

PRICE DYNAMICS WITH CUSTOMER MARKETS*

BY LUIGI PACIELLO, ANDREA POZZI, AND NICHOLAS TRACHTER¹

Einaudi Institute for Economics and Finance and CEPR, Italy; Einaudi Institute for Economics and Finance and CEPR, Italy; Federal Reserve Bank of Richmond, U.S.A.

Using microdata from a U.S. retailer we document that customer turnover responds to pricing. We study the optimal price setting of a firm when its demand has an extensive margin that is elastic to price due to customers' opportunity to search for an alternative supplier. The price pass-through of idiosyncratic productivity shocks is incomplete, with the most productive firms passing through more. Firm demand is more persistent than price. Higher demand is associated with lower markups due to higher search intensity, despite flexible prices. We find empirical support for these predictions in microdata from the retail industry.

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. Several authors have emphasized that accounting for the influence of customer markets on firm pricing has relevant implications for the propagation of aggregate shocks to prices and output. These studies do not typically microfound customer reallocation across firms (Rotemberg and Woodford, 1991), or they rely on consumption habit formation abstracting from consumer flows (Ravn et al., 2006).

In this article, we study firm pricing with customer retention concerns in a model with endogenous customer dynamics and heterogeneous firm productivity. We show that the interaction of endogenous customer turnover and heterogeneous productivity delivers two main sets of results. The mechanism at the heart of the model allows us to match two important features of price and demand dynamics. To retain customers, firms have to absorb part of the productivity shocks, causing incomplete price pass-through; as firms with different productivity face endogenously different demand schedules, the price pass-through will be heterogeneous. Inertia in the customer base of a firm induces a greater persistence in firm demand than in firm productivity (Foster et al., 2016). Furthermore, we show that these predictions hold true in microdata from U.S. retailers. Finally, our study offers a methodological contribution by building a framework to study the link between firm pricing and demand that features both customer turnover and price dispersion of identical products. Hence, our setup lends itself naturally to quantification of

*Manuscript received October 2016; revised December 2017.

¹ Luigi Paciello acknowledges the financial support of the European Research Council under ERC Starting Grant 676846. We benefited from comments on earlier drafts at numerous conferences, workshops, and seminars. We thank Fernando Alvarez, Lukasz Drozd, Huberto Ennis, Mike Golosov, Bob Hall, Christian Hellwig, Hugo Hopenhayn, Eric Hurst, Pat Kehoe, Philipp Kircher, Francesco Lippi, Erzo Luttmer, Kiminori Matsuyama, Guido Menzio, Dale Mortensen, Jaromir Nosal, Ezra Oberfield, Nicola Pavoni, Facundo Piguillem, Valerie Ramey, Leena Rudanko, Katja Seim, the editor, and two anonymous referees. The views expressed in this article are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Richmond or the Federal Reserve System. Please address correspondence to: Nicholas Trachter, Research Department, Federal Reserve Bank of Richmond, 701 E Byrd St, Richmond, VA 23219. Phone: 8044050266. E-mail: trachter@gmail.com.

the key margins shaping the benefit and cost of searching by matching these observable statistics from microdata.

The link between pricing and customer base is the central tenet of our model and of the literature on customer markets more generally. Yet the existing evidence on this mechanism consists mostly of anecdotes and surveys (Blinder et al., 1998; Fabiani et al., 2007). Therefore, we begin our analysis by presenting what we believe is the first direct evidence linking firm prices with customer base evolution. We exploit scanner data documenting pricing and customer base evolution for a major U.S. retailer. The data contain information on all the shopping trips each household makes to the chain. This allows us to infer when customers leave the retailer by looking at prolonged spells without purchases at the chain. Combining these data with detailed information on the prices posted by the retailer, we are able to study the relation between a customer's decision to abandon the firm and the price of the goods she consumes there. We show that an increase in the price significantly raises the probability that the customer leaves the firm. This implies that the customer base is elastic to prices: a 1% change in the price of the goods typically consumed by the customers would raise the firm yearly customer turnover from 14% to 21%.

Next, we introduce a microfounded model of firm pricing with customer markets and focus on the interaction of the pricing response to idiosyncratic productivity shocks with customers' search intensity. The distinctive feature of our setting is that we endogenize customer dynamics. We do so by explicitly modeling the game between a firm and its customers. 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, each customer can decide to pay an idiosyncratic search cost² to observe the state of another randomly selected firm, compare it to that of her 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 the common invest\harvest trade-off (Galenianos and Gavazza, 2017): 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 framework to study how the relationship between customer and price dynamics is shaped, in equilibrium, by idiosyncratic production and search costs. Even though the pool of customers matched with each firm is characterized by the same distribution of search costs, the threshold to search varies across firms. In particular, since higher productivity is associated to lower prices on average and thus a higher value of staying in the match, the threshold to search decreases with firm productivity. As a result of the heterogeneity in customer retention concerns, we obtain heterogeneity in price pass-through of cost shocks. Although the most productive firms' pricing decision is unaffected by customer retention concerns, the remaining firms face competition for customers resulting in incomplete price pass-through of cost shocks. As more productive firms charge lower prices, they are net gainers of customers and grow faster. Less productive firms are net losers of customers, despite charging lower markups. However, given that the dynamics in the customer base are inertial, firm demand is substantially more persistent than firm prices and productivity. We find support both for the incomplete price pass-through of cost shocks and the persistence of firm demand in retailer level data.

Finally, we use our calibrated model as a laboratory to study the effect of a preference shock that shifts the utility from consumption. A positive preference shock raises customers' search intensity since it is more valuable for them to be matched with sellers offering a lower price. There are more consumers looking to switch, which incentivizes firms to lower their markup to retain them. The result is that markups drop, magnifying the effect of the demand shock on

² Modeling the market friction as a search cost suits well our application since search costs have been found to importantly affect price dispersion in clustered retail markets (Sorensen, 2000) similar to those to which our empirical application refers.

consumption, and ties to a recent but very active area of research that emphasizes the importance of consumer shopping behavior for macroeconomic dynamics (Bai et al., 2012; Coibion et al., 2015; Nevo and Wong, 2015; Kaplan and Menzio, 2016;).

Related Literature. Our article 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 that arise as the outcome of customers' optimal search decisions in response to firms' pricing. Since we model the product market friction as a search cost, we relate to other studies looking at the firm price-setting problem in models where search costs prevent customers from freely moving to the lowest price supplier (Fishman and Rob, 2003; Alessandria, 2004; Menzio, 2007). Our model has the distinctive feature of delivering both price dispersion and customer reallocation in equilibrium. This is a particularly desirable property, as it offers the chance to quantify the model, matching available statistics on price dispersion and customer turnover.

We share with the literature on deep habits (Ravn et al., 2006) the interest for the impact of aggregate shocks on markups through their effect on the elasticity of demand. The main difference is that in deep habits models there is, typically, no *extensive* margin of demand, as each consumer buys from all firms at any point in time, albeit with different habit and expenditure, and all the adjustment in demand takes place along the *intensive* margin. In our model, instead, the extensive margin plays a key role.

Several studies look at the implications of product market frictions for business cycle fluctuations. Bai et al. (2012) analyze a demand-driven business cycle model where preference shocks affect consumers' search incentives and consumption by directly impacting production efficiency, so to show up as shocks to the Solow residual. Petrosky-Nadeau and Wasmer (2015) and Kaplan and Menzio (2016) study the interaction of labor and product market frictions, linking unemployment dynamics to consumer search effort. In these models, whether consumer search amplifies or dampens the recessionary implications of higher unemployment depends on how consumer search intensity comoves with unemployment. Coibion et al. (2015) document the relationship between the household expenditure allocation across retailers and unemployment and find that households pay on average a lower price when unemployment is higher. Although we share with these papers the interest on the transmission of aggregate shocks to search intensity, we explore an additional channel through which aggregate shocks affect the opportunity cost of searching.

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. 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. Kleshchelski and Vincent (2009) examine the impact of customer markets on the pass-through of idiosyncratic cost shocks to prices in an economy where firms are identical and there is no price dispersion. 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 (2016) studies a setting where firms cannot price discriminate across customers and use sales to attract new customers. Burdett and Coles (1997) study the role of firm size for pricing when firms use the price to attract new customers. The industrial organization literature has also studied the implications of customer markets for a variety of subjects. For instance, Foster et al. (2016) stress their role in affecting firm survival, and Einav and Somaini (2013) and Cabral (2016) focus on their effect on the competitive environment.

Finally, our article relates to the literature studying equilibrium price dispersion. Although Burdett and Judd (1983), Burdett and Menzio (2018), and Menzio and Trachter (2015, 2018) obtain price dispersion as an equilibrium outcome without relying on firm heterogeneity, price

dispersion in our model is driven by the combination of linear pricing and heterogeneity in firm productivity along the lines of Reinganum (1979).

The rest of the article 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, characterize the equilibrium, and discuss its calibration. In Section 4 we present some quantitative predictions of the model and compare them with empirical evidence from our data. In Section 5 we introduce an application of the model with the goal of studying the implications of customer markets for the propagation of aggregate shocks. Section 6 concludes.

2. THE LINK BETWEEN PRICES AND CUSTOMER DYNAMICS

We use novel microdata to provide direct evidence that firm prices have an effect on the evolution of their customer base and to obtain statistics allowing us to calibrate the key parameters of our model in Section 3. In particular, we document that changes in the price of the basket of goods typically bought by a customer at a large U.S. retailer affect the probability of that customer abandoning the retailer. This result provides a compelling motivation to modeling the link between customer dynamics 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; Fabiani et al., 2007). To the best of our knowledge, we are the first to document this fact using microdata based on actual customers' decisions.

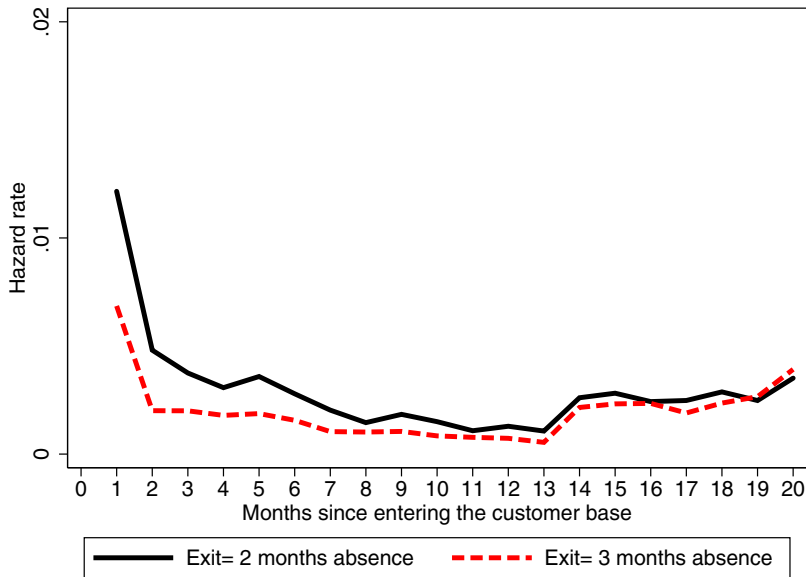
2.1. Data Sources and Variable Construction. The empirical study of the interaction of consumer shopping behavior and retail prices presents two challenges. First, we have to define what it means to exit the customer base; second, we need to identify the price to which customers react. Below we briefly describe our approach; the details are left to the Appendix.

To identify customer base evolution we rely on a data set (henceforth, "retailer consumers panel") consisting of cashier register records on purchases by a panel of households carrying a loyalty card of a large U.S. supermarket chain.³ The most important feature of this data set is that it allows us to keep track of a set of loyal customers of the chain. In fact, 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. We consider every customer shopping at the retailer in a given week as belonging to the chain's customer base in that week and assume that a household has exited the customer base when she has not shopped at the chain for eight or more consecutive weeks. We assume that the decision to exit matured in response to the prices observed the last time the customer visited the chain.

The definition of exit takes into account that brief spells without purchases do not necessarily imply that a shopper has left the chain: she may just be consuming her inventory or being on vacation. However, a regular customer is unlikely to experience a long spell without shopping for reasons other than having switched to a different chain. In fact, for the average household in our sample, only four days elapse between consecutive grocery trips, and the 99th percentile of this statistic is 28 days. This implies that the period of absence we require before determining that a customer has exited the customer base is probably a conservative choice.

