Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Inflation, Price Dispersion, and Welfare: The Role of Consumer Search

Francisca Sara-Zaror

2024-047

Please cite this paper as:
Sara-Zaror, Francisca (2024). "Inflation, Price Dispersion, and Welfare: The Role of Consumer Search," Finance and Economics Discussion Series 2024-047. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2024.047.

Inflation, Price Dispersion, and Welfare: The Role of Consumer Search

Francisca Sara-Zaror*†‡

May 21, 2024

Abstract

In standard macroeconomic models, the costs of inflation are tightly linked to the price dispersion of identical goods. Therefore, understanding how price dispersion empirically relates to inflation is crucial for welfare analysis. In this paper, I study the relationship between steady-state inflation and price dispersion for a cross section of U.S. retail products using scanner data. By comparing prices of items with the same barcode, my measure of relative price dispersion controls for product heterogeneity, overcoming an important challenge in the literature. I document a new fact: price dispersion of identical goods increases steeply around zero inflation and becomes flatter as inflation increases, displaying a ϒ-shaped pattern. Current sticky-price models are inconsistent with this finding. I develop a menu-cost model with idiosyncratic productivity shocks and sequential consumer search that reproduces the new fact and exhibits realistic price-setting behavior. In the model, inflation-induced price dispersion increases shoppers' incentives to search for low prices and thus competition among retailers. The positive welfare-maximizing inflation rate optimally trades off the efficiency gains from lower markups and the resources spent on search.

JEL Codes: E31, E50, L11, L16.

---

*Federal Reserve Board. E-mail: francisca.sara-zaror@frb.gov

†I am greatly indebted to my co-chairs Fernando Alvarez and Mike Golosov, and my dissertation committee members Greg Kaplan and Rob Shimer, for invaluable advice, support, and encouragement at every stage of this project. I am grateful to Agustin Hurtado, Josh Morris-Levenson, Hyejin Park, and Marta Prato for insightful feedback; to my discussants Kieran Larkin, Gaetano Gaballo, and Jane Ryngaert, as well as participants of the ASSA 2023, Midwest Macro Meetings 2023, SED 2023, Inflation: Drivers and Dynamics Conference 2023, XXV Workshop in International Economics and Finance, and multiple seminars for helpful comments.

‡Researcher(s) own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

A salient feature of micro-level data is that the prices of identical goods vary across sellers. In current macroeconomic models, inflation is a crucial determinant of price dispersion, and the costs of inflation mainly arise from inflation-induced price dispersion. In this paper, I address three questions: What is the empirical relationship between inflation and price dispersion? What does this relationship imply for current monetary theories? What do we learn for the welfare analysis of inflation?

Studying how price dispersion empirically relates to inflation imposes several challenges. The first is to measure price dispersion accurately. For this, we need to observe the prices different sellers charge for *identical* goods. Typically, granular datasets which have been essential to establish facts on micro-price rigidity do not satisfy this requirement. For instance, quotes in the micro-data underlying the CPI are grouped into product categories with limited information about the specific product being measured or the ability to link identical products across sellers. Thus, the comparison of prices across sellers, even within narrow product categories, would not take product heterogeneity into account (Nakamura et al., 2018).

I overcome this challenge using highly detailed scanner data for retail products in the U.S. The dataset contains prices and quantities of products sold in over 35,000 stores across the country between 2006 and 2020. The observations are identified by the product barcode, the week, and the retailer where the transaction was carried out. More than 3 million barcodes are available, each of them sold by 50 stores on average. Each retailer belongs to one of over a thousand geographically dispersed counties. With these data, I can compare prices of products with identical barcodes sold across different retailers on the same date and county. Hence, my measure of relative price dispersion controls for several sources of heterogeneity which a priori are unrelated to inflation.

A shortcoming of this dataset is that it is only available for a relatively low and stable aggregate inflation period. Therefore, the variation in aggregate inflation is insufficient to statistically identify its time-series comovement with price dispersion. In contrast, the variation in product-level inflation rates across markets and product categories is substantial. While aggregate inflation fluctuated between -1% and 7% over the sample period, product-level inflation ranged between -30% and 30%. Thus, my approach is to study the relationship between inflation and price dispersion by exploiting cross-sectional variation and derive implications of aggregate inflation using a multi-product model calibrated to match the disaggregate evidence.[1]

I contribute to the literature by documenting a new fact: in the cross section, small deviations from zero inflation sharply increase price dispersion of identical goods. As inflation increases, price dispersion becomes a flatter function of inflation. When plotting price dispersion against inflation, the relationship resembles the Greek capital letter upsilon; thus, I refer to it as Υ-shaped. The Υ-shaped pattern is prevalent in the data and robust to various econometric specifications. Moreover, the cross-sectional behavior of additional pricing moments such as the average frequency and the absolute size of price changes is consistent with time-series evidence in the literature.

Standard sticky-price models cannot reproduce the Υ-shaped pattern while accounting for the other pricing moments. Thus, I develop a menu-cost model with idiosyncratic productivity shocks and sequential consumer search that can. In the model, the representative household has a worker who supplies labor and a continuum of shoppers, each purchasing a continuum of goods. A continuum of retailers produces and sells each homogeneous good using labor. Retailers set nominal prices and face a menu cost; they adjust because relative prices erode at a deterministic inflation rate – which varies among products because of differing productivity growth rates – and idiosyncratic shocks hit their productivity. The non-degenerate and time-invariant cross-sectional distribution of relative prices results from both product-level inflation and idiosyncratic shocks.

Every period, a continuum of shoppers enter, search, purchase, and leave the market for each product. They take the relative-price distribution as given and search sequentially for the lowest price. For each additional price draw, shoppers pay a heterogeneous search cost. Their strategy is to search until finding an offer lower than their reservation price. Each buyer ends up purchasing the homogeneous good from one of many nearly identical retailers. Therefore, shopping behavior determines the equilibrium demand curve. Product-level inflation, on the other hand, affects the returns to search by directly impacting price dispersion.

A demand curve that changes endogenously with inflation through search is key to replicate the Υ-shaped relationship between inflation and price dispersion in the data. The intuition is as follows. At zero inflation, the only source of price dispersion is the idiosyncratic productivity shocks. Thus, searching might be profitable only for low-search-cost buyers. In this case, retailers with high productivity draws have incentives

[1] In what follows, I will use the terms "product-level inflation" and "inflation" indistinctively and refer to economy-wide inflation as "aggregate inflation."

to set a price low enough to attract searchers. It turns out that, since relative prices are fixed, sellers optimally bunch at the reservation price of searchers. Only retailers with low productivity draws, serving high-search-cost shoppers who do not search, set higher prices. Because of this bunching behavior, price dispersion at zero inflation is relatively low.

A small deviation from zero inflation makes relative prices continuously drift downward; bunching at any price would require retailers to pay the menu costs every period, which is suboptimal. Therefore, retailers let their prices erode before adjusting, and price dispersion increases. More dispersed prices increase the returns to search, producing a feedback effect on price dispersion. On the one hand, highly productive sellers have incentives to charge lower prices, attracting more customers and increasing their sales volume. On the other hand, because searchers flee from high prices, the least productive retailers end up serving a larger fraction of non-searchers or captive shoppers, even having incentives to increase their prices. As a result, price dispersion rises sharply.

Furthermore, I provide empirical evidence on shopping behavior and inflation that supports the theory. The model predicts the higher is absolute inflation in a given market, the lower are the prices shoppers visiting several stores pay. I test this prediction using consumer panel data for the same retail products in the scanner data. These data contain the barcodes of the items households purchased on each shopping trip, the quantities they bought, and the prices they paid. In addition, we can identify the number of distinct retailers a household visited to purchase the good. I merge shopping-behavior measures from these data to product-level inflation from the retail scanner data. I find that a household visiting ten stores when absolute inflation is 10% pays 1% less than when inflation is zero and 2.4% less than a household visiting only one store.

I jointly calibrate the parameters of the model to match the average markup, the absolute size and frequency of price changes, shopping behavior at zero inflation, and the Υ-shaped pattern of product-level price dispersion in inflation. The latter is crucial to identify the parameters of the search-cost distribution. When the density of shoppers with a positive but negligible marginal search cost is relatively large, at zero inflation, more productive retailers bunch at the highest price that low-search-cost shoppers will accept because there are many of these searchers. For a small positive level of inflation, dispersion at the lowest prices increases because of real price erosion, prompting the large density of low-search-cost shoppers to search more, which increases price dispersion further, producing the data pattern.

In my model, the welfare effects of product-level and aggregate inflation are ambiguous. The costs of inflation stem from adjustment costs, as in standard menu-cost models, and from search costs: inflation-induced price dispersion increases shoppers' returns to search, and thus the resources spent on the search for the best prices. Through price dispersion and costly search, inflation can be beneficial for welfare. The calibrated model suggests that it is. Given that search activity is limited at zero inflation, retailers set high markups and extract the surplus from consumers. A low positive inflation rate increases price dispersion and the returns to search considerably. Highly productive retailers charge lower prices to attract more shoppers, particularly those who have a lower search cost and search more. As a consequence, markups decrease, generating significant efficiency gains. These gains dissipate for large levels of inflation: the least productive retailers, whose customers become mostly captive shoppers, optimally charge higher prices, increasing average markups.

Related literature Extensive research on the theoretical relationship between price dispersion and inflation exists. My paper is directly related to two types of models in the literature. First, monetary models in which sticky prices and inflation generate price dispersion. Second, models of monetary exchange in which price dispersion arises from buyers' incomplete information on the prices charged by each seller.

In models where nominal price changes are costly and no aggregate or idiosyncratic uncertainty exists (Barro, 1972; Sheshinski and Weiss, 1977; Benabou, 1988), the optimal policy is an (S, s) pricing rule: the firm keeps the nominal price fixed while the real price drifts continuously from the initial level S to the terminal level s, at which point it jumps back to S. The higher the expected inflation, the larger the distance between such bounds. If inflation is constant and firms follow a common (S, s) rule, the cross-sectional distribution of real prices is log-uniform on $[s, S]$.[2] In this case, the coefficient of variation is increasing and concave in the max-min price ratio, S/s.[3] Because this ratio generally increases with inflation, price dispersion tends to be increasing and concave in inflation. Such a prediction on cross-sectional pricing behavior

[2]Caplin and Spulber (1987) show a log-uniform on $[s, S]$ is the only cross-sectional distribution consistent with: (i) time-invariant common (S, s) rules, and (ii) a price index that grows at the aggregate inflation rate.

[3]The coefficient of variation of a log-uniform distribution is $\sqrt{\left(\frac{S/s+1}{S/s-1}\right)\frac{\log(S/s)}{2} - 1}$.

is in line with my findings. However, the predictions regarding firms' dynamic price-setting behavior are at odds with micro-level empirical evidence: at low to moderate inflation rates, price decreases are as common as price increases, and the absolute size of adjustments is significantly larger than aggregate inflation.

Golosov and Lucas (2007) replicate those features of firms' dynamic pricing behavior by introducing idiosyncratic productivity shocks in a menu-cost model. In their setting, firms adjust not only because the real price is declining but also because the real cost of production is changing stochastically. For low to moderate inflation levels, the large idiosyncratic shocks induce firms to adjust before inflation erodes their real prices too much; cross-sectional price dispersion is practically constant and increases smoothly with inflation. When inflation is high, the main reason for firms to adjust is to catch up with aggregate inflation, as in Sheshinski and Weiss (1977). Therefore, price dispersion is U-shaped in inflation.

In contrast, New Keynesian models typically assume prices are sticky but adjusting opportunities arrive at an exogenous rate, as in Calvo (1983). Because firms do not have the option to adjust before their prices drift far from optimal levels, cross-sectional price dispersion rises rapidly with inflation. Nevertheless, idiosyncratic shocks tend to make price dispersion smooth at zero inflation. Thus, both models predict a U-shaped relationship between inflation and price dispersion, at odds with my findings.

Head and Kumar (2005), on the other hand, study the effects of inflation on price dispersion by embedding the price posting environment of Burdett and Judd (1983) in a model of monetary exchange. In their setting, buyers hold fiat money and search non-sequentially for a seller. The equilibrium price distribution is non-degenerate if some buyers observe a single price quote, whereas others observe more than one. The model predicts a positive relationship between inflation and price dispersion which, as in my model, is tightly related to market power. Inflation erodes the purchasing power of fiat money, so the fraction of buyers observing a single price increases. In response, sellers pricing at the upper end of the distribution raise their prices by a relatively large amount: since a higher share of their customers are captive buyers, the decline in sales will be small. Conversely, sellers pricing at the lower end of the distribution are constrained in their price increases by the fact that they can lose a significant volume of sales to competitors. An important shortcoming of this framework is that prices are fully flexible, contrary to what the data shows.[4]

[4] In the data, prices stay fixed for at least four months on average. See Nakamura and Steinsson (2008) for a representative characterization of price changes in the U.S.

Head et al. (2012) embed Burdett and Judd (1983) pricing into a dynamic New Monetarist model that delivers price dispersion even without inflation or idiosyncratic productivity shocks. In a similar setting, Burdett and Menzio (2018) introduce a menu cost of adjustment: for high-enough values of this cost, the cross-sectional distribution of real prices converges to a log-uniform, and for values close to zero, it converges to Burdett and Judd (1983). The main drawback of these models is that they are unable to generate price decreases for positive inflation. Burdett and Menzio (2017) overcome this issue introducing idiosyncratic shocks to sellers. However, unless the menu cost is sufficiently high, there is a negative relationship between inflation and price dispersion.[5]

In my theory, I assume sellers face idiosyncratic productivity shocks and a menu cost of adjusting nominal prices. These assumptions generate realistic firm pricing behavior and allow us to take the model to the data. I borrow the firm price-setting block from Golosov and Lucas (2007), as is standard in the literature with micro-founded sticky prices. On the consumer's side, I assume heterogeneous shoppers search sequentially for the lowest price. Sequential search facilitates the introduction of firm-level heterogeneity and the mapping between model and data on the firms' side. The consumer search block is primarily based on Benabou (1992). As we will see, both price stickiness and incomplete information of consumers are key to generate the Υ-shaped relationship between inflation and price dispersion in the data.

My paper also contributes to the scarce literature studying the empirical relationship between inflation and price dispersion.[6] Tommasi (1993) finds a positive correlation between absolute inflation and price dispersion of identical goods, but the scope of the findings is limited: the sample is for 15 products sold by five supermarkets in Argentina. Reinsdorf (1994) uses the micro-level data underlying the U.S. CPI to compute price dispersion of *similar* goods. Assuming price dispersion is linear in inflation – not absolute inflation – he finds the variables are negatively correlated.

Nakamura et al. (2018) extend the dataset used by Reinsdorf (1994) back to 1977;

[5] In this model, sellers follow a (Q, S, s) rule: they let inflation erode real prices until they reach s, at which point they pay a menu cost and adjust to a random price in $[S, Q]$. As inflation increases, S increases but Q and s remain constant. As a result, more sellers price in $[s, S]$, where the distribution is less dispersed than in $[S, Q]$.

[6] Mostly due to data limitations, earlier research focused on price *change* dispersion – or relative price variability (RPV) – instead of price *level* dispersion. Van Hoomissen (1988), Lach and Tsiddon (1992), and Beaulieu and Mattey (1999) study the relationship between inflation and intra-market RPV; Parsley (1996) and Debelle and Lamont (1997) focus on inter-market RPV and inflation. Although some of these papers use RPV as a proxy for price dispersion, the relationship between both variables is not straightforward in models or data (Nakamura et al., 2018).

the resulting data exhibit significant variability in aggregate inflation. Nevertheless, products are identified by narrow categories (e.g., “carbonated drinks”), not by barcodes or brands. Therefore, as the authors state, much of the within-category dispersion likely results from differences in product size and quality. To overcome this data limitation, they analyze the relationship between inflation and the absolute size of price changes instead. They argue that such a relationship should inform about inflation and price dispersion because that is the case in current sticky-price models (i.e., New Keynesian and Golosov and Lucas style menu-cost models). They show that, although annual inflation has fluctuated between -2% and 12% since 1977, the mean absolute size of price changes has been practically constant. They conclude that the main costs of inflation in the models they study are absent in the data.

Alvarez et al. (2019) use the micro-level price data underlying the Argentinian CPI between 1988 and 1997, a period in which monthly inflation ranged from 200% to less than zero. These data allow comparing the prices of goods, with the same brand and package, across stores every two weeks. The authors measure aggregate price dispersion as the residual variance in a regression of prices on a rich set of fixed effects. They find the elasticity of price dispersion with respect to absolute inflation is zero for inflation below 10% per year and close to one-third at high inflation rates.

Sheremirov (2020) studies the cross-sectional relationship between inflation and price dispersion using retailer scanner data. The main differences with the data I use, aside from being gathered by a different company, are the period (2001-2011), the coverage (31 categories comparable to the 1,000 I have), and the existence of flags for temporary price reductions (i.e., sales). He documents a negative relationship between inflation and dispersion of prices including sales. After removing sales, this correlation becomes positive. Although my measures of price dispersion include sales, Sheremirov’s findings suggest removing them should not affect my results qualitatively.

Because I observe products at the barcode level, I overcome the challenge Nakamura et al. (2018) and earlier research faced. In addition, by comparing prices of identical goods before computing aggregates, my measure of price dispersion is closer to the models than the measure in Alvarez et al. (2019). Unlike Reinsdorf (1994) and Sheremirov (2020), I do not impose linearity when studying the comovement of inflation and price dispersion; the flexibility of a non-parametric specification allows me to uncover the Υ-shaped relationship between both variables.

Finally, I contribute to the literature on the costs and benefits of inflation. The costs of inflation are typically associated with price dispersion of identical goods. In current sticky-price models (i.e., New Keynesian or menu-cost models with idiosyncratic shocks and without consumer search), inflation-induced price dispersion tends to decrease aggregate labor productivity and welfare. The intuition is that as nominal prices stay fixed, relative prices drift away from their optimal levels under a positive inflation rate. Hence, relative prices no longer reflect the relative costs of production, negatively affecting efficiency. The welfare losses – even for low to moderate inflation rates – are substantial in the New Keynesian model because price dispersion increases significantly with inflation (Burstein and Hellwig, 2008; Nakamura et al., 2018).

Menu-costs models, on the other hand, predict the negative effects of low to moderate inflation on welfare are negligible because: (i) price dispersion is essentially flat in inflation, and (ii) the physical cost of changing prices is relatively small (Nakamura et al., 2018; Alvarez et al., 2019).[7] Moreover, in the presence of a zero lower bound on nominal interest rates, the benefits from a positive level of inflation more than offset such costs (Blanco, 2021).[8]

In monetary models with consumer search (Benabou, 1988, 1992; Diamond, 1993; Head and Kumar, 2005), inflation-induced price dispersion can be welfare-improving. In particular, if monopolistic competition arises from costly consumer search instead of imperfect substitutability of the goods, consumers could *exclusively* buy from sellers charging the lowest prices. Then, by increasing price dispersion, inflation can increase the returns to search and decrease firms' market power, potentially increasing welfare.

The article proceeds as follows. Section 2 presents evidence from scanner data on the Υ-shaped relationship between price dispersion and inflation. Section 3 develops a menu-cost model with consumer search that reproduces the Υ-shaped pattern. Section 4 explains the intuition behind this result. Section 5 presents evidence on shopping behavior and inflation that supports the mechanism. Section 6 shows the model fit to the data and Section 7 discusses the welfare implications of inflation in the calibrated model. Section 8 concludes.

[7] Burstein and Hellwig (2008) consider, in addition, the effects of inflation on the opportunity cost of holding real money balances. They find the welfare costs arising from price dispersion are negligible compared to those from this extra channel.

[8] Danziger (1988) shows that in a menu-cost model without idiosyncratic shocks, a low positive inflation rate can be better than zero. Intuitively, in the absence of inflation, firms charge the static profit-maximizing price at all times; under positive inflation, lower real prices in the periods preceding the adjustment make consumers better off.

2 Evidence on inflation and price dispersion

How does price dispersion empirically relate to aggregate inflation? The first requirement to answer this question is to precisely measure the price dispersion of identical goods. To do so, we need to observe the prices that several sellers in a particular geographic area charge for the same good on the same date. The NielsenIQ Retail Scanner dataset for the U.S. satisfies this condition, as I explain in the next subsection.

Second, we would require these granular price data over different economy-wide inflation regimes: we are interested in the relationship between *aggregate* inflation – which can be influenced by the monetary authority – and price dispersion. Nonetheless, the data is only available for a relatively short period with low and stable aggregate inflation, implying the variation for a time-series analysis is insufficient. Conversely, the variation in county- and product-level inflation rates is substantial. Thus, I identify the relevant relationship by exploiting cross-sectional variation and derive aggregate implications using a model calibrated to match the disaggregate evidence.

2.1 Data

The NielsenIQ Retail Scanner Data are available for the 2006-2020 period and have non-durable items at the barcode level (Universal Product Code – UPC), sold mostly by grocery, drugstore, and mass-merchandise chains. These products can be categorized at the module level (i.e., highly substitutable products that differ only in their brand), then at the group level (i.e., modules serving similar purposes), and finally at the department level. An example of a module would be "Ground and Whole Bean Coffee" in the product group "Coffee" from the department "Dry Grocery". The data contain 10 departments, 125 groups, and 1,078 modules, approximately.

For each week and UPC, stores report total units sold and total revenues. At the UPC level, we observe product description, brand, multi-pack, size, and additional characteristics in some cases (e.g., flavor).[9] Over 35,000 stores are located across the U.S., classified into around 1,460 counties. Each store can be associated to a retail chain.[10] I use the Retail Scanner dataset to compute disaggregate inflation measures

[9] I exclude private-label products because NielsenIQ alters the barcode so a particular store cannot be identified: for generics, the UPC does not represent a unique good.

[10] The retail chain identifier might refer to (1) the Corporate Parent if the company centralizes their data release to NielsenIQ across all of their retail brands, or (2) the Retail Banner of the store if the company decentralizes their data release to NielsenIQ through multiple data centers representing each of their retail banners.

and seller-level statistics on price dispersion.

Price dispersion is defined as contemporaneous discrepancies between the prices offered by different sellers of the same good around an average price. Scanner data allow us to be consistent with such a definition because a product is defined by its barcode, so we compare prices of the same good across stores.[11] Therefore, we can compute price dispersion controlling for product heterogeneity. Moreover, each product is linked to a store, and because each store is associated to a geographic market, we can control for systematic regional differences. Lastly, weekly-level data allow us to compare the prices across stores at a given instant of time.

2.2 Variable construction

I define products as items that have the same barcode and thus the same brand, size, and characteristics. I compute inflation and price dispersion at the product × county × month (k, m, t) level. I use monthly instead of weekly prices to minimize the incidence of missing values. For the same reason, I only keep product × store pairs that appear in every month of a given year. To compute meaningful measures of price dispersion, I require that each product is sold by at least ten stores in each county and month.

Inflation I define the price store i in county m and month t charges for product k as

$$P_{ikm,t} = \frac{\sum_{\tau=1}^{T_t} Rev_{ikm,\tau}}{\sum_{\tau=1}^{T_t} q_{ikm,\tau}} = \frac{Rev_{ikm,t}}{q_{ikm,t}},$$

where $Rev_{ikm,\tau}$ and $q_{ikm,\tau}$ are total revenues and total units sold, respectively, in each week τ of month t; $T_t \in \{4, 5\}$, depending on the month. Price-level inflation is the annualized monthly average price change across stores selling product k in county m,

$$\pi_{km,t} = 12 \times \sum_{i=1}^{N_{km,t}} \frac{(\ln P_{ikm,t} - \ln P_{ikm,t-1})}{N_{km,t}},$$

where $N_{km,t}$ is the total number of stores. To validate this inflation measure, I compute a sales-weighted average over all food-related product × county pairs and compare it with the official statistics. Figure A.1 shows aggregate inflation for food categories

[11] Since the retail chain identifier does not distinguish between parent companies and retail banners, I focus my analysis on stores. In the robustness checks below, I compute price dispersion across retail chains. The qualitative results remain unchanged.

in these data closely tracks the food-at-home CPI reported by the BLS for the same period.

Price dispersion The main measure of price dispersion that I use is the unweighted standard deviation of log prices across stores for each product, county, and month:

$$\sigma_{km,t} = \sqrt{\frac{1}{N_{km,t}-1}\sum_{i=1}^{N_{km,t}}\left(\ln P_{ikm,t} - \frac{\sum_{i=1}^{N_{km,t}} \ln P_{ikm,t}}{N_{km,t}}\right)^2}.$$

In addition, I compute alternative measures of price dispersion (interquartile range, 90-10, 90-50, and 50-10 ratios of the price distribution) for each (k, m, t).[12] [13]

2.3 Estimation and results

A comparative-statics analysis assumes inflation is constant and price dispersion is computed using an invariant distribution. Therefore, to analyze the data, we need to choose a time horizon within which we assume retailers' behavior completely adjusts to changes in inflation. Because the data show retailers adjust their prices every four to six months on average, I conjecture that the adjustment period is 12 months. Annual inflation and price dispersion for each product $\times$ county correspond to the average across months for each year:

$$\pi_{km,y} = \frac{\sum_{t=1}^{12} \pi_{km,t}}{12}; \quad \sigma_{km,y} = \frac{\sum_{t=1}^{12} \sigma_{km,t}}{12}.$$

Table A.1 shows descriptive statistics for these product-level measures and for a sales-weighted average inflation. The final sample for estimation includes more than 9 million product $\times$ county pairs for 268,149 unique products and 904 counties. Each pair is present in the data for an average of 4.4 years and contains information for an average of 26 stores. Disaggregate inflation exhibits significantly larger variation than aggregate inflation, with a coefficient of variation (CV) of 8.15 and 1.13, respectively.

[12] Weighting prices using annual store sales when constructing product-level inflation or price dispersion measures does not affect the main results.

[13] Other than computing sales-weighted measures, one could control for sellers' heterogeneity by removing seller-fixed effects from the price before obtaining price dispersion. However, eliminating persistent price-level differences generates downward-biased estimates of price dispersion, thus causing a downward shift in its relationship with product-level inflation. Despite the level shift, the qualitative results regarding the shape of the relationship remain unchanged.

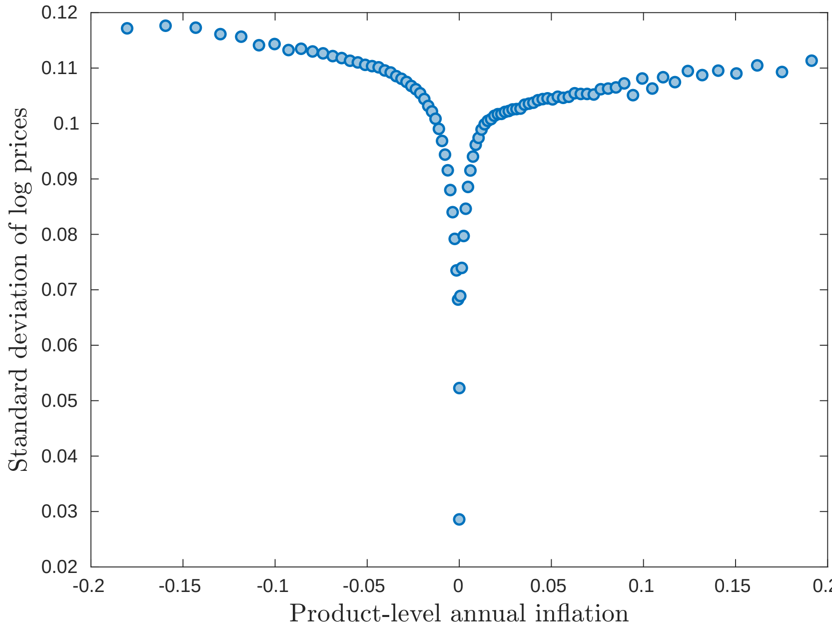

Figure 1: Product-level price dispersion and inflation, raw data

Each dot corresponds to average price dispersion for each of 100 equally sized product-level annual inflation bins. The unit of observation is a product $\times$ county $\times$ year, for 40.1 million in total.

To understand the relationship between disaggregate price dispersion and inflation, I start by constructing a binned scatterplot: I divide annual product-level inflation $\pi_{km,y}$ into 100 equally sized bins and average price dispersion $\sigma_{km,y}$ within each bin. In this way, we can analyze the data without imposing any parametric structure. As Figure 1 shows, price dispersion is at its lowest average levels when inflation is close to zero; as inflation deviates from zero, price dispersion increases steeply, flattening at 2%. Hence, the data suggest the relationship between both variables is non-differentiable at zero. Because this pattern resembles the Greek capital letter upsilon, I hereafter refer to it as upsilon-shaped, or Υ-shaped.

Next, I assess how the results change when controlling for observable and unobservable factors in the product, county, or time dimensions. In particular, I estimate the following non-parametric regression of price dispersion on inflation:

$$\sigma_{km,y} \quad = \quad \sum_{n=1}^{100} \beta_n \mathbf{1}_{\{\pi_{km,y} \in B_n\}} + a_{k,y} + b_{m,y} + \alpha \log N_{km,y} + \epsilon_{km,y}. \tag{1}$$

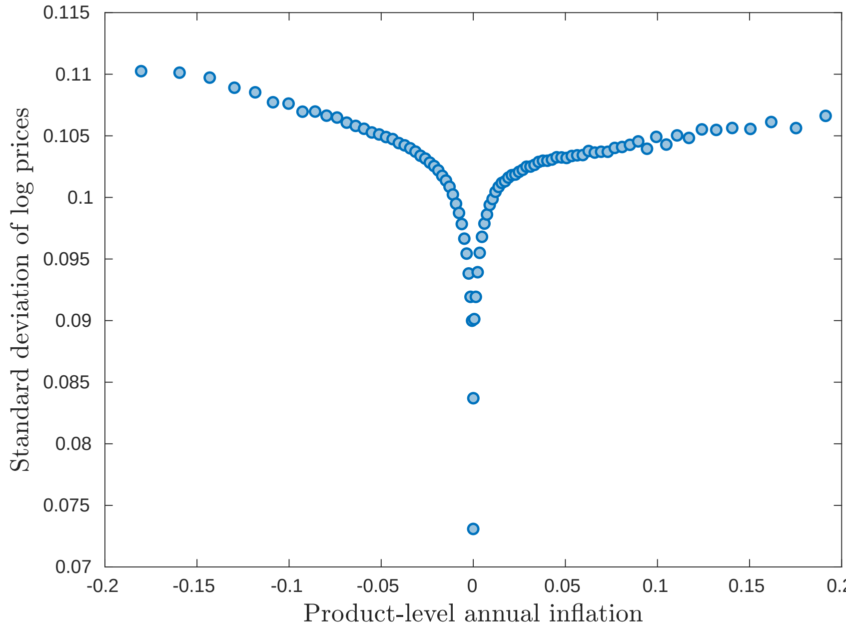

Figure 2: Product-level price dispersion and inflation, predicted

The dots correspond to average price dispersion conditional on covariates for each of 100 equally sized product-level annual inflation bins as predicted by equation (1). The unit of observation is a product $\times$ county $\times$ year, for 40.1 million in total.

The coefficients $\{\beta_n\}_{n=1}^{100}$ correspond to average price dispersion at each equally sized inflation bin $\{B_n\}_{n=1}^{100}$, conditional on covariates: $a_{k,y}$ and $b_{m,y}$ are product-year and county-year fixed effects, respectively, and $N_{km,y}$ denotes the total number of stores. This specification implies we identify $\{\beta_n\}_{n=1}^{100}$ by exploiting the cross-product inflation variation within each county and year, taking into account unobservable product-level differences that might be changing over time. In addition, we control for the number of stores, because research has shown this variable explains a substantial degree of the cross-sectional price-dispersion variation in the data (Hitsch et al., 2019).

Figure 2 shows the price dispersion predicted by inflation using estimates from equation (1). The results confirm the Υ-shaped pattern between price dispersion and inflation is prevalent in the data, even when comparing different products in a given county-year, or a product across counties for the same year (Figure A.3). In particular, including fixed effects increases average price dispersion around zero inflation and makes price dispersion a steeper function of absolute inflation for values larger than 2%.

Moreover, the Υ-shaped relationship is robust to controlling non-parametrically (Figure A.5) for the number of stores; that is, binned scatterplots of $\pi_{km,y}$ and $\sigma_{km,y}$ keeping products sold by at least 10, 25, or 50 stores display the same pattern.

On the other hand, price dispersion is approximately symmetric around zero. Figure A.6 plots estimates of equation (1) for 100 equally sized absolute inflation bins and the curves in Figure 2 on the same x-axis. The figure shows that the relationship between price dispersion and absolute inflation is increasing and concave around zero; it becomes linear for values larger than 2%, and steeper for negative than positive inflation.

2.3.1 Robustness checks

Category-level statistics With 268,149 unique products classified into 1,078 product modules/categories, a few overrepresented categories could be driving the product-level results. To show this is not the case, I aggregate product-level statistics into categories using total annual sales as weights.[14] As Figure A.4 shows, the category-level analysis qualitatively reproduces the main product-level findings of this paper. I also provide the robustness checks described below at the category level.

Statistical significance To assess the statistical significance of the estimated relationship, I compute confidence bands that reflect the underlying variance of the data. I construct these bands with standard errors clustered by product $\times$ county pair to control for within-pair error correlation over time. Because the confidence band in Figure A.7 covers the entire function with probability 0.95 we cannot reject, visually, that price dispersion is Υ-shaped in inflation.

Year-by-year estimates I test whether the relationship holds for each county and year separately. I start by assuming price dispersion is a symmetric, continuous, and differentiable function of inflation, which I identify exploiting within-county-year variation across products[15]:

$$\sigma_{km,y} = f_{m,y}\left(|\pi_{km,y}|\right) + \alpha \log N_{km,y} + \epsilon_{km,y}.$$

[14]I consider all k sold by at least five stores in each m and t. The category-level dispersion measure is the squared root of the sales-weighted average product-level standard deviation of log prices.

[15]I estimate this function non-parametrically by fitting a second-degree polynomial within each bin and forcing the curves to be smoothly connected at the boundaries of the bins.

Under these assumptions, I test two hypotheses for each county-year function $f_{m,y}$. First, that the function is monotonically non-decreasing; second, that it is concave[16]:

$$\inf_{x>0} f'_{m,y}(x) \geq 0; \quad \sup_{x>0} f''_{m,y}(x) \leq 0.$$

The top panel in Table A.4 shows the first hypothesis is rejected for 25% and the second for 35% of the total county-year combinations. When we restrict absolute inflation to values lower than 2%, the same hypotheses are rejected for only 6% and 9%, respectively, of the county-year pairs. The bottom panel of the table shows similar results when we estimate the relationship for each product and year, exploiting cross-county variation. Table A.5 shows that the results are even stronger when using categories instead of products.

Sales and weekly-level statistics Temporary price reductions, or sales, are not flagged in the data, so I use three different filters to remove them from weekly posted prices and construct a weekly series of regular prices. Defining prices as total revenues over quantity sold each week instead of month would allow for better identification of price changes and, therefore, sales.[17] To construct the monthly-level statistics, I keep the regular price in the third week of the month and product $\times$ store pairs that appear in every month of a given year.[18] Figure A.9 shows the Υ-shaped pattern holds for each of the regular-price series, having only a level difference with their posted-price counterpart (Figure A.10).

Future inflation Up to this point, we have analyzed the relationship between current inflation and price dispersion. One could argue, nonetheless, that price-setting behavior depends on the inflation rate retailers expect between nominal price adjustments, not the one they observe. Define future inflation as the average realized inflation rate for the expected duration of the nominal price. As before, I choose 12 months as an upper bound for price duration. Figure A.8 shows the Υ-shaped pattern also holds for price dispersion and future inflation – an expected result in an environment with low and stable aggregate inflation.

[16]Both tests require differentiability of $f_{m,y}$. The estimates using pooled data suggest this function might be non-differentiable at zero. I assume $f_{m,y}$ is symmetric around zero to overcome this issue.

[17]See Section B for further discussion, and details on data cleaning and the filters used.

[18]Although I chose the third week of the month, using any other week between the 1st and 4th yields the same results.

Alternative measures of price dispersion To verify that the results are not specific to the standard deviation of log prices, I repeat the estimation procedure using different measures of price dispersion. Figure A.12 shows the relationship between the 90-10 ratio of the price distribution and inflation, which, again, is Υ-shaped. The same pattern holds for the variance of log prices (Figure A.11), the interquartile range (Figure A.13), and the 90-50 and 50-10 ratios of the price distribution (Figure A.14).

Retail-chain pricing To make sure that the results are not driven by particular pricing strategies of retail chains/parent companies across their stores, I construct a chain-level price.[19] For each product and month, I define the monthly price as total revenues over quantity sold by a given retail chain/parent company. I compute inflation and price dispersion statistics among chains/parent companies. Figure A.15 shows a clear Υ-shaped relationship between product-level inflation and price dispersion among retail chains.

3 Model

The empirical evidence in the previous section speaks to the relationship between *product-level* inflation and price dispersion. Product-level inflation has a product-specific real component – due to, for example, productivity or cost trends – and a nominal component. Because the inflation variation I exploit is not purely nominal, we cannot directly conclude the relationship between *aggregate* inflation and price dispersion is Υ-shaped.

In this section, I develop a multi-product monetary model that generates a Υ-shaped relationship between product-level price dispersion and allows us to study the effects of *aggregate* inflation. We start by describing the general structure of the model, which is illustrated in Figure 3.

Time is continuous and no aggregate uncertainty exists. A representative household has a worker who supplies labor and a measure-one continuum of heterogeneous shoppers, each purchasing a continuum of products.

Each homogeneous product is sold by a measure-one continuum of monopolistically competitive retailers. Retailers are infinitely lived, set nominal prices, and have

[19] For example, DellaVigna and Gentzkow (2019) find retail chains charge nearly-uniform prices across stores.

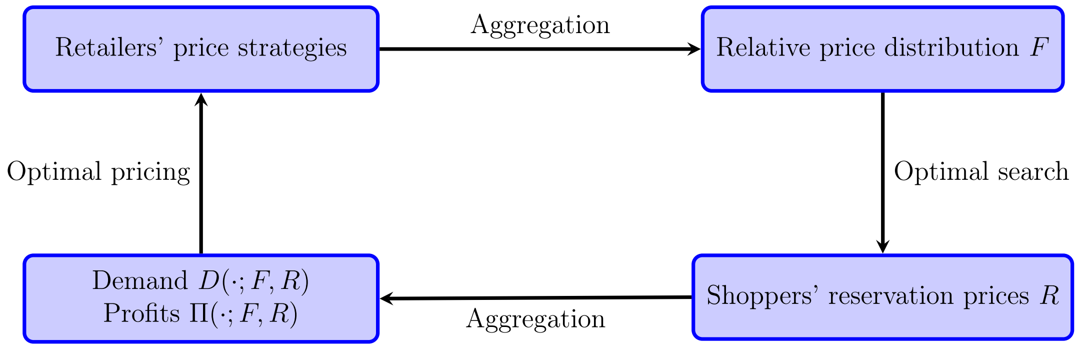

Figure 3: General model structure

a production technology linear in labor. A fixed cost is incurred when changing nominal prices, and retailers would like to adjust them for three reasons: (i) they face idiosyncratic transitory shocks to their productivity, (ii) product-specific productivity changes at a deterministic rate, and (iii) the nominal wage is increasing at a deterministic rate. These elements change retailers' production costs and therefore their desired prices.[20] At the same time, both wage inflation and idiosyncratic shocks are key to produce a non-degenerate and stationary cross-sectional distribution of relative prices, while product-specific productivity growth generates variation in product-level inflation rates. Retailers make their price-setting decisions taking into account the shopping behavior of buyers.

Each instant of time, a continuum of shoppers enter, search, purchase, and leave the market for each product. Shoppers take the relative-price distribution as given and search for the lowest price, paying a cost for each new price draw. The search cost is heterogeneous across buyers. As in a McCall (1970) type of search, buyers follow a reservation-price strategy: they accept offers up to a relative reservation price at which they are indifferent between buying and searching again. Each buyer ends up purchasing a product from one of many nearly identical retailers.

Monopolistic competition is an outcome of the model. Because search is costly, each retailer sells the good to a positive share of buyers. In equilibrium, the magnitude of this share depends inversely on the retailer's price. The aggregation of consumer search rules generates a downward-slopping demand curve for the retailer. Through

[20] Equivalently, we could assume the existence of a perfectly competitive sector of manufacturers who produce one unit of the good using one unit of labor. Retailers purchase the good from manufacturers and face idiosyncratic shocks to the cost of selling – rather than producing – the good. As a result of perfect competition, the price that retailers pay for one unit of the good equals the nominal wage.

shopping behavior, the equilibrium demand curve (thus, the profit function) depends on the inflation rate. This feature will be at the center of the Υ-shaped pattern between inflation and price dispersion.

In equilibrium, product-level inflation is determined by both nominal wage and product-specific productivity growth; aggregate inflation is the average of product-level inflation rates. We assume the nominal wage growth rate is the policy parameter that the monetary authority can control to influence aggregate inflation, and that productivity growth rates are the only difference among products.

3.1 Retailers

A measure-one continuum of monopolistically competitive retailers indexed by i_k sell the homogeneous product k. These retailers are infinitely lived, with discount rate ρ. Money is used as the unit of account, so retailers set their prices $P_{i_k,t}$ in nominal terms. The production technology of a retailer selling good k is given by

$$y_{i_k,t} = A_{k,t} v_{i_k,t} l_{i_k,t},$$

where $A_{k,t}$ is product-specific productivity, $v_{i_k,t} \in [\underline{v}, \overline{v}]$ is the retailer-specific productivity, and $l_{i_k,t}$ is labor. The retailer's productivity follows the mean-reverting process:

$$d \log v_{i_k,t} = -\rho_v \log v_{i_k,t} dt + \sigma_v d\mathcal{Z}_{i_k,t}, \quad \rho_v > 0, \tag{2}$$

where $\mathcal{Z}_{i_k,t}$ is a standard brownian motion with zero drift and unit variance, distributed independently across retailers.

The labor market is competitive and retailers hire labor at a nominal wage W_t. We assume the nominal wage and product-specific productivity grow at constant rates λ_W and $\lambda_{A,k}$, respectively. Because relative prices are defined with respect to the nominal wage and product-specific productivity, that is, $p_{i_k,t} \equiv P_{i_k,t} A_{k,t}/W_t$, a fixed nominal price implies a relative price eroding at a rate $\pi_k \equiv \lambda_W - \lambda_{A,k}$. Each nominal price adjustment costs the retailer $\kappa > 0$ labor units.

A central aspect of the model is that the demand curve that a single retailer faces (and more importantly, its elasticity) is endogenously determined by the optimal behavior of shoppers and other retailers. Denote this downward-sloping demand as a function of the relative price by $\check{D}_{k,t}(p)$. The instantaneous nominal profit function for a retailer charging a nominal price $P_{i_k,t}$, given the nominal wage W_t, product-specific productivity $A_{k,t}$, and the stochastically determined labor productivity $v_{i_k,t}$, is

$$\check{\Pi}_{k,t}\left(P_{i_k,t}, v_{i_k,t}\right) = \left(P_{i_k,t} - \frac{W_t}{A_{k,t} v_{i,t}}\right) \times \check{D}_{k,t}\left(\frac{P_{i_k,t} A_{k,t}}{W_t}\right).$$

At any date t, retailers are characterized by a pair $\left(P_{i_k,t}, v_{i_k,t}\right)$. Because the nominal wage and product-specific productivity grow at constant rates, and the retailer-specific productivity follows a stationary process, real aggregates are expected to be time invariant. Therefore, we can express the problem in terms of relative prices and conjecture an equilibrium in which the current joint distribution of relative prices and productivity, $\phi_{k,t}(p, v)$, is time invariant. The equilibrium instantaneous real profit function is

$$\Pi_k\left(p_{i_k,t}, v_{i_k,t}\right) = \left(p_{i_k,t} - \frac{1}{v_{i_k,t}}\right) \times D_k\left(p_{i_k,t}\right),$$

where $D_k(p) \equiv A_{k,t}^{-1} \check{D}_{k,t}(p)$ is the time-invariant component of demand.

Let $\psi_k\left(p_{i_k,0}, v_{i_k,0}\right)$ denote the present value of a retailer i that begins at $t = 0$ with the relative price $p_{i_k,0}$ and idiosyncratic productivity $v_{i_k,0}$. Given the strategies of shoppers and other retailers, i chooses a time $T \geq 0$ to adjust and a price p' to solve

$$\psi_k\left(p_{i_k,0}, v_{i_k,0}\right) = \max_T \mathbb{E}\left\{\int_0^T e^{-\rho t} \Pi_k\left(p_{i_k,t}, v_{i_k,t}\right) dt + e^{-\rho T} \max_{p'}\left[\psi_k\left(p', v_{i_k,T}\right) - \kappa\right]\right\}.$$

This time-invariant Bellman equation is standard in the literature, and its solution is well known.[21] The optimal policy of a retailer with productivity level v is to leave its nominal price unchanged if p is between $p_{L,k}(v)$ and $p_{U,k}(v)$. If the relative price hits any of these bounds, the retailer pays the menu cost κ and adjusts to $\hat{p}_k(v)$.

Before stating the problem of the representative household, defining the stationary product-specific posted-price distribution F_k as the marginal of $\phi_k(p, v)$ over v is useful:

$$F_k(p) = \int_{\underline{p}_k}^{p} \int_v \phi_k(x, v)\, dv\, dx, \quad p \in \left[\underline{p}_k, \overline{p}_k\right],$$

where $\underline{p}_k$ and $\overline{p}_k$ are determined in equilibrium by retailers' optimal pricing behavior.[22]

[21] In fact, firms in Golosov and Lucas (2007) solve the same problem. The main difference is that in the present paper, monopolistic competition is a result of costly consumer search.

[22] Retailer pricing strategies are monotonic in productivity $v \in [\underline{v}, \overline{v}]$, implying $\underline{p}_k = p_{L,k}(\overline{v})$ and $\overline{p}_k = p_{U,k}(\underline{v})$ in equilibrium.

3.2 Representative household

A representative household has a worker and a measure-one continuum of shoppers, each buying a continuum of goods. The worker supplies labor and shares income equally across shoppers. At instant dt, the shoppers enter, search, purchase, and leave the market for each product. The worker and the shoppers are replaced by a new household an instant later. Because the problem of the household is static, I drop the time subscripts in what follows.[23] [24]

Shoppers have incomplete information about the prices: they know the posted-price distributions F_k, but not the price that each seller charges. Because each good is homogeneous and price dispersion exists, buyers have incentives to search for the lowest price. Shoppers j search sequentially for a retailer i_k. They receive a first price quote for free but need to pay a disutility cost for each subsequent draw. We assume the search cost is heterogeneous.[25] The shopper searches S times to find a retailer i_k from whom he buys $\check{q}_{j,k}$ units of the good k at the nominal price P.

Given the search outcomes of shoppers, P and S, the household head chooses labor supply L and quantities of the consumption goods to solve

$$\begin{aligned}
&\max_{\left\{\check{q}_{j,k}(\cdot)\right\}_{j,k},L} \quad \frac{Q^{1-\sigma}}{1-\sigma} - \tau L - \alpha \bar{C} \\
&Q = \left(\int_0^1 Q_k^{1-\frac{1}{\theta}} \, dk \right)^{\frac{1}{1-\frac{1}{\theta}}} \\
&Q_k = \left(\int_0^1 Z_k^{\frac{1}{\eta}} \mathbb{E}_{j,k} \left[\check{q}_{j,k}\left(P\right)^{1-\frac{1}{\eta}} \right] dj \right)^{\frac{1}{1-\frac{1}{\eta}}} \\
&\bar{C} = \int_0^1 \int_0^1 \mathbb{E}_{j,k} \left[\gamma C\left(S\right) \right] dj \, dk \\
s.t. \quad &\int_0^1 \int_0^1 \mathbb{E}_{j,k} \left[P \check{q}_{j,k}\left(P\right) \right] dj \, dk = WL + \mathcal{D}
\end{aligned}$$

[23] Under this assumption, we rule out complex issues such as learning and intertemporal arbitrage. Then, we can focus on the main links between search and price-setting behavior.

[24] The assumption of instantaneous consumers also implies search happens quickly compared with relative-price erosion. This implication is reasonable for the U.S. economy, where inflation is low and sellers take between 4 and 10 months (Bils and Klenow, 2004; Nakamura and Steinsson, 2008) to change the prices of goods that households purchase every week.

[25] Heterogeneity in the search cost is key to generate a link between inflation and search activity. If buyers were identical, they would search at most once for any inflation level. Therefore, higher levels of inflation would not increase the resource cost of search and affect welfare through this margin.

where utility comes from the consumption aggregate Q, and disutility from labor L and search expenses $\bar{C}$. The consumption aggregate Q nests product-level consumption, Q_k, which at the same time nests shoppers' consumption, $\check{q}_{j,k}$.[26] The latter depends on the nominal price to be paid by shopper j to retailer i_k for product k, P, which is made explicit for clarity. The total disutility that shopping for any product generates is given by $\alpha\gamma C(S)$, where γ is the heterogeneous component of the search cost and $C(\cdot)$ is a function relating the total search cost to the total number of draws. I assume that γ is randomly distributed across shoppers within the household according to G. Both P and S are random outcomes of the search process, and the operator $\mathbb{E}_{j,k}$ indicates that such process varies among shoppers j and products k.[27] A deterministic demand shifter Z_k exists for each product. Finally, the household receives labor income, WL, and dividends from retailers, $\mathcal{D}$. We assume $\theta > 0$ and $\eta > 1$.

The problem's solution implies each shopper receives the same demand schedule,

$$\check{q}_k(P) = \left(\frac{\tau P Z_k^{-\frac{1}{\eta}}}{W}\right)^{-\eta} Q^{\frac{\eta}{\theta}-\eta\sigma} Q_k^{1-\frac{\eta}{\theta}}. \tag{3}$$

The associated indirect utility of the household is given by

$$U = \sigma\frac{Q^{1-\sigma}}{1-\sigma} + \tau\frac{\mathcal{D}}{W} - \alpha\bar{C}. \tag{4}$$

Therefore, the marginal change in utility as a particular price changes is

$$\frac{\partial U}{\partial P} = -Q^{\left(\frac{\eta}{\theta}-\sigma\eta\right)} Q_k^{\left(1-\frac{\eta}{\theta}\right)} \frac{\tau Z_k}{W}\left(\frac{\tau P}{W}\right)^{-\eta} \equiv U_k'(P)$$

and the surplus that each shopper derives from buying $\check{q}_k(P)$ units of k at P is

$$\check{V}_k(P) \equiv \int_P^{\infty} -U_k'(x)\,dx = Q^{\left(\frac{\eta}{\theta}-\sigma\eta\right)} Q_k^{\left(1-\frac{\eta}{\theta}\right)} \frac{1}{\eta-1}\left(\frac{\tau P Z_k^{\frac{1}{1-\eta}}}{W}\right)^{1-\eta}. \tag{5}$$

[26] We can interpret preferences for shopper's j consumption as preferences for differentiated varieties of k. While this nested CES structure facilitates the comparison with the literature, imperfect substitutability among goods is not necessary to generate monopoly power or replicate the Υ-shaped relationship between product-level price dispersion and inflation. A one-product model where shoppers have, e.g., quadratic or iso-elastic preferences and household's surplus is the simple average of shoppers' surpluses can produce the same qualitative results.

[27] If searching is so costly that all shoppers accept the first free draw, this setting is equivalent to Golosov and Lucas (2007).

As in the literature, we assume $Z_k = A_k^{1-\eta}$ and rewrite (3) and (5) in terms of p:

$$\begin{aligned} q_k(p) &= A_k (\tau p)^{-\eta} Q^{\frac{\eta}{\theta}-\eta\sigma} Q_k^{1-\frac{\eta}{\theta}}; \\ V_k(p) &= Q^{\left(\frac{\eta}{\theta}-\sigma\eta\right)} Q_k^{\left(1-\frac{\eta}{\theta}\right)} \frac{(\tau p)^{1-\eta}}{\eta - 1}. \end{aligned}$$

Shoppers take as given retailers' pricing strategies F_k, the utility they contribute to the household $V_k(p)$, and the utility cost of the marginal search,

$$\gamma \times c(S) \equiv \gamma \times \alpha \left[C(S) - C(S-1)\right].$$

Their objective is to find a stopping rule that maximizes

$$\mathbb{E}_{j,k}\left[V_k(p_S) - \gamma_j \times \sum_{s=2}^{S} c(s)\right],$$

where p_S is the random price at which they end up buying after S costly searches. The common component of the marginal search cost, $c(S)$, is monotonic in the integer S, and the maximum number of draws, $\overline{S}$, can be or not be finite. To explain the main insights of the model, I will assume a constant marginal search cost (i.e., $c(S) = \alpha$) and an unlimited number of draws (i.e., $\overline{S} \to \infty$), and relax both assumptions in the calibration section.

Under these assumptions, the optimal strategy of the shoppers is to search for a retailer until they find an offer below their reservation price.[28] To compute the optimal stopping rule, we first define the value that an offer p to purchase k has for a shopper j, $B_k(p, \gamma)$. This value is the maximum between accepting such an offer and continuing to search after paying the marginal search cost:

$$B_k(p,\gamma) = \max\left\{V_k(p), -\gamma \times \alpha + \int_{\underline{p}_k}^{\overline{p}_k} B_k(u,\gamma)\, dF_k(u)\right\}.$$

Then, we define the reservation price r_k as the relative price that makes shopper j

[28]Because utility is linear in search expenditures and the number of searches is not limited, the solution to this problem is the same with or without recall of previous offers.

indifferent between buying and searching again:

$$V_k(r_k) = \int_{\underline{p}_k}^{\overline{p}_k} [\max\{V_k(p), V_k(r_k)\}] \, dF_k(p) - \gamma \times \alpha.$$

If this equation has a solution, the optimal policy for the shopper is to accept any offer p if $p \leq r_k$, and to continue searching otherwise.[29] Alternatively, the reservation price equates the expected benefit and the cost of the marginal search:

$$\Gamma_k(r_k) \equiv \alpha^{-1} \int_{\underline{p}_k}^{r_k} [V_k(p) - V_k(r_k)] \, dF_k(p) \tag{6}$$

$$\Rightarrow \Gamma_k(r_k) = \frac{\tau}{\alpha A_k} \int_{\underline{p}_k}^{r_k} q_k(p) F_k(p) \, dp = \gamma. \tag{7}$$

A buyer with $\gamma < \tilde{\gamma}_k \equiv \lim_{r_k \to \infty} \Gamma_k(r_k)$ would rather search than accept a zero-surplus offer: the expected benefit of rejecting the first free offer and searching is greater than the marginal search cost. Because the marginal return to search $\Gamma_k(r_k)$ is increasing and continuously differentiable on $[\underline{p}_k, \infty)$, the reservation price is well defined and given by $r_k = R_k(\gamma) \equiv \Gamma_k^{-1}(\gamma)$, with $\underline{p}_k \leq r_k < \infty$.[30] Buyers with $\gamma > \tilde{\gamma}_k$ accept the first free offer as long as it does not exceed their maximum willingness to pay, so we let $R_k(\gamma) \to \infty$.

The reservation price $R_k(\gamma)$ is an increasing function of the search cost. This result implies that all types γ for which $p \leq R_k(\gamma)$ accept a given offer p. An equivalent and convenient statement is that buyers whose search cost is larger than the marginal return to search at offer p, that is, $\gamma \geq \Gamma_k(p)$, are the ones that accept such an offer.

3.3 Retailer-level demand function

Assume $\gamma \in [\underline{\gamma}, \overline{\gamma}]$ and that the distribution G is continuous and differentiable, with an associated probability density function g. Therefore, $g(\gamma)$ buyers exist with search cost γ. They search at random until finding one of the $F_k(R_k(\gamma))$ retailers charging $p \leq R_k(\gamma)$. Each of these retailers retains $g(\gamma)/F_k(R_k(\gamma))$ buyers type γ, and each of them purchases $q_k(p)$ units of the good. From the previous section, we know the buyers with $\gamma \geq \Gamma_k(p)$ would accept the offer p. Aggregating all the shoppers that a

[29] Lippman and McCall (1976) present a detailed discussion and proof.

[30] By integrating the left-hand side of equation (6) by parts, it is possible to obtain equality (7). Therefore, $\Gamma_k'(r_k) = \frac{\tau}{\alpha A_k} q_k(r_k) F_k(r_k) \geq 0$.

seller serves by setting price p gives:

$$N_k(p) \equiv \int_{\Gamma_k(p)}^{\overline{\gamma}} \frac{g(\gamma)}{F_k(R_k(\gamma))} d\gamma.$$

This function indicates the total number of transactions at a given price p, and, using equation (7), we can show it is decreasing in the price:

$$N_k'(p) \quad = \quad -\frac{\tau}{\alpha A_k} q_k(p) \, g(\Gamma_k(p)) \leq 0. \tag{8}$$

Intuitively, shoppers can take advantage of price dispersion and flee from higher prices. Nonetheless, in an equilibrium where $\gamma > \tilde{\gamma}_k$, a fraction of the shoppers will accept the first free draw, independently of the price. In this case, the number of transactions at each price is

$$N_k(p) = \int_{\Gamma_k(p)}^{\tilde{\gamma}_k} \frac{g(\gamma)}{F_k(R_k(\gamma))} d\gamma + 1 - G(\tilde{\gamma}_k),$$

where $1 - G(\tilde{\gamma}_k)$ is the fraction of *captive* shoppers. On the other hand, the number of shoppers that a retailer can attract by lowering its price is limited. Denote by $\underline{r}_k$ the minimum reservation price in the product k market, $R_k(\underline{\gamma})$. A seller who sets $p \leq \underline{r}_k$ serves all shoppers and reaches the maximum number of transactions, $N_k(\underline{r}_k)$.

Because the quantity each buyer purchases is independent of its type γ, we can express the equilibrium demand function for a retailer charging p as the product of two components:

$$\check{D}_k(p) = q_k(p) \times N_k(p). \tag{9}$$

The first term on the right corresponds to the *intensive margin* of demand, or the units that each buyer purchases at price p, $q_k(p)$. The second and novel component for menu-cost models represents the *extensive margin* of demand, $N_k(p)$. The extensive margin plays a crucial role in determining the price elasticity of demand that a retailer faces. For $p \geq \underline{r}_k$, the extensive margin tends to increase the overall elasticity of the demand curve:

$$\epsilon_{\check{D}_k}(p) \quad = \quad \begin{cases} \eta & p < \underline{r}_k \\ \eta + \epsilon_{N_k}(p) & p \geq \underline{r}_k \end{cases} \tag{10}$$

$$\epsilon_{N_k}(p) \quad = \quad \frac{p}{N_k(p)} \times \frac{\tau}{\alpha A_k} q_k(p) \, g(\Gamma_k(p)) \geq 0. \tag{11}$$

An extensive-margin elasticity that is determined endogenously is at the center of the relationship between inflation, price dispersion, and efficiency: optimal search rules in equation (6) depend on the price distribution F_k, which in turn depends on product-level inflation. Thus, the demand elasticity at each price changes with product-level inflation through shopping behavior, affecting markups and efficiency. This feature is absent in the standard Dixit-Stiglitz or Kimball demand systems typically used in the literature, where the demand elasticity as a function of the price remains constant for any level of inflation.

3.4 Equilibrium

In a stationary equilibrium, for each product k:

1. Worker chooses buying strategy $\check{q}_k(\cdot)$ (with associated product-level aggregate Q_k) given shoppers' search outcomes implied by their reservation prices $R_k(\gamma)$;

2. Shoppers $\gamma \in \left[\underline{\gamma}, \overline{\gamma}\right]$ choose $R_k(\gamma)$ given $\check{q}_k(\cdot)$ and the time-invariant relative posted-price distribution $F_k(p)$;

3. Market demand $\check{D}_k(p)$ results from aggregating individual demand from shoppers and is consistent with $R_k(\gamma)$ and $F_k(p)$;

4. Retailers choose optimal pricing strategies $\Psi_k(v) = \{p_{L,k}(v), p_{U,k}(v), \hat{p}_k(v)\}$ given $\check{D}_k(p)$;

5. Joint distribution of relative prices and productivity $\phi_k(p, v)$ and relative posted-price distribution $F_k(p)$ are consistent with $R_k(\gamma)$ and $\Psi_k(v)$.

Moreover, product-level consumption aggregates Q_k must be consistent with the household-level consumption aggregate Q.

The equilibrium is solved as a fixed-point problem. We start by guessing Q in an outer loop, then solve the equilibrium for each product – equivalently, each π_k – in an inner loop. Given a guess for Q_k, $N_k(p)$, and the constant component of the market demand,

$$D_k^0\left(p; N_k^0, Q_k^0\right) = (\tau p)^{-\eta} Q^{\frac{\eta}{\theta}-\eta\sigma} \left(Q_k^0\right)^{1-\frac{\eta}{\theta}} N_k^0(p),$$

we solve the retailers' problem to get the pricing strategies, $\Psi_k^0(v; D_k^0)$, and the associated relative posted-price distribution, $F_k^0(p; D_k^0)$. Given F_k^0, we compute the search strategies of shoppers and aggregate them to update the extensive margin of demand $N_k^1(p; F_k^0, Q_k^0)$. Using both F_k^0 and N_k^1, we update the product-level aggregate Q_k^1 and obtain $D_k^1(p; N_k^1, Q_k^1)$. We stop iterating the inner loop when $D_k^1(p) \approx D_k^0(p)$ for every p. After finding an equilibrium for each product k, we update Q by aggregating every Q_k, and repeat this process until Q converges.

In equilibrium, the total number of transactions equates the total number of buyers:

$$\int_{\underline{p}_k}^{\bar{p}_k} N_k(p)\, dF_k(p) = 1.$$

The posted-price distribution weighted by the number of transactions, $N_k dF_k$, can be interpreted as the distribution of *transaction* prices – the distribution of prices paid by shoppers. This distribution, which assigns a higher probability than F_k to lower prices, is the one we use to express the product-level consumption aggregate in terms of prices:

$$Q_k = Q^{1-\theta\sigma}\left(\int_0^1 \mathbb{E}_{j,k}(\tau p)^{1-\eta}\, dj\right)^{-\frac{1}{1-\eta}} = Q^{1-\theta\sigma}\left(\int_{\underline{p}_k}^{\bar{p}_k} (\tau p)^{1-\eta} N_k(p)\, dF_k(p)\right)^{-\frac{\theta}{1-\eta}}.$$

4 Inflation and price dispersion in the theory

The model in the previous section produces a ϒ-shaped relationship between price dispersion of identical goods and product-level inflation. A parametrization achieving this result has a search cost distribution g where $\underline{\gamma} = 0$, $g(0) = 0$, and $\lim_{\gamma \to 0} g'(\gamma) = +\infty$, meaning that, although search is costly for everyone, the density of shoppers with a positive but negligible marginal search cost is relatively large.[31] To understand the mechanism, starting to analyze the equilibrium at zero inflation is useful.

Zero inflation Assume a trivial starting condition for our algorithm such as $N_k^0(p) = 1\ \forall p$. Without inflation, the problem of the retailer reduces to picking a price that will stay fixed while idiosyncratic productivity changes. A retailer with initial productivity v sets the nominal (and relative) price $\hat{p}_k(v)$, and adjusts to $\hat{p}_k(v')$ whenever future

[31] Figure 6c plots such a search cost distribution. Table A.7 shows the exact parameters used in this section.

productivity v' is too high or too low.[32] Therefore, the support of the relative price distribution is determined by the optimal prices conditional on adjustment: $\underline{p}_k = \hat{p}_k(\underline{v})$ and $\overline{p}_k = \hat{p}_k(\overline{v})$.

Let $\check{\phi}(v)$ denote the stationary distribution of productivity implied by the bounded mean-reverting process (2). With $\check{\phi}(\underline{v}) > 0$, the probability of a retailer charging the lowest price is strictly positive. Optimal search rules (equation (7)) and $F_k(\underline{p}_k) > 0$ imply shoppers in the *interval* $(0, \Gamma_k(\underline{p}_k + \epsilon)]$ will search until finding $p \leq \underline{r}_k = \underline{p}_k + \epsilon$, as seen in Figure 6b. Under the assumption that $\lim_{\gamma \to 0} g'(\gamma) = +\infty$, $g(\Gamma_k(\underline{r}_k)) > 0$: a retailer setting $p \leq \underline{r}_k$ will attract a strictly positive share of shoppers. From equation (8), it is clear that the slope of the updated transaction function, $N_k^1(p)$, changes discretely at $p = \underline{r}_k$, generating a sharp kink in the demand at this price.

In an equilibrium with a sharply kinked demand, as Figure 4 shows, $\underline{r}_k$ maximizes static profits of retailers with higher productivity draws. If they charge a price slightly lower than $\underline{r}_k$, they will not attract more customers because they are already serving the maximum possible, $N_k(\underline{r}_k)$. On the other hand, if they charge a price slightly higher than $\underline{r}_k$, low-search-cost shoppers will reject the offer and search: the extensive margin of demand activates and demand elasticity jumps.

Moreover, pricing at the kink is not only optimal but also feasible for those retailers: relative prices are fixed, because zero inflation implies that relative and nominal prices are equal. Therefore, the more productive retailers bunch at $\underline{r}_k$, which generates a point mass in the posted-price distribution as in Figure 5a. This point mass expands the interval of buyer types searching for $p \leq \underline{r}_k$, sustaining an equilibrium demand function with a sharp kink at $\underline{r}_k$. Since high-search-cost shoppers accept prices greater than $\underline{r}_k$, retailers with low productivity draws can set higher prices, explaining the relatively low but positive price dispersion in an equilibrium with zero inflation.[33]

Non-zero inflation To illustrate the role of search in explaining our main result, suppose inflation increases by a very small amount while the demand function stays fixed at the zero-inflation one. For a positive level of inflation, relative prices are continuously drifting downward. To stay at the kink, retailers would need to pay the

[32] That is, whenever productivity reaches the upper or lower boundaries, $p_{U,k}^{-1}(\hat{p}_k(v))$ and $p_{L,k}^{-1}(\hat{p}_k(v))$, respectively.

[33] The kinked demand complicates the analytical proof of existence and uniqueness of the equilibrium at $\pi_k = 0$ considerably. If an equilibrium exists, to evaluate whether it is unique, I solve the model for different guesses of the transaction function, $N_k^0(p)$. For the parameters in the present paper, I find convergence to the same equilibrium always exists.

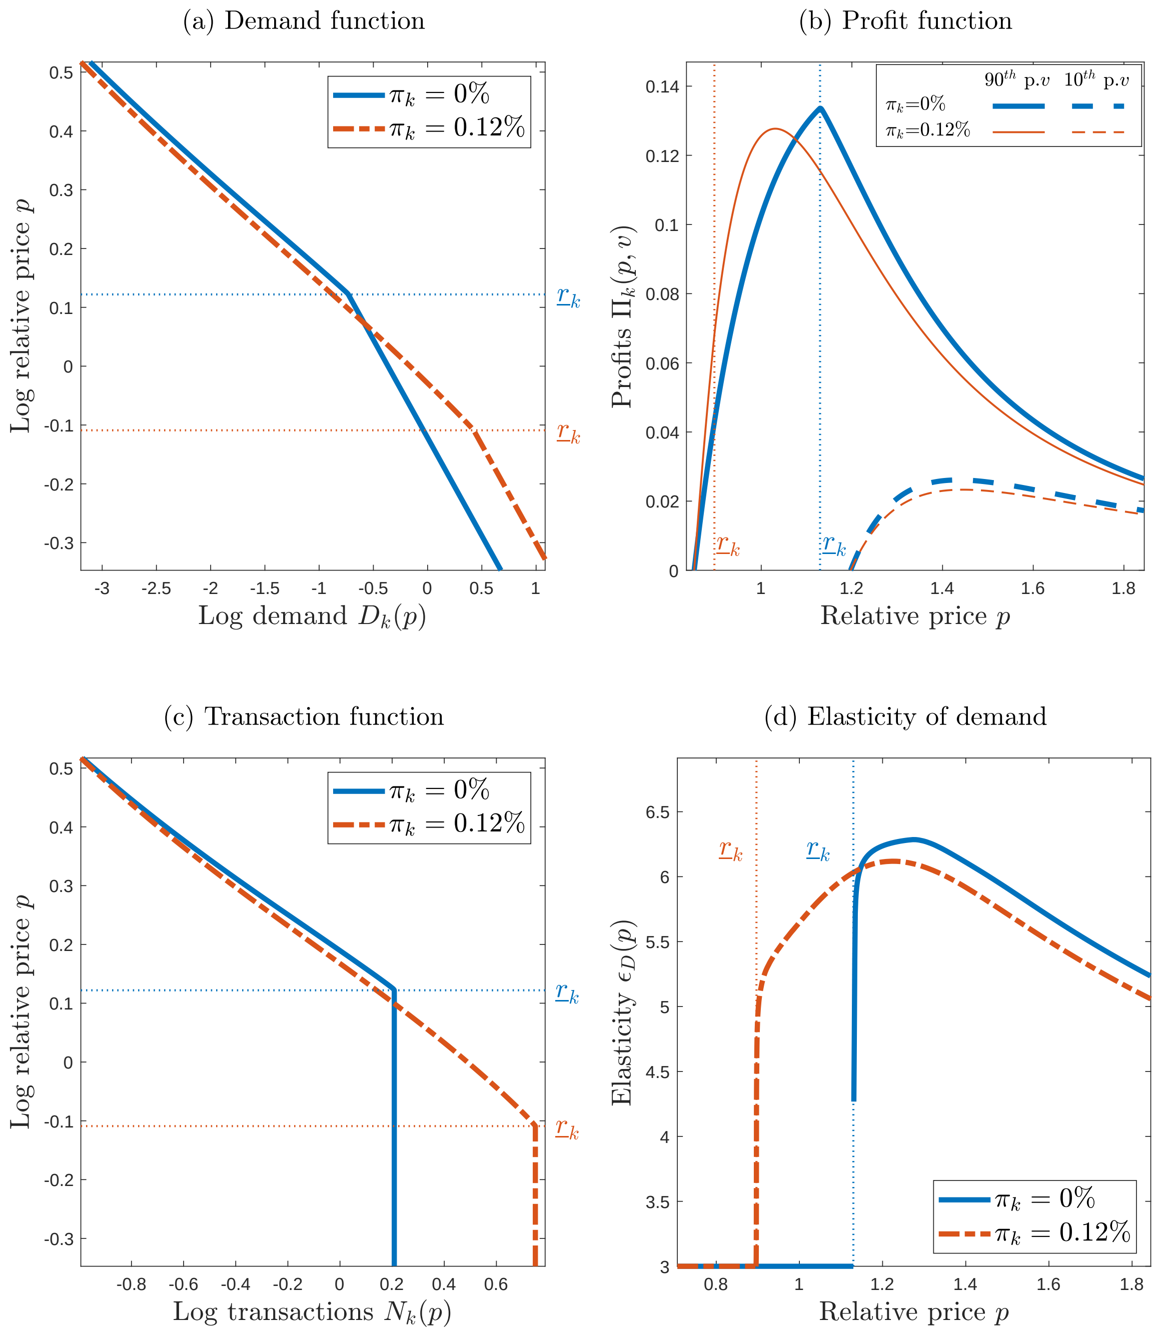

Figure 4: Retailer-level equilibrium functions at $\pi_k = 0\%$ and $\pi_k = 0.12\%$

The figures describe the retailer-level equilibrium functions at $\pi_k = 0\%$ and $\pi_k = 0.12\%$ (annualized). The value $\underline{r}_k$ indicates the lowest reservation price among shoppers. The upper-right panel shows the profit function for retailers with 90th and 10th percentile productivity draws.

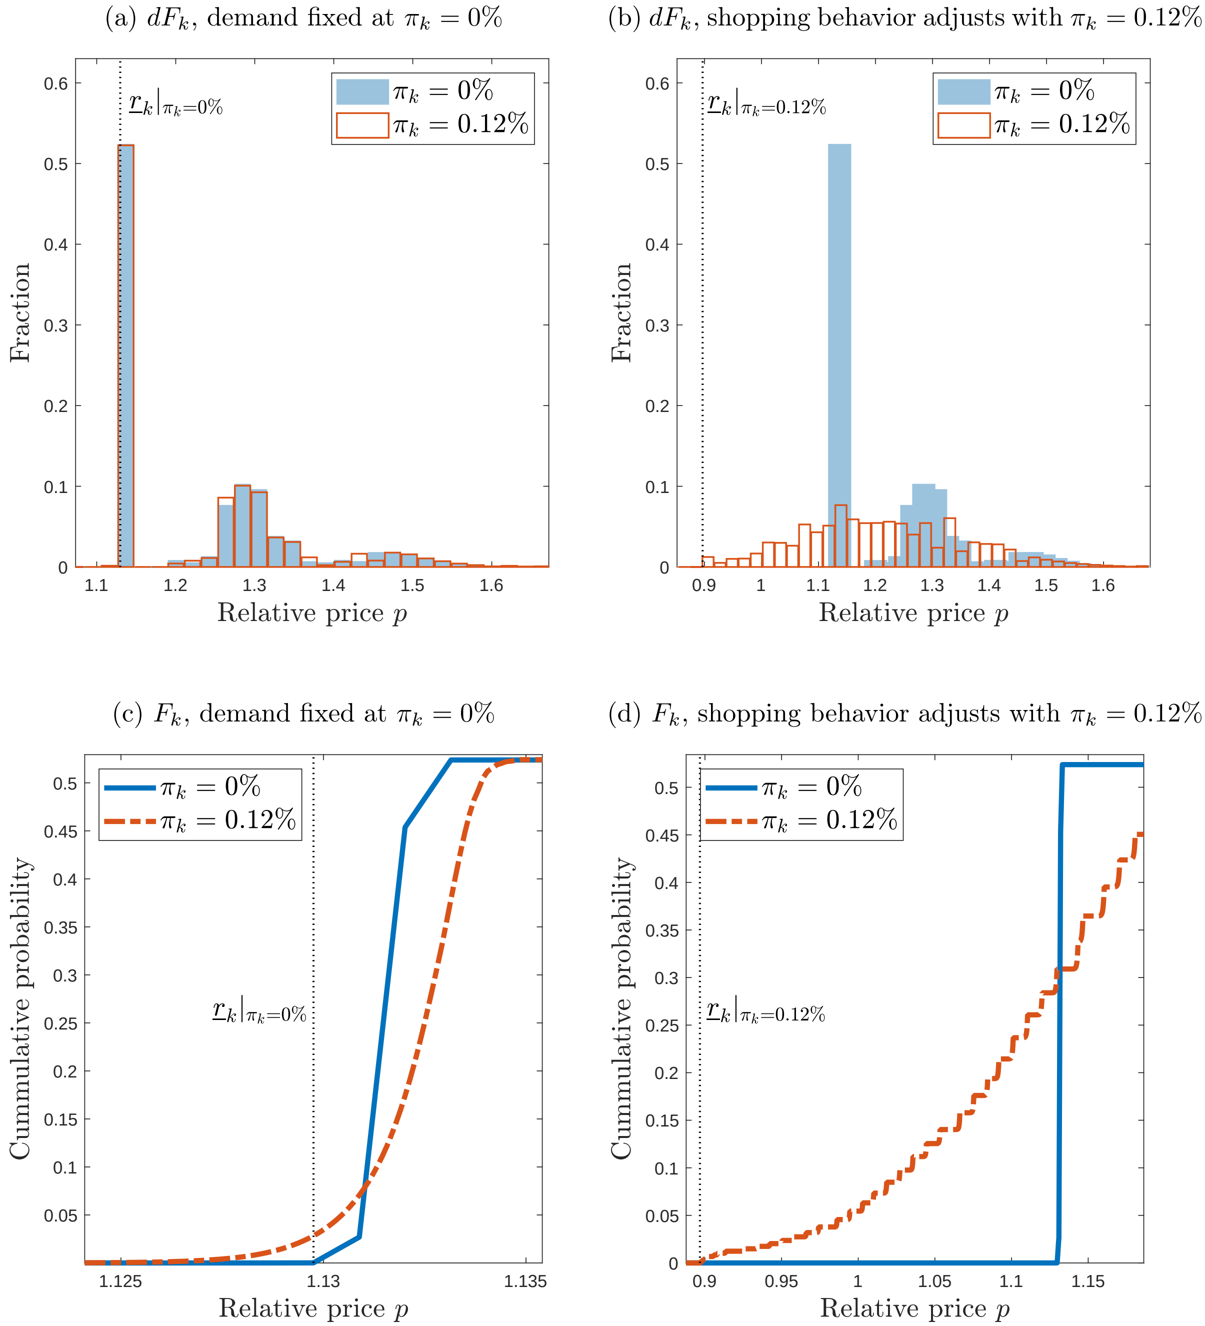

Figure 5: Equilibrium price distributions at $\pi_k = 0\%$ and $\pi_k = 0.12\%$

The figures show the posted-price distribution for different levels of annualized product-level inflation, π_k. The top figures show the probability dF_k, while the bottom ones show the cumulative probability F_k truncated just above 0.5. The left panels show the posted-price distribution at $\pi_k = 0\%$ and at $\pi_k = 0.12\%$ keeping the demand function fixed at $\pi_k = 0\%$. The value $\underline{r}_k$ indicates the lowest reservation price among shoppers.

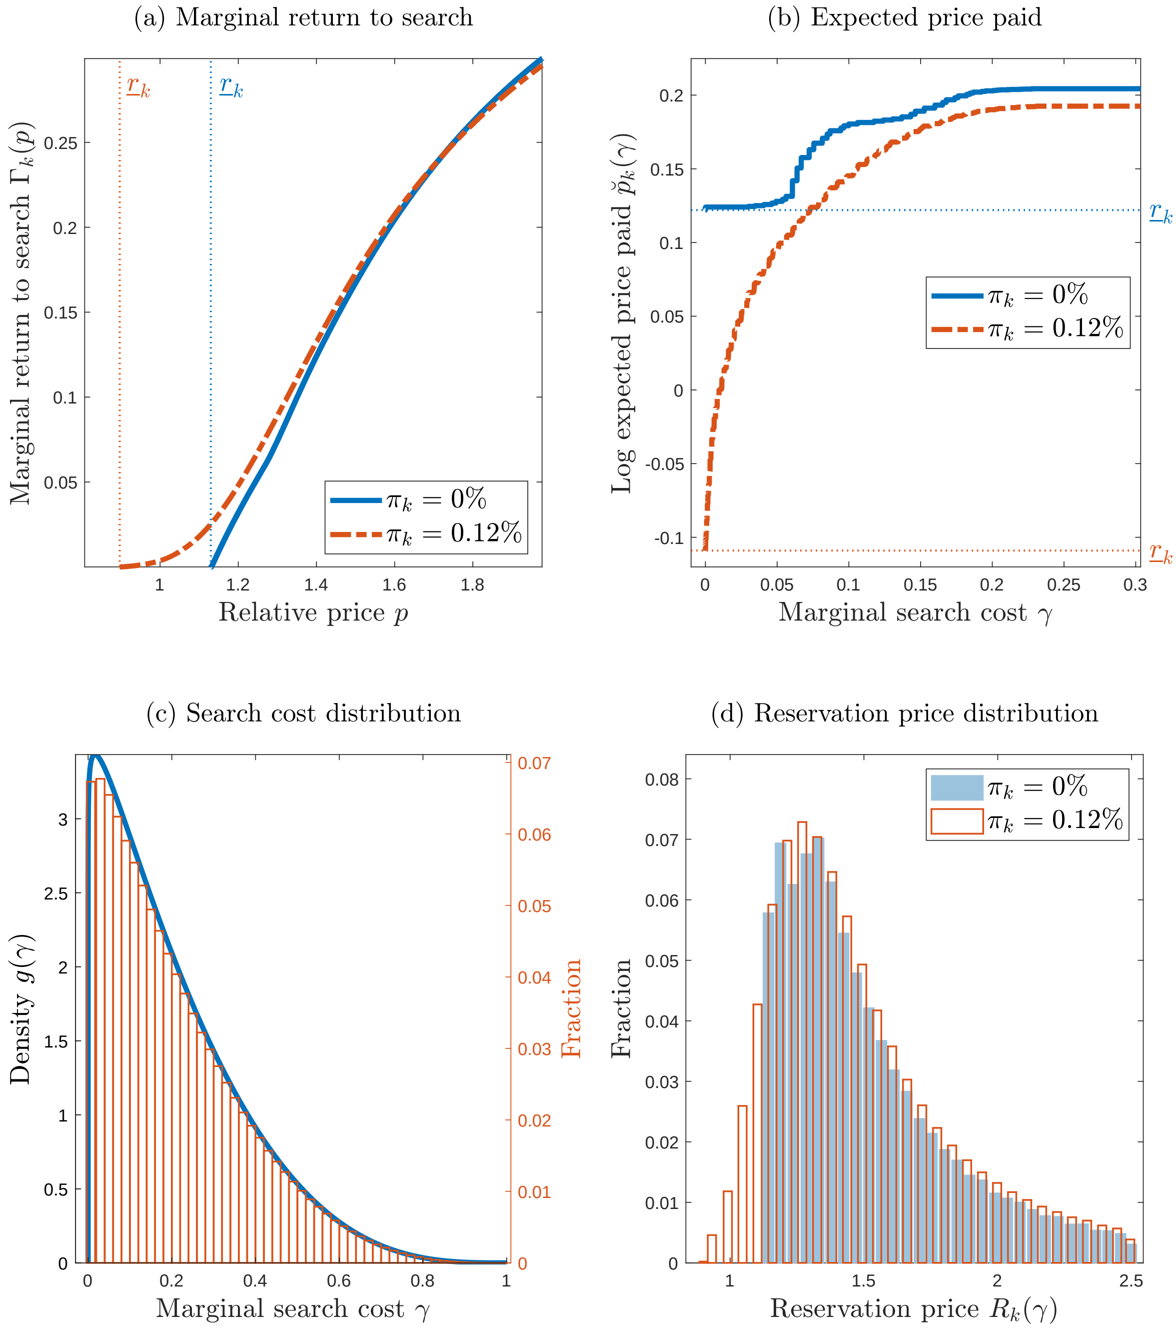

Figure 6: Equilibrium search at $\pi_k = 0\%$ and $\pi_k = 0.12\%$

The lower-right panel shows the search cost distribution $g(\gamma)$, where $\gamma \sim Beta(g_a, g_b)$, $g_a = 1.05$, $g_b = 4$, and $\gamma \in [0, 1]$. The other panels show equilibrium search strategies for $\pi_k = 0\%$ and $\pi_k = 0.12\%$ annualized, with the x-axis truncated above $\overline{p}_k$ or $\Gamma_k(\overline{p}_k)$ (upper-left plot) for clarity. The expected relative price paid conditional on the type is defined as $\breve{p}_k(\gamma) \equiv \mathbb{E}_k[p|p \leq R_k(\gamma)]$. The value $\underline{r}_k$ indicates the lowest reservation price among shoppers.

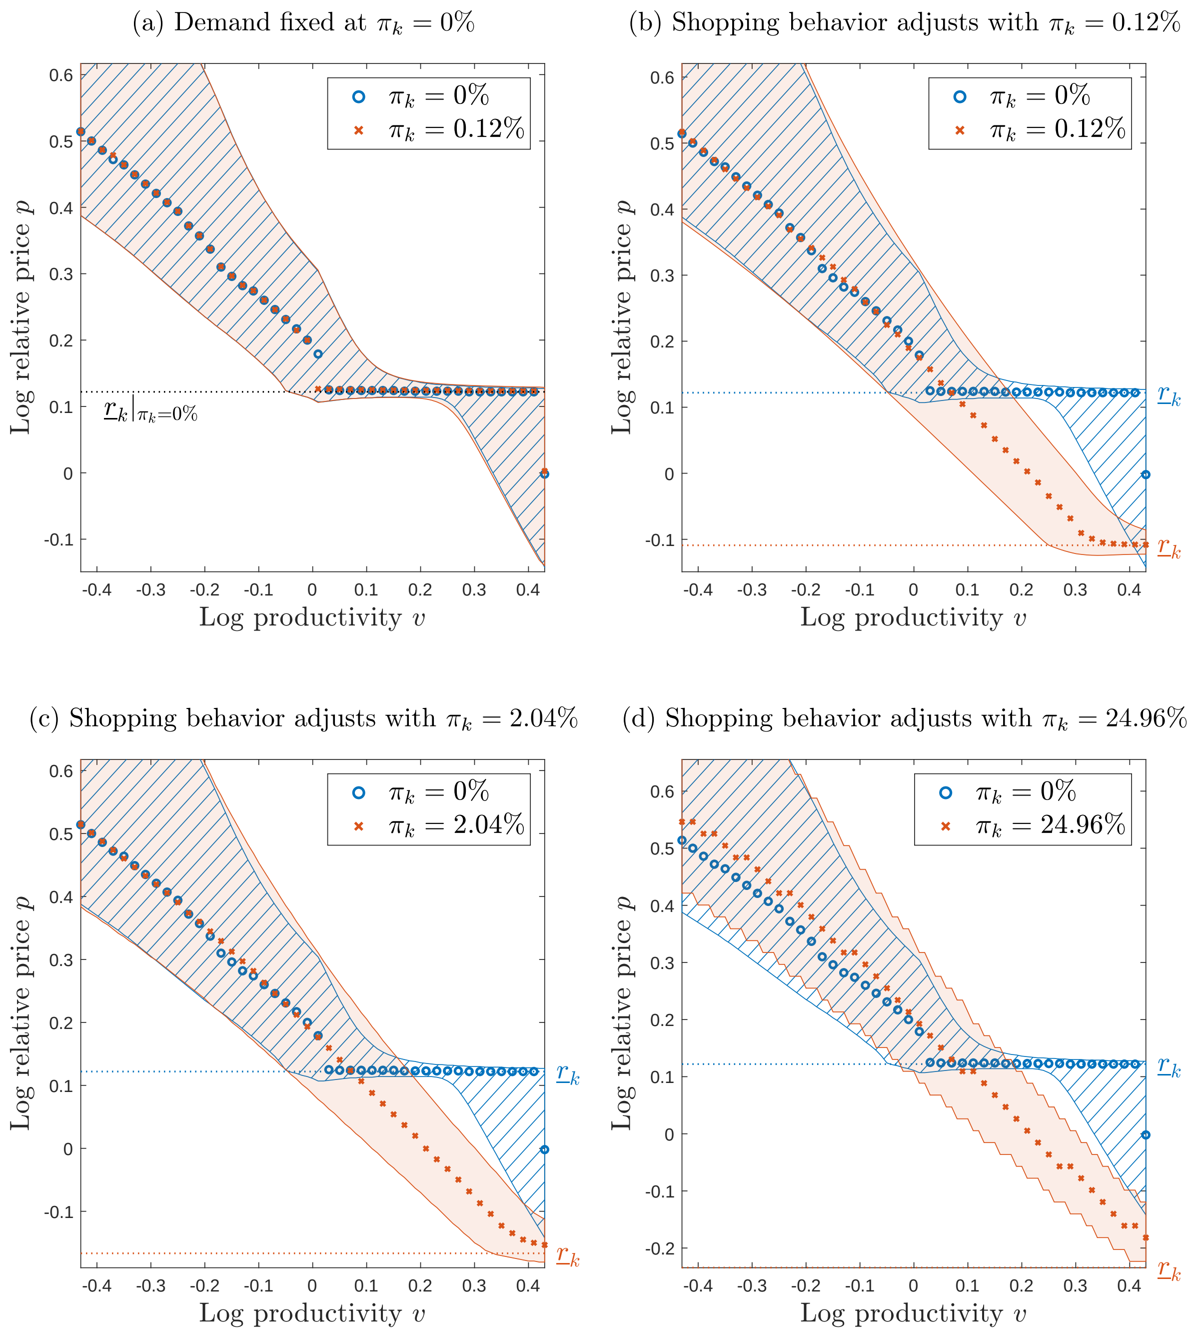

Figure 7: Optimal pricing strategies

The figures show the pricing bounds for different levels of annualized product-level inflation, π_k. The solid lines represent the upper and lower bounds $p_{U,k}(v)$ and $p_{L,k}(v)$; the dots and crosses, $\hat{p}_k(v)$. The upper-right panel shows the optimal pricing strategies at $\pi_k = 0\%$ and at $\pi_k = 0.12\%$ keeping the demand function fixed at $\pi_k = 0\%$. The other panels show equilibrium pricing strategies for different levels of π_k when shopping behavior adjusts. The value $\underline{r}_k$ indicates the lowest reservation price among shoppers.

menu cost every period: bunching at any price is not optimal. Thus, retailers let their relative prices erode to save on menu costs, generating dispersion around $\hat{p}_k(v)$ and decreasing the minimum market price to $\underline{p}_k = p_{L,k}(\underline{v})$. Nevertheless, as Figure 5c shows, the probability of a retailer charging the lowest price tends to zero: because of the relatively large idiosyncratic productivity shocks, most retailers would have already adjusted before reaching the lower bound of the inaction region. For the same reason, and as in standard menu-cost models, the price distribution (thus price dispersion) stays virtually constant around zero inflation (see Figure 5a).

Now, we let shoppers adjust their behavior to these inflation-induced changes in the relative price distribution. A lower minimum market price and higher dispersion at the bottom of the price distribution increase the returns to search of low-search-cost shoppers. However, only those at the very bottom of the search cost distribution set the lowest reservation price: as $\underline{p}_k$ becomes harder to find (i.e., $F_k(\underline{p}_k) \to 0$), only buyers with $\gamma = \Gamma_k(\underline{p}_k + \epsilon) \to 0$ will search until finding $p \leq \underline{r}_k = \underline{p}_k + \epsilon$ (Figure 6a). As dispersion at the left tail of the reservation price distribution increases (Figure 6d), the kink in demand smooths out.[34] Therefore, the bunching behavior is no longer optimal for more productive retailers: they can attract shoppers by decreasing their prices, increasing dispersion at low prices further (Figures 5b and 5d) and generating the jump in price dispersion around $\pi_k = 0$ we observe in the data. The retailer-level equilibrium functions at a small positive level of inflation are represented in Figure 4.[35]

In sum, the feedback between inflation and shopping behavior is key to reproduce the ϒ-shaped relationship between price dispersion of identical goods and product-level inflation. When inflation equals zero, price dispersion – especially at the lowest prices – is relatively low: its only source is idiosyncratic productivity shocks since retailers' relative prices stay fixed. Given that the lowest prices are relatively easy to find, a large share of low-search-cost shoppers will be looking for these prices. Therefore, more productive retailers bunch around the lowest prices to attract those shoppers.

From zero to positive inflation, price dispersion increases through two channels. The first is the menu-cost channel: relative prices drift downward continuously, so more productive retailers allow suboptimal price levels to pay the menu cost less often.

[34] From equation (8), we have that $N_k'(\underline{r}_k) \to 0$ if $\Gamma_k(\underline{r}_k) \to 0$.

[35] The mechanism is asymmetrical around zero inflation. When inflation is small but negative, relative prices drift upward, generating dispersion around $\hat{p}_k(v)$ and decreasing $\underline{p}_k$. Nevertheless, $F_k(\underline{p}_k) > 0$ since retailers reach the lower bound of the inaction region because of productivity shocks, not inflation as in the $\pi_k > 0$ case. Therefore, there is a sharp kink in the demand at $p = \underline{r}_k$ but such a price is below the optimum for most retailers, so price dispersion jumps but less than when $\pi_k > 0$.

Absent search, price dispersion barely changes since the region of inaction is mainly determined by idiosyncratic productivity shocks (Figure 7a). The second is the search channel: lower prices and higher dispersion at the bottom of the price distribution increase the returns to search, and to attract more shoppers, more productive retailers charge even lower prices. Figure 7b shows that, through the search channel, inflation changes the optimal price of these retailers considerably, increasing price dispersion discretely around $\pi_k = 0$. As inflation departs from zero, price dispersion increases smoothly: the search channel becomes less relevant since only shoppers at the very bottom of the search cost distribution change their behavior (Figures 7c and 7d).

5 Evidence on price dispersion and search

The model predicts that, within a market, buyers who visit more sellers pay lower prices. Moreover, the expected price reduction from visiting an additional seller is more significant the larger the absolute inflation.

To test this prediction and provide supporting evidence for the theory, I use the NielsenIQ Consumer Panel Data. These data contain the barcodes of the items that households purchased on each shopping trip, the quantities they bought, the prices they paid, and whether they used coupons to pay. Although information about the physical store is limited, every transaction can be associated to a retail chain. Therefore, we observe from how many different retailers a household purchased a particular product in a given time period. I assume the number of visits to distinct retailers is positively correlated with the number of prices observed, where the latter is inversely correlated with the search cost of shoppers in the model.

Households are sampled from 53 geographically dispersed Scantrack markets (each roughly corresponding to an MSA), and detailed demographic information about them is available.[36] The product categories are the same as in the Retail Scanner Data.

A unique feature of NielsenIQ's datasets is that they can be merged through product categories and geographic markets, so household shopping patterns can be linked to product $\times$ market variables such as inflation and price dispersion. In this way, we can study how shopping behavior interacts with inflation to affect the prices buyers pay for a given good while controlling for several sources of heterogeneity.

[36] NielsenIQ provides projection factors to make the sample demographically representative of the Scantrack market population.

5.1 Variable construction

The unit of analysis is a household × department × quarter. For each household in the Consumer Panel Data, I compute the relative price paid for a basket of goods belonging to a department in a given quarter, and the number of distinct retailers visited to purchase such a basket in the same quarter. Then, I merge these data with inflation and price dispersion at the department-quarter level in the geographic market to which the household belongs.

I aggregate purchases and shopping trips at the department level for two reasons. First, in an average shopping trip to a particular retailer, households purchase 7.4 distinct items. Thus, households likely consider the price of a *bundle* of goods when deciding which store to visit. Working at the department level takes this element into account. Second, this level of aggregation maximizes the amount of variation in the variables of interest. At the household × category level, for instance, shoppers visit, on average, 1.1 retailers each quarter (CV of 0.20), whereas at the household × department level, they visit three retailers (CV of 0.70). Working at the household level produces an even larger variation in the number of retailers visited (8.2 on average, with a CV of 1.9) but limits the inflation variation required for the analysis.

Relative price paid For each product with barcode k in the market m and quarter t, we define the average price paid among households as

$$\bar{P}_{km,t} = \frac{\sum_{h=1}^{H_{km,t}} \sum_{l=1}^{L^h_{km,t}} P^{h,l}_{km,t} q^{h,l}_{km,t}}{\sum_{h=1}^{H_{km,t}} \sum_{l=1}^{L^h_{km,t}} q^{h,l}_{km,t}},$$

where $P^{h,l}_{km,t}$ is the price household h paid for $q^{h,l}_{km,t}$ units of good k in shopping trip l; $H_{km,t}$ denotes the number of households buying product k in market m and quarter t, and $L^h_{km,t}$ is total shopping trips by household h. Following Kaplan and Menzio (2015), I construct the relative price that a household pays for the goods in a department as total expenditure over hypothetical expenditure at market-average prices:

$$p^h_{dm,t} = \frac{X^h_{dm,t}}{\bar{X}^h_{dm,t}} = \frac{\sum_{k=1}^{D_{dm,t}} \sum_{l=1}^{L^h_{km,t}} P^{h,l}_{km,t} q^{h,l}_{km,t}}{\sum_{k=1}^{D_{dm,t}} \sum_{l=1}^{L^h_{km,t}} \bar{P}_{km,t} q^{h,l}_{km,t}},$$

where $D_{dm,t}$ is the total number of products in department d for each m and t.

Number of retailers visited The expenditure by households in a department d and quarter t, $X^h_{dm,t}$, can also be expressed in terms of shopping trips to distinct retailers:

$$X^h_{dm,t} = \sum_{i=1}^{N^h_{dm,t}} \sum_{l=1}^{L^{h,i}_{km,t}} X^{h,l}_{dm,t},$$

where $X^{h,l}_{dm,t}$ is expenditure in shopping trip l; $L^{h,i}_{km,t}$ denotes the total number of shopping trips to retailer i; $N^h_{dm,t}$ is the number of distinct retailers household h visited to purchase items in department d and quarter t.

Table A.6 shows descriptive statistics for the sample. The total number of household × department × quarter observations is 24,468,527, with 149,324 households, 10 departments, 56 quarters, and 53 geographic markets.

5.2 Estimation and results

Let $\bar{\pi}_{dm,t}$ denote department × market-level quarterly inflation, corresponding to a sales-weighted average of product × market-level quarterly inflation.[37] I study the effects of inflation through shopping behavior on the relative prices paid by households by estimating

$$\begin{aligned} \log p^h_{dm,t} = & \sum_{S=1}^{10} \beta_{0,S} \mathbf{1}_{\{N^h_{dm,t}=S\}} + \sum_{S=1}^{10} \beta_{1,S} \mathbf{1}_{\{N^h_{dm,t}=S\}} \times |\bar{\pi}_{dm,t}| \\ & + \mu' X^h_{dm,t} + a_{d,t} + b_{m,t} + c_{B,t} + \epsilon^h_{dm,t}. \end{aligned} \quad (12)$$

The indicator variables take the value of 1 when household h visits $S \in [1, 10]$ distinct stores in a given quarter to purchase items in department d; the coefficients $\{\beta_{0,S}\}_{S=1}^{10}$ correspond to average relative prices paid by households visiting S retailers when inflation is zero; more importantly, the coefficients $\{\beta_{1,S}\}_{S=1}^{10}$ indicate how and by how much absolute inflation affects the relative prices paid by households according to their shopping behavior. The theory predicts $\beta_{1,10} < 0$ and that the sequence $\{\beta_{1,S}\}_{S=1}^{10}$ is decreasing: when absolute inflation is high, shoppers searching more find the lowest prices.

As in the model, the unobservable search cost determines the number of retailers

[37] Quarterly inflation is computed as the average of the 12-month differences for the three months in the quarter.

visited and the price paid, so its omission might generate biased estimates of $\{\beta_{0,S}\}_{S=1}^{10}$ and $\{\beta_{1,S}\}_{S=1}^{10}$. In an effort to control for the search cost, I first divide households into bins according to demographic observables.[38] Under the assumption that households within a bin B share a similar search cost, I then include bin $\times$ quarter fixed effects, $c_{B,t}$, in the estimation.[39] In this way, we alleviate endogeneity concerns while exploiting variation across buyers at the same date, which is the relevant variation in the model.

Note that, because we identify the coefficients $\{\beta_{0,S}\}_{S=1}^{10}$ using within-bin variation, their interpretation changes: they indicate the effect that visiting S retailers has on prices paid by buyers with a similar search cost. If, for example, the marginal search cost is constant, the reservation prices will not change with S, and $\beta_{0,S} = \beta_{0,S'}$ $\forall S \neq S'$ since a given buyer expects to pay the same price for any number of draws. Therefore, these coefficients will inform about the marginal search cost structure, so we will use them to calibrate the model.

In addition, I include department $\times$ quarter and market $\times$ quarter fixed effects, so we identify $\{\beta_{1,S}\}_{S=1}^{10}$ by exploiting the cross-department inflation variation within each geographic market and quarter, taking into account unobservable department-level differences that might be changing over time. The vector $X^h_{dm,t}$ contains a set of shopping behavior controls (number of shopping trips and transactions, and fraction of transactions paid with coupons) that might also affect prices paid. The observations are weighted using NielsenIQ sampling weights, and standard errors are two-way clustered by household and product department $\times$ market $\times$ quarter combination.

Figure 8 shows the average relative prices paid by households depending on the number of distinct retailers they visit and the level of absolute inflation. The findings support the theory: a household that visits 10 stores when absolute inflation is 10% pays 1.0% less than when inflation is zero and 2.4% less than a household that visits only one store. Moreover, for shoppers with a similar search cost, the average relative price paid decreases with the number of visits when inflation is zero. Through the lens of the model, this evidence suggests that reservation prices, thus the marginal search cost, are decreasing in the number of visits.

[38] The variables are household size, household income bin, presence of children, household head age and education, employment and marital status, number of household heads, and race. For a given household, these characteristics can change over time.

[39] There are 49,600 households on average per quarter and 8,620 bins, so around 6 distinct households per bin and quarter.

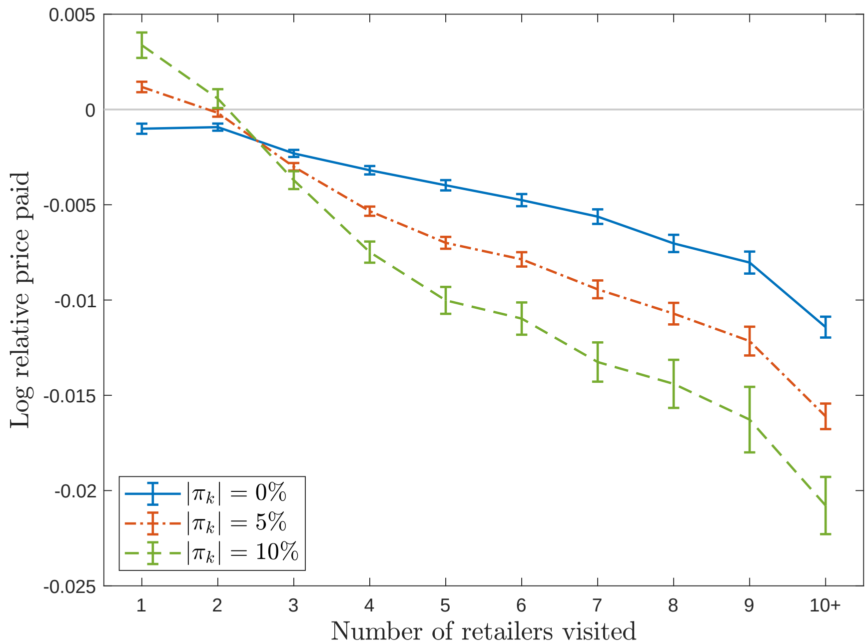

Figure 8: Shopping behavior and absolute inflation

The figure shows the relationship between log relative prices paid, number of retailers visited, and department × market-level quarterly absolute inflation as predicted by equation (12). The unit of observation is a household × product department × quarter, and the total number of observations is 24,468,527. Observations are weighted using NielsenIQ sampling weights. The confidence intervals (bars) