer reallocation leads to a countercyclical response of markups to demand shocks.

JEL classification: E30, E12, L16

Keywords: customer markets, price setting, product market frictions

1 Introduction

The customer base of a firm -the set of customers buying from it at a given point in time- is an important determinant of firm performance. Its effects are long lasting, as customer-supplier relationships are subject to a certain degree of stickiness (Hall (2008)). Starting with Phelps and Winter (1970), a large literature has stressed that the price is an important instrument to attract and retain customers, and firms seek to maintain and grow their customer base through their pricing decisions. However, our only evidence on the link between pricing and customer base dynamics rests on anecdotes or surveys (Blinder et al. (1998), and Fabiani et al. (2007)). Furthermore, little is known about the quantitative relevance of the implications of this link for pricing. The aim of this paper is to address both of these gaps.

We exploit novel data from the retail industry to provide direct evidence that firm prices influence households decision to stop buying from their current supplier. A natural implication of this result is that firms should take the effect on the customer base into account when they set prices. To understand whether this is an important mechanism, we build and estimate a microfounded model of firm pricing with customer markets. We find that the predictions of the model match well with the cross-sectional distribution of prices and the pass-through of cost shocks implied by our data. This suggests that customer markets play a relevant role in price formation, which is important to acknowledge given the recent interest of the macro literature on retail prices (see, for example, Burstein and Jaimovich (2012), Coibion et al. (forthcoming), and Kaplan and Menzio (forthcoming,a)). Furthemore, our findings lend support to a growing literature that uses customer markets to explain a wide array of phenomena, from business cycle dynamics (Kaplan and Menzio (forthcoming,b)) to the movements in international relative prices (Drozd and Nosal (2012)), and firm performance (Gourio and Rudanko (2014)).

Our empirical analysis relies on scanner data documenting pricing and customer base evolution for a major U.S. retailer. The data contain information on grocery purchases for a large sample of customers between 2004 and 2006. Household-level scanner data are particularly well suited to study customer base dynamics. First, we observe a wealth of details on all the shopping trips each household makes to the chain (list of goods purchased, prices, quantities, etc.). More importantly, we can infer when customers leave the retailer by looking at prolonged spells without purchases at the chain. Hence, these data allow us to study the relation between a customer's decision to abandon the firm and the price of the bundle of goods she consumes there. We show that an increase in the price significantly raises the probability that the household leaves the firm. This implies that the customer base is elastic to prices: a 1% change in the price of the customer's typical basket of grocery

1

goods would raise the firm yearly customer turnover from 14% to 21%.

To asses the implications of customer markets for pricing, we setup a model where firms choose prices taking into account their effects on the dynamics of their customer base. Customers respond to price changes but their ability to reallocate across suppliers is impaired by the presence of search frictions. Modeling the market friction as search costs suits well our application since search costs have been found to importantly affect price dispersion in clustered retail markets (Sorensen (2000)) similar to those our data refer to. The distinctive feature of our setting is that it delivers endogenous customer dynamics, arising as a consequence of customers search and exit decisions.

Customers start each period in the customer base of the firm from which they bought in the previous period. Every period, firms draw a new idiosyncratic productivity level, and post a price. Then, customers can decide to pay an idiosyncratic search cost to observe the state of another randomly selected firm, compare it to that of their old supplier, and decide where to buy (*extensive* margin of demand). After these decisions have been made, each customer decides her purchased quantity of the good (*intensive* margin of demand). In this setting, firms face a trade-off between charging a higher price and extracting more surplus from customers, versus posting a lower price to extract a lower surplus but from a larger mass of customers.

While being tractable, the model provides a rich laboratory to study how the relationship between customer and price dynamics is shaped, in equilibrium, by idiosyncratic production and search costs. The price posted by the firm and its current level of productivity determine the value that customers obtain if they remain in the firm's customer base. Customers of firms in the left tail of the distribution of such values are more likely to search and leave. As a result, firms have incentives to avoid being in the left tail of this distribution, which introduces an element of strategic complementarity in prices affecting the shape of the price distribution. The same mechanism also affects firms incentive to pass-through cost shocks.

We use the estimated price elasticity of the customer base, jointly with additional moments from our data, to identify the key objects of the model: the distribution of search costs and the productivity process. We use the estimated model to quantify the relevance of customer market in explaining the patterns we observe in the data. Our model delivers a yearly average turnover of 9%, nearly two-thirds of what we measure in the data (14%). This implies that customer dynamics triggered by variation in prices due to idiosyncratic cost shocks can explain a large fraction of the overall turnover. The relevance of customer markets in explaining salient features of real world pricing is also witnessed by two other distinctive predictions of the model: i) a price distribution with higher kurtosis and smaller dispersion than the distribution of production cost, ii) incomplete pass-through of cost shocks. Not only

do both of these features match well with the evidence from our pricing data but the high kurtosis of the price distribution has also been independently documented for homogeneous packaged good by Kaplan and Menzio (forthcoming,a).

Finally, we use our model to study the effect of an aggregate demand shock on markups. In ad-hoc models of customer markets firms react to strengthened demand by raising margins, generating procyclical markups. This result can be overturn by introducing habits into the utility function of consumers (Ravn et al. (2006)). We show that countercyclical behavior can also occur by endogenizing the evolution of the customer base. In our setting, a positive demand shock does not only make customers more eager to consume the good but it also pushes them to search harder for better sellers. Therefore, the demand shock increases the mass of customers who are potentially up for grab, motivating firms to lower their prices, and compress their markups, to retain them.

Related Literature.   Our paper relates to the seminal work by Phelps and Winter (1970) who study the pricing problem of a firm facing customer retention concerns. In their paper, the response of the firm's customer base to a change in the firm's price is modeled with an ad hoc function. We instead endogenize customer dynamics which arise as the outcome of customers' optimal search decisions in response to firms' pricing. Fishman and Rob (2003), Alessandria (2004), and Menzio (2007) also study the firm price-setting problem in models where search costs prevent customers from freely moving to the lowest price supplier. Fishman and Rob (2003) study the implications of customer markets for firm dynamics. Alessandria (2004) shows that such a model can generate large and persistent deviations from the law of one price, consistent with the empirical evidence on international prices. Menzio (2007) looks at the role of asymmetric information and commitment in the optimal pricing decision of the firm. Differently from our paper, in these studies customers face a homogeneous search cost and, as a result, optimal pricing is such that no endogenous customer dynamics occur in equilibrium.

Unlike the literature cited above, we exploit micro data to discipline our model and provide a quantitative assessment of the relevance of customer markets for pricing. This relates our findings to contributions that aim at documenting empirical stylized facts on the behavior of prices and markups. Our evidence on the shape of the price distribution ties in to the recent empirical work by Kaplan and Menzio (forthcoming,a). While their focus is on customers and the price they pay for the same good (or bundle of goods), we are interested in the point of view of sellers and the price they charge.

Another set of related contributions uses customer markets to address questions different from the ones we study here. Gourio and Rudanko (2014) explore the relationship between

the firm's effort to capture customers and its performance. They show that customer markets have nontrivial implications for the relationship between investment and Tobin's q. Drozd and Nosal (2012) introduce in a standard international real business cycle model the notion that, when producers want to increase sales, they must exert effort to find new customers. This extension helps to rationalize a number of empirical findings on the dynamics of international prices and trade. Dinlersoz and Yorukoglu (2012) focus on the importance of customer markets for industry dynamics in a model where firms use advertising to disseminate information to uninformed customers. Shi (2011) studies a setting where firms cannot price discriminate across customers and use sales to attract new customers. Kleshchelski and Vincent (2009) examine the impact of customer markets on the pass-through of idiosyncratic cost shocks to prices but focus on a case with no heterogeneity in firm productivity and markups. Burdett and Coles (1997) study the role of firm size for pricing when firms use the price to attract new customers. Their work complements ours: price and customer dynamics in their setting are shaped by the heterogeneity in firm size (age). For us, the driving force is the heterogeneity in productivity.

Several other studies analyze the implications of product market frictions for business cycle fluctuations (Petrosky-Nadeau and Wasmer (forthcoming), Bai et al. (2012), and Kaplan and Menzio (forthcoming,b)). The industrial organization literature has also studied the implications of customer markets for a variety of subjects. For instance, Foster et al. (2013) stress their role in affecting firm survival and Einav and Somaini (2013) and Cabral (2014) focus on their effect on the competitive environment.

Finally, we relate to the literature on deep habits where persistence in demand is due to the preferences of the agents rather than to costly search. The pricing and aggregate implications of this setup have been studied by Ravn et al. (2006), Ravn et al. (2010), and Nakamura and Steinsson (2011).

The rest of the paper is organized as follows. Section 2 presents the data and descriptive evidence of the relationship between customer dynamics and prices. In Section 3 we lay out the model and in Section 4 we characterize the equilibrium. In Section 5 we discuss identification and estimation of the model. In Section 6 we present some quantitative predictions of the model and compare them with empirical evidence from our data. In Section 7 we introduce an application of the model with the goal of studying the implications of customer markets for the dynamics of markups. Section 8 concludes.

2 The link between prices and customer dynamics

We use new micro data to provide direct evidence that firm prices have an effect on the evolution of the customer base. In particular, we document that changes in the price of the goods included in their regular consumption basket at a large US supermarket chain affect customers' decision to abandon the chain. This result provides a compelling motivation to modeling the link between customer base and pricing policy, lending support to the central tenet of the growing literature on customer markets. Pre-existing evidence of this relationship is based on survey data where firms report concerns about customer retention as the main reason for their reluctance to adjust prices (see Blinder et al. (1998), and Fabiani et al. (2007)). To the best of our knowledge, we are the first to document this fact using micro data based on actual customers' decisions. Besides supporting the mechanism at the heart of our model, the descriptive evidence presented in this section provides statistics useful to estimate the model and help us quantify the importance of customer markets in shaping firm price setting.

2.1 Data sources and variable construction

We exploit cashier register data from a large U.S. supermarket chain on purchases by a panel of households carrying a loyalty card of the chain.[1] The chain operates over a thousand stores across 10 states, and the data reflect this geographical dispersion. For every trip made at the chain between June 2004 and June 2006 by customers in the sample, we have information on the date of the trip, store visited, and list of goods purchased (as identified by their Universal Product Code, UPC), as well as quantity and price paid. The customers in our sample make an average of 150 shopping trips at the chain over the two years; if those trips were uniformly distributed, that would imply visiting a store of the chain six times per month. The average expenditure per trip is $69 for the average household. There is a great deal of variation (the 10th percentile is $29; the 90th is $118) explained, among other things, by income and family size of the different households.[2]

This brief description highlights how our data are well suited to study customer base dynamics. Not only are they rich in detail about households' grocery consumption (products purchased, expenditure, prices, etc.) but they also include a panel dimension which is cru-

---

[1]The chain is able to associate the loyalty cards belonging to different members of a same family to a single household identifying number, which is the unit of observation in our data. Therefore, in the analysis we use the terms "customer" and "household" interchangeably.

[2]Our data do not include information on purchases by customers not carrying a loyalty card of the chain. The focus of this study, however, is on "regular" customers who can be meaningfully said to be part of the customer base of a firm. This seems to be more the case for individuals who sign up for a loyalty card than for occasional shoppers who do not do that.

cial to observe the evolution of the customer base of the supermarket chain. Furthermore, given the significant market share held by the chain -it ranks consistently among the top supermarket chains countrywide, its wide geographical spread and the fact that it offers all the mainstream packaged good products, we can think of the behavior of its customers as representative of the population of retail shoppers at large. We use the information included in the data to construct the variables needed to measure the comovement between the customer's decision to exit the customer base and the price of her typical basket of goods posted at the chain: (i) an indicator signaling when the household is exiting the chain's customer base, and (ii) the price of the household basket. Below we briefly describe the procedure followed to obtain them; the details are left to Appendix A.

We consider every customer shopping at the retailer in a given week as belonging to the chain's customer base in that week. We assume that a household has exited the customer base when she has not shopped at the chain for eight or more consecutive weeks. The decision to exit is imputed to the last time the customer visited the chain. Although brief spells without purchases can be justified with alternative explanations (e.g. consuming inventory or going on vacation), the typical customer is unlikely to experience an eight-week spell without shopping for reasons other than having switched to a different chain. In fact, for the average household in our sample, four days elapse between consecutive grocery trips and the 99th percentile of this statistic is 28 days, half the length of the absence we require before inferring that a household is buying its groceries at a competing chain. This suggests that the eight-week window is a conservative choice.[3]

We construct the price of the basket of grocery goods usually purchased by the households in the following fashion. We identify the goods belonging to a household's basket using scanner data on items the household purchased over the two years in the sample. In a particular week t, the price paid by customer i, shopping at store j for its basket, represented by the collection of UPC's in K_i, is

$$p_{ijt} = \sum_{k \in K_i} \omega_{ik} \, p_{kjt} \;, \qquad \omega_{ik} = \frac{\sum_t E_{ikt}}{\sum_{k \in K_i} \sum_t E_{ikt}} \;, \tag{1}$$

where p_{kjt} is the price of UPC k in week t at the store j where customer i shops, and E_{ikt} is the expenditure (in dollars) by customer i in UPC k in week t. Note that the price of the basket is household specific because households differ in their choice of grocery products (K_i) and in the weight such goods have in their budget (ω_{ik}). We face the common problem

---

[3]We experimented with 4 weeks and 12 weeks as alternative lengths of the period of absence required to infer the exit from the customer base. In both cases the results are qualitatively similar. However, in the 12-week case the number of exit events becomes too small and we do not have the power to detect significant effects.

that household scanner data only contain information on prices and quantities of UPCs when they are actually purchased. Therefore, we complement them with store level data on weekly revenues and quantities sold.[4] This data allows us to back out weekly prices of each UPC in the sample by dividing total revenues by total quantity sold as in Eichenbaum et al. (2011).

2.2 Evidence on customer base dynamics

The availability of individual level scanner data allows us to study the determinants of a customer's decision to exit the customer base of the firm she is currently shopping from. We estimate a linear probability model where the dependent variable is an indicator for whether the customer has left the customer base of the chain in a particular week. Our aim is to capture the effect of the price posted by the chain for the basket of goods purchased by the customer on her decision to exit. In Table 1, we report results of regressions of the following form,

$$Exit_{it} = b_0 + b_1 \log(p_{it}) + b_2 \log(p_{it}^{mkt}) + b_3 tenure_{it} + X_i'c + \varepsilon_{it} \ . \tag{2}$$

Our main interest is on the coefficient of the retailer price of the basket, b_1. Existing theories on the link between prices and customer dynamics (Phelps and Winter (1970)) stress that a firm's ability to retain its customers should be influenced by its idiosyncratic price variations but not from aggregate shocks that move the competitors' prices as well. To isolate idiosyncratic price variations, we control for the prices posted by the competitors in the same market of the chain using information from the IRI Marketing data set.[5] This data allows us to compute the average price of the basket in a market for every customer (p_{it}^{mkt}). To further control for sources of aggregate variation, we include in the regression year-week fixed effects that account for time-varying drivers of the decision of exiting the customer base common across households (e.g., disappearances due to travel during holiday season).

The coefficient b_1 is then identified by *UPC-chain* specific shocks as those triggered, for example, by the expiration of a contract between the chain and the manufacturer of a UPC. We also observe the price of a same good moving differently in different stores within the chain, for instance due to variation in the cost of supplying the store linked to logistics (e.g. distance from the warehouse). Since these shocks can hit differently goods with different intensity in delivery cost (e.g. refrigerated vs. nonrefrigerated goods), *UPC-store* specific shocks also contribute to our identification. It is important to notice that we do not need

---

[4]The retailer changes the price of the UPCs at most once per week, hence we only need to construct weekly prices to capture the entire time variation.

[5]A detailed description of the data can be found in Bronnenberg et al. (2008). All estimates and analyses in this paper based on Information Resources Inc. data are by the authors and not by Information Resources Inc. We provide additional details on the IRI data and on the construction of the price index for the competitors of the chain in Appendix A.

to assume that such shocks will make a supermarket uniformly more expensive than the competition. Shocks that affect the convenience of a chain with respect to a subset of goods suffice to induce the customers who particularly care about those goods to leave. Kaplan and Menzio (forthcoming,a) use independent scanner data to provide ample evidence for this source of variation. They report that the bulk of price dispersion arises not from the difference from high-price and low-price stores but from dispersion in the price of a particular good (or product category) even among stores with similar overall price level.

The retailer price in equations (2) can be endogenous if the chain conditions to variables unobserved to the econometrician that also influence the customer's decision to leave. We use a measure of cost provided by the retailer to construct an instrument for the price of the basket. This measure represents a replacement cost, i.e. the cost for the retailer of restocking the product. It includes the wholesale price but also other costs associated with logistics (delivery to the store, etc.). Eichenbaum et al. (2011) use the same measure and treat it as a good approximation of the retailer's marginal cost. We use it to calculate the cost of the basket, obtained as the weighted average of the replacement cost of the UPCs included in it. Its calculation is analogous to that described in equation (1) to obtain the price of the basket.

We include in our specification observable characteristics (age, income, and education) matched from Census 2000 and also consider location as a potential driver of the decision to exit. We control for the number of supermarket stores in the zip code of residence of the customer and factor her convenience in shopping by calculating the distance in miles between her residence and the closest store of the chain and the closest alternative supermarket. To pick up the heterogeneity in the type of goods different customers include in their basket, we control for the price volatility of the customer-specific basket and for its price in the first week in the sample, as a scaling factor. Finally, we calculate customer tenure, defined as the number of consecutive weeks the customer has spent in the customer base of the chain, and include it in the regression to account for the fact that long-term customers of the chain may be less willing to leave it *ceteris paribus*. All of these controls aim at absorbing heterogeneity across customers and ensure that it does not bias our estimate of the elasticity of the probability of exit to prices. This effort is important as we will later use this estimate to quantify a model where customers are heterogeneous only in their search cost.

The main specification in column (1) shows that the basket price posted by the retailer significantly impacts the probability of leaving. The effect is also quantitatively important. The average probability of exiting the customer base (0.3% weekly) implies a yearly turnover of 14%; if the retailer's prices were 1% higher, its yearly turnover would jump to 21%. The coefficient on the competitors' price, which we would expect to enter with a negative sign, is

Table 1: Effect of the price of the basket on the probability of exiting the customer base

Exiting: Missing at least 8 consecutive weeks

	(1)	(2)	(3)	(4)
$\log(p)$	0.14**	-0.01	0.16*	0.15**
	(0.071)	(0.030)	(0.089)	(0.070)
Walmart entry			0.019*	
			(0.011)	
$\log(p^{mkt})$	0.001	0.000		0.001
	(0.001)	(0.001)		(0.001)
Tenure	-0.002***	-0.003***	-0.004***	-0.002***
	(0.001)	(0.000)	(0.000)	(0.000)
Observations	52,670	52,670	66,182	52,101

Notes: An observation is a household-week pair. The results reported are calculated through two-stages least squares where we use the logarithm of the cost of the basket (constructed based on the replacement cost provided for each UPC by the retailer) as instrument for the logarithm of the price of the basket. In column (2), the price of the household basket is substituted with a price index for the store overall. In column (4), the exit of the customer is assumed to have occurred in the first week of absence in the eight (or more) weeks spell without purchase at the chain rather than the week of the last shopping trip before the hiatus. We trim from the sample households in the top and bottom 1% in the distribution of the number of trips over the two years. Coefficients on a series of variables are not reported for brevity: demographic controls matched from Census 2000 (ethnicity, family status, age, income, education, and time spent commuting) as well as distance from the closest outlet of the supermarket chain and distance from the closest competing supermarket (provided by the retailer). The logarithm of the price of the household basket in the first week in the sample and the standard deviation of changes in the log-price of the household basket over the sample period are included as a controls in all specifications. Week-year fixed effects are also always included. Robust standard errors are in parenthesis. ***: Significant at 1% **: Significant at 5% *: Significant at 10%.

not significant. This may be due to the fact that the IRI data only allow us to imperfectly capture competitors' behavior. In fact, the IRI dataset contains price information only on a subset of the goods included in a customer's basket, although it arguably covers all the major product categories. Furthermore, the IRI data do not contain detailed information on the location of the outlets. This introduces measurement error in our construction of the set of stores a customer considers as options for her shopping. The negative coefficient on tenure confirms the intuition that the longer the relationship between a firm and a customer, the less likely they are to be interrupted. Among the several individual characteristics we control for, it is worth mentioning that distance from stores of the chain and distance from the closest competing store enter with the expected sign. Customers living in proximity of a store of the chain are less likely to leave it, and those living closer to competitors' stores are more inclined to do so.

In columns (2)-(4) we assess the robustness of these findings. We start by replacing the

price of the individual basket with a price index for the store basket, defined as the average of the prices of the UPC's sold by the store where the customer buys, weighted for their sales. Such price is, by construction, equal for all the customers shopping in the same store. Column (2) shows that this results in a price coefficient that is negative and not significant, confirming the importance of being able to construct individual specific baskets in order to make inference on customer's behavior.

In column (3), we experiment with an alternative way to control for the effect of competition: we exploit episodes of entry by Walmart, a major retailer with which our chain is in direct competition. We use data from Holmes (2011) to identify the date of entry by a Walmart supercenter, i.e. a store selling groceries on top of general discount goods, in a zip code where the retailer we study also operates a supermarket. The resulting event study allows us to measure the effect of the retailer price on the probability of exit controlling for the most relevant change in the competitive environment. The coefficient obtained falls in the same ballpark as the estimate in the main specification, which is reassuring on the effectiveness of the IRI price in measuring the competitors' behavior.

In column (4), we change the assumption on the imputation of the date of exit. Rather than assuming that the customer left on the occasion of her last trip to the store, we posit that the exit occurred in the first week of her absence. Even in this case, the main result stays unaffected. Finally, we performed a placebo test to investigate whether it is possible to obtain results with the same level of significance of our main specification out of pure chance. We estimated our main specification 1,000 times each time with a different dependent variable where exits from the customer base, while kept constant in number, are randomly assigned to customers. We find that only in 2.8% of the cases the simulation yields and price coefficient are positive and significant at 5%.

3   The model

The economy is populated by a measure one of firms producing an homogeneous good and a measure Γ of customers who consume it. The economy is in steady state and there are no foreseen aggregate shocks.

3.1   The problem of the firm

In this section we describe the pricing problem of the firm. Firms produce a homogenous good and are indexed by j. They choose each period the price p of the good they produce. We assume a linear production technology $y^j = z^j \ell^j$ where ℓ is the production input, and

z^j is the firm-specific productivity. Idiosyncratic productivity is distributed according to a conditional cumulative distribution function $F(z'|z)$ with bounded support $[\underline{z}, \bar{z}]$. We also assume that $F(z'|z_h)$ first order stochastically dominates $F(z'|z_l)$ for any $z_h > z_l$ to induce persistence in firm productivity. The profit per customer accrued to the firm is given by $\pi(p, z) \equiv d(p)(p - w/z)$, where the constant $w > 0$ denotes the marginal cost of the input ℓ, and the function $d(\cdot)$ is a downward sloping demand function.[6] We assume that profits per customer are single-peaked in p.

We denote by m_{t-1}^j the *customer base* of firm j at the beginning of period t, which consists of the mass of customers who bought from firm j in period $t-1$. The customer base of each firm j evolves according to a law of motion determined in equilibrium and that we conjecture, and later verify, is given by $m_t^j = F(m_{t-1}^j, p_t^j, z_t^j)$. The mass of customers buying from firm j in the current period depends on the price and productivity of the firm in the current period, as well as on the mass of customers who bought from it in the previous period.

We study a stationary Markov Perfect equilibrium where pricing strategies are a function of the current state. As there are no aggregate shocks, the aggregate state is constant and the relevant state for firm j problem in period t is the pair $\{z_t^j, m_{t-1}^j\}$. The firm pricing problem in its recursive form solves

$$\tilde{W}(z_t^j, m_{t-1}^j) = \max_p \; m_t^j \, \pi(p, z_t^j) \; + \; \beta \int_{\underline{z}}^{\bar{z}} \tilde{W}(z', m_t^j) \, dF(z'|z_t^j) \tag{3}$$

$$\text{s.t.} \quad m_t^j = F(m_{t-1}^j, p_t^j, z_t^j) \, , \tag{4}$$

where $\tilde{W}(z_t^j, m_{t-1}^j)$ denotes the firm value at the optimal price. The price impacts firm value through two channels. First, it affects the level of profits per customer as in standard models of firm pricing. Given our assumption of single-peakedness of the profit function $\pi(p, z)$, there is a unique level of p that maximizes the profits per customer. Second, the price p affects the dynamics of the customer base. It does so by influencing the mass of customers buying from the firm in the current period m_t^j, and, if there is persistence in the evolution of the customer base, by affecting the mass of customers buying from the firm in future periods. As a result, the pricing problem of the firm is dynamic in nature.

We study an environment where there is persistence in the customer base, as in Phelps and Winter (1970) and Rotemberg and Woodford (1999). In these models an ad-hoc functional form for the evolution of the customer base is assumed, where the mass of customers served by a firm is given by the product of its original customer base and a growth rate, which depends on its (relative) price. Our law of motion for the customer base preserves this

---

[6]In Section 5, we extend this framework adding a model of the labor market to endogenize the wage w.

standard structure and is given by:

$$F(m_{t-1}^j, p_t^j, z_t^j) \equiv m_{t-1}^j \, \Delta(p_t^j, z_t^j) \tag{5}$$

This similarity notwithstanding, there are two important innovations that we introduce. First, we generalize the law of motion so that it can depend not only on the price the firm sets but also on its productivity. This extension allows to use our model to study the equilibrium price pass-through of idiosyncratic cost shocks. It also proves useful when we bring the model to the data, since having heterogeneity in productivity helps us matching the cross-sectional variation in prices. Second, while the customer base evolution is typically characterized with ad-hoc functional form assumptions, the form of our $\Delta(\cdot)$ function depends on the equilibrium distribution of prices as well as on the distributions of productivity and search costs. Accounting for this dependence matters for the estimation where we will match micro moments obtained from customers' decisions. Moreover, it has important implications when using the model for policy experiments, as we will illustrate with the application in Section 7. However, we do share an important feature with classic customer market models: the growth rate of the customer base does not depend on the initial mass of customers. This property allows for a substantial simplification of the firm's problem. In particular, it can be obtained that the value function of a firm is homogeneous of degree one in m, i.e. $\tilde{W}(z,m) = m \, \tilde{W}(z,1) \equiv m \, W(z)$, where using equation (3) and $m_t^j = m_{t-1}^j \, \Delta(p_t^j, z_t^j)$, it is immediate to show that $W(z_t^j)$ solves[7]

$$W(z_t^j) \; = \; \max_p \; \Delta(p, z_t^j) \, \pi(p, z_t^j) + \beta \, \Delta(p, z_t^j) \int_{\underline{z}}^{\bar{z}} W(z') dF(z' \,|\, z_t^j) \,. \tag{6}$$

The relevant state to the firm pricing problem is its productivity, as the level of the customer base affects multiplicatively the firm value. The solution to the firm problem in equation (6) gives an optimal pricing strategy that depends on productivity and we denote by $\hat{p}(z)$. We emphasize that, while the initial level of the customer base does not affect the optimal price, its evolution does. This follows as a change in the price affects the growth rate of the customer base, i.e., the value of $\Delta(p,z)$, and given the persistence of the customer base, this affects the firm value in the current period as well as in future periods. As a result, the optimal price of a firm will be affected by customer base growth considerations, which are driven by variation in idiosyncratic productivity.

Specifically, the objective of the firm maximization problem can be expressed as the

---

[7]Under the assumption that the discount rate β is low enough so that the maximization operator in equation (6) is a contraction, by the contraction mapping theorem we can conclude that our conjecture about homogeneity of $\tilde{W}(z,m)$ is verified.

product of two terms, $W(z) \equiv \Delta(\hat{p}(z), z) \Pi(\hat{p}(z), z)$, where $\Pi(p, z)$ denotes the expected present discounted value of each customer to the firm. Under the assumption that the functions $\Delta(p, z)$ and $\pi(p, z)$ are differentiable in p, the first order condition to the firm problem is given by

$$\frac{\partial \Pi(p, z)}{\partial p} \frac{p}{\Pi(p, z)} = -\frac{p}{\Delta(p, z)} \frac{\partial \Delta(p, z)}{\partial p} \ . \tag{7}$$

We will discuss conditions under which equation (7) is necessary and sufficient in Section 4. The function $\Pi(p, z)$ is maximized at the static profit maximizing price,

$$p^*(z) \equiv \frac{\varepsilon_d(p)}{\varepsilon_d(p) - 1} \frac{w}{z} \ , \tag{8}$$

where $\varepsilon_d(p) \equiv \partial \log(d(p))/\partial \log(p)$ denotes the *intensive margin* elasticity of demand. Equation (7) implies however that, due to concerns about customer dynamics, the optimal price is in general different from the one that maximizes static profits. The following lemma discusses some key properties of the optimal price.

Lemma 1 *Let $\Delta(p, z)$ be continuous in p, and let $\varepsilon_m(p, z) \equiv \partial \log(\Delta(p, z))/\partial \log(p)$. If a price $\bar{p}(z)$ exists such that $\varepsilon_m(p, z) > 0$ for all $p > \bar{p}(z)$, and $\varepsilon_m(p, z) = 0$ for all $p \leq \bar{p}(z)$, then we have $\hat{p}(z) \in [\bar{p}(z), p^*(z))$ if $\bar{p}(z) < p^*(z)$, and $\hat{p}(z) = p^*(z)$ otherwise.*

The proof is an immediate implication of equation (7). If the growth in the customer base is non-decreasing in the price, equation (7) implies that setting a price above the static profit maximizing price is never optimal. Hence, $\hat{p}(z) \leq p^*(z)$ for all z. Moreover, if the growth in the customer base is strictly decreasing in the price in a neighborhood of the static profit-maximizing price $p^*(z)$, the optimal price is distorted downwards with respect to it, i.e. $\hat{p}(z) < p^*(z)$. The first order condition in equation (7) illustrates the trade-off the firm faces when setting the price in a region where customer retention is a concern. When $\bar{p}(z) < p^*(z)$, the optimal price balances the marginal benefit of an increase in price (more profit per customer) with the cost (decrease in the customer base). The requirement that the solution to the firm problem must satisfy the first order condition implies that we study equilibria where the firm objective, and in particular $\Delta(p, z)$, is differentiable in p. In Section 4, we will derive the necessary equilibrium properties that guarantee that the assumptions of Lemma 1 are satisfied.

3.2 The problem of the customer

In this section we evaluate the problem of the customers.

Let $v(p)$ denote the customer surplus as a function of the price p of the good, associated to the demand function $d(p)$. We assume that $v(p)$ is continuously differentiable with $v'(p) < 0$, and bounded above with $\lim_{p \to 0^+} v(p) < \infty$. These properties are satisfied in standard models of consumer demand.

Each customer starts period t in the customer base of the firm she bought from in period $t-1$. The customer observes perfectly the state of the firm she is matched to, in particular she observes its productivity. Given the equilibrium pricing function of the firm, this allows her to assess the probability distribution of the path of prices of that firm. After observing the state of her current match, the customer decides whether she wants to pay a search cost to draw another firm. The search cost $\psi^i \geq 0$ is measured in units of customer surplus, it is idiosyncratic to each customer i and it is drawn each period from a cumulative distribution $G(\psi)$, with an associated density $g(\psi)$. For tractability, we restrict our attention to density functions that are continuous on all the support. Heterogeneity, albeit transitory, in search costs makes the customer base a continuous function of the price and allows us to study firms' pricing decisions that are not necessarily knife-edge in the trade-off between maximizing demand and markups.

The customer can search at most once per period. Search is random, with the probability of drawing a particular firm j^{new} being proportional to its customer base, i.e. $m_{t-1}^{j^{new}}/\Gamma$. As in Fishman and Rob (2005), this assumption captures the idea that consumers search new suppliers not by randomly sampling firms but by randomly sampling other consumers, and then visiting the stores that these frequent. On the technical side, this is the key assumption that will allow us to obtain $m_t^j = m_{t-1}^j \Delta(p_t^j, z_t^j)$. Conditional on searching, the customer observes the state of the new match and takes then another decision concerning whether to exit the customer base of her initial firm and match to the new firm. In particular, the customer compares the distribution of the path of current and future prices at the two firms and buys from the firm offering higher expected value. Finally, we assume that a customer cannot recall a particular store once she exits its customer base. Figure 1 summarizes timing and payoffs of the problem of the customer.

We next characterize the customer problem. Let $V(p_t^j, z_t^j, \psi_t^i)$ denote the value function of a customer i who has drawn a search cost ψ_t^i and is matched to firm j, which has current productivity z_t^j and posted price p_t^j. This value function solves the following problem,

$$V(p_t^j, z_t^j, \psi_t^i) = \max \left\{ \bar{V}(p_t^j, z_t^j), \, \hat{V}(p_t^j, z_t^j) - \psi_t^i \right\} , \tag{9}$$

where $\bar{V}(p, z)$ is the customer's value if she does not search, and $\hat{V}(p, z) - \psi$ is the value if

14

Figure 1: The problem of a customer matched with a firm with productivity z in period t
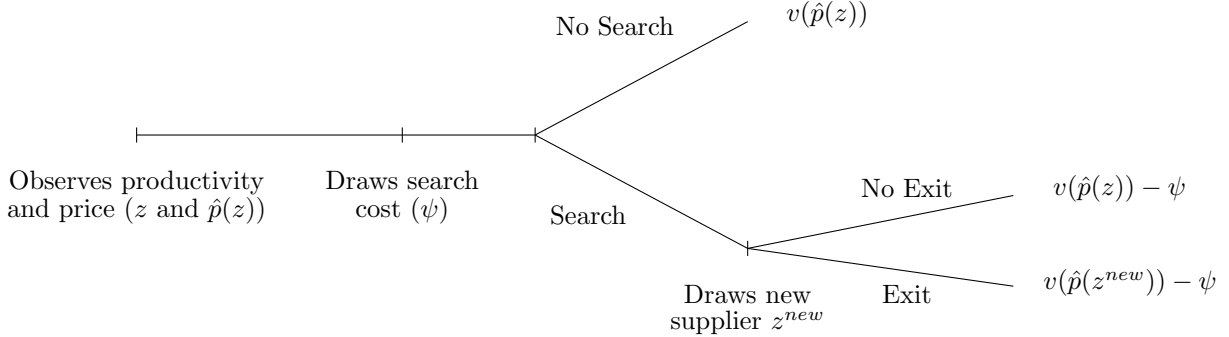
she does search. The value in the case of not searching is

$$\bar{V}(p_t^j, z_t^j) = v(p_t^j) + \beta \int_0^\infty \int_{\underline{z}}^{\bar{z}} V(\hat{p}(z'), z', \psi') \, dF(z'|z_t^j) \, dG(\psi') \ . \tag{10}$$

We notice that the state of the firm problem depends on the on the productivity z because the pricing function $\hat{p}(\cdot)$ mapping future productivity into prices in the Markov equilibrium makes productivity z a sufficient statistic for the distribution of future prices at the firm.[8] The value when searching is given by

$$\hat{V}(p_t^j, z_t^j) = \int_{-\infty}^{+\infty} \max\left\{\bar{V}(p_t^j, z_t^j), \, x\right\} dH(x) \ , \tag{11}$$

where the customer takes expectations over all possible draws of potential new firms, and where $H(\cdot)$ is the equilibrium cumulative distribution of continuation values from which the firm draws a new potential match when searching. For instance, $H(\bar{V}(p_t^j, z_t^j))$ is the probability of drawing a potential match offering a continuation value smaller than or equal to the current match. The following lemma describes the customer's optimal search and exit policy rules.

Lemma 2 *The customer matched to a firm with productivity z_t^j charging price p_t^j: i) searches if she draws a search cost $\psi_t \le \hat{\psi}(p_t^j, z_t^j)$, where $\hat{\psi}(p, z) \equiv \int_{\bar{V}(p,z)}^{\infty} \left(x - \bar{V}(p, z)\right) dH(x) \ge 0$ is the threshold to search; ii) conditional on searching, exits if she draws a new firm promising a continuation value $\bar{V}^{new}$ larger than the current match, i.e. $\bar{V}^{new} \ge \bar{V}(p_t^j, z_t^j)$.*

---

[8]We also notice that the state of the firm problem includes the current price p, despite in equilibrium productivity is enough to determine the current price, as this notation is needed to study the game between the firm and its customers where the firm could, in principle, deviate from the equilibrium price.

The proof of the lemma follows immediately from equations (9)-(11). The lemma states that, as search is costly, not all customers currently matched to a given firm exercise the search option, only those with a low search cost do so. Notice that the threshold $\hat{\psi}(p,z)$ depends both on the price of the firm, p, and its productivity, z. The dependence on the price is straightforward, following from its effect on the surplus $v(p)$ that the customer can attain in the current period. The intuition behind the dependence on the firm's productivity is that, as searching is a costly activity, the decision of which firm to patronize is a dynamic one, and involves comparing the value of remaining in the customer base of the current firm with the value of searching. Because of the Markovian structure of prices, the customer's expectation about future prices is completely determined by the firm's current productivity.

We are now ready to describe the equilibrium dynamics of the customer base as a function of price and productivity, given the optimal search and exit strategy of the customers. The customer base of firm j evolves as follows:

$$m_t^j = m_{t-1}^j - \underbrace{m_{t-1}^j G\left(\hat{\psi}(p_t^j, z_t^j)\right)\left(1 - H(\bar{V}(p_t^j, z_t^j))\right)}_{\text{customers outflow}} + \underbrace{\frac{m_{t-1}^j}{\Gamma} Q\left(\bar{V}(p_t^j, z_t^j)\right)}_{\text{customers inflow}}, \quad (12)$$

where $G(\hat{\psi}(p_t^j, z_t^j))$ is the fraction of customers searching from firm j customer base, a fraction $1 - H(\bar{V}(p_t^j, z_t^j))$ of which actually finds a better match and exits the customer base of firm j. The ratio m_{t-1}^j/Γ is the probability that searching customers in the whole economy draw firm j as a potential match. The function $Q(\bar{V}(p_t^j, z_t^j))$ denotes the equilibrium mass of searching customers currently matched to a firm with continuation value smaller than $\bar{V}(p_t^j, z_t^j)$. Therefore, the product of the two amounts to the mass of new customers entering the customer base of firm j. We can express the dynamics of the customer base as $m_t^j = m_{t-1}^j \Delta(p_t^j, z_t^j)$, where the function $\Delta(\cdot)$ is given by

$$\Delta(p, z) \equiv 1 - G\left(\hat{\psi}(p, z)\right)\left(1 - H(\bar{V}(p, z))\right) + \tfrac{1}{\Gamma} Q\left(\bar{V}(p, z)\right). \quad (13)$$

Notice that the growth of a firm is independent of its customer base and, therefore, of its size. This verifies the conjecture about the customer base dynamics made in Section 3.1.

4 Equilibrium

In this section we define an equilibrium, discuss its existence, and characterize its general properties. We start by defining the type of equilibrium we study.

Definition 1 *Let $\mathcal{V}(z) \equiv \bar{V}(\hat{p}(z), z)$ and $p^*(z)$ be given by equation (8). We study stationary Markovian equilibria where $\mathcal{V}(z)$ is non-decreasing in z. A stationary equilibrium is then*

(i) *a search and an exit strategies that solve the customer problem for given equilibrium pricing strategy $\hat{p}(z)$, as defined in Lemma 2;*

(ii) *a firm pricing strategy $\hat{p}(z)$ that solves equation (7) for each z, given customers dynamics summarized by the function $\Delta(p, z)$ in equation (13);*

(iii) *an invariant distribution $K(\cdot)$ determining the mass of customers matched to firms with productivity smaller than z at the beginning of the period, where $K(z)$ solves*

$$K(z) = \int_{\underline{z}}^{z} \int_{\underline{z}}^{\bar{z}} \Delta(\hat{p}(x), x) \, dF(s|x) \, dK(x) \;, \tag{14}$$

for each $z \in [\underline{z}, \bar{z}]$ with boundary condition $\int_{\underline{z}}^{\bar{z}} dK(x) = 1$;

(iv) *two invariant distributions, $H(\cdot)$ and $Q(\cdot)$, that solve*

$$H(x) = K(\hat{z}(x)) \quad and \quad Q(x) = \Gamma \int_{\underline{z}}^{\hat{z}(x)} G(\hat{\psi}(\hat{p}(z), z)) \, dK(z) \;,$$

for each $x \in [\mathcal{V}(\underline{z}), \mathcal{V}(\bar{z})]$, where $\hat{z}(x) = \max\{z \in [\underline{z}, \bar{z}] : \mathcal{V}(z) \leq x\}$.

We study equilibria where the continuation value to customers is non-decreasing in productivity, implying that customers' rank of firms coincides with their productivity. This is a natural outcome as more productive firms are better positioned to offer lower prices and therefore higher values to customers.

The following proosition provides results regarding the endogenous evolution of the customer base.

Proposition 1 *Search behavior of customers and its implications for customer base growth:*

(1) *The value function $\bar{V}(p, z)$ is strictly decreasing in p and increasing in z. The threshold $\hat{\psi}(p, z)$ is strictly increasing in p and decreasing in z.*

(2) *Let $\bar{p}(z)$ denote the price p that solves $\bar{V}(p, z) = \mathcal{V}(\bar{z})$. We have that $\Delta(p, z)$ is continuous, strictly decreasing in p for all $p > \bar{p}(z)$, and constant for all $p \leq \bar{p}(z)$. Also, $\Delta(p, z)$ is increasing in z.*

Part (1) of the proposition characterizes the equilibrium search behavior of customers, while part (2) showcases its implications for the evolution of the customer base. The proof of part (1) follows directly from the assumption of $v(p)$ being strictly decreasing in p.[9] This part of the proposition states that customers obtain strictly higher value from firms offering a lower current price and, if $\mathcal{V}(z)$ is increasing in z, also from firms characterized by higher current productivity. Notice that, under persistence in the productivity process, a sufficient condition for the latter is that equilibrium prices are decreasing in productivity. As a result, customers are not only more likely to search and exit from firms charging higher prices, but they are also more likely to do so from firms with lower productivity. Part (2) of the proposition follows directly from part (1) and states that the growth of the customer base is decreasing in the current price because a higher price reduces the current surplus and therefore the value of staying matched to the firm. When the price is low enough that no firm in the economy offers a higher value to the customer, the customer base is maximized and a further decrease in the price has no impact on the customer growth. The growth of the customer base increases with firm productivity, as a larger z is associated to higher continuation value which increases the value of staying matched to the firm. Also, notice that the second part of the proposition verifies that the assumptions of Lemma 1 are satisfied in equilibrium, so that optimal prices diverge from those obtained from a setup without customer base dynamics.

The next proposition states conditions under which the equilibrium that we evaluate exists and characterizes its properties.

Proposition 2 *Let productivity be i.i.d. with $F(z'|z_1) = F(z'|z_2)$ continuous and differentiable for any z' and any pair $(z_1, z_2) \in [\underline{z}, \bar{z}]$, and let $G(\psi)$ be differentiable for all $\psi \in [0, \infty)$, with $G(\cdot)$ differentiable and not degenerate at $\psi = 0$. There exists an equilibrium as in Definition 1 where $\hat{p}(z)$ satisfies equation (7), and*

(i) $\hat{p}(z)$ is strictly decreasing in z, with $\hat{p}(\bar{z}) = p^*(\bar{z})$ and $\hat{p}(\bar{z}) < \hat{p}(z) < p^*(z)$ for $z < \bar{z}$, implying that $\mathcal{V}(z)$ is strictly increasing. Moreover, the optimal markups are given by

$$\mu(p, z) \equiv \frac{p}{w/z} = \frac{\varepsilon_d(p)}{\varepsilon_d(p) - 1 + \varepsilon_m(p, z)\,\Pi(p, z)/(d(p)\,p)}\,, \tag{15}$$

where $p = \hat{p}(z)$ for each z.

(ii) $\hat{\psi}(\hat{p}(z), z)$ is strictly increasing in z, with $\hat{\psi}(\hat{p}(\bar{z}), \bar{z}) = 0$ and $\hat{\psi}(\hat{p}(z), z) > 0$ for $z < \bar{z}$, implying that $\Delta(\hat{p}(z), z)$ is strictly increasing, with $\Delta(\hat{p}(\bar{z}), \bar{z}) > 1$ and $\Delta(\hat{p}(\underline{z}), \underline{z}) < 1$.

---

[9]From the assumption we obtain that $\bar{V}(p, z)$ is decreasing in p. If $\mathcal{V}(z)$ is increasing in z, $\bar{V}(p, z)$ increases with z because of persistence in the productivity process. Finally, notice that $\frac{\hat{\psi}(p,z)}{\partial p} = -v'(p)\,(1 - H(V(p, z))) \geq 0$ and that $\frac{\partial \hat{\psi}(p,z)}{\partial z} = -\frac{\partial \bar{V}(p,z)}{\partial z}\,(1 - H(V(p, z))) \leq 0$.

The proof of the proposition can be found in Appendix B.1. Here we just point out that, while the results of Proposition 2 refer to the case of i.i.d. productivity shocks, numerical results in Section 6 show they hold even in the case of persistent productivity processes.

We now comment on the properties of the equilibrium highlighted in the proposition. The equilibrium is characterized by price dispersion: this is important, as price dispersion is what motivates customers to search. Price dispersion is driven by heterogeneity in firm productivity, as in Reinganum (1979), and by the level and dispersion of search frictions.[10] More productive firms charge lower prices and, therefore, offer higher continuation value to customers. If all the firms had the same productivity, Proposition 2 would imply a unique equilibrium where the price is that maximizing static profits, $p^*(\tilde{z})$, and as a result the customer base of every firm would be constant.[11] The equilibrium is also characterized by dispersion in customer base growth: more productive firms grow faster, and there is a positive mass of lower productivity firms that have a shrinking customer base and a positive mass of higher productivity firms that are expanding their customer base.

Optimal markups in equation (15) depend on three distinct terms: $\varepsilon_d(p)$, $\varepsilon_m(p, z)$, and $x(p, z) \equiv \Pi(p, z)/(d(p) \, p)$. The terms $\varepsilon_d(p)$ and $\varepsilon_m(p, z)$ represent the price elasticities of quantity purchased (per-customer) and of customer growth, respectively. We notice that the elasticity of total firm demand to the price, i.e. $m \, \Delta(p, z) \, d(p)$, is given by $\varepsilon_d(p) + \varepsilon_m(p, z)$. An increase in price reduces total current demand both because it reduces quantity per customer (*intensive margin effect*) and because it reduces the number of customers (*extensive margin effect*). Moreover, the optimal markup solves a dynamic problem as a loss in customers has persistent consequences for future demand due to the inertia in the customer base. This dynamic effect is captured by the term $x(p, z)$, which measures the firm present discounted value of a customer scaled by the current revenues. It follows that active customer markets are associated with a strictly lower markup than the one that maximizes static profit; the lower, the larger the product $\varepsilon_m(p, z) \, x(p, z)$.

Finally, the next remark explores two interesting limiting cases of our model and showcases the effect of search frictions on price dispersion.

Remark 1 *Let search costs be scaled as $\psi \equiv n \, \tilde{\psi}$, where $n > 0$. That is, let the value*

---

[10]For tractability we will abstract from (possible) equilibria where symmetric firms charge different prices.

[11]This special case is useful to understand our relation to Diamond (1971), which shows in a simple search model that the resulting equilibrium when search cost is positive exhibits no price dispersion and firms behaving as monopolists. Our model delivers a different outcome because we allow firms to differ in idiosyncratic productivity but can generate the Diamond (1971) results if heterogeneity in productivity is shut down.

function in equation (9) be

$$V(p, z, \psi) = \max \left\{ \bar{V}(p, z), \, \hat{V}(p, z) - n\tilde{\psi} \right\} .$$

Two limiting cases of the equilibrium stated in Definition 1:

(1) *Let* $n \to \infty$. *Then, in equilibrium: (i) the optimal price maximizes static profits, i.e.* $\hat{p}(z) = p^*(z)$ *for all* $z \in [\underline{z}, \bar{z}]$, *and (ii) there is no search in equilibrium.*

(2) *Let* $\pi(p^*(\bar{z}), \underline{z}) > 0$ *and let the assumptions of Proposition 2 be satisfied. Then,* $\hat{p}(\bar{z}) = p^*(\bar{z})$ *and* $\max\{\hat{p}(z)\} = \hat{p}(\underline{z})$ *approaches* $p^*(\bar{z})$ *as* $n \to 0$. *As a result, in the limit, there is no price dispersion in equilibrium and customers do not search.*

A proof of the remark can be found in Appendix B.2. The first limiting case explores the equilibrium when we let search costs diverge to infinity. The model then reduces to one where customer base concerns are not present. Because the customer base is unresponsive to prices, the firm problem reverts to a standard price-setting problem commonly studied in the macroeconomics literature: the firm sets the price p, taking into account only its impact on static demand $d(p)$. In equilibrium, optimal prices maximize static profits, i.e. $\hat{p}(z) = p^*(z)$ for all $z \in [\underline{z}, \bar{z}]$. There is price dispersion, and there is no search in equilibrium. The second limiting case explores the equilibrium when search costs become arbitrarily small. We restrict attention to the model that satisfies the assumptions of Proposition 2, so that the first order condition presented in equation (7) is necessary for optimality.[12] In this case, as the scale of search costs becomes arbitrarily small, equilibrium prices approach the static profit maximizing price of the most productive firm, $\hat{p}(\bar{z})$. As a result, there is no price dispersion and customers do not search.

5 Parametrization and analysis of the model

In this section we discuss the procedure followed to estimate the model. More details are provided in Appendix C.

To parametrize our model, we need to choose the discount factor (β) as well as four functions: the demand function, $d(p)$, the surplus function $v(p)$, the distribution of search costs $G(\psi)$, and the conditional distribution of productivity $F(z'|z)$. We assume that a period in the model corresponds to a week to mirror the frequency of our data. We fix the firm

---

[12]The assumption $\pi(p^*(\bar{z}), \underline{z}) > 0$ is purely technical, and it ensures that the first derivative of the profit function is bounded in the relevant range.

discount rate to $\beta = 0.995$. In the set of parameters that we consider, this level of β ensures that the max-operator in equation (6) is a contraction.[13]

We assume that customers have logarithmic utility in consumption. Consumption is defined as a composite of two types of goods $c \equiv \left[\alpha^{\frac{1}{\theta}} d^{\frac{\theta-1}{\theta}} + (1-\alpha)^{\frac{1}{\theta}} n^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}$, with $\theta > 1$ governs the demand elasticity and $\alpha \in (0, 1)$ is the weight of good d in the utility function.[14] The first good (that we label d) is supplied by firms facing product market frictions as described in Section 3; the other good (n) acts as a numeraire and it is sold in a frictionless centralized market. The sole purpose of good n is to microfound a downward sloping demand $d(p)$ and, therefore, to allow for an intensive margin of demand. We set $\alpha = 1/2$ and choose the parameter θ so that the implied average intensive margin elasticity of demand is 7, a value in the range of those used in the macro literature.[15] The customer budget constraint is given by $p\,d + n = I$, where I is the agent's nominal income, which we normalize to one.[16] We set the nominal wage equal to the price of the numeraire good, so that $w = 1$.[17]

While we fix the parameters listed above using external sources, our data allow us to estimate the key parameters of the model: those characterizing the idiosyncratic productivity process and the search cost distribution using a minimum-distance estimator. The productivity process influences the variability of prices, which is necessary for customers to obtain any benefit from search. The parameters of the search cost distribution, on the other hand, directly determine how costly it is to search. Below, we select functional forms for these objects, and explain the moments we choose in our data to identify the associated parameters. The discussion on identification is provided only for the sake of intuition; given the nonlinearity of the model, all the moments contribute to the identification of all the parameters.

We assume that the productivity evolves according to a process of the following form:

$$
\log(z_t^j) = \begin{cases} \log(z_{t-1}^j) & \text{with probability } \rho \\ \log(z') \sim N(0, \sigma) & \text{with probability } 1 - \rho \end{cases}.
$$

---

[13]This level of β reflects that the effective discount rate faced by the firm is the product of the usual time preference discount factor and a rescaling element which takes into account the time horizon of the decision maker, as for instance the average tenure of CEOs in the retail food industry reported in Henderson et al. (2006).

[14]Moving from these assumptions we can derive a demand function, $d(p)$, and a customer surplus function, $v(p)$, consistent with the assumptions made in Section 3.

[15]The choice of α affects the benefits of searching for a supplier of good d. As we will choose the distribution of observation costs ψ so to target the sensitivity of the searching decision to price variation, the choice of α is not important on this dimension.

[16]In Appendix D we show that I can be derived based on a model of the labor market.

[17] This is equivalent to assume that the numeraire good n is produced by a competitive representative firm with linear production function and unitary labor productivity. See Appendix D for details.

Our theoretical model describes how persistence and volatility of productivity (ρ and σ, respectively) determine autocorrelation and volatility of the resulting firm prices. We therefore estimate ρ and σ by matching the autocorrelation and the volatility of the logarithm of firm prices observed in the data. Here we face the complication that, while the theoretical model described the behavior of customers buying from firms producing a single homogeneous good, our application relates to supermarket stores selling bundles of goods.[18] We address this issue focusing on price indeces of composite bundles of goods as it is routinely done in the literature (Smith (2004)). Since customers baskets are in large proportion composed of packaged goods, which are standardized products available with identical characteristics at many different supermarket stores, the assumption that the basket is a homogenous good is not unwarranted. We have already introduced a consumer specific price index for the basket of goods each customer typically buys at the supermarket and used it to provide evidence on the elasticity of exit from the customer base to prices. To study the properties of the pricing process, we rely on an analogous price index constructed, however, at the store level.[19] We find that the autocorrelation of log-prices in the data is equal to 0.58, while the unconditional measure of standard deviation is 0.02, on a weekly basis.

The other important parameters in our model are those governing the search cost distribution. We assume that customers draw their search cost from a Gamma distribution with shape parameter ζ, and scale parameter λ. The Gamma is a flexible distribution and fits the assumptions we made over the G function in the specification of the model. In particular, for $\zeta > 1$, we obtain that the distribution of search costs is differentiable at $\psi = 0$.[20]

To identify the parameters of the search cost distribution we exploit the estimates of the relationship between the price and the probability of exiting the customer base discussed in Section 2. We identify the scale parameter λ by matching the average effect of log-prices on the exit probability predicted by the model to its counterpart in the data, measured by the

---

[18]The choice of focusing on the customer base of the store rather than that of one of the branded product it sells is data driven. With data from a single chain we cannot track the evolution of the customer base of a single brand. In fact, if we observed customers no longer buying a particular brand we could only infer that they are not buying it at the chain we analyze, but we could not rule out that they are buying it elsewhere.

[19]We perform the analysis at the store level and focus on the property of a price index for the store (p_{jt}). This index is computed in a fashion analogous to the customer price index. It is the average of the price of all the UPCs sold in store j, weighted for the share of revenues they represent. To estimate persistence and volatility of prices in the data, we exploit the first year of the sample span to obtain the store-level average of the price index and use it to demean the variables so to remove store fixed effects. We then estimate on the second year of data the equation $(\log(p_{jt}) - \overline{\log(p_j)}) = k_0 + k_1(\log(p_{jt-1}) - \overline{\log(p_j)}) + \tau_t + \epsilon_{jt}$ pooling all stores. Time fixed effects are included to purge the data from aggregate effects and isolate the variation in price driven by the idiosyncratic component. The estimate of $k_1 = 0.58$ provides us a measure of the autocorrelation of prices. The estimate of the standard deviation of ϵ, $\sigma(\epsilon) = 0.0168$, gives us an estimate of price volatility.

[20]In our estimation procedure we do not impose any constraints on the values the parameter ζ can take. Our unconstrained point estimate lies in the desired region.

parameter b_1 in equation (2). It is important to notice that this parameter was estimated controlling for a number of household-specific characteristics (demographics, distance from the store, etc.) which affect the decision on the store where to shop.

The parameter ζ measures the inverse of the coefficient of variation of the search cost distribution. In the model, higher dispersion of search costs (i.e., lower ζ) implies more mass on the tails of the distribution of search costs. The latter is associated to larger variation in the sensitivity of the exit probability to the price. In the data, we measure this variation by fitting a spline to equation (2), allowing for the marginal effect of price on the probability of exit to vary for different terciles of price levels. We find that higher prices are associated to higher value of b_1 as predicted by the model. The dispersion in the estimates of b_1 is 0.03. The parameter ζ is estimated by matching this number to an equivalent statistic generated by the model.

Table 2: Parameter estimates

	Estimates
Persistence of productivity process, ρ	0.60 [0.58, 0.61]
Volatility of productivity innovations, σ	0.11 [0.07, 0.12]
Scale parameter of search cost distribution, λ	0.033 [0.012, 0.041]
Shape parameter of search cost distribution, ζ	3.00 [2.60, 3.20]

Notes: 95% confidence intervals reported in parenthesis are computed by block bootstrap.

We define $\Omega \equiv [\zeta \ \lambda \ \rho \ \sigma]'$ as the vector of parameters of interest and estimate it with a minimum-distance estimator. Denote by $v(\Omega)$ the vector of the moments predicted by the model as a function of parameters in Ω, and by v_d the vector of their empirical counterparts. The n^{th} iteration of the estimation procedure unfolds according to the following steps:

1. Pick values for the parameters ρ_n, σ_n, λ_n and ζ_n from a given grid,

2. Solve the model and obtain the vector $v(\Omega_n)$,

3. Evaluate the objective function $(v_d - v(\Omega_n))'\Sigma \ (v_d - v(\Omega_n))$. Where Σ is a weighting matrix that we assume to be the identity matrix.

We select as estimates the parameter values from the proposed grid that minimize the objective function. Implementing step 2 requires solving a fixed point problem in equilibrium

prices. In particular, given our definition of equilibrium and the results of Proposition 2, we look for equilibria where prices are in the interval $[p^*(\bar{z}), p^*(\underline{z})]$. In principle, our model could have multiple equilibria; however, numerically we always converge to the same equilibrium despite starting from different initial conditions.

The estimation results are summarized in Table 2. Persistence and volatility of the productivity process are easy to interpret. To provide context for the parameters of the search cost distribution, we perform a back-of-envelope calculation of the dollar value of the search costs the estimated distribution implies. We find that the average search cost paid, conditional on searching, by customers in our model accounts for 0.3% of their weekly income.

6   Price and customer dynamics

In this section we use the parameter estimates reported in Table 2 to solve for the equilibrium pricing and search policies implied by our model (henceforth "baseline economy"). We check these predictions against our data on retail prices to assess whether encompassing customer market in a model of price formation helps explaining the available evidence. Data on retail prices are routinely used by both micro and macroeconomist for a variety of purposes. Improving our understanding of the forces that contribute to shape them is important to the progress of this growing literature.

We also contrast results from our model with those implied by the limiting case where competition for customers is shut down by raising the scale of search costs (λ) to infinity. We refer to this benchmark as "counterfactual economy." Since in the counterfactual economy search costs are infinite, customers do not change firm and firms do not compete for customers. It follows that in our baseline specification firms face both an intensive and an extensive margin elasticity of demand; whereas only the latter is present in the counterfactual economy. To make the comparison meaningful, we fix θ in the counterfactual economy so that the resulting average total elasticity of demand is the same as in our baseline economy (i.e. $\varepsilon_q = \varepsilon_d^{baseline} + \varepsilon_m^{baseline} = \varepsilon_d^{counterfactual}$). We also estimate σ and ρ targeting the same volatility and autocorrelation of prices as in the baseline estimates. Hence, the two economies are observationally equivalent with respect to the (average) price elasticity of demand and the equilibrium price process.

Customer dynamics.   A first order difference between our model and the counterfactual economy is the presence of customer dynamics in equilibrium. In the counterfactual economy, a firm's customer base is constant; whereas in our model firms with high productivity experience positive net growth of their customers base and lower productivity firms are net losers

of customers. On this dimension, the comparison with the empirical evidence is obviously partial to our baseline model since the fact that customers move across competing firms is observed in virtually every industry. Quantitatively, the model predicts a yearly customer turnover of about 9% in front of an unconditional frequency of exit from the customer base in our data of 14% on a yearly basis. Thus, price variation arising from idiosyncratic cost shocks explains almost two-thirds of customer dynamics observed in the data.

The distribution of prices. As we have shown in Section 3, the presence of customer markets also affects firms' pricing strategies. Here we document this fact analyzing the distribution of prices implied in equilibrium by our model. Since data on prices are typically easily available to researchers, we believe that assessing the predictions of the model in this respect makes for a particularly relevant empirical test. Indeed, in a recent paper Kaplan and Menzio (forthcoming,a) explore the features of the price distribution with the explicit goal to provide evidence against which the empirical relevance of price setting models can be tested.

Figure 2: The distribution of standardized prices: model and data
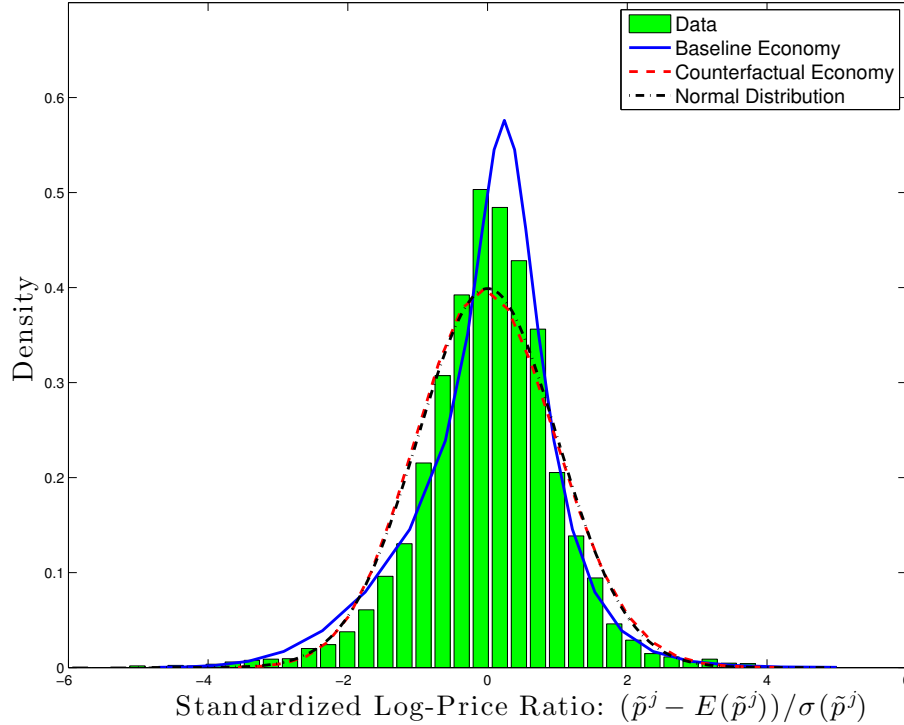
Notes: In the figure, we plot the distribution of the prices implied by our basline (blue line) and counterfactual economy (red line), as well as the empirical distribution of the price index of each store of the retail chain (green histogram). All prices are appropriately normalized to allow comparability: for details, see Appendix E.

In Figure 2, we compare the distribution of prices in the baseline economy with that of the counterfactual economy and with the empirical distribution emerging from the data. The shape of the price distribution generated by the model with customer markets matches quite closely the empirical distribution of prices in our data. Like in the data, the baseline model shows a high concentration of prices around the mean: the fraction of prices within half a standard deviation is 46% (45.5.% in the data), and excess kurtosis is 4.2 (4.6 in the data). Kaplan and Menzio (forthcoming,a) report a similar shape for the price distribution of homogeneous goods in the grocery sector; a finding we replicate in Appendix E.

The good fit achieved by our model does not result from targeting the shape of the price distribution in our estimation or from ad-hoc assumptions to introduce excess kurtosis. The underlying productivity innovations in the model are drawn from a normal distribution. In fact the counterfactual economy, which is characterized by prices roughly equal to a constant markup over marginal cost, displays a nearly normal price distribution, and therefore cannot explain the excess kurtosis found in the data. Instead, it is the very mechanism at the core of the model to induce this shape. The high clustering of prices around the mean derives from the higher extensive margin elasticity faced by firms in this region. In fact, a given variation in production cost is associated to smaller variation in prices because stronger customer retention concerns reduce the pass-through of cost shocks in this region. High productivity firms instead face weaker competition for customers and therefore decrease their price more in response to a reduction in production cost, contributing to a fatter left tail of the price distribution. Low productivity firms have small markups and hence do not have room to absorb cost variation in their markup. They adjust their price more in response to a cost increase, contributing to a fatter right tail of the price distribution.

Pass-through of cost shocks.   An alternative way to compare the model predictions on pricing with the empirical evidence is to analyze the pass-through of idiosyncratic shocks. Our calculations for the baseline model imply an average pass-through of idiosyncratic cost shocks equal to 13%, well below the 79% predicted in the counterfactual economy.[21] In the presence of competition for customers, a price increase leads to a persistent loss of customers. Instead, dynamic in the customer base is not a factor in the counterfactual economy. This implies that, when experiencing an increase in production cost, a firm in a customer market economy has an extra incentive to reduce the pass-through to price by compressing its margin.

---

[21]Note that the pass-through is incomplete even in the counterfactual economy because, with CES preferences, the demand of good i depends on the relative price p_i/P. With a finite number of goods in the basket of the customer, an increase in p_i also increases the price of the basket, P, thus reducing the overall increase in p_i/P and effect on demand. The effect on P is larger, the higher the weight of good i in the basket, that is the lower the price p_i and the higher its demand. Therefore, the elasticity of demand $\varepsilon_d(p)$ increases in p.

The measure of cost provided by the retailer allows to check the model's predictions regarding pass-through by regressing the log-price index of each store in a given week on its log-cost index. To avoid inflating the short-term (weekly) pass-through due to the persistence of both price and cost variables, we include in the specification lagged values of the independent variable. We experiment with an alternative way to deal with the persistence of the dependent variable by measuring the short-term pass-through using first differences. Finally, we include time and market fixed effects to control for aggregate trend (e.g. demand shocks) that can move prices independently from cost shifts.

Table 3: Pass-through of idiosyncratic shocks

Dep. variable	(1) $\log(p_t^j)$	(2) $\log(p_t^j)$	(3) $\Delta \log(p_t^j)$
$\log(cost_t^j)$	0.17***	0.24***	
	(0.04)	(0.09)	
$\log(cost_{t-1}^j)$	0.04	0.06	
	(0.03)	(0.04)	
$\log(cost_{t-2}^j)$	-0.01	0.02	
	(0.05)	(0.07)	
$\log(cost_{t-3}^j)$	0.06***	0.05*	
	(0.02)	(0.03)	
$\log(cost_{t-4}^j)$	0.07	0.07	
	(0.05)	(0.05)	
$\Delta \log(cost_t^j)$			0.13*
			(0.07)
Observations	12,915	8,295	8,295
MSA f.e.	No	Yes	Yes
Time f.e.	No	Yes	Yes

**Notes:** An observation is a store(j)-week(t) pair. The dependent variable is the price index of the store and the independent variables are the cost index of the store and its lags. Standard errors are in parenthesis and are clustered at the store level. ***: Significant at 1% **: Significant at 5% *: Significant at 10%.

Once again we find that data and theory seem to align. The results from this exercise are reported in Table 3: the pass-through measured in the data is in line with the predictions of a customer markets model and much lower than what the counterfactual economy would imply.[22]

---

[22]This result is not inconsistent with evidence of complete pass-through presented by Eichenbaum et al. (2011) using the same data. First, they measure pass-through conditional on price adjustment; whereas we look at the unconditional correlation between prices and costs. Second, they deal with UPC-level pass-through while we measure pass-through of a basket of goods. If retailers play strategically with the pricing

Illustrating the role of search costs. To further illustrate the mechanics of our model, in Figure 3 we plot the results from a comparative static exercise where we simulate it for different values of the scale of the search cost, λ. Unlike in Figure 2, we are not forcing dispersion and persistence of the price process to be the same across the different simulations; therefore all the parameters, except λ, are kept constant at the values reported in Table 2. This implies that the distribution of production cost is not changing and therefore any change in the distribution of prices is explained by an equivalent variation in the distribution of markups.

Figure 3: Equilibrium prices as a function of the scale of search costs, λ
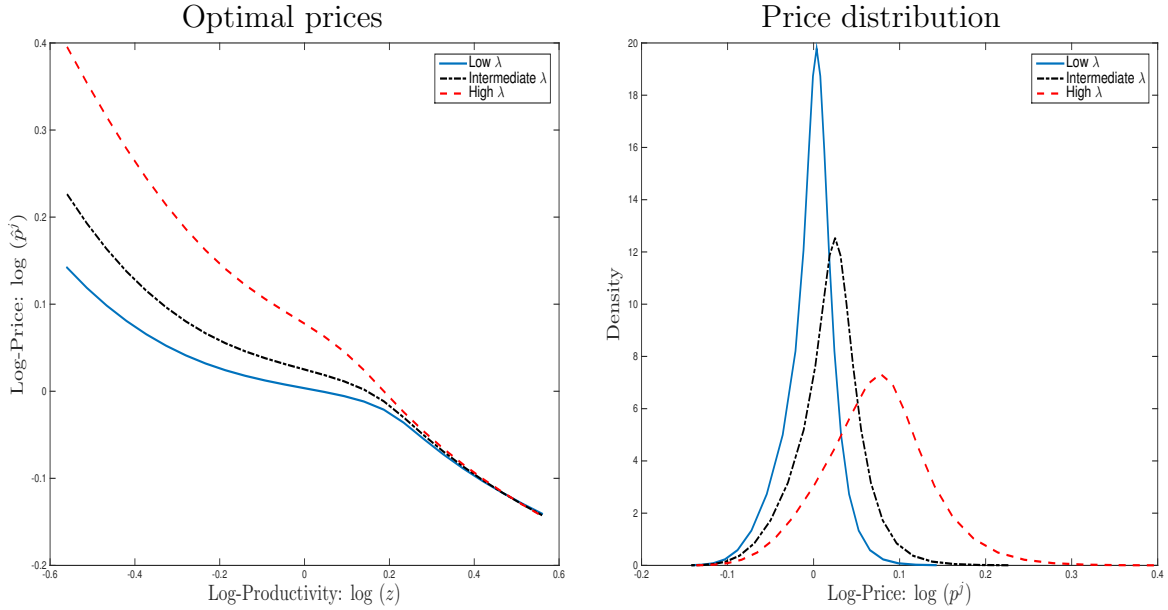
Notes: In the left panel, we plot the optimal log-prices as a function of productivity. In the right panel we plot the cross-sectional distribution of log-prices. The blue solid line refers to our baseline economy with customer markets at parameters estimated in Section 5 ($\lambda = 0.03$). The black dotted line refers to our baseline economy where we set $\lambda = 0.08$. The red dashed line refers to our baseline economy where we raise λ to 0.12.

When competition for customers is higher (lower λ) there is a greater risk for a firm to lose customers to competitors that are more productive. This increases the incentive for each firm to price closer to firms with higher productivity, resulting in a shift of mass from the right to the left of the price distribution. The incentive to reduce prices in response to increased competition for customers is weaker for more productive firms. In fact, in the extreme, the firm with the highest productivity always charges the same price (i.e. the price maximizing profits per customer), independently of λ. As displayed in the left panel of Figure 3, lower

---

of different products, for example lowering margins on some UPC to compensate the cost increase they experienced on others, we can obtain both high UPC-level pass-through and low basket-level pass-through.

search costs imply a flatter pricing policy as a function of productivity, and therefore lower pass-through of productivity shocks on average.

7   Demand shocks and markups cyclicality

Starting with Phelps and Winter (1970), there has been an active debate on the role of customer markets in shaping aggregate markup dynamics, in particular for the propagation of transitory demand shocks (Rotemberg and Woodford (1991, 1999)). Although it was commonly conjectured that customer markets could be a source of countercyclical markups, this result proved hard to achieve. In ad-hoc models with customer markets, in fact, a positive demand shock only makes customers more keen to buy in the current period relative to future periods, leading to a pro-cyclical response of markups. In this section we show that once we endogenize customer dynamics, demand shocks generate countercyclical aggregate markups. In our setting a positive demand shock not only increases current demand, but also affects customers incentives to search and hence firms' probability of losing customers. This increases competition for customers, providing incentives to lower markups.[23]

To illustrate this point, we perform the following experiment. We consider the economy in steady state as calibrated in Section 5. We then increase by 10% the weight of good d -the one traded in the frictional market- in the consumption basket. This means that the parameter α in the utility function of the customer jumps on impact from 0.5 to 0.55. The shock is unforeseen, it realizes before pricing and customer's exit decisions are taken, and it dies out according to an AR(1) process with autoregressive coefficient 0.9, which implies a half-life of around a quarter.[24]
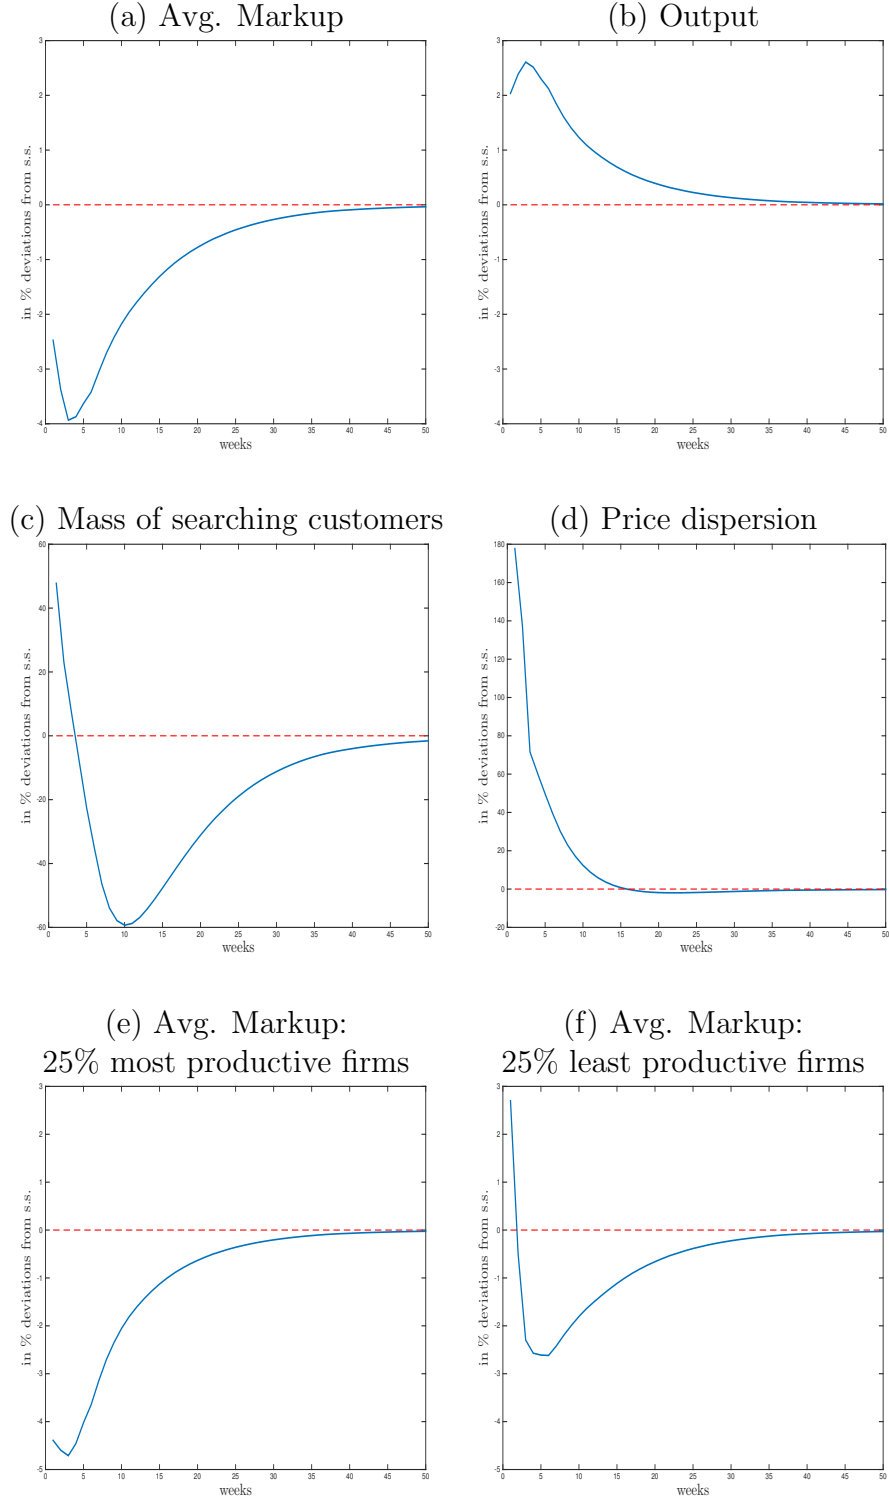
Figure 4 (plot (a)) shows that our model generates a countercyclical response of the average markup to demand shocks, which implies a cyclical movement of output (panel (b)). Panel (c) highlights why this result can arise in our model by displaying how the mass of customers who decide to search for a new firm responds to the demand shock. Panel (d) shows that price dispersion substantially increases following the demand shock. In our setting, the fact that customers value the good more does not only increase their demand for it but it also motivates them to search harder for firms that sell it cheap. The jump in the number of customers searching provides an opportunity for firms to snatch customers away from competitors which incentivize them to offer good prices by lowering their markups. This

---

[23]The habit formation literature (Ravn et al. (2006)) provides another example of countercyclical markups arising from individual behavior.

[24]Since the shock implies aggregate dynamics, we augment our economy with a simple equilibrium model of the labor market to capture the general equilibrium effects of the shock on wages and income. More details on this extension are provided in Appendix D.

Figure 4: The response to a temporary 10% demand shock

(a) Avg. Markup

(b) Output

(c) Mass of searching customers

(d) Price dispersion

(e) Avg. Markup:
25% most productive firms

(f) Avg. Markup:
25% least productive firms

Notes: Plot (a) displays the impulse response of the average markup (weighting firms by their customer base) to a 10% increase in consumer preferences for the good traded in the frictional market (i.e. an increase in the parameter α in the CES utility function of the customer. Plots (b) and (c) show the responses to the same shock of output and the mass of customers who decide to search. Plot (d) depicts the evolution of price dispersion after the demand occurs. Panels (e) and (f) plot the average response of average markup for firms in the top and bottom quartile of the productivity distribution. The dashed red line denotes the pre-shock steady state level.

30

effect is not present in original models of customer markets where the shock to preferences does not impact directly the probability that customers search.

Panel (e) and (f) provide additional insights on our main result by revealing that the markup response is heterogeneous in the level of productivity, thus accounting for the response of price dispersion. Productive firms respond along the lines we already described: they have the potential to snatch away searching customers and try to do it offering good deals. Therefore, the average markup for this group falls on impact. Low productivity firms find themselves in a very different situation. Their customers always had incentives to search, which the preferences shock has now strengthened. Furthermore, their high cost leaves them little scope for price cuts which could help them hold on to their customers. Therefore, they choose to raise their markups to extract more profits from the customers they can retain. This means that the demand shock triggers a reallocation of customers from low to high productivity firms. A result of such movement is that a firm that wanted to attract customers a few periods after the shock has hit must offer a better price (i.e. set a lower markup) than before. In fact, the average searching customer at this point belongs to the customer base of a more productive firm than it was the case on impact. This explains why markups of high productivity firms keep falling for a few periods after the initial drop. Eventually, as the shock dissipates, the opportunity of attracting large masses of customers vanishes and firms bring their markups back to the steady state level. This allows low productivity firms to regain the lost customers by lowering their markup.

The process of customer reallocation also explains why during the transition the mass of searching customers falls below the steady state level. At a certain stage, more customers will be matched with highly productive firms with respect to the initial period. Hence, their incentive to search will be weaker than it was then. As low and high productivity firms respectively lower and raise their markups, the distribution of customers across firms by productivity level reverts to the initial one with the associated baseline search intensity.

8 Concluding remarks

This paper provides an assessment of the importance of customer dynamics in shaping firm pricing strategies. We exploit novel data to document the premise of the customer market literature that prices influence customers' decision to exit the customer base of their firm. We then develop and estimate a rich yet tractable model to assess the role that customer markets play in determining price setting. We find that this mechanism is important to explain salient features of the data. These results bring new evidence on the forces involved in price formation which should be of great interest for the increasing number of both micro

and macroeconomists analyzing price data for an assortment of purposes. Finally, we use our framework to study aggregate dynamics and show that the conventional wisdom according to which customer market models cannot generate countercyclical markups in response to demand shocks relates only to ad-hoc model of customer base evolution. Our microfoundation linking customer dynamics to search and exit decision by individual customers can results in a countercyclical response of markup to a temporary demand shock.

Our study relies on a number of simplifying assumptions, whose relaxation seems of interest for future research. First, for tractability we refrain from explicitly modeling persistent heterogeneity in customers search/opportunity costs (although we control for these factors in the empirical analysis) and we do not allow for price discrimination. The presence of customers heterogeneity in shopping behavior is well documented (Aguiar and Hurst (2007)), which makes studying its implications for optimal pricing and customer dynamics an important topic. Due to lack of data, we do not consider the role of advertising in generating demand dynamics (Hall (2014)). While our conjecture is that the analysis of the pricing incentives presented in this paper would still apply, we think that extending the analysis to advertising, as well as to other strategies to attract and retain customers, and confront the results with direct firm level evidence, could provide new insights about firms behavior.

References

Aguiar, M. and Hurst, E. (2007). Life-cycle prices and production, *American Economic Review* **97**(5): 1533–1559.

Alessandria, G. (2004). International Deviations From The Law Of One Price: The Role Of Search Frictions And Market Share, *International Economic Review* **45**(4): 1263–1291.

Bai, Y., Rull, J.-V. R. and Storesletten, K. (2012). Demand Shocks that Look Like Productivity Shocks, *working paper*.

Blinder, A. S., Canetti, E. R. D., Lebow, D. E. and Rudd, J. B. (1998). *Asking About Prices: A new Approach to Understanding Price Stickiness*, New York: Russell Sage Foundation.

Bronnenberg, B. J., Kruger, M. W. and Mela, C. F. (2008). The IRI Marketing data set, *Marketing Science* **27**(4): 745–748.

Burdett, K. and Coles, M. G. (1997). Steady state price distributions in a noisy search equilibrium, *Journal of Economic Theory* **72**(1): 1 – 32.

Burstein, A. and Jaimovich, N. (2012). Understanding movements in aggregate and product-level real exchange rates, *working paper.*

Cabral, L. (2014). Dynamic pricing in customer markets with switching costs, *Working paper.*

Coibion, O., Gorodnichenko, Y. and Hong, G. H. (forthcoming). The cyclicality of sales, regular and effective prices: Business cycle and policy implications, *American Economic Review* .

Diamond, P. A. (1971). A model of price adjustment, *Journal of Economic Theory* **3**(2): 156–168.

Dinlersoz, E. M. and Yorukoglu, M. (2012). Information and Industry Dynamics, *American Economic Review* **102**(2): 884–913.

Drozd, L. A. and Nosal, J. B. (2012). Understanding international prices: Customers as capital, *American Economic Review* **102**(1): 364–95.

Eichenbaum, M., Jaimovich, N. and Rebelo, S. (2011). Reference prices and nominal rigidities, *American Economic Review* **101**(1): 234–262.

Einav, L., Leibtag, E. and Nevo, A. (2010). Recording discrepancies in Nielsen Homescan data: Are they present and do they matter?, *Quantitative Marketing and Economics* **8**(2): 207–239.

Einav, L. and Somaini, P. (2013). A model of market power in customer markets, *Journal of Industrial Economics* **61**(4): 938–986.

Fabiani, S., Loupias, C., Martins, F. and Sabbatini, R. (2007). *Pricing decisions in the euro area: how firms set prices and why*, Oxford University Press, USA.

Fishman, A. and Rob, R. (2003). Consumer inertia, firm growth and industry dynamics, *Journal of Economic Theory* **109**(1): 24 – 38.

Fishman, A. and Rob, R. (2005). Is bigger better? customer base expansion through word-of-mouth reputation, *Journal of Political Economy* **113**(5): 1146–1175.

Foster, L., Haltiwanger, J. and Syverson, C. (2013). The slow growth of new plants: Learning about demand?, *working paper.*

Gourio, F. and Rudanko, L. (2014). Customer capital, *The Review of Economic Studies* **81**: 1102–1136.

Hall, R. E. (2008). General equilibrium with customer relationships: A dynamic analysis of rent-seeking, *working paper*.

Hall, R. E. (2014). What the cyclical response of advertising reveals about markups and other macroeconomic wedges, *working paper*.

Henderson, A. D., Miller, D. and Hambrick, D. C. (2006). How quickly do CEOs become obsolete? Industry dynamism, CEO tenure, and company performance, *Strategic Management Journal* **27**: 447–460.

Holmes, T. (2011). The diffusion of Wal-Mart and economies of density, *Econometrica* **79**: 253–302.

Kaplan, G. and Menzio, G. (forthcoming,a). The morphology of price dispersion, *International Economic Review* .

Kaplan, G. and Menzio, G. (forthcoming,b). Shopping externalities and self-fulfilling unemployment fluctuations, *Journal of Political Economy* .

Kleshchelski, I. and Vincent, N. (2009). Market share and price rigidity, *Journal of Monetary Economics* **56**(3): 344–352.

Menzio, G. (2007). A search theory of rigid prices, *working paper*.

Nakamura, E. and Steinsson, J. (2011). Price setting in forward-looking customer markets, *Journal of Monetary Economics* **58**(3): 220–233.

Petrosky-Nadeau, N. and Wasmer, E. (forthcoming). Macroeconomic Dynamics in a Model of Goods, Labor and Credit Market Frictions, *Journal of Monetary Economics* .

Phelps, E. S. and Winter, S. G. (1970). Optimal price policy under atomistic competition, *in* G. C. Archibald, A. A. Alchian and E. S. Phelps (eds), *Microeconomic Foundations of Employment and Inflation Theory*, New York: Norton.

Ravn, M. O., Schmitt-Grohe, S. and Uribe, M. (2010). Incomplete cost pass-through under deep habits, *Review of Economic Dynamics* **13**: 317–332.

Ravn, M., Schmitt-Grohe, S. and Uribe, M. (2006). Deep habits, *Review of Economic Studies* **73**(1): 195–218.

Reinganum, J. F. (1979). A simple model of equilibrium price dispersion, *Journal of Political Economy* **87**(4): 851–58.

Rotemberg, J. J. and Woodford, M. (1991). Markups and the business cycle, *NBER Macroeconomics Annual 1991, Volume 6*, NBER Chapters, National Bureau of Economic Research, Inc, pp. 63–140.

Rotemberg, J. J. and Woodford, M. (1999). The cyclical behavior of prices and costs, *working paper*.

Shi, S. (2011). Customer relationship and sales, *working paper*.

Smith, H. T. (2004). Supermarket choice and supermarket competition, *Review of Economic Studies* **71**: 235–263.

Sorensen, A. T. (2000). Equilibrium price dispersion in retail markets for prescription drugs, *Journal of Political Economy* **108**(4): 833–850.

Appendix - Not for publication

A Data sources and variables construction

A.1 Data and selection of the sample

The empirical evidence presented in Section 2 is based on two data sources provided by a large supermarket chain that operates over 1,500 stores across the United States. This implies that we can observe our agents behavior only when they shop with the chain; on the other hand, cash register data contain significantly less measurement error than databases relying on home scanning (Einav et al. (2010)).

   The main data source contains information on grocery purchases at the chain between June 2004 and June 2006 for a panel of over 11,000 households. For each grocery trip made by a household, we observe date and store where the trip occurred, the collection of all the UPCs purchased with quantity and price paid. The data include information on the presence and size of price discounts but do not generally report redemption of manufacturer coupons. Data are collected through usage of the loyalty card; purchases made without using the card are not recorded. However, the chain ensures that the loyalty card has a high penetration by keeping to a minimum the effort needed to register for one. Furthermore, nearly all promotional discount are tied to ownership of a loyalty card, which provides a strong incentive to use it.

   Household-level scanner data report information on the price paid conditional on a certain item having been bought by the customer. Therefore, if we do not observe at least one household in our sample buying a given item in a store in a week, we would not be able to infer the price of the item in that store-week. This has important implications as our definition of basket requires us to be able to attach a price to each of the item composing it in every week, even when the customer does not shop. The issue can be solved using another dataset with information on weekly store revenues and quantities between January 2004 and December 2006 for a panel of over 200 stores. For each good (identified by its UPC) carried by the stores in those weeks, the data report total amount grossed and quantity sold. Exploiting this information, we can calculate unit value prices every week for every item in stock in a given store, whether or not that particular UPC was bought by one of the households in our main data. Unit value prices are computed using data on revenues and quantities sold as

$$UVP_{stu} = \frac{TR_{stu}}{Q_{stu}} \ ,$$

where TR represent total revenues and Q the total number of units sold of good u in week t in store s.

As explained in Eichenbaum et al. (2011), this only allows us to recover an average price for goods that were on promotion. In fact the same good will be sold to loyalty card carrying customers at the promotional price and at full price to customers who do not have or use a loyalty card. Without information on the fraction of these two types of customers it is not possible to recover the two prices separately. Furthermore, since prices are constructed based on information on revenues, missing values can originate even in this case if no unit of a specific item is sold in a given store in a week. This is, however, an unfrequent circumstance and involves only rarely purchased UPCs, which are unlikely to represent important shares of the basket for any of the households in the sample. For the analysis, we only retain UPCs with at most two nonconsecutive missing price observations and impute price for the missing observation interpolating the prices of the contiguous weeks.

On top of reporting revenues and quantities for each store-week-UPC triplet, the store-level data also contain a measure of cost. This variable is constructed on the basis of the estimated markup imputed by the retailer for each item and includes more than the simple wholesale cost of the item (the share of transportation cost, etc.). Eichenbaum et al. (2011) suggest to think about it as a measure of replacement cost, i.e. the cost of placing an item on the shelf to replace an analogous one just grabbed by a consumer. We use this measure to construct our instrument of the basket price.

It is important to notice that the retail chain sets different prices for the same UPC in different geographic areas, called "price areas." The retailer supplied store-level information for 270 stores, ensuring that we have data for at least one store for each price area. In order to use unit value prices calculated from store-level data to compute the price of the basket of a specific household, we need to determine to which price area the store(s) at which she regularly shops belong. This information is not supplied by the retailer that kept the exact definition of the price areas confidential. A possible solution is to infer in which price areas the store(s) visited by a household are located by comparing the prices contained in the household panel with those in the store data. In principle the household data should give information on enough UPC prices in a given week to identify the price area representative store whose pricing they are matching. However, even though two stores belonging in the same price area should have the same prices, they may not have the same unit value prices if the share of shoppers using the loyalty card differs in the two stores. Therefore, we choose to restrict our analysis to the set of customers shopping predominantly (over 80% of their grocery expenditure at the chain) in one of the 270 stores for which the chain provided complete store-level data. This choice is costly in terms of sample size: only

1,336 households (or 12% of the original sample) shop at one of the 270 stores for which we have store-level price data. However, since the 270 representative stores were randomly chosen, the resulting subsample of households should not be subject to any selection bias.

A final piece of the data is represented by the IRI-Symphony database. We use store-level data on quantities and revenues for each UPC in 30 major product categories for a large sample of stores (including small and mom & pop ones) in 50 Metropolitan Statistical Areas in the United States. The data allow to construct unit value prices for all the stores competing with the chain who provided the main dataset. However, the coarse geographic information prevents us from matching each customer with the stores closer to her location (in the same zip code, for instance) and forces us to adopt the MSA as our definition of a market.

A.2 Variables construction

Exit from customer base. The dependent variable in the regression presented in equation (2) is an indicator for whether a customer is exiting the customer base of the chain. With data on grocery purchases at a single retail chain it is hard to definitively assess whether a household has abandoned the retailer to shop elsewhere or is simply not purchasing groceries in a particular week, for instance because she is just consuming its inventory. In fact, we observe households when they buy groceries at the chain but do not have any information on their shopping at competing grocers. Our choice is to assume that a customer is shopping at some other store when she has not visited any supermarket store of the chain for at least eight consecutive weeks. The *Exit* dummy is then constructed so that it takes value of one in correspondence to the last visit at the chain before a spell of eight or more weeks without shopping there. Table 4 summarizes shopping behavior for households in our sample. It is immediate to notice that an eight-week spell without purchase is unusual, as customers tend to show up frequently at the stores. This strengthens our confidence that customers missing for an eight-week period have indeed switched to a different retailer.

Composition of the household basket and basket price. The household scanner data deliver information on all the UPCs a household has bought through the sample span. We assume that all of them are part of the household basket and, therefore, the household should care about all of those prices. Some of the items in the household's basket are bought regularly, whereas others are purchased less frequently. We take this into account when constructing the price of the basket by weighting the price of each item by its expenditure share in the household budget. The price of household i's basket purchased at store j in week t is computed as:

Table 4: Descriptive statistics on customer shopping behavior

	Mean	Std.dev.	25th pctile	75th pctile
Number of trips	150	127	66	200
Days elapsed between consecutive trips	4.2	7.5	1	5
Expenditure per trip ($)	69	40	40	87
Frequency of exits	0.003	0.065		

$$p_{ijt} = \sum_{k \in K_i} \omega_{ik}\, p_{kjt} \;, \qquad \omega_{ik} = \frac{\sum_t E_{ikt}}{\sum_k \sum_t E_{ikt}} \;,$$

where K^i is the set of all the UPCs (k) purchased by household i during the sample period, p_{kjt} is the price of a given UPC k in week t at the store j where the customer shops. E_{ikt} represents expenditure by customer i in UPC k in week t and the ω_{ik}'s are a set of household-UPC specific weights. There is the practical problem that the composition of the consumer basket cannot vary through time; otherwise basket prices for the same customer in different weeks would not be comparable. This requires that we drop from the basket all UPCs for which we do not have price information for every week in the sample. However, the price information is missing only in instances where the UPC registered no sales in a particular week. It follows that only low market-share UPCs will have missing values and, therefore, the UPCs entering the basket computation will represent the bulk of each customer's grocery expenditure. The construction of the cost of the basket follows the same procedure where we substitute the unit value price with the measure of replacement cost provided by the retailer.

We choose to calculate the weights using the total expenditure in the UPC by the household over the two years in the sample. This can lead to some inaccuracy in identifying the goods the customer cares for at a given point in time. For example, if a customer bought only Coke during the first year and only Pepsi during the second year of data, our procedure would have us give equal weight to the price of Coke and Pepsi throughout the sample period. If we used a shorter time interval, for example using the expenditure share in the month, we would correctly recognize that she only cares about Coke in the first twelve months and only about Pepsi in the final 12 months. However, weights computed on short time intervals are more prone to bias induced by pricing. For example, a two-weeks promotion of a particular UPC may induce the customer to buy it just because of the temporary convenience; this would

give the UPC a high weight in the month. The effect of promotion is instead smoothed when we compute weights using expenditure over the entire sample period.

The construction of the price of the competitors occurs in two steps. First, we use the IRI data and the same procedure described above to obtain a price for the basket of each consumer at every store located in her same MSA. Next, we average those prices across stores to obtain the average market price of the consumer basket. In particular, the price is computed as:

$$p_{it}^{mkt} = \sum_{z \in M^i} s_z \sum_{k \in K_i} \omega_{ik} p_{kzt}, \quad \omega_{ik} = \frac{\sum_t E_{ikt}}{\sum_k \sum_t E_{ikt}}, \quad s_z = \frac{\sum_t R_{zt}}{\sum_{z' \in M} \sum_t R_{z't}}$$

where M^i is the MSA of residence for customer i and R_{zt} represents revenues of store z in week t. In other words, in the construction of the competitors' price index, stores with higher (revenue-based) market shares weight more.

Composition of the store basket and basket price. The construction of the price (and cost) index for the store is conceptually analogous to that described above for the household basket. In principle, we would want to compute the store price index including all the UPCs sold at a store throughout the sample period, weighted by the share of revenues they generated. However, to keep the composition of the store basket constant through time, we must restrict ourselves to the UPCs for which we have no missing price information in any of the weeks in the sample span. This severely reduces the size of the store basket. At the same time, goods without missing information are the best sellers, which are those the store is likely to care more about.

B   Proofs

B.1   Proof of Proposition 2

Monotonicity of prices. Monotonicity of optimal prices follows from an application of Topkis' theorem. In order to apply the theorem to the firm problem in equation (6) we need to establish increasing differences of the firm objective $\Delta(p, z) \Pi(p, z)$ in $(p, -z)$. Under the standard assumptions we stated on $\pi(p, z)$, it is easy to show that $\Pi(p, z)$ satisfies this property. The customer base growth function does not in general verify the increasing difference property. However, under the assumption of i.i.d. productivity, $\Delta(p, z)$ is independent of z, which, together with Proposition 1, is sufficient to obtain the result. We first show that optimal prices $\hat{p}(z)$ are non-increasing in z. Given, that productivity is i.i.d. and that we look for equilibria where $\hat{p}(z) \geq p^*(\bar{z})$, we have that $\bar{p}(z) = p^*(\bar{z})$ for each z. From Lemma 1 we know

that, for a given z, the optimal price $\hat{p}(z)$ belongs to the set $[p^*(\bar{z}),\ p^*(z)]$. Over this set, the objective function of the firm,

$$W(p,z) = \Delta(p,z)\left(\pi(p,z) + \beta\,\text{constant}\right), \tag{16}$$

is supermodular in $(p, -z)$. Notice the i.i.d. assumption implies that future profits of the firm do not depend on current productivity as future productivity, and therefore profits, are independent from it. Similarly, $\Delta(p,z)$ does not depend on z, as the expected future value to the customer does not depend on the productivity of the current match as future productivity is independent from it. Abusing notation, we replace $\Delta(p,z)$ by $\Delta(p)$. To show that $W(p,z)$ is supermodular in $(p, -z)$ consider two generic prices p_1, p_2 with $p_2 > p_1 > 0$ and productivities $z_1, z_2 \in [\underline{z}, \bar{z}]$ with $-z_2 > -z_1$. We have that $W(p_2, z_2) - W(p_1, z_2) \le W(p_2, z_1) - W(p_1, z_1)$ if and only if

$$\Delta(p_2)d(p_2)(p_2 - w/z_2) - \Delta(p_1)d(p_1)(p_1 - w/z_2) \le \Delta(p_2)d(p_2)(p_2 - w/z_1) - \Delta(p_1)d(p_1)(p_1 - w/z_1),$$

which, since $\Delta(p_2)d(p_2) < \Delta(p_1)d(p_1)$ as $d(p)$ is strictly decreasing and $\Delta(p)$ is non-increasing, is indeed satisfied if and only if $z_2 < z_1$. Thus, $W(p,z)$ is supermodular in $(p, -z)$. By application of the Topkis Theorem we readily obtain that $\hat{p}(z)$ is non-increasing in z.

Existence of equilibrium. Next we prove existence of an equilibrium. The fixed point problem is a mapping from candidate functions of equilibrium prices, $\hat{p}(z)$, to the firm's optimal pricing strategy, $\hat{p}(z)$, where an equilibrium is one where $\hat{p}(z) = \hat{p}(z)$ for each z. Notice that $W(p,z)$ in equation (16) is continuous in (p,z). By the theorem of maximum, $\hat{p}(z)$ is upper hemi-continuos and $W(\hat{p}(z), z)$ is continuos in z. Given that $\hat{p}(z)$ is non-increasing in z it follows that $\hat{p}(z)$ has a countably many discontinuity points. We thus proceed as follows. Let $\hat{\mathcal{P}}(z)$ be the set of prices that maximize the firm problem. Whenever a discontinuity arises at some $\tilde{z}$ (so that $\hat{\mathcal{P}}(\tilde{z})$ is not a singleton), we modify the optimal pricing rule of the firm and consider the convex hull of the $\hat{\mathcal{P}}(\tilde{z})$ as the set of possible prices chosen by the firm with productivity $\tilde{z}$. The constructed mapping from $\mathcal{P}(z)$ to $\hat{\mathcal{P}}(z)$ is then upper-hemicontinous, compact and convex valued. We then apply Kakutani's fixed point theorem to this operator and obtain a fixed point. Finally, notice that since the convexification procedure described above has to be applied only a countable number of times, the set of convexified prices has measure zero with respect to the density of z. Hence, they do not affect the fixed point.

It is important to point out that differentiability of the distribution of productivity F is not needed for the existence of an equilibrium. We assume it to ensure that $H(\cdot)$ and $Q(\cdot)$ are almost everywhere differentiable so that equation (7) is a necessary condition for

optimal prices (see below). However, even when F is not differentiable and the first order condition cannot be used to characterize the equilibrium, an equilibrium with the properties of Proposition 2 exists where $\hat{p}(z)$ and $\hat{\psi}(\hat{p}(z), z)$ are monotonic in z but not necessarily strictly monotonic for all z.

Necessity of the first order condition. We show that Q and H are almost everywhere differentiable, so that Lemma 1 implies that equation (7) is necessary for an optimum. We guess that $\hat{p}(z)$ is strictly decreasing and almost everywhere differentiable. It immediately follows that $\mathcal{V}(z)$ is strictly increasing in z and almost everywhere differentiable. Then, given the assumption that F is differentiable, we have that K is differentiable. From $H(x) = K(\mathcal{V}^{-1}(x))$ it follows that H is also almost everywhere differentiable. Given that G and H are differentiable, so is Q. Then the first order condition in equation (7) is necessary for an optimum, which indeed implies that $\hat{p}(z)$ is strictly decreasing and differentiable in z in any neighborhood of the first order condition. Finally, given that $\hat{p}(z)$ has a countably many discontinuity points, it has countably many points where it is not differentiable, and the first order condition does not apply at those points, but applies everywhere else. These points have measure zero with respect to the density of z and therefore $\hat{p}(z)$ is almost everywhere differentiable.

Point (i). We already proved that $\hat{p}(z)$ is non-increasing in z. The proof that $\hat{p}(z)$ is strictly decreasing follows by contradiction. Consider that $\hat{p}(z_1) = \hat{p}(z_2) = \tilde{p}$ for some $z_1, z_2 \in [\underline{z}, \bar{z}]$. Also, without loss of generality, assume that $z_1 < z_2$. Given that we already established the necessity of the first order condition presented in equation (7) when prices are monotonic, suppose that it is satisfied at the pair $\{z_2, \tilde{p}\}$. Notice that, because of the assumed i.i.d. structure of productivity shocks together with $\pi_z(p, z) < 0$, it is not possible that the first order condition is also satisfied at the pair $\{z_1, \tilde{p}\}$. Moreover, because the first order condition is necessary and we already established that $\hat{p}(z)$ cannot be increasing at any z, we conclude that the optimal price at z_1 is strictly larger than at z_2. That is, $\hat{p}(z_1) > \hat{p}(z_2)$. Notice that this verifies the conjecture used to prove the necessity of the first order condition, which in turn validates the use of equation (7) here.[25]

Notice that, because $\hat{p}(z)$ is strictly decreasing in z, the fact that $v'(p) < 0$ together with i.i.d. productivity, implies, through an application of the contraction mapping theorem, that

---

[25]If prices are not strictly decreasing, this argument cannot be used as the first order condition is not necessary. However, it is possible to prove that $\hat{p}(z)$ is strictly decreasing in z for some region of z. The argument follows by contradiction. Suppose that $\hat{p}(z)$ is everywhere constant in z at some $\tilde{p}$. Then $\bar{p}(z) = \tilde{p}$ for all z. If $\tilde{p} > p^*(\bar{z})$, then $\tilde{p}$ would not be optimal for firm with productivity $\bar{z}$, which would choose a lower price. If $\tilde{p} = p^*(\bar{z})$, then continuous differentiability of G together with $H = G = Q = 0$ at the conjectured constant equilibrium price imply that the first order condition is locally necessary for an optimum, and a firm with productivity $z < \bar{z}$ would have an incentive to deviate according to equation (7), and set a strictly higher price than $\tilde{p}$. Finally, the result that $\hat{p}(z) < p^*(z)$ for all $z < \bar{z}$ and that $\hat{p}(\bar{z}) = p^*(\bar{z})$ follows from applying Lemma 1, and using that $\hat{p}(z) \geq \hat{p}(\bar{z})$ and $\bar{p}(z) = \hat{p}(\bar{z})$ for all z.

$\mathcal{V}(z) = \bar{V}(\hat{p}(z), z)$ is increasing in z.

Point (ii). $\hat{\psi}(p, z) \geq 0$ immediately follows its definition. The fact that $\mathcal{V}(z)$ is strictly increasing in z, together with Proposition 1 , immediately implies that $\hat{\psi}(\hat{p}(\bar{z}), \bar{z}) = 0$ and that $\hat{\psi}(\hat{p}(z), z)$ is strictly increasing in z. Proposition 1 also implies that $\Delta(\hat{p}(z), z)$ is increasing in z. Because of price dispersion, some customers are searching, which guarantees that $\Delta(\hat{p}(\bar{z}), \bar{z}) > 1$. Likewise, $\Delta(\hat{p}(\underline{z}), \underline{z}) < 1$.

B.2   Proof of Remark 1

Part (1). Start by noticing that, because the mean of $G(\psi)$ is positive, the expected value of searching diverges to $-\infty$ as n diverges to infinity. Because prices are finite for all $z \in [\underline{z}, \bar{z}]$, the value of not searching is bounded. As a result, customers do not search so that firms do not face customer base concerns. Formally, $\bar{p}(z) \to \infty$ for all $z \in [\underline{z}, \bar{z}]$. Because $p^*(z)$ is finite for all $z \in [\underline{z}, \bar{z}]$, it follows immediately that $p^*(z) < \bar{p}(z)$ for all $z \in [\underline{z}, \bar{z}]$. Then, using Lemma 1 we obtain that $\hat{p}(z) = p^*(z)$ for all $z \in [\underline{z}, \bar{z}]$.

Part (2). From Proposition 2 we have that, in equilibrium, the highest price is $\hat{p}(\underline{z})$. Moreover, under the assumptions of Proposition 2, the first order condition is a necessary condition for optimality of prices. We use this to show that, as n approaches zero, $\hat{p}(\underline{z})$ has to approach $\hat{p}(\bar{z}) = p^*(\bar{z})$.

In equilibrium, it is possible to rewrite equation (7), evaluated at $\{\hat{p}(\underline{z}), \underline{z}\}$, as $LHS(\hat{p}(\underline{z}), n) = RHS(\hat{p}(\underline{z}), n)$, where

$$
\begin{aligned}
LHS(\hat{p}(\underline{z}), n) &\equiv G'\left(\frac{\hat{\psi}(\hat{p}(\underline{z}), \underline{z})}{n}\right) \frac{\hat{\psi}_p(\hat{p}(\underline{z}), \underline{z})}{n} \\
&+ \left(G\left(\frac{\hat{\psi}(\hat{p}(\underline{z}), \underline{z})}{n}\right) H'(\bar{V}(\hat{p}(\underline{z}), \underline{z})) + \frac{1}{\Gamma} Q'(\bar{V}(\hat{p}(\underline{z}), \underline{z}))\right) \bar{V}_p(\hat{p}(\underline{z}), \underline{z}) , \\
RHS(\hat{p}(\underline{z}), n) &\equiv -\frac{\pi_p(\hat{p}(\underline{z}), \underline{z})}{\Pi(\hat{p}(\underline{z}), \underline{z})} \left(1 - G\left(\frac{\hat{\psi}(\hat{p}(\underline{z}), \underline{z})}{n}\right)\right) ,
\end{aligned}
$$

given that $H(\bar{V}(\hat{p}(\underline{z}), \underline{z})) = Q(\bar{V}(\hat{p}(\underline{z}), \underline{z})) = 0$.

Suppose that as $n \downarrow 0$, $\hat{\psi}(\hat{p}(\underline{z}), \underline{z})$ does not converge to zero. Then, $G\left(\frac{\hat{\psi}(\hat{p}(\underline{z}), \underline{z})}{n}\right) \uparrow 1$ as $n \downarrow 0$. This implies that $\lim_{n \downarrow 0} RHS(\hat{p}(\underline{z}), n) > 0$.

Consider now the function $LHS(\hat{p}(\underline{z}), n)$. Again, suppose that as $n \downarrow 0$, $\hat{\psi}(\hat{p}(\underline{z}), \underline{z})$ does not converge to zero. Notice that the second term of the function approaches a finite number as $\bar{V}_p(\hat{p}(\underline{z}), \underline{z})$ is bounded by assumptions on $v(p)$ and $H'(\bar{V}(\hat{p}(\underline{z}), \underline{z}))$ and $Q'(\bar{V}(\hat{p}(\underline{z}), \underline{z}))$ being bounded as a result of Proposition 2. Moreover, as long as $\hat{p}(\underline{z}) > \bar{p}(z) = p^*(\bar{z})$, we have

that $\hat{\psi}_p(\hat{p}(\underline{z}), \underline{z}) > 0$ so that $\hat{\psi}_p(\hat{p}(\underline{z}), \underline{z})/n$ diverges as n approaches zero. This means that $G'\left(\frac{\hat{\psi}(\hat{p}(\underline{z}), \underline{z})}{n}\right)\frac{\hat{\psi}_p(\hat{p}(\underline{z}), \underline{z})}{n}$ is divergent, and therefore the first order condition cannot be satisfied.

This analysis concluded that, if $\hat{\psi}(\hat{p}(\underline{z}), \underline{z})$ does not converge to zero as n becomes arbitrarily small, the first order condition, i.e. equation (7), cannot be satisfied. This occurs because $LHS(\hat{p}(\underline{z}), n)$ would diverge to infinity, while $RHS(\hat{p}(\underline{z}), n)$ would remain finite. It then follows that, as n approaches zero, a necessary condition is that $\hat{\psi}(\hat{p}(\underline{z}), \underline{z})$ also approaches zero. This condition can be restated as requiring that $\hat{p}(\underline{z})$ approaches $\bar{p}(z)$ as n approaches zero. Moreover, given the assumptions of Proposition 2, $\bar{p}(z) = \hat{p}(\bar{z}) = p^*(\bar{z})$.

In the end, if $\hat{p}(\underline{z})$ approaches $p^*(\bar{z})$ as n becomes arbitrarily small (so that $\hat{\psi}(\hat{p}(\underline{z}), \underline{z}) \to 0$ and $\hat{\psi}_p(\hat{p}(\underline{z}), \underline{z}) \to 0$), we have that $\lim_{n\downarrow 0} LHS(\hat{p}(\underline{z}), n) < \infty$ and $\lim_{n\downarrow 0} RHS(\hat{p}(\underline{z}), n) < \infty$ as $\pi_p(p^*(\bar{z}), \underline{z})$ is bounded as $\pi(p^*(\bar{z}), \underline{z}) > 0$. However, if $\hat{p}(\underline{z})$ does not approach $p^*(\bar{z})$ as n becomes arbitrarily small, we have that $LHS(\hat{p}(\underline{z}), n)$ diverges as n approaches zero, while $LHS(\hat{p}(\underline{z}), n)$ remains finite. As the first order condition has to be satisfied in equilibrium, a necessary condition is that, as n approaches zero, the highest price in the economy, i.e. $\hat{p}(\underline{z})$, has to approach the lowest price in the economy, i.e. $p^*(\bar{z})$.

C   Numerical solution of the model

In order to solve the model, we start by setting the parameters. The parameters β, w, q, and I are constant throughout the numerical exercises. For the set of estimated parameters $\Omega_n = [\lambda_n, \zeta_n, \rho_n, \sigma_n]'$, we set a search grid. The grid is different for each parameter, as they differ both in their levels and in the sensitivity of the statistics of interest to their variation. We consider a grid with an interval of 0.01 for σ, 0.05 for ρ, 0.5 for ζ, and 0.01 for λ. Each Ω_n corresponds to a particular combination of parameters among these grids. For each Ω_n we set θ to obtain $E[\varepsilon_d(z)] = 7$.

We next describe how we solve for the equilibrium of the model for a given combination of parameters. We start by discretizing the AR(1) process for productivity to a Markov chain featuring $N = 25$ different productivity values. We then conjecture an equilibrium function $\hat{p}(z)$. Given our definition of equilibrium and the results of Proposition 2, we look for equilibria where $\hat{p}(z) \in [p^*(\bar{z}), p^*(\underline{z})]$ for each z, and $\hat{p}(z)$ is decreasing in z. Our initial guess for $\hat{p}(z)$ is given by $p^*(z)$ for all z. We experiment with different initial guesses and found that the algorithm always converges to the same equilibrium.

Given the guess for $\hat{p}(z)$, we can compute the continuation value of each customer as a function of the current price and productivity, i.e. $\bar{V}(p, z)$, and solve for the optimal search and exit thresholds as described in Lemma 2. Given $\hat{p}(z)$ and the customers' search and exit thresholds we can solve for the distributions of customers $Q(\cdot)$ and $H(\cdot)$ as defined in

Definition 1. Notice that the latter also amounts to solve for a fixed point in the space of functions. Here, standard arguments for the existence of a solution to invariant distribution for Markov chains apply. Therefore, the assumption that $F(z'|z) > 0$ and $\Delta(\hat{p}(z), z) > 0$ ensure the existence of a unique $K(z)$ that solves equation (14). Finally, given $Q(\cdot)$, $H(\cdot)$, $\hat{p}(z)$ and $\bar{V}(p, z)$, we solve the firm problem and the obtain optimal firm prices given by the function $\hat{p}(z)$. We use $\hat{p}(z)$ to update our conjecture about equilibrium prices $\hat{p}(z)$, and iterate this procedure until convergence to a fixed point where $\hat{p}(z) = \hat{p}(z)$ for all $z \in [\underline{z}, \bar{z}]$.

Once we have solved for the equilibrium of the model at given parameter values, we construct the statistics to be matched to their data counterpart as follows.

- Log-price dispersion:

$$\hat{\sigma}_p \equiv \sqrt{\sum_j K(z_j)(\log(\hat{p}(z_j)) - M_p)^2}$$

  where $M_p = \sum_j K(z_j) \log(\hat{p}(z_j))$ and $K(z_j)$ is the equilibrium fraction of customers buying from firms with productivity z_j.

- Average comovement between the probability of exiting the customer base and the price:

$$\hat{b}_1 = Cov(E(z), \log(\hat{p}(z)))/(\sigma_p)^2$$

  where $E(z) \equiv G(\hat{\psi}(\hat{p}(z), z))(1 - H(\bar{V}(\hat{p}(z), z)))$, and $Cov(E(z), \log(\hat{p}(z))) = \sum_i K(z_j)(\log(\hat{p}(z_j)) - M_p)(E(z_j) - M_E)$ and $M_E = \sum_j K(z_j) E(z_j)$.

- Dispersion in the marginal effect of the price on the probability of exiting the customer base:

$$\hat{\sigma}_{b_1} = \sqrt{\sum_j K(z_j)(\hat{b}_1(z_j) - M_{\hat{b}_1})^2}$$

  where $\hat{b}_1(z_j) = G'(\hat{\psi}(\hat{p}(z), z))/G(\hat{\psi}(\hat{p}(z), z))(1 - H(\bar{V}(\hat{p}(z), z)))^2$ and $M_{\hat{b}_1} = \sum_i K(z_j)\hat{b}_1(z_j)$.

The autocorrelation of prices, $\hat{\rho}_p$ coincides with the parameter ρ governing the persistence and autocorrelation of productivity. Thus the model-predicted statistics used to estimate the parameters are given by the vector $v(\Omega_n) = [\hat{\rho}_p, \sigma_p, \hat{b}_1, \hat{\sigma}_{b_1}]'$. We then evaluate the objective function $(v_d - v(\Omega_n))'\Sigma (v_d - v(\Omega_n))$ at each iteration. We assume the weighting matrix Σ to be the identity matrix. We select as estimates the parameter values from the proposed grid that minimize the objective function and check that the optimum in the interior of the assumed grid.

D A simple model of the labor market

In this appendix we provide details on how the model can be extended to use it to evaluate the role of aggregate shocks.

We assume that each household is divided into a mass Γ of shoppers/customers and a representative worker. The preferences of the household are given by

$$E_t \left[\int_0^\Gamma V_t(p_t^i, z_t^i, \psi_t^i) \, di - J_t \right] \quad , \tag{17}$$

where $V_t(p_t^i, z_t^i, \psi_t^i)$ is defined as in equation (9) and it is the value function that solves the customer problem in Section 3.2. We denote the disutility from the sequence of labor ℓ_T as $J_t \equiv \phi \sum_{T=t}^\infty \beta^{T-t} \ell_T$ with $\phi > 0$. The aggregate state for the household includes the distribution of prices, the distribution of customers over the different firms and the levels of income, the wage, and their laws of motion. Given that we allow for aggregate shocks, we have to consider the possibility that the aggregate state varies over time. We index dynamics in the aggregate state through the time subscript t for the value function.

The worker chooses the path of ℓ_t that maximizes household preferences in equation (17). The search problem of each customer is as described in Section 3.2. As for the consumption decision, each customer allocates her income across consumption of the good sold in the local market, the demand of which we denote by d, and another supplied in a centralized market by a perfectly competitive firms, the demand of which we denote by n, to solve the following problem

$$v_t(p_t) = \max_{d, \, n} \frac{\left(\alpha^{\frac{1}{\theta}} d^{\frac{\theta-1}{\theta}} + (1-\alpha)^{\frac{1}{\theta}} n^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}(1-\gamma)}}{1-\gamma} \tag{18}$$

$$\text{s.t.} \quad p_t \, d + q_t \, n \leq I_t, \tag{19}$$

where $\theta > 1$ and $I_t \equiv (w_t \ell_t + D_t)/\Gamma$ is nominal income, which the customer takes as given. Nominal income depends on the household labor income $(w_t \ell_t)$ and dividends from firms ownership (D_t). The first order condition to the problem in equations (18)-(19) delivers the following standard downward sloping demand function for variety d

$$d_t(p_t) = \alpha \frac{I_t}{P_t} \left(\frac{p_t}{P_t} \right)^{-\theta} . \tag{20}$$

where $P_t = (\alpha \, (p_t)^{1-\theta} + (1-\alpha) \, (q_t)^{1-\theta})^{\frac{1}{1-\theta}}$ is the price of the consumption basket. Without loss of generality we use the price q_t as the numeraire of the economy. From the first order

conditions for the household problem, we obtain that the stochastic discount factor is given by $\beta \, \Lambda_{t+1}/\Lambda_t$, where $\Lambda_{t+s} = \int_0^\Gamma \left(c_{t+s}^i\right)^{-\gamma}/P_{t+s}^i \, di$ is the household marginal increase in utility with respect to nominal income; c_{t+s}^i denotes customer i's consumption basket in period $t+s$.

The production technology of the perfectly competitively sold good (good n) is linear in labor, so that its supply is given by $y_t^n = Z_t \, \ell_t^n$, where Z_t is aggregate productivity, and ℓ_t^n is labor demand by this firm. The production technology of the other good (good d) is also linear in labor, so that its supply is given by $y_t^j = Z_t \, z_t^j \, \ell_t^j$, where Z_t is aggregate productivity, and ℓ_t^n is labor demand by this firm, where j indexes one particular producer. Perfect competition in the market for variety n and in the labor market implies that workers are paid a wage equal to the marginal productivity of labor so that $w_t = q_t \, Z_t$. Equilibrium in the labor markets requires $\ell_t = \ell_t^n + \int_0^1 \ell_t^j \, dj$.

There are two exogenous driving processes in our economy: aggregate productivity Z and the numeraire q. We consider an economy in steady state at period t_0 where expectations are such that $Z_t = 1$ and $q_t = 1$ for all $t \geq t_0$. Notice that in this case the economy coincides with the economy described in Section 3.

E   Price distribution

Standardization of prices. In Figure 2 we compare the empirical distribution of store-level prices with the one predicted from the model. To allow for this comparison, we have to perform an appropriate standardization to remove from the data sources of heterogeneity across stores which may influence pricing but are not modeled in our framework. We construct a price index for store j in each week t as the average of the prices of the UPCs sold by the outlet, weighted for the share of total revenues they generate in the entire sample. We then use the IRI data to obtain the average market price, $p_t^{j,mkt}$, given by the period t average price index for the same basket of goods posted by retailers operating in the same Metropolitan Statistical Area where store j is located. The procedure requires the price index of the store to be computed on the subset of UPCs for which we have price information both in the retailer's data and for each store in the IRI data for every week in the sample. This substantially reduces the size of the store basket: in our data, the average store price index is computed using about 1,000 UPCs. On the upside, the procedure naturally selects the best-selling products (for which price information is more likely to appear continuously for all stores).

The statistics plotted in the histogram in Figure 2 is the (standardized) ratio of the store and the market price index where we remove permanent cross-market heterogeneity normalizing both the numerator and the denominator of the ratio by their averages computed

over time:

$$\tilde{p}^j \equiv \log(\frac{p^j/\overline{p^j}}{p^{j,mkt}/\overline{p^{mkt}}})$$

Stores whose coefficient of variation for the ratio exceeds 1 are trimmed.

Although in the stationary equilibrium of the model, the market price $(p_t^{j,mkt})$ does not vary over time (and there is only one market), the price distributions implied by the baseline and the counterfactual models plotted in Figure 2 are also normalized. This is to make them comparable with the data, where the market price can vary through time and controlling for it is necessary to isolate the price variation driven by idiosyncratic shocks. Furthemore, the normalized price is reported in deviation from its mean and divided by its standard deviation, to allow for comprability between the prices implied by the baseline and the counterfactual models, which have different variance.

Price distribution of single UPCs. Kaplan and Menzio (forthcoming,a) perform a thorough study of the properties of the distribution of prices in the grocery sector which is closely related to ours. They use Nielsen data to document the features of the price distribution of both single UPCs and bundles of goods bought by the consumers. Whereas the former displays high kurtosis; for the latter they find that the distribution is nearly normal. Since, our analysis in section Section 6 focuses on a normalized store-level price index; whereas theirs considers an index of dispersion of households' expenditure in grocery stores (not necessarily at a same store or chain), our results on store baskets and their evidence on price distribution and dispersion for bundles of goods cannot be directly compared.

However, Kaplan and Menzio (forthcoming,a) also present evidence at the single good (UPC) level. Although this is not the relevant level of observation for our study, we use our data to replicate their findings and establish that any difference between our and their results on bundles of goods comes from the choice of a different object of interest, rather than from some dishomogeneity in the underlying data.

In Figure 5 we plot a distribution comparable to the one Kaplan and Menzio (forthcoming,a) report in their Figure 2, panel (a). In particular, we take the set of all the UPCs used to compute the store-level price index whose distribution we depict in Figure 2. For each UPC k sold in store j belonging to Metropolitan Statistical Area m, we take the price posted by the store in week t $(P_{kt}^{j(m)})$ and normalize it dividing it by the mean of the prices posted in the same week for the same UPC by the stores active in the same MSA $(\overline{P_{kt}^m})$. In computing the MSA average, we weight the different stores by their market shares. Formally, we define the normalized price as follows

$$p_{kt}^{j(m)} = \frac{P_{kt}^{j(m)}}{\overline{P_{kt}^m}}$$

In Figure 5 we plot the distribution of the normalized price across UPCs, stores and weeks. Just as the comparable figure in Kaplan and Menzio (forthcoming,a), the distribution exhibits excess kurtosis. It is unimodal, has a peak close to the mean, and thicker tails than a normal distribution with the same mean and variance.

Figure 5: Distribution of normalized UPC prices
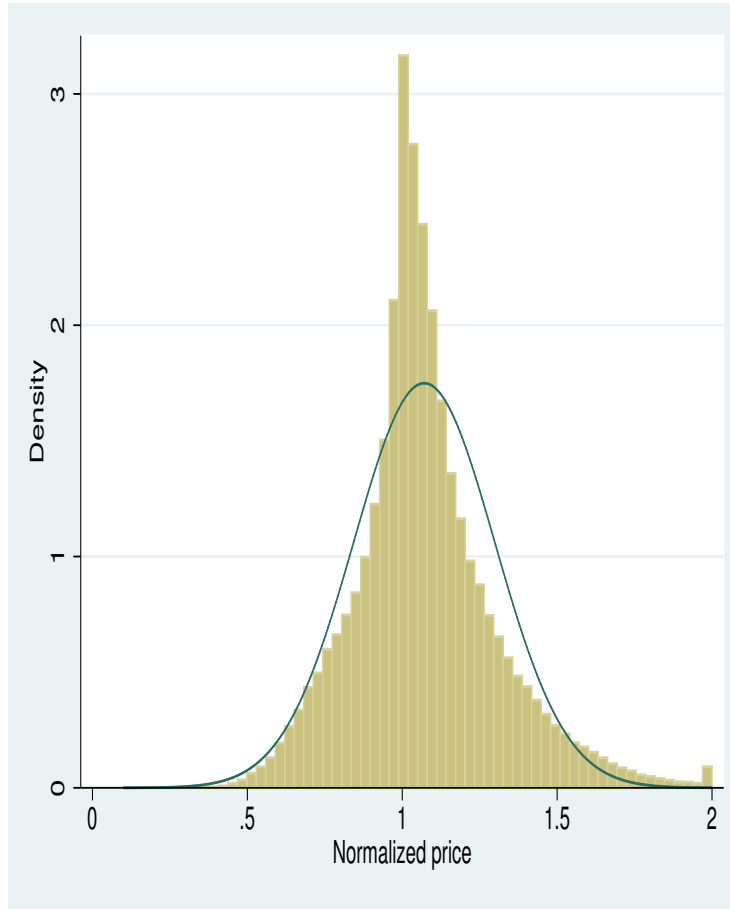
Notes: The histogram plots the distribution of normalized prices across UPCs, stores, and weeks. The normalized price of a UPC in a week is defined as the ratio of the weekly price of the UPC at a store of the chain that shared data with us to the average price of the same UPC in the Metropolitan Statistical Area where the store is located. The latter is computed using the IRI Marketing database. The set of UPCs considered is that used to compute the store-level price index whose distribution is presented in Figure 2; we discard UPCs whose coefficient of variation is larger than 1. The solid line plots the density of a normal with the same mean and variance as the empirical distribution of the normalized prices.