In Figure 1 we plot the monthly hazard rate of exiting the customer base for our sample of customers, that is the probability of exiting the customer base of the firm as a function of the time elapsed since the first time we observe the customer shopping at the firm (tenure of

³ 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. The household identifier also allows us to track members of a same household when they lose (and replace) their individual loyalty card.



NOTES: The figure plots the hazard rate for our sample of households, where failure is defined as exit from the customer base. We report two hazard rates for different criteria to determine whether the customer has exited: two consecutive months without shopping at the chain (our baseline definition, continuous black line), and three consecutive months without shopping at the chain (dashed red line). To ensure that all individuals have similar potential length for their spells, we only consider the first spell as customer for those having multiple ones and we only retain households whose first trip at the chain occurs within the first 40 weeks in our sample.

FIGURE 1

HAZARD RATE OF EXIT FROM THE CUSTOMER BASE
[COLOR FIGURE CAN BE VIEWED AT WILEYONLINELIBRARY.COM]

the customer). We explore the sensitivity of customer base evolution to our definition of exit by displaying two hazard rates. The solid line refers to our baseline definition; the dashed line represents the hazard rate if we extend to three months the absence spell required to determine that the customer has exited. The first thing to notice is that the hazard rate is not overly sensitive to the definition of exit. The hazard rate is higher at very short tenure lengths since there we include also households who were just casual shoppers at the chain. At tenures larger than a month it is much lower, oscillating between 0.2% and 0.4%. The second noteworthy fact emerging from the plot is that the customer base is quite sticky: The implied probability of never exiting the customer base in 20 months is between 70% and 80%. These values lend support to industry estimates on the customer attrition rate surveyed in Gourio and Rudanko (2014).

The second object we need to construct is the price that affects each customer's exit decision. In our data, households shop at the retailer's store for a number of different goods, defined at the UPC level. In the consumer data, however, we only observe a price of a good when it is purchased by some household in the sample. To construct a price for the basket of grocery goods usually purchased by each households we need to observe the price of each item belonging to the basket of each household whether or not the UPC is actually purchased. We do so by exploiting data, previously used and documented by Eichenbaum et al. (2011), on store level weekly⁴ revenues and quantities for the full set of UPCs purchasable at stores of the chain (henceforth, "retailer price data"). We recover the weekly price for each UPC by dividing the revenues by the quantity sold in the week. Then, we construct the price paid by customer i , shopping at store j in week t , for its basket as the average price of the goods included in

⁴ The retailer changes the price of each good at most once per week; hence the frequency of the retailer price data captures the entire time variation.

TABLE 1
EFFECT OF THE PRICE OF THE BASKET ON THE PROBABILITY OF EXITING THE CUSTOMER BASE

	Exiting: Missing at least Eight Consecutive Weeks					
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(p_{it})$	0.013*	0.14**	0.16**	0.15**		0.034*
	(0.008)	(0.066)	(0.080)	(0.070)		(0.020)
$\log(p_{it}) \times \text{Walmart entry}$			0.018**			
			(0.009)			
$\log(\bar{p}_{it})$	0.001	0.001		0.001	0.000	0.001
	(0.001)	(0.001)		(0.001)	(0.001)	(0.001)
$\log(p_t^{j(t)})$					0.01	
					(0.03)	
Tenure	-0.002***	-0.002***	-0.003***	-0.003***	-0.003***	-0.001***
	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	52,670	52,670	66,182	52,101	52,670	52,670

NOTES: An observation is a household–week pair. The results reported are calculated through two-stage 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 an instrument for the logarithm of the price of the basket. Column (1) reports OLS estimates of Equation (2); columns (2)–(6) are instrumental variable estimates where the price of the basket is instrumented with the replacement cost of the basket. In column (4), the exit of the customer is assumed to have occurred in the first week of absence in the eight- (or more) week spell without purchase at the chain instead of the week of the last shopping trip before the hiatus. In column (5), the price of the household basket is substituted with a price index for the store where the customer shops (identical for all the customers shopping at the same store). In column (6), we allow the weights used in the construction of the household basket price to vary between the first and the second year in the sample. 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. Standard errors are in parentheses and account for within-household correlation through a two-steps feasible-GLS estimator. *** Significant at 1%; ** significant at 5%; and * significant at 10%.

the basket, weighted for the share of grocery expenditure of the household she represents. Namely:

$$(1) \quad p_{ijt} = \sum_{k \in K_i} \omega_{ik} p_{kjt}, \quad \omega_{ik} = \frac{\sum_t E_{ikt}}{\sum_{k \in K_i} \sum_t E_{ikt}},$$

where p_{kjt} is the price of UPC k in week t at the store j where customer i shops, K_i represents the collection of UPCs belonging to household i 's basket, and E_{ikt} is the expenditure (in dollars) by customer i in UPC k in week t . The latter two objects are measured using the retailer consumers' panel. It is important to notice 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}).

2.2. *Evidence on Customer Base Dynamics.* 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:

$$(2) \quad Exit_{it} = b_0 + b_1 \log(p_{it}) + b_2 \log(\bar{p}_{it}) + b_3 tenure_{it} + X'_i c + \varepsilon_{it}.$$

Our main interest is on the coefficient of the retailer price of the basket, b_1 , which is a measure of the elasticity of the exit decision to price.⁵ 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., fluctuations in the price of gas affect differently stores at different distances 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. We do not need 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 (2015) use the Kielts–Nielsen scanner data to provide ample evidence for this type of variation. They report that the bulk of price dispersion arises not from the difference between high-price and low-price stores but from dispersion in the price of a particular good (or product category) even among stores with a similar overall price level.

Since the retailer price in Equation (2) can be endogenous if the chain conditions to variables unobserved to the econometrician that also influence the customer's decision to leave, we instrument it using information on replacement cost for each UPC included in the retailer price data.⁶ We use the UPC level replacement cost to construct the cost of the customer's basket with a procedure analogous to the one we followed to obtain the price of the basket: We calculate it as the weighted average of the replacement cost of the UPCs included in the basket.

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.⁷ These data allows us to compute for every customer the average (cross retailers) price of her basket in the market where she lives (\bar{p}_{it}). 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 the holiday season).

The additional controls in our specification account for sources of customer heterogeneity that could influence their exit decision. The limited number of exits occurring in our sample implies that the within unit variation in the dependent variable is low. Therefore, we cannot control nonparametrically for cross-household heterogeneity using household or store fixed effects. Instead, we include in our specification a rich set of covariates that control for the main characteristics affecting store choice: demographics, location, market characteristics, and tenure. The demographic variables (age, income, and education) are matched from Census 2000. We calculate, using data on grocery shop location by Reference US, both the distance (in miles) between a household's residence and the closest store of the chain and that to the closest supermarket of a competing brand. We account for market structure by controlling for the total number of supermarket stores in the zip code of residence of the customer. To pick up the heterogeneity in the type of goods different customers include in their basket, we control

⁵ To ease notation, we have dropped the j superscript: it is implicit that p_{it} is the price of the basket purchased by consumer i in week t at the store j where she usually shops.

⁶ This represents the replacement cost for the chain, that is, 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) treat this measure as a good approximation of the retailer's marginal cost.

⁷ A detailed description of the data can be found in Bronnenberg et al. (2008). All estimates and analyses in this article 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 the Appendix.

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*.

In column (1) we report the OLS estimates of the linear probability model. The effect of the household's basket price on the probability of exiting is positive and significant. However, the basket price could be endogenous: Stores with better unobserved characteristics will both set higher prices and suffer fewer exits. As a result, the elasticity of exit to the price of the basket is likely to be biased toward zero.

Consistent with this view, the baseline IV specification in column (2) shows the basket price posted by the retailer to have a much more significant impact on the probability of leaving. A weekly price elasticity of the customer base equal to 0.14 implies that if the retailer's prices were 1% higher for a full year, the customer base would decrease 7%. The coefficient on the competitors' price, which we would expect to enter with a negative sign, is 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 data set contains price information only on a subset of the goods purchased by households in our sample, 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 is 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.

Additional columns in Table 1 present robustness checks of our main result. 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, that is, 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 our 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 modify the assumption on the imputation of the date of exit. Instead of 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 is unaffected. In column (5), we replace the price of the individual basket with a price index for the store where she buys ($p_t^{j(i)}$). The store price is a price index of a composite bundles of goods for each store so to accommodate the multiproduct nature of grocery retailing (Smith, 2004), and its construction resembles that of the price index for the customer individual baskets. It is calculated as the average of the prices of the goods sold by the store, weighted by the amount of revenue they generate.⁸ Formally the price index for store j in week t is

$$(3) \quad p_t^j = \sum_{u \in A_j} \omega_u^j p_{ut}^j, \quad \omega_u^j = \frac{\sum_t R_{ut}^j}{\sum_{u \in A_j} \sum_t R_{ut}^j},$$

⁸ In principle, we would want to include the prices of all the UPCs carried by the store. In practice, this is not possible because the information on price is missing for some UPCs in certain store-weeks. Therefore, the price index for the store is computed using a constant set of UPCs for which we have a complete time series of prices at the store during our sample.

where p_{ut}^j is the price of UPC u in week t at store j , A_j is the set of goods in assortment at store j , and R_{ut}^j are revenues from UPC k to store j in week t . By construction, $p_t^{j(i)}$ is identical for all the customers shopping in the same store. The coefficient on $p_t^{j(i)}$ is negative and not significant, confirming the importance of being able to construct individual specific baskets in order to make inference on customers' behavior.

Since we employ time-invariant weights in constructing the household basket price, we are introducing measurement error if the set of products purchased by a household changes. To gauge the impact of changes in the composition of the basket, in column (6) we estimate a specification where the basket price is calculated with time-varying weights. In particular, for all the households shopping at least 20 times in each of the two years in our sample, we construct separate weights for the first and the second year. For the remaining 15% of the households in our data, we still compute the basket price using a single set of weights. The qualitative result aligns with the finding in our baseline specification—the basket price positively and significantly affects the probability of leaving the chain—although the implied elasticity is lower.⁹

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 does the simulation yield a price coefficient that is positive and significant at 5%.

3. THE MODEL

The economy is populated by a measure one of firms producing a homogeneous good and by a measure one 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. Firms produce the same homogeneous good. We assume a linear production technology $y = z\ell$ where ℓ is the production input, and z is the firm-specific productivity. Idiosyncratic productivity is distributed according to a conditional cumulative distribution function $F(z'|z)$ with bounded support $[z, \bar{z}]$. 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 p denotes the price, the constant $w > 0$ denotes the marginal cost of the input ℓ , and the function $d(\cdot)$ is a downward sloping demand function.¹⁰ We assume that profits per customer are single-peaked in p .

Firms differ not only in their idiosyncratic productivity but also in the mass of customers buying from them. In particular, we denote by m the firm's *customer base* that is defined as the mass of customers who bought from that firm in the previous period, adjusted for an exogenous attrition rate δ . Starting from a given customer base m , the mass of customers actually buying from the firm in the current period is determined in equilibrium, and we conjecture, and later verify, that it is given by the function $\mathcal{M}(m, p, z)$ depending on the price and productivity of the firm in the current period as well as on the customer base.

We assume a constant probability κ of a firm exiting the market. Once a firm exits the market it loses all of its customers and its value is zero. An exiting firm is replaced by a new firm

⁹ We note that time-varying weights open the door to another set of concerns. Namely, unobserved shocks affecting households' decisions on the composition of their basket could also influence their decision to stay in the customer base of the firm. We have no clear way of establishing if the quantitative differences between this specification and baseline one is due to this potential endogeneity channel.

¹⁰ In Section 5, we extend this framework, adding a model of the labor market to endogenize the wage w .

that starts with a customer base m_0 and draws a productivity z_0 from the invariant productivity distribution $\bar{F}(z)$ associated to the conditional distribution $F(z'|z)$.¹¹

We study a stationary Markov Perfect equilibrium where pricing strategies are a function of the current state. Firms set prices every period without commitment and without discriminating across customers.¹² As there are no aggregate shocks, the aggregate state is constant and the relevant state for the firm problem in period t is the pair $\{z, m\}$. The firm pricing problem in its recursive form solves

$$(4) \quad \begin{aligned} \tilde{W}(z, m) = \max_p \mathcal{M}(m, p, z) \pi(p, z) + \beta(1 - \kappa) \int_{\underline{z}}^{\bar{z}} \tilde{W}(z', m') dF(z'|z) \\ \text{s.t. } m' = (1 - \delta) \mathcal{M}(m, p, z), \end{aligned}$$

where $\tilde{W}(z, m)$ 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. In fact, it influences the mass of customers buying from the firm in the current period and, if there is persistence in the evolution of the customer base, 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). These models assume a functional form for the evolution of the customer base 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 conjectured law of motion for customers preserves this standard structure and is given by

$$(5) \quad \mathcal{M}(m, p, z) \equiv m \Delta(p, z).$$

This similarity notwithstanding, there are two important innovations that we introduce. First, although the customer evolution is typically characterized with ad hoc functional form assumptions, our $\Delta(\cdot)$ function is endogenous and results from the solution to a game between the firm and its customers. It 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 5.

Second, 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 us to study the mapping from the distribution of productivities to the distribution of prices. It also proves useful when we bring the model to the data, since having heterogeneity in productivity helps us to match the cross-sectional variation in prices. Our formulation does, however, 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 1 in m ,

¹¹ The invariant distribution is obtained by solving $\bar{F}(z) = \int_{\underline{z}}^{\bar{z}} F(z|x) d\bar{F}(x)$ for all $z \in [\underline{z}, \bar{z}]$.

¹² See Nakamura and Steinsson (2011) for a model of pricing with customer markets where a commitment to a price path can be sustained in equilibrium.

that is, $\tilde{W}(z, m) = m \tilde{W}(z, 1) \equiv m W(z)$, where using Equation (4) and $\mathcal{M}(m, p, z) = m \Delta(p, z)$, it is immediate to show that $W(z)$ solves¹³

$$(6) \quad W(z) = \max_p \Delta(p, z) \pi(p, z) + \Delta(p, z) \beta (1 - \eta) \int_z^{\bar{z}} W(z') dF(z' | z),$$

where $\eta \equiv \kappa + \delta(1 - \kappa)$ is the probability of exogenous dissolution of the firm–customer match due to either firm or customer random exit. The relevant state to the firm pricing problem is its productivity, as the level of the customer base affects the firm value multiplicatively. The solution to the firm problem in Equation (6) gives an optimal pricing strategy that depends on productivity. We denote it by $\hat{p}(z)$.

We emphasize that, although the initial level of the customer base does not affect the optimal price, its evolution does. A change in the price affects the growth rate of the customer base, that is, the value of $\Delta(p, z)$, and given the persistence of the customer base, it affects the firm value in the current period as well as in future periods. Our framework is well suited to capture the relationship between firm prices and customer dynamics when this is driven by variation in idiosyncratic productivity; extending it to encompass how firm size affects this relationship is an interesting direction for future research.

The objective of the firm maximization problem can be expressed as the 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

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

where we define $\varepsilon_m(p, z) \equiv \partial \log(\Delta(p, z)) / \partial \log(p)$ as the *extensive margin* elasticity of demand. The function $\Pi(p, z)$ is maximized at the static profit maximizing price,

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

where we define $\varepsilon_d(p) \equiv \partial \log(d(p)) / \partial \log(p)$ as the *intensive margin* elasticity of demand. 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. Due to concerns about customer dynamics, the optimal price is in general different from the one that maximizes static profits. The optimal price balances the marginal benefit of an increase in price (more profit per customer) with the cost (decrease in the customer base). If the growth in the customer base is nonincreasing 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 pushed downwards with respect to it, that is, $\hat{p}(z) < p^*(z)$. 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 .¹⁴

3.2. The Problem of the Customer. Customers value the good sold by the firms described in the previous section according to the function $v(p)$, denoting the customer surplus associated to the demand function $d(p)$. We assume that $v(p)$ is continuously differentiable with $v'(p) < 0$,

¹³ 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.

¹⁴ In Section 3.3, we derive the necessary equilibrium properties that guarantee that these conditions are satisfied.

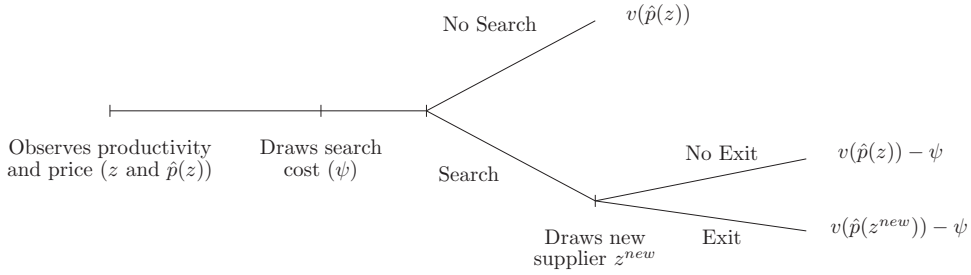


FIGURE 2

THE PROBLEM OF A CUSTOMER MATCHED TO A FIRM WITH PRODUCTIVITY z

and bounded above with $\lim_{p \rightarrow 0^+} v(p) < \infty$. These properties are satisfied in standard models of consumer demand.

Each customer starts the period in the customer base of the firm she bought from in previous period. At the beginning of every period, a customer can be randomly reallocated to a new entrant because either the firm she was matched with exited (with probability κ) or with probability δ the customer herself leaves for random reasons (for instance, she moved to a different city). We allow for random exit to acknowledge that price dynamics are unlikely to account for all the exits observed in the data. Conditional on a firm surviving, random exit is i.i.d. across customers of that firm.

After random reallocation has taken place, 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 $\psi \geq 0$ is measured in units of customer surplus, it is idiosyncratic to each customer, 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 being proportional to its customer base m . As in Burdett and Vishwanath (1988) and Fishman and Rob (2005), this assumption captures the idea that consumers search for new suppliers not by randomly sampling firms but by randomly sampling other consumers. On the technical side, this is the key assumption that will allow us to solve for an equilibrium where the value of a firm scales up multiplicatively with its customer base. Conditional on searching, the customer observes the state of the new match and then makes 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 firm once she exits its customer base. Figure 2 summarizes timing and payoffs of the problem of the customer.

We next characterize the customer problem. Let $V(p, z, \psi)$ denote the value function of a customer i who has drawn a search cost ψ and is matched to firm j , which has current productivity z and posted price p . This value function solves the following problem:

$$(9) \quad V(p, z, \psi) = \max\{\bar{V}(p, z), \hat{V}(p, z) - \psi\},$$

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 she does search. The value in the case of not searching is

$$(10) \quad \bar{V}(p, z) = v(p) + \beta \mathbb{E}_G \left[(1 - \eta) \int_{\underline{z}}^{\bar{z}} V(\hat{p}(x), x, \psi') dF(x|z) + \eta \int_{\underline{z}}^{\bar{z}} V(\hat{p}(x), x, \psi') d\bar{F}(x) \right].$$

We notice that the state of the firm problem depends 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. We also notice that the state of the firm problem includes the current price p , despite the fact that 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. Finally, the expectation operator $\mathbb{E}_G[\cdot]$ refers to the realization of future search costs that are drawn from the i.i.d. distribution G . The value function $\bar{V}(p, z)$ is strictly decreasing in p and increasing in z . Given the specifics of the search technology, the value to the searching customer is given by

$$(11) \quad \hat{V}(p, z) = \int_{-\infty}^{+\infty} \max \{ \bar{V}(p, z), x \} dH(x),$$

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.

We are now ready to describe the customer's optimal search and exit policy rules. Such policies are characterized by simple cut-off rules. A customer matched to a firm with productivity z charging price p searches if she draws a search cost $\psi \leq \hat{\psi}(p, z)$, where

$$\hat{\psi}(p, z) \equiv \int_{\bar{V}(p, z)}^{\infty} (x - \bar{V}(p, z)) dH(x) \geq 0$$

is the threshold to search. Conditional on searching, the customer exits if she draws a new firm promising a continuation value \bar{V}^{new} larger than the current match, that is, $\bar{V}^{new} \geq \bar{V}(p, z)$. Note that the threshold $\hat{\psi}(p, z)$ is strictly increasing in p . 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 notice that if the continuation value is increasing in z (a sufficient condition is that $\hat{p}(z)$ is decreasing), then the threshold $\hat{\psi}(p, z)$ is decreasing in z .

3.3. Equilibrium. In this section we define an equilibrium, discuss its existence, and characterize its general properties. First we derive the equilibrium dynamics of the customer base as a function of price and productivity, given the optimal search and exit strategy of customers. Given customers' optimal decision rule, the mass of customers buying from a firm with productivity z and charging price p is given by $\mathcal{M}(m, p, z) = m \Delta(p, z)$, with

$$(12) \quad \Delta(p, z) \equiv 1 - \underbrace{G(\hat{\psi}(p, z))(1 - H(\bar{V}(p, z)))}_{\text{customer outflow}} + \underbrace{Q(\bar{V}(p, z))}_{\text{customer inflow}},$$

where $G(\hat{\psi}(p, z))$ is the fraction of customers searching from the firm customer base, a fraction $1 - H(\bar{V}(p, z))$ of which actually finds a better match and exits the customer base of the firm. The mass m is the probability that searching customers in the whole economy draw the firm as a potential match. The function $Q(\bar{V}(p, z))$ denotes the equilibrium mass of searching customers currently matched to a firm with continuation value smaller than $\bar{V}(p, z)$. Therefore, the product $m Q(\bar{V}(p, z))$ amounts to the mass of new customers entering the customer base. Equation (12) verifies the conjecture about the equilibrium customer dynamics made in Section 3.1.

We study equilibria where the continuation values to customers are nondecreasing 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, therefore offering higher values to customers.

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 nondecreasing in z and $\hat{p}(z) \geq p^*(\bar{z})$ for all $z \in [z, \bar{z}]$. A stationary equilibrium is then

- (i) search and exit strategies that solve the customer problem in Equations (9)–(11);
- (ii) a firm pricing strategy $\hat{p}(z)$ that solves Equation (7) for each z ;
- (iii) a customer base for new entrant firms $m_0 = \eta/\kappa$, with $\eta = \kappa + \delta(1 - \kappa)$;
- (iv) a dynamic of the customer base at a surviving firm with productivity z given by $m' = (1 - \delta) \Delta(\hat{p}(z), z)m$, where $\Delta(\cdot)$ is given by Equation (12);
- (v) an invariant distribution of customers $K(\cdot)$ over productivities that for each z solves

$$K(z) = (1 - \eta) \int_{\underline{z}}^z \int_{\underline{z}}^{\bar{z}} \Delta(\hat{p}(x), x) dF(s|x) dK(x) + \eta \bar{F}(z);$$

- (vi) two invariant distributions, $H(\cdot)$ and $Q(\cdot)$, that solve

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

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

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

PROPOSITION 1. Let productivity be i.i.d. with $F(z'|z_1) = F(z'|z_2)$ continuous and differentiable for any z' and any pair $(z_1, z_2) \in [z, \bar{z}]$, 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 $p^*(\bar{z}) < \hat{p}(z) < p^*(z)$ for $z < \bar{z}$, implying that $\mathcal{V}(z)$ is strictly increasing. Moreover, the optimal markups are given by

$$(13) \quad \mu(p, z) \equiv \frac{p}{w/z} = \frac{\varepsilon_d(p)}{\varepsilon_d(p) - 1 + \varepsilon_m(p, z) \Pi(p, z)/(d(p) p)},$$

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

- (ii) $\hat{\psi}(\hat{p}(z), z)$ is strictly decreasing 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 $\Delta(\hat{p}(z), z) < 1$.

The proof of the proposition can be found in Appendix A.3.1. Here we just point out that, although the results of Proposition 1 refer to the case of i.i.d. productivity shocks, numerical results in Section 4 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 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.¹⁵ More productive firms charge lower prices and, therefore, offer higher continuation value to customers. If all the firms had the same productivity, Proposition 1 would imply a unique equilibrium where the price is that maximizing static profits, $p^*(z)$, and as a result the customer base of every firm would be constant.¹⁶ 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 (13) depend on three distinct terms: $\varepsilon_d(p)$, $\varepsilon_m(p, z)$, and $\bar{\pi}(p, z) \equiv \Pi(p, z)/(d(p)p)$. 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, that is, $m \Delta(p, z) d(p)$, 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 $\bar{\pi}(p, z)$, 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 markup, the larger the product $\varepsilon_m(p, z) \bar{\pi}(p, z)$.

An important assumption in Proposition 1 is that the search cost distribution $G(\cdot)$ is not degenerate at $\psi = 0$. We introduce the restriction in order to simplify the solution of the model. First, the restriction implies that the first order condition of the problem as described in Equation (7) is necessary for an equilibrium. If the restriction was not in place, the growth rate of the customer base of a firm charging price p and with productivity z , that is, $\Delta(p, z)$ (see Equation (12)), would not be a differentiable function of the price set by the firm (through the exit threshold $\hat{\psi}(p, z)$). In fact, there would be a region on the price space where a small change in price would have a discrete change in the customer base of a firm. The implied kink in the customer base of the firm translates into a kink on firm's profits, which in turn makes the optimal price of a firm $\hat{p}(z)$ not necessarily characterized by the first-order condition of profits with respect to prices. Second, the restriction implies that no customer samples new firms for free. This restriction guarantees that we abstract from equilibria as the one studied in Burdett and Mortensen (1998).

Finally, it is useful to discuss two interesting limiting cases of our model reported in the following corollary (see Appendix A.3.2 for the proofs).

COROLLARY 1. *Let search costs be scaled as $\psi \equiv \iota \tilde{\psi}$, where $\iota > 0$. That is, let the value function in Equation (9) be $\max\{\tilde{V}(p, z) \hat{V}(p, z) - \iota \tilde{\psi}\}$. Let $\pi(p^*(\bar{z}), \bar{z}) > 0$ and the assumptions of Proposition 1 be satisfied.*

- (1) *Let $\iota \rightarrow \infty$. Then, (i) the optimal price maximizes static profits, that is, $\hat{p}(z) \rightarrow p^*(z)$ for all $z \in [\underline{z}, \bar{z}]$, and (ii) there is no search in equilibrium.*
- (2) *Let $\iota \rightarrow 0^+$. Then, (i) there is no price dispersion, that is, $\hat{p}(z) \rightarrow p^*(\bar{z})$ for all $z \in [\underline{z}, \bar{z}]$, and (ii) there is no search in equilibrium.*

¹⁵ For tractability we abstract from (possible) equilibria where symmetric firms charge different prices. We also notice that allowing firms to use two-part tariffs would still generate price dispersion as customers are heterogeneous in their search costs. A pricing strategy that would eliminate price dispersion in our setup is a customer-specific two-part tariff.

¹⁶ This special case is useful to understand our relation to Diamond's (1971) results. Our model delivers equilibrium price dispersion as a result of heterogeneity in productivity. If productivity was homogeneous as in Diamond (1971) the monopoly price would be the only equilibrium price.

These two limiting cases highlight the tight relationship among the size of the search cost, competition for customers, and price dispersion. The first limiting case concerns 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, that is, $\hat{p}(z) = p^*(z)$ for all $z \in [z, \bar{z}]$. There is price dispersion, and there is no search in equilibrium.

The second limiting case concerns the equilibrium when search costs become arbitrarily small. As long as search costs are not equal to zero, equilibrium pricing in the model is characterized by the first-order condition presented in Equation (7), and the firm with highest productivity, that is, \bar{z} , sets the lowest possible, price $p^*(\bar{z})$. So when search costs are arbitrarily close to zero, all customers would search for a new supplier so long as they are matched to a firm with $p > p^*(\bar{z})$. Firms have then incentives to cut their price until they set $p = p^*(\bar{z})$, as there is a discrete gain in customers. Hence the only possible equilibrium is one where all firms charge $p = p^*(\bar{z})$. As a result, there is no price dispersion and customers do not search. We notice that there is a discontinuity in equilibrium prices at $\iota = 0$. When search costs are actually equal to zero, the first-order condition is no longer necessary and sufficient to determine optimal prices, and the equilibrium price is the competitive price of the most productive firm, $p = w/\bar{z}$, for the same reasons present in Diamond (1971).

3.4. Parametrization of the Model. We assume that a period in the model corresponds to a month and fix the discount rate to $\beta = 0.995$. We assume that households have logarithmic preferences over consumption given by $\log(c)$. Consumption is defined as a composite of two types of goods, $c \equiv (d^{\frac{\theta-1}{\theta}} + n^{\frac{\theta-1}{\theta}})^{\frac{\theta}{\theta-1}}$, with $\theta > 1$ governing the demand elasticity. The first good (that we label d) is supplied by firms facing product market frictions as described above; the other good (n) acts as a numeraire and is sold in a frictionless centralized market. We set $\theta = 3.62$ so that the cross-firms average price-elasticity of demand (including both extensive and intensive margins) is equal to 2.7, which is the elasticity of the logarithm of the dollar grocery expenditure to the basket price estimated for households using data from our consumer retailer panel.¹⁷ The customer budget constraint is given by $p d + n = I$, where I is the household's total expenditure, which we normalize to one.¹⁸ We set $\delta = 0.004$ corresponding to a yearly customer turnover of 5% that is independent of price variation. This number is chosen to match the average U.S. cross-counties migration rate estimated by the Census Bureau for the period 2004–5. We set $\kappa = 0.008$ corresponding to a yearly firm exit rate of 10%.¹⁹

We assume that the logarithm of idiosyncratic firm productivity evolves according to an AR(1) process, $\log(z') = \rho \log(z) + \varepsilon$, where ε is i.i.d. normally distributed, $\varepsilon \sim N(0, \sigma)$.²⁰ We estimate ρ and σ by matching the autocorrelation and cross-firms dispersion of log-prices as estimated by Kaplan and Menzio (2015). In particular, they report a probability of about 0.4 that the average price of a store remains in the same quartile of the price distribution after one year, which allows us to estimate a monthly autocorrelation of $\rho = 0.95$;²¹ we use σ to target the component of price dispersion of identical products (UPCs) due to the store-specific and

¹⁷ The average elasticity of demand is obtained in the model by summing over the intensive and extensive margins at the firm level and then aggregating over firms: $\int_{\bar{z}}^z [\varepsilon_m(\hat{p}(z), z) + \varepsilon_d(\hat{p}(z))] dK(z)$.

¹⁸ In Section 5 we show that I can be derived based on a model of the labor market where the numeraire good n is produced by a competitive representative firm with linear production function and unitary labor productivity.

¹⁹ Such a number is on the high end of the range of estimates reported by Dunne et al. (1988) for U.S. firms. We use a relatively high firm exit rate to compensate for the absence in our model of customer returning to a previously visited store that, everything else being equal, would reduce the incentives to retain a customer.

²⁰ Operationally, we approximate the AR(1) through a discrete Markov chain with the methodology proposed by Tauchen (1986) and we discretize the values of $\log(z)$ over 25 points.

²¹ See Table 4 in Kaplan and Menzio (2015).

TABLE 2
PARAMETERS CALIBRATION

	Estimates	Target
Random customer attrition, δ	0.004	Yearly migration rate, 5%
Firm death rate, κ	0.008	Yearly firm death, 10%
Elasticity of substitution, θ	3.62	Price elasticity of demand, 2.7
Productivity persistence, ρ	0.95	Price quartile yearly persistence, 0.4
Productivity shocks volatility, σ	0.13	Log-price dispersion, 0.057
Scale of search cost, λ	1.33	Yearly customer attrition, 0.22
Coeff. variation search cost, $1/\sqrt{\zeta}$	0.667	Price sensitivity of exit b_1 , 0.14

store-good-specific components, that is, 0.057.²² Finally, we assume that customers draw their search cost from a Gamma distribution with shape parameter ζ and scale parameter λ . The parameter λ governs the scale of the search cost distribution. A higher λ implies higher search cost on average and, for given dispersion in prices, lower propensity to search of customers (see the discussion in Section 3.3). We estimate it by matching an average yearly customer attrition rate of 22% per year, as obtained from our analysis in Section 2.²³ We calibrate ζ by matching the average effect of log-prices on the exit probability predicted by the model to its counterpart in the data, measured by $b_1 = 0.14$ in column (2) of Table 1. The larger ζ , the more concentrated distribution of search costs, so that a change in prices causes a larger change in the fraction of customers searching. More details on our estimation algorithm for ρ , σ , λ , and ζ are provided in Appendix A.4. Table 2 collects the calibrated parameters and associated targets. Search costs incurred in the entire cross section of households in a given period are equivalent to a 0.08% reduction in aggregate expenditure I . The average search cost incurred conditional on searching is equivalent to a 4% reduction in monthly grocery expenditure.²⁴ The estimated process for productivity implies a cross-sectional dispersion in log-production costs of 0.25. We notice that such cross-sectional dispersion in costs is about 5 times the cross-sectional dispersion in prices, which is the mirror image of the incomplete pass-through of cost shocks to prices in our model: On average, the pass-through is 26%. Such a prediction is consistent with estimates of low price pass-through from our data, discussed later in Section 4.²⁵

4. PRICE PASS-THROUGH AND DEMAND PERSISTENCE

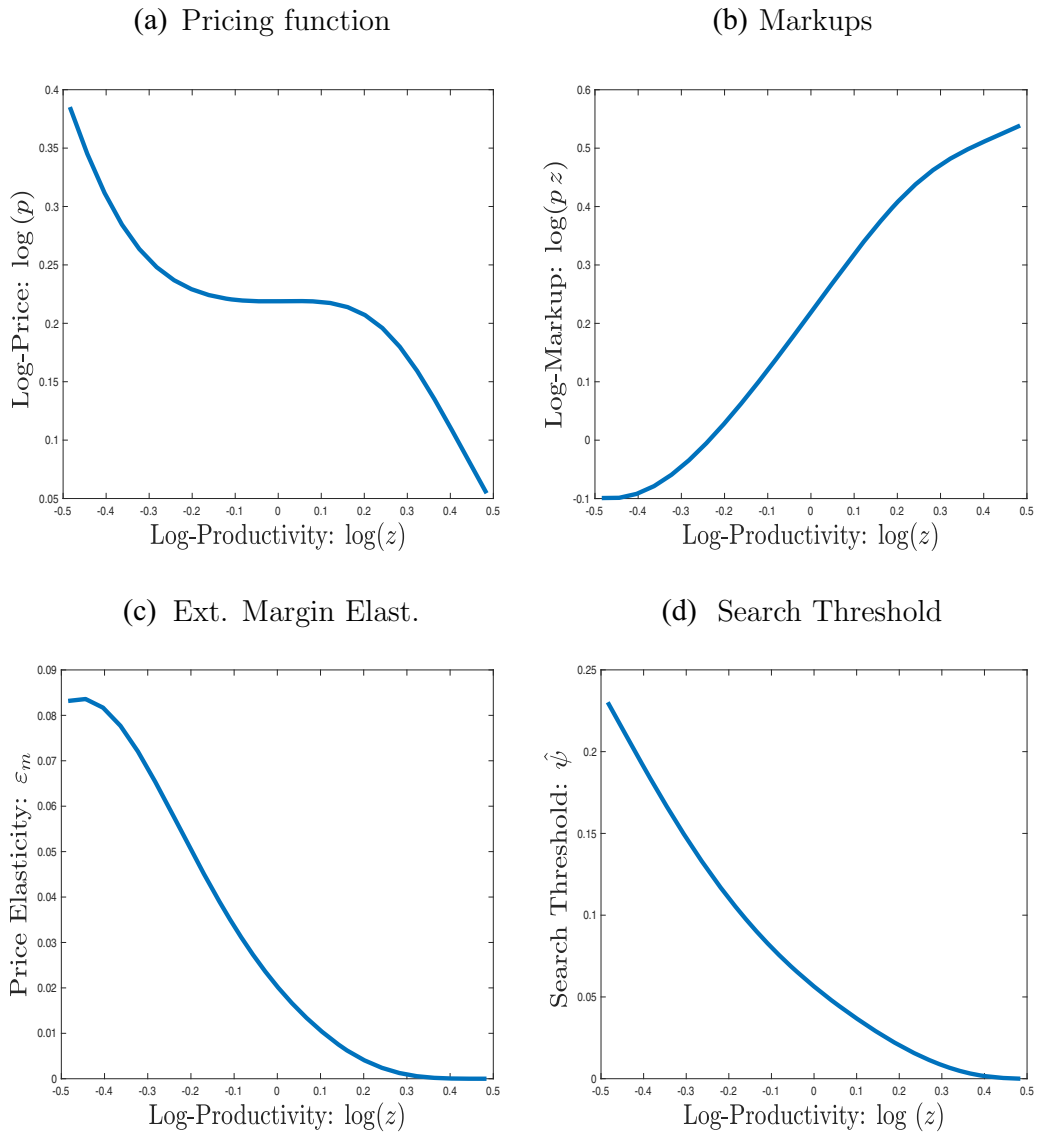
In this section we illustrate the properties of our model using the parametrization introduced in Section 3.4. We highlight the implications for the two main aspects of interest of our analysis—the distribution of prices and the evolution of the customer base across firms—and explain how they derive from the presence of the extensive margin of demand associated to consumers searching and how it affects the price pass-through of idiosyncratic shocks. We compare the qualitative model prediction with the available empirical evidence to validate our mechanism.

²² This is the type of variation that we have used to estimate the customer exit probability in Equation (2). Kaplan and Menzio (2015) estimate the cross-sectional price dispersion for identical products measured as UPCs to be 0.19, on average, across different products; 30% of such dispersion is due to the store-specific and store-good-specific components.

²³ Such a number is in the range of estimates reviewed by Gourio and Rudanko (2014) for different sectors—the range is 10%–25%.

²⁴ Given that in our data average monthly expenditure is \$424, a back of the envelope calculation provides that the average search cost is \$17. This number is consistent with the empirical industrial organization literature that aims to estimate switching costs. For example, Hong and Shum (2006) find that average search costs for books can be as high as \$51 (for the sequential search model that maps closer to our approach), Honka (2014) finds that average search cost range from \$35 to \$170 for auto insurance, and Giulietti et al. (2014) find that the average search cost for the electricity market in Britain ranges from £20 to £50.

²⁵ We notice that the interdecile of the log-productivity distribution implied by our calibration is 0.9, below the value estimated across U.S. retailers by Decker et al. (2016).



NOTES: The histogram in panel (a) plots the optimal log-prices as a function of productivity. In panel (b), we plot the optimal markups as a function of productivity. In panel (c), we plot the extensive margin elasticity. In panel (d), we plot the threshold below which consumers search. All the plots were obtained using the parametrization of the model summarized in Table 2.

FIGURE 3

EQUILIBRIUM PRICE DYNAMICS
[COLOR FIGURE CAN BE VIEWED AT WILEYONLINELIBRARY.COM]

Price dynamics. In panel (a) of Figure 3, we plot the optimal price as a function of productivity. The relationship is flat at intermediate levels of productivity and steep at low and high levels of productivity. It follows that our model delivers two implications related to the pass-through of idiosyncratic productivity shocks. First, the pass-through is incomplete, with firms passing through an average of 26% of cost shocks to prices. Second, the pass-through is heterogeneous in productivity: high for firms in the right and the left tail of the productivity distribution and low for firms of average productivity.

The incomplete price pass-through is due to the endogeneity of optimal markups, which positively comove with firm productivity (panel (b)). As more productive firms charge persistently lower prices, they are more attractive to customers and face lower demand elasticity. The heterogeneity in pass-through is explained by the heterogeneity in demand elasticity (panel (c)).²⁶ The most productive firms face low risk that customers will leave since they offer high expected value to their customers relative to the average firm, so that the threshold to search is lower (panel (d)). As a result, a low fraction of their customers is searching and an even lower fraction ends up exiting the customer base: Only customers drawing tiny search costs will search, and among those the ones that exit need to find a better match. This means that variations in the threshold to search are associated to small variations in the mass of customers searching. Therefore, these firms face low extensive margin elasticity and can afford nearly complete pass-through, enjoying high markups.

As productivity decreases, the threshold to search, as well as the probability of drawing a better match, increases. Price variations (which affect the search threshold) are associated with significant changes in the mass of customers searching. Therefore, firms with productivity around the average face high extensive margin elasticity and, as Equation (13) dictates, will offset increases in production cost with reduction in markups. This explains the flatness of the pricing function in that region. Finally, as productivity approaches the left tail of the distribution, the extensive margin elasticity flattens again. This happens because customers paying higher prices at lower productivity firms substitute toward the numeraire good (good n). Therefore, everything else being equal, variations in the price of the good with product market friction (good d) have less of an impact on the utility of these customers.²⁷ We next provide microeconomic evidence on the incomplete and heterogeneous price pass-through predicted by our model. Our estimates are based on the retailer price data introduced in Section 2 that include both the price and the replacement cost for every item. To measure pass-through of idiosyncratic shocks, we regress the log-price index of each store in a given week on its log-cost index. The price index of store j , p_t^j , is constructed as described in Section 2, and the cost index χ_t^j is analogously computed using the data on replacement cost provided by the retailer. To avoid inflating the short-term pass-through due to the persistence of both price and cost variables, we include in the specification lagged values of the independent variable. 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. In column (1) of Table 3, we report the estimates of our baseline specification, indicating a point estimate of 0.24, which is not statistically different from the average pass-through predicted by our model, that is, 0.26. As a robustness check, in column (2) we report estimates without time and market fixed effects. In column (3), 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. In both cases, the estimates imply an incomplete pass-through of cost shocks to prices, even smaller than our baseline specification. In column (4), we test the model prediction that firms with higher productivity (i.e., lower cost) should display higher pass-through. We consider the distribution of store price index cost for each store and construct an indicator variable “High productivity” that takes value 1 when a store’s cost is in the bottom tercile of the distribution.²⁸ We include the “High productivity” dummy and its interaction with the store cost variable in the specification to estimate short-run pass-through; we cannot replicate the same exercise for the long-run pass-through specification of column (1) because the lagged realization of costs may lie in different terciles of the cost distribution

²⁶ Equation (13) implies that optimal prices also depend on the intensive margin elasticity (ε_d) and the value of a customer ($\bar{\pi}$). Their role is, however, quantitatively small. Therefore, here we concentrate only on the role of the extensive margin of demand.

²⁷ Note that the threshold to search at low productivity firms is in the increasing part of the density of search cost, so the flattening of the extensive margin elasticity at low productivity is not due to a smaller mass of customers exiting at the margin.

²⁸ This is an imperfect test of our model implication, which refers to the productivity of a firm with respect to the cross-sectional distribution in a given week and not to the cross-time distribution for a given firm.

TABLE 3
PRICE PASS-THROUGH OF IDIOSYNCRATIC COST SHOCKS

Dependent Variable	(1) $\log(p_t^j)$	(2) $\log(p_t^j)$	(3) $\Delta \log(p_t^j)$	(4) $\Delta \log(p_t^j)$
$\log(\chi_t^j)$	0.24*** (0.09)	0.17*** (0.04)		
$\log(\chi_{t-1}^j)$	0.06 (0.04)	0.04 (0.03)		
$\log(\chi_{t-2}^j)$	0.02 (0.07)	-0.01 (0.05)		
$\log(\chi_{t-3}^j)$	0.05* (0.03)	0.06*** (0.02)		
$\log(\chi_{t-4}^j)$	0.07 (0.05)	0.07 (0.05)		
$\Delta \log(\chi_t^j)$			0.13* (0.07)	0.11* (0.06)
$\Delta \log(rc_t^j)$ × High productivity				0.40** (0.15)
High productivity			(0.00)	-0.00
Observations	8,295	12,915	8,295	8,295
MSA f.e.	Yes	No	Yes	Yes
Time f.e.	Yes	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 parentheses and are clustered at the store level. *** Significant at 1%; ** Significant at 5%; and * Significant at 10%.

making the characterization a firm as being high productivity ambiguous. The positive large and significant coefficient on the interaction term means that cost pass-through is much higher when firms are more productive. In particular, the pass-through of a high productivity store is still incomplete, but it is over three times larger than the average pass-through of stores in the top two terciles of their cost distribution.

Demand dynamics. In this section we compare the persistence of total sales and prices. The latter can be linked to the intensive margin of demand $d(p)$. The total demand for the firm is the sum of the demands of all of its customers, $m \times d(p)$. Even firms with declining productivity, which will post higher prices and, therefore, see their demand per customer shrink, can have strong total demand if they have a large base of customers.

To evaluate the predictions of our model about the persistence of demand against available data, we compare the persistence of an AR(1) process fitted to actual data on store level prices and revenues to that of an AR(1) estimated on data simulated using our model. For each of the 1,727 stores in the IRI data in the years from 2004 to 2006, we construct a monthly store level price as the average of the weekly store prices computed following the same formula as in Equations (3). We also sum revenues from sales of all the UPCs sold at each store during a month to obtain monthly store revenues. For both variables, we estimate an AR(1) process that includes month×year and market fixed effects leading to an estimated coefficient of 0.82 for prices and 0.93 for revenues (bottom row of Table 4). Hence an idiosyncratic shock affecting demand has a half-life of 10 months, whereas a shock to prices has a half-life of only four months. When performing the same exercise on artificial data simulated from the model, we find that firm level demand is very persistent and substantially more so than the underlying process of the price.²⁹ The monthly autoregressive coefficients are 0.98 for demand and 0.75 for

²⁹ To make the estimates comparable to the ones obtained with the IRI data, we simulate 36 months of observations for 1,727 firms in our model.

TABLE 4
ESTIMATE OF AN AR(1) PROCESS FOR REVENUES AND PRICES

	Revenues: $p \times d(p) \times m$	Price: p
Model	0.98 (0.00001)	0.75 (0.00021)
Data	0.93 (0.005)	0.82 (0.007)

NOTES: We report the estimates of the coefficient ρ_i from the process $\log(x_{it}) = a + \rho_i \log(x_{it-1}) + \epsilon_{it}$, where x_i stands for revenues in the first column and price in the second column. The top row reports estimates on the data for revenues and prices obtained by simulating 1,727 firms for 36 months in our baseline economy with parameters described in Section 3.4. The bottom row reports estimates on data on store level revenues and prices from IRI. Standard errors in parentheses.

prices, so that the model correctly predicts that firm revenues are substantially more persistent than prices. However, we notice that the gap in persistence between the demand and the price processes in the model is larger than in the IRI data. This can be explained by the fact that, while the model allows for idiosyncratic cost shocks, it does not include idiosyncratic demand shocks, which could water down the serial correlation of the demand process in the IRI data.

5. THE PROPAGATION OF DEMAND SHOCKS

In this section we study the effect of a demand shock that shifts the utility of consumption without affecting the disutility of search. The equations of the model will now be indexed by a time subscript t , capturing the aggregate state. As we want to study the effects of aggregate shocks in general equilibrium, we also need to endogenize household income. We do so through a simple model of a perfectly competitive labor market. The household takes the wage as given, and the wage is determined in a centralized market to clear labor demand. We assume that the representative household is divided into a mass one of shoppers and a representative worker. The worker takes care of supplying labor in the perfectly competitive labor market and then shares labor income equally across the shoppers who instead buy goods according to the model described in Section 3.³⁰ The expected discounted utility of the household is given by

$$(14) \quad \sum_{t=t_0}^{\infty} \beta^{t-t_0} \left\{ \int_0^1 [\xi_t \log(c_t(i)) - \psi(i) \mathcal{S}_t(z(i), \psi(i))] di - \phi \ell_t \right\},$$

where $c_t(i)$ is the consumption basket of shopper i matched to a firm with productivity $z(i)$, after search decisions are taken; $\mathcal{S}_t(z(i), \psi(i))$ is an indicator variable equal to one if shopper i , with search cost $\psi(i)$ and matched to a firm with productivity $z(i)$, decides to search and equal to zero otherwise; ξ_t is an aggregate shock to the utility of consumption; ϕ captures the disutility from working and is set equal to one so that income I is normalized to one in steady state, as in our partial equilibrium calibration. The worker chooses the path of labor supply ℓ_t that maximizes household utility in Equation (14) under perfect foresight. In particular the worker trades off higher disutility of labor ℓ_t with higher labor income $w_t \ell_t$ to be distributed equally across all shoppers, so that total income available to shoppers is given by $I_t = w_t \ell_t + D_t$, where D_t are firms profits rebated to the households. The worker internalizes the impact that higher labor income will have on the shoppers' decisions both in terms of search activity and consumption allocation but cannot discriminate across shoppers, so that she has to divide labor proceeds equally across shoppers. The mappings from income I_t to the distributions of consumption $c_t(i)$

³⁰ We are implicitly assuming that the worker cannot discriminate across the different shoppers. This assumption reduces the dimensionality of the problem, removing heterogeneity in income across consumers. This is a common shortcut in the literature (see Shi, 1997).

TABLE 5
COMPARATIVE STATICS

	Income	Average Markup	Consumption/Income
Low demand economy, $\xi = 1.00$	1.00	0.29	0.91
High demand economy, $\xi = 1.28$	1.28	0.26	0.92

NOTES: Two economies are simulated in steady state at the parameter values of Table 2. In economy A we set $\xi_t = 1$, in economy B we set $\xi_t = 1.28$ for all t .

and search activity $\mathcal{S}_t(z(i), \psi(i))$ are obtained from the solution of the model in Section 3. For simplicity we assume that individual shoppers are not allowed to save, whereas representative households do not save in equilibrium given assets are in zero net supply. The production technology of the good sold in the perfectly competitively market (good n) is linear in labor, with unitary productivity. Perfect competition in the market for good n and in the labor market implies that workers are paid a wage equal to the marginal productivity of labor so that $w_t = q_t$ for all t , where q_t is the price of good n . We use good n as a numeraire and set $q_t = 1$.

Steady state comparative statics. We compare two economies in steady state, denoted by A and B . These economies are identical but for the value of ξ_t . In particular, all parameters are as in Table 2 and described in Section 3.4, but demand in economy A is lower than in economy B because consumers in economy A value consumption less than in B . Operationally, we set $\xi_t = 1$ for all t in economy A and $\xi_t = \bar{\xi} > 1$ for all t in economy B . Given the log-utility preferences for consumption and the linear disutility in labor, equilibrium steady state income is $I = 1$ in economy A and $I = \bar{\xi}$ in economy B , as labor supply increases one for one with the shift in marginal utility of consumption. Thus, we choose $\bar{\xi} = 1.28$ so that the log-difference in income between economies A and B matches the standard deviation in the logarithm of monthly income per capita across U.S. counties in 2005, which indeed equals 0.28.³¹

In Table 5, we report statistics from simulating economies A and B in steady state. Economy A is characterized by higher average markups and, as a result, lower average consumption to income ratio (defined as $\int_0^1 c(i)di/I$) than economy B . In particular, markups are lower by 3 percentage points on average in economy B , allowing for consumption to be higher by 1% for each unit of income spent on consumption. Lower average markups in economy B are the outcome of the higher propensity to search, as households value demand more relative to the disutility of search, resulting in higher elasticity of demand.

We test the prediction of our model about the relationship between markups and aggregate expenditure using county level data. First, we compute median county level markups using the price and replacement cost reported in the “retailer price data” by all available stores. The data report the markup for each UPC-week at each of the 270 stores of the chain representative of the price areas. We construct weekly markup for each store by averaging the markups on all the UPCs on sale at a store in a given week and then obtain yearly store markup as the median of the weekly markups for each store. Finally, our county-level measure of markup in a given year is the average of the markups across all the retailer’s store located in the county in that year. Then, we regress the computed cross section of markups on the log of income per capita in the county obtained from the Bureau of Economic Analysis. Since shifts in the opportunity cost of search may affect firm pricing (Kaplan and Menzio, 2016) and are correlated with our measure of demand, we control for the county unemployment rate. The estimated regression coefficient is equal to 0.084 and is statistically significant, whereas the standard deviation of log of income per capita across counties is 0.28. Although we do not give a causal interpretation to such estimate, we notice that the implied correlation between income and markups is not only qualitatively but also quantitatively similar to the one produced by our model: When we consider two counties that differ in the log of income per capita by one standard deviation

³¹ We focus on 2005 because that is the only full calendar year for which we have markup data.

(28 percentage points) the difference in associated markups amounts to 2.35 percentage points in the data against 2.8 in the model.

Impulse responses to a transitory shock. We consider our baseline economy in steady state at $t = t_0$, calibrated as described in Section 3.4. We consider the dynamics following an unforeseen aggregate shock that takes the economy temporarily away from the steady state and assume that after the shock has hit, there is perfect foresight on the path of the aggregate state. In particular, we assume that $\xi_t = 1$ at all $t < t_0$, when agents expect $\xi_t = 1$ to hold also in the future. At $t = t_0$, ξ_{t_0} increases unexpectedly and then reverts back to steady state following an AR(1) process, that is, $\log(\xi_t) = \rho_\xi \log(\xi_{t-1})$ for $t > t_0$.³² We calibrate ρ_ξ and ξ_{t_0} by fitting an AR(1) process on the log of yearly per capita income at the U.S. county level, obtained from the Bureau of Economic Analysis for the period 1990–2015, controlling for cross-county heterogeneity including county fixed effect and for aggregate shocks including a full set of year dummies.³³ The AR(1) process in our model is at the monthly frequency, whereas the data come at the yearly frequency. Thus, we obtain $\xi_{t_0} = 1.007$ and $\rho_\xi = 0.96$ so that our monthly AR(1) implies the same volatility and persistence of the yearly AR(1) process estimated in the data, corresponding respectively to 2% standard deviation of year-on-year income and half-life of 18 months.

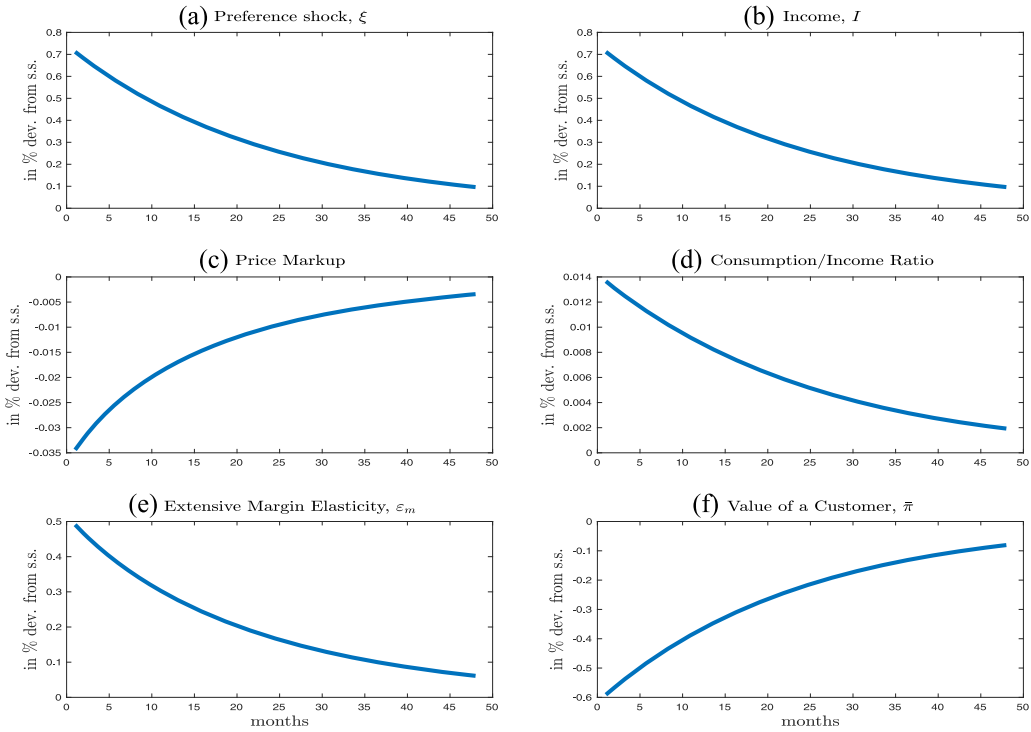
Figure 4 plots the impulse responses of several variables of interest to the transitory demand shock (panel (a)). As in standard macro models, a higher utility of consumption relative to the utility of leisure causes higher labor supply and higher income (panel (b)). In particular, we have $I_t = \xi_t$ for all t in our model. Aggregate consumption goes up more than income (panel (d)) because of the fall in average markups (panel (c)). The lower markups are the result of the increased benefits of searching. The fact that more customers may be looking for a new supplier increases customer retention concerns for firms, manifesting into a higher average extensive margin elasticity of demand (panel (e)), inducing firms to lower their markups. We notice that the fall in markups associated to the increase in expenditure is smaller than in the case of the permanent shift in preferences of Table 5. In particular, the implied partial elasticity of average markups to ξ is about 5% for this transitory shock, against 10% in the case of a permanent change. The smaller sensitivity of markups is explained by the response of the average customer value, $\bar{\pi}(p, z)$, in panel (f) that falls in response to a transitory demand shock, reducing the overall effect of the shock on markups. In particular, the mean reverting path of income implies that demand from a customer today is higher than tomorrow, so that the firm has an incentive to set higher prices today to take advantage of the higher demand.

6. CONCLUDING REMARKS

This article provides novel empirical evidence on the elasticity of a firm customer base to its price. Using data from a large U.S. supermarket chain, we showed that an increase in the prices is associated to a higher customer attrition. Motivated by this finding we have developed a rich model to study the implications of the extensive margin of demand for price setting. It emerged from our analysis that customer retention concerns substantially reduce the price pass-through of idiosyncratic cost shocks. Hence firm demand is substantially more persistent than the underlying price and productivity processes. These predictions find empirical support in data from the U.S. grocery sector. We used our calibrated model as a laboratory to study the effects of both transitory and permanent shifts in aggregate demand. Higher demand is associated to higher customers' propensity to search, resulting in higher demand elasticity.

³² Operationally, we guess the path of the distribution of customers over productivity $K_t(z)$ and then solve backward for the optimal pricing and customer decisions at each time $t \geq t_0$ starting from the final steady state (which is identical to the initial one) at $t = t_0 + T$ for some large T . Then we update our guess about $K_t(z)$ and iterate until convergence.

³³ We do not have a time series for markups, as we only observe markups for a full calendar year in 2005. However, we estimate the process on the same set of counties used for the cross-sectional analysis which exploited markup information.



NOTES: The values are measured in percentage deviations from the steady state. All plots report the impulse response to a $\xi_0 = 1.007$ shift in the utility of consumption with an half-life of 18 months.

FIGURE 4

IMPULSE RESPONSES TO A TRANSITORY PREFERENCE SHOCK
 [COLOR FIGURE CAN BE VIEWED AT WILEYONLINELIBRARY.COM]

Increases in the extensive margin elasticity of demand lead firms to lower markups, amplifying the effect of demand shocks on consumption.

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 the limitations in our data, we do not consider the role of advertising in generating demand dynamics (Hall, 2014), or of firm pricing in attracting customers (Dinlersoz and Yorukoglu, 2012). Although our conjecture is that the analysis of the pricing incentives presented in this article would still apply, this would be an interesting dimension in which to extend our work.

APPENDIX: DATA SOURCES AND VARIABLES CONSTRUCTION

This appendix provides additional information on the data sources presented in Section 2. We also document more in depth the procedure used to construct the main variables used to empirically assess the relevance of the extensive margin of demand.

A.1. Data and Selection of the Sample. The retailer that provided both the price data and the consumer panel is a large supermarket chain that operates over 1,000 stores across the United States. It is a high/low supermarket chain selling grocery goods as well as household supplies; it could be compared to Kroger or Tesco.

TABLE A.1
DESCRIPTIVE STATISTICS ON CUSTOMER SHOPPING BEHAVIOR

	Mean	Standard deviation	25th percentile	75th percentile
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		

A.1.1. *Sampling and representativeness.* The Consumer Panel data include complete purchase data for over 11,000 customers of the chain sampled from the major markets for the retailer, excluding those where it operates under acquired brands. Households are tracked through usage of the supermarket loyalty card; purchases made without using the card are not recorded. However, the chain ensures that the loyalty card has a high penetration, for instance, by keeping to a minimum the effort needed to register. Furthermore, nearly all promotional discounts are tied to ownership of a loyalty card, which provides a strong incentive to sign up and use it. Therefore, we can consider the customers in our sample as representative of the population of noncasual shoppers at the chain.

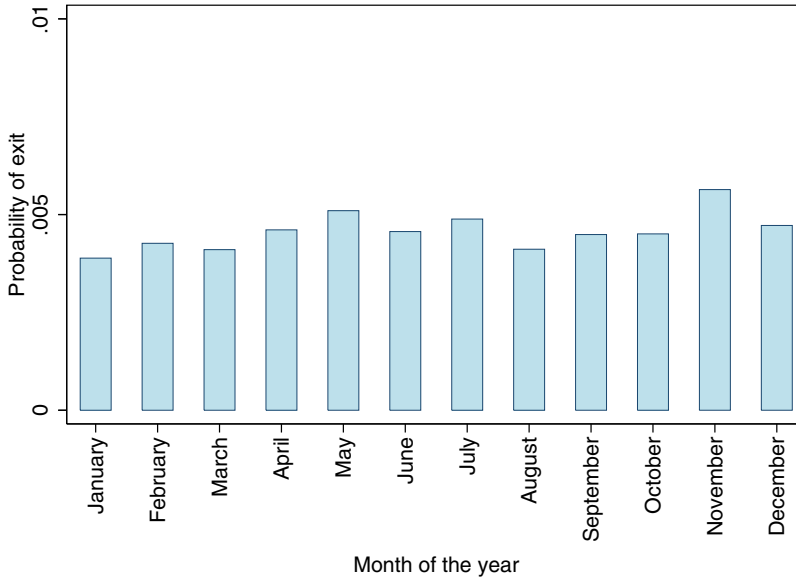
The Price Data cover 270 stores. This is about one-fifth of the stores operated by the retailer; however, the chain sets different prices for the same UPC in different geographic areas, called “price areas.” The set of stores for which the retailer provided information was designed so that at least one store for each price area would be included.

A.2. Variables Construction.

A.2.1. *Exit from the 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 customer has abandoned the retailer to shop elsewhere or she is simply not purchasing groceries in a particular week, for instance because she is just consuming her inventory. In fact, we observe households when they buy groceries at the chain but do not have any information on their shopping at competing firms. We choose 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 a value of one in correspondence to the last visit at the chain before a spell of eight or more weeks without shopping there. Table A.1 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.

In Figure A.1 we document the seasonality in exit rates and find that the probability of exit is roughly stable across months.

Finally, in Table A.2 we show that the qualitative result of our reduced form analysis on the elasticity of the exit decision to price is not sensitive to the definition of exit. The table replicates the specification in the first column of Table 1 for different definitions of “exit.” We experiment with the number of consecutive weeks without a visit to a store of the chain required to declare that the customer has exited the customer base of the retailer both shortening our baseline of eight to four weeks and tightening the requirement for exit asking for 12, 16, or even 20 consecutive weeks without trips. To estimate each of these models on the same sample, we can only use data up to January 2006 so that we leave five months at the end of the sample to assess occurrence of an exit in the specification where 20 weeks of absence are required. As a consequence, the number of observations and the point estimate for the baseline specification are not exactly identical to those in Table 1.



NOTES: The figure plots the unconditional probability of exit, computed as the ratio of the number of exits and the number of shopping trips by customers of the chain, by month. The definition of exit adopted for the plot is our baseline one: lack of shopping trips at the chain for eight consecutive weeks or more.

FIGURE A.1

SURVIVAL IN THE CUSTOMER BASE
 [COLOR FIGURE CAN BE VIEWED AT WILEYONLINELIBRARY.COM]

TABLE A.2
 ROBUSTNESS TO EXIT DEFINITION

	<i>Consecutive Weeks without Trips Needed to Observe "Exit"</i>				
	4 weeks	8 weeks	12 weeks	16 weeks	20 weeks
$\log(p_{it})$	0.79*** (0.229)	0.15** (0.023)	0.08 (0.059)	0.09* (0.047)	0.10** (0.042)
$\log(\bar{p}_{it})$	0.003 (0.078)	0.001 (0.618)	0.001 (0.001)	0.000 (0.000)	0.000 (0.000)
Tenure	-0.001*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)
Observations	44,647	44,647	44,647	44,647	44,647

NOTES: The table reports replications of the specification in column (1) of Table 1 with different definition of the dependent variable. 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. Standard errors are in parenthesis and account for within-household correlation through a two-steps feasible-GLS estimator. *** Significant at 1%; ** Significant at 5%; and * Significant at 10%.

For more stringent definitions (i.e., 12, 16, or 20 weeks), the estimated parameters are significant at the conventional levels, with the exception of the specification where exit occurs after a three-month absence from the chain. Moreover, even though the point estimates go down slightly, none of them is significantly different from our baseline estimate when testing at 1%.

Moving to a shorter horizon to identify exit introduces a different set of considerations. When we consider as exits also absence of four weeks from the chain, we obtain a positive and significant elasticity of the probability of exit to the price of the basket but its size is larger than

our baseline estimate. The fact that this estimate departs significantly from the range of values we find for all the other definitions implies that there is a number of customers who disappear from the chain for at least four weeks but return before eight weeks. These could be bargain hunting households willing to move repeatedly across firms to take advantage of temporary price changes. Evidence in favor of this interpretation comes from the fact that the probability of returning to the retailer (i.e., purchasing there again after having exited the customer base) is significantly related to posted prices when the definition of exit requires a hiatus of four weeks, whereas there is no correlation for the other definitions. However, customers who have truly left the customer base of their firm could not be that responsive to the price of the firm they used to shop at as they would have to pay the search cost to rejoin it. This indicates that a four-week window tends to pick up households who did not truly abandoned the customer base of the firm.

A.2.2. Weekly UPC prices. The Consumer Panel reports 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 cannot infer the price of the item in that store-week. However, our definition of basket requires us to be able to attach a price to each of the items composing it in every week, even when the customer does not shop. The issue can be solved using the price data that report information on weekly store revenues and quantities regardless of the shopping decisions of the households in our Consumer Panel. We use data on store level revenues and quantities sold in the price data to compute unit value prices as

$$UVP_{tu}^j = \frac{TR_{tu}^j}{Q_{tu}^j},$$

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

As explained in Eichenbaum et al. (2011), the unit value price delivers the exact price paid for items sold at full price. For goods on promotion, it delivers an average between the full and the discounted price. 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, a rare occurrence and involves only infrequently 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.

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

TABLE A.3
COMPARISON OF MULTISTORE AND SINGLE STORE SHOPPING HOUSEHOLDS

	Mean for Multistore Households	Mean for Single Store Households	Difference
<i>Shopping Behavior</i>			
Net expenditure	433.10	424.23	(8.87) ^{***}
Number of unique UPCs bought	105.76	105.56	(0.20)
Number of unique categories bought	83.17	82.94	(0.23)
Number of shopping trips	7.99	7.66	(0.34) ^{***}
<i>Demographic characteristics</i>			
% male	49.26	48.97	(0.30)
% black	4.66	4.80	(-0.14)
% Hispanic	11.80	9.78	(2.03) ^{***}
% age 18–24	7.35	6.96	(0.39)
% age 25–34	15.22	14.56	(0.66) [*]
% age 35–44	18.37	18.70	(-0.33)
% age 45–54	15.19	14.94	(0.25)
% age 55–64	9.01	8.93	(0.09)
% age 65 or more	11.25	11.14	(0.11)
% HS degree	40.71	40.05	(0.67)
% college degree	48.94	50.37	(-1.42) [*]
% employed	65.47	66.07	(-0.60)
Per capita income	34,880.80	34,600.98	(279.82)
Distance from retailer	1.36	1.41	(-0.05)
<i>Probability of exit</i>			
Probability of exit	0.0018	0.0017	(0.0001)

NOTES: The table reports the means of different characteristics of households who shop at multiple stores of the retail chain that shared the microdata with us and compares them with those for households who predominantly shop at a same store of the chain. The latter group is the one we use in our analysis. The shopping behavior variables are computed at monthly frequency; the probability of exit is instead weekly. The third column of the table displays the difference between the mean and reports the level of significance at which no difference between the means of the two groups can be rejected. *** Significant at 1% and * significant at 10%.

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. Since the 270 representative stores were randomly chosen, the selection of the stores should not induce selection in the subsample of households we analyze.

The need for considering only households that predominantly shop at the same store of the chain could instead generate selection if the probability of switching store within the same retailer is correlated to the propensity of the household to change the retailer at which it shops. To address this concern, in Table A.3 we compare the behavior of the sample of households selected for our analysis with that of households we discarded because they shopped at multiple outlets of the chain. We look at three main dimensions: shopping behavior (frequency and size of grocery shopping trips), demographic characteristics, and probability of exit. The demographic characteristics are not reported by the households themselves but come from Census data matched at the block group level, which corresponds roughly to a city block. This is the finest level of geographical detail available in the Census and ensures that, while aggregate, the figures reported originate from a group of highly homogenous respondents. The probability of exiting the customer base of the retailer is computed using our baseline definition (i.e., eight consecutive weeks without shopping trips at the retailer).

As far as shopping behavior is concerned, the differences in the total expenditure in grocery and in the number of trips are statistically significant between the two groups but not economically large. Multistore shoppers spend 8 extra dollars per month in grocery and take an extra

trip to a store of the chain every three months. Very few of the demographic variables are significantly different between the two samples and the gaps are small. Finally, households in the two samples have basically the same weekly probability of exit.

A.2.3. 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

$$p_{it}^j = \sum_{u \in K_i} \omega_{iu} p_{ut}^j, \quad \omega_{iu} = \frac{\sum_t E_{iut}}{\sum_u \sum_t E_{iut}},$$

where K^i is the set of all the UPCs (u) purchased by household i during the sample period and p_{ut}^j is the price of a given UPC u in week t , at the store $costj$ where the customer shops. E_{iut} represents expenditure by customer i in UPC u in week t , and the ω_{iu} '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.

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 Metropolitan Statistical Area. Next, we average those prices across stores to obtain the average market price of the consumer basket. In particular, the price is computed as

$$\bar{p}_{it} = \sum_{j \in m(i)} s^j \sum_{u \in K_i} \omega_{iu} p_{ut}^j, \quad \omega_{iu} = \frac{\sum_t E_{iut}}{\sum_u \sum_t E_{iut}}, \quad s^j = \frac{\sum_t R_t^j}{\sum_{j' \in M} \sum_t R_t^{j'}},$$

where $m(i)$ is the market of residence for customer i , and R_t^j represents revenues of store j in week t . In other words, in the construction of the competitors' price index stores with higher (revenue-based) market shares weight more.

To make the price index of the basket at the Retail comparable with the same figure calculated for the competitors, we have to use the same set of UPCs. This implies that the price indices are computed based only on UPCs that appear both in the Retailer data and in the IRI ones. The main limitation here is that IRI covers only a subset of the product categories on which we have shopping information from the Retailer data: For instance, nail products are not covered in IRI. However, IRI covers all the major product categories, reducing concerns on the effect of this limitation on the calculation of the price of the basket. In theory, it is also possible that this criterion would make us discard UPCs referring to identical products that have different codes at different retailers for administrative reasons. This issue seems quantitatively negligible. Conditional on the set of product categories covered by IRI, 87% of the UPCs in the Retailer data have a match in the IRI data, and the matched UPCs account for 97% of the total household expenditure.

A.3. *Proofs.*

A.3.1. *Proof of Proposition 1.* The following lemma discusses some key properties of the optimal price useful to prove Proposition 1.

LEMMA A.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 of the lemma is an immediate implication of Equation (7). We next prove the results of Proposition 1.

Monotonicity of prices. Monotonicity of optimal prices follows from an application of Topkis' theorem. 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, 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 $p \leq \bar{p}(z)$. Under the assumption of i.i.d. productivity, $\Delta(p, z)$ is independent of z , which is sufficient to obtain the result. We first show that optimal prices $\hat{p}(z)$ are nonincreasing 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 A.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,

$$(A.1) \quad W(p, z) = \Delta(p, z) (\pi(p, z) + \beta \text{ constant}),$$

is supermodular in $(p, -z)$. Note 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 $z_1, z_2 \in [z, \bar{z}]$ with $-z_2 > -z_1$. We have that $W(p_2, z_2) - W(p_1, z_2) \leq W(p_2, z_1) - W(p_1, z_1)$ if and only if

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

which, since $\Delta(p_2)d(p_2) < \Delta(p_1)d(p_1)$ as $d(p)$ is strictly decreasing and $\Delta(p)$ is nonincreasing, 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 nonincreasing in z .

Existence of equilibrium. Next we prove the 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, $\hat{p}(z)$. Note that $W(p, z)$ in Equation (A.1) is continuous in (p, z) . By the theorem of maximum, $\hat{p}(z)$ is upper hemicontinuous and $W(\hat{p}(z), z)$ is continuous in z . Given that $\hat{p}(z)$ is nonincreasing in z it follows that $\hat{p}(z)$ has 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 \bar{z} (so that $\hat{\mathcal{P}}(\bar{z})$ is not a singleton), we modify the optimal pricing rule of the firm and consider the convex hull of the $\hat{\mathcal{P}}(\bar{z})$ as the set of possible prices chosen by the firm with productivity \bar{z} . The constructed mapping from $\mathcal{P}(z)$ to $\hat{\mathcal{P}}(z)$ is then upper hemicontinuous, compact, and convex valued. We then apply Kakutani's fixed point theorem to

this operator and obtain a fixed point. Finally, note 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 1 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 A.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 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.

Proof of Point (i). We already proved that $\hat{p}(z)$ is nonincreasing in z . The proof that $\hat{p}(z)$ is strictly decreasing follows by contradiction. Consider that $\hat{p}(z_1) = \hat{p}(z_2) = \bar{p}$ 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 $\{z_2, \bar{p}\}$. Note 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 $\{z_1, \bar{p}\}$. 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, $\hat{p}(z_1) > \hat{p}(z_2)$. Note 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.³⁴

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

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

³⁴ 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 \bar{p} . Then $\bar{p}(z) = \bar{p}$ for all z . If $\bar{p} > p^*(\bar{z})$, then \bar{p} would not be optimal for firm with productivity \bar{z} , which would choose a lower price. If $\bar{p} = p^*(\bar{z})$, 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 \bar{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 A.1, and using that $\hat{p}(z) \geq \hat{p}(\bar{z})$ and $\bar{p}(z) = \hat{p}(\bar{z})$ for all z .

A.3.2 Proof of Corollary 1.

Part (1): Start by noticing that, because the mean of $G(\psi)$ is positive, the expected value of searching diverges to $-\infty$ as ι 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, $\bar{p}(z) \rightarrow \infty$ 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 A.1 we obtain that $\hat{p}(z) = p^*(z)$ for all $z \in [\underline{z}, \bar{z}]$.

Part (2): From Proposition 1 we have that, in equilibrium, the highest price is $\hat{p}(z)$. Moreover, under the assumptions of Proposition 1, the first-order condition is a necessary condition for optimality of prices. We use this to show that, as ι approaches zero, $\hat{p}(z)$ has to approach $\hat{p}(\bar{z}) = p^*(\bar{z})$. In equilibrium, it is possible to rewrite Equation (7), evaluated at $\{\hat{p}(z), \underline{z}\}$, as $LHS(\hat{p}(z), \iota) = RHS(\hat{p}(z), \iota)$, where

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

given that $H(\bar{V}(\hat{p}(z), \underline{z})) = Q(\bar{V}(\hat{p}(z), \underline{z})) = 0$. Suppose that as $\iota \downarrow 0$, $\hat{\psi}(\hat{p}(z), \underline{z})$ does not converge to zero. Then, $G(\frac{\hat{\psi}(\hat{p}(z), \underline{z})}{\iota}) \uparrow 1$ as $\iota \downarrow 0$. This implies that $\lim_{\iota \downarrow 0} RHS(\hat{p}(z), \iota) > 0$. Consider now the function $LHS(\hat{p}(z), \iota)$. Again, suppose that as $\iota \downarrow 0$, $\hat{\psi}(\hat{p}(z), \underline{z})$ does not converge to zero. Note that the second term of the function approaches a finite number as $\bar{V}_p(\hat{p}(z), \underline{z})$ is bounded by assumptions on $v(p)$ and $H'(\bar{V}(\hat{p}(z), \underline{z}))$ and $Q'(\bar{V}(\hat{p}(z), \underline{z}))$ being bounded as a result of Proposition 1. Moreover, as long as $\hat{p}(z) > \bar{p}(z) = p^*(\bar{z})$, we have that $\hat{\psi}_p(\hat{p}(z), \underline{z}) > 0$ so that $\hat{\psi}_p(\hat{p}(z), \underline{z})/\iota$ diverges as ι approaches zero. This means that $G'(\frac{\hat{\psi}(\hat{p}(z), \underline{z})}{\iota})\frac{\hat{\psi}_p(\hat{p}(z), \underline{z})}{\iota}$ is divergent, and therefore the first-order condition cannot be satisfied.

This analysis concluded that if $\hat{\psi}(\hat{p}(z), \underline{z})$ does not converge to zero as ι becomes arbitrarily small, $LHS(\hat{p}(z), \iota)$ would diverge to infinity, whereas $RHS(\hat{p}(z), \iota)$ would remain finite, and the first-order condition, that is, Equation (7), cannot be satisfied. It then follows that, as ι approaches zero, a necessary condition is that $\hat{\psi}(\hat{p}(z), \underline{z})$ also approaches zero. This condition can be restated as requiring that $\hat{p}(z)$ approaches $\bar{p}(z)$ as ι approaches zero. Moreover, given the assumptions of Proposition 1, $\bar{p}(z) = \hat{p}(\bar{z}) = p^*(\bar{z})$.

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

A.4. Numerical Solution of the Model. To solve the model, we start by setting the parameters. The parameters β, κ, δ , 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 , 0.025 for ρ , 0.25 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 $E[\varepsilon_d(z) + \varepsilon_m(z)] = 2.7$.

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 1, we look for equilibria where $\hat{p}(z) \in [p^*(\bar{z}), p^*(z)]$ 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 find 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, that is, $\bar{V}(p, z)$, and solve for the optimal search and exit thresholds. 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. Note 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 $K(z)$. Finally, given $Q(\cdot)$, $H(\cdot)$, $\hat{p}(z)$, and $\bar{V}(p, z)$, we solve the firm problem and 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 $z \in [z, \bar{z}]$.

Once we have solved for the equilibrium of the model at given parameter values. 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 is in the interior of the assumed grid.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Data S1

REFERENCES

- AGUIAR, M., AND E. HURST, "Life-Cycle Prices and Production," *American Economic Review* 97(5) (2007), 1533–59.
- ALESSANDRIA, G., "International Deviations from the Law of One Price: The Role of Search Frictions and Market Share," *International Economic Review* 45(4) (2004), 1263–91.
- BAI, Y., J.-V. RIOS RULL, AND K. STORESLETTEN, "Demand Shocks that Look Like Productivity Shocks," mimeo, 2012.
- BLINDER, A. S., E. R. D. CANETTI, D. E. LEBOW, AND J. B. RUDD, *Asking About Prices: A New Approach to Understanding Price Stickiness* (New York: Russell Sage Foundation, 1998).
- BRONNENBERG, B. J., M. W. KRUGER, AND C. F. MELA, "The IRI Marketing Data Set," *Marketing Science* 27(4) (2008), 745–8.
- BURDETT, K., AND M. G. COLES, "Steady State Price Distributions in a Noisy Search Equilibrium," *Journal of Economic Theory* 72(1) (1997), 1–32.
- , AND K. L. JUDD, "Equilibrium Price Dispersion," *Econometrica* 51(4) (1983), 955–69.
- , AND G. MENZIO, "(Q,S,s) Pricing Rules," *Review of Economic Studies* 85 (2018), 892–928.
- , AND D. T. MORTENSEN, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review* 39(2) (1998), 257–73.
- , AND T. VISHWANATH, "Balanced Matching and Labor Market Equilibrium," *Journal of Political Economy* 96(5) (1988), 1048–65.
- CABRAL, L., "Dynamic Pricing in Customer Markets with Switching Costs," *Review of Economic Dynamics* 20 (2016), 43–62.
- COIBION, O., Y. GORODNICHENKO, AND G. H. HONG, "The Cyclicity of Sales, Regular and Effective Prices: Business Cycle and Policy Implications," *American Economic Review* 105(3) (2015), 993–1029.
- DECKER, R. A., J. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA, "Declining Business Dynamism: Implications for Productivity," Mimeo No. 23, Brookings Institution Hutchins Center, 2016.
- DIAMOND, P. A., "A Model of Price Adjustment," *Journal of Economic Theory* 3(2) (1971), 156–68.
- DINLERSOZ, E. M., AND M. YORUKOGLU, "Information and Industry Dynamics," *American Economic Review* 102(2) (2012), 884–913.
- DROZD, L. A., AND J. B. NOSAL, "Understanding International Prices: Customers as Capital," *American Economic Review* 102(1) (2012), 364–95.
- DUNNE, T., M. J. ROBERTS, AND L. SAMUELSON, "Patterns of Firm Entry and Exit in US Manufacturing Industries," *RAND Journal of Economics* (1988), 495–515.

- EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO, "Reference Prices and Nominal Rigidities," *American Economic Review* 101(1) (2011), 234–62.
- EINAV, L., AND P. SOMAINI, "A Model of Market Power in Customer Markets," *Journal of Industrial Economics* 61(4) (2013), 938–86.
- FABIANI, S., C. LOUPIAS, F. MARTINS, AND R. SABBATINI, *Pricing Decisions in the Euro Area: How Firms Set Prices and Why* (New York: Oxford University Press, 2007).
- FISHMAN, A., AND R. ROB, "Consumer Inertia, Firm Growth and Industry Dynamics," *Journal of Economic Theory* 109(1) (2003), 24–38.
- , AND ———, "Is Bigger Better? Customer Base Expansion through Word-of-Mouth Reputation," *Journal of Political Economy* 113(5) (2005), 1146–75.
- FOSTER, L., J. HALTIWANGER, AND C. SYVERSON, "The Slow Growth of New Plants: Learning about Demand?" *Economica* 83(329) (2016), 91–129.
- GALENIANOS, M., AND A. GAVAZZA, "A Structural Model of the Retail Market for Illicit Drugs," *American Economic Review* 107(3) (2017), 858–96.
- GIULIETTI, M., M. WATERSON, AND M. WILDENBEEST, "Estimation of Search Frictions in the British Electricity Market," *Journal of Industrial Economics* 62(4) (2014), 555–90.
- GOURIO, F., AND L. RUDANKO, "Customer Capital," *The Review of Economic Studies* 81 (2014), 1102–36.
- HALL, R. E., "General Equilibrium with Customer Relationships: A Dynamic Analysis of Rent-Seeking," mimeo, 2008.
- , "What the Cyclical Response of Advertising Reveals about Markups and Other Macroeconomic Wedges," mimeo, 2014.
- HOLMES, T., "The Diffusion of Wal-Mart and Economies of Density," *Econometrica* 79 (2011), 253–302.
- HONG, H., AND M. SHUM, "Using Price Distributions to Estimate Search Costs," *RAND Journal of Economics* 37(2) (2006), 257–75.
- HONKA, E., "Quantifying Search and Switching Costs in the US Auto Insurance Industry," *RAND Journal of Economics* 45(4) (2014), 847–84.
- KAPLAN, G., AND G. MENZIO, "The Morphology of Price Dispersion," *International Economic Review* 56(4) (2015), 1165–206.
- , AND ———, "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," *Journal of Political Economy* 124(3) (2016), 771–825.
- KLESCHCHESKI, I., AND N. VINCENT, "Market Share and Price Rigidity," *Journal of Monetary Economics* 56(3) (2009), 344–52.
- MENZIO, G. "A Search Theory of Rigid Prices," mimeo, 2007.
- , AND N. TRACHTER, "Equilibrium Price Dispersion with Sequential Search," *Journal of Economic Theory* 160(C) (2015), 188–215.
- , AND ———, "Equilibrium Price Dispersion across and within Stores," *Review of Economic Dynamics* 28 (2018), 205–20.
- NAKAMURA, E., AND J. STEINSSON, "Price Setting in Forward-Looking Customer Markets," *Journal of Monetary Economics* 58(3) (2011), 220–33.
- NEVO, A., AND A. WONG, "The Elasticity of Substitution between Time and Market Goods: Evidence from the Great Recession," mimeo, Northwestern University, 2015.
- PETROSKY-NADEAU, N., AND E. WASMER, "Macroeconomic Dynamics in a Model of Goods, Labor, and Credit Market Frictions," *Journal of Monetary Economics* 72 (2015), 97–113.
- PHELPS, E. S., AND S. G. WINTER, "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, 1970).
- RAVN, M., S. SCHMITT-GROHE, AND M. URIBE, "Deep Habits," *Review of Economic Studies* 73(1) (2006), 195–218.
- REINGANUM, J. F., "A Simple Model of Equilibrium Price Dispersion," *Journal of Political Economy* 87(4) (1979), 851–58.
- ROTEMBERG, J. J., AND M. WOODFORD, "Markups and the Business Cycle," *NBER Macroeconomics Annual 1991*, Volume 6 (Chicago: University of Chicago Press, 1991) 63–140.
- , AND ———, "The Cyclical Behavior of Prices and Costs," mimeo, 1999.
- SHI, S., "A Divisible Search Model of Fiat Money," *Econometrica* 65(1) (1997), 75–102.
- , "Customer Relationship and Sales," *Journal of Economic Theory* 166 (2016), 483–516.
- SMITH, H. T., "Supermarket Choice and Supermarket Competition," *Review of Economic Studies* 71 (2004), 235–63.
- SORENSEN, A. T. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs," *Journal of Political Economy* 108(4) (2000), 833–50.
- TAUCHEN, G., "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters* 20(2) (1986), 177–81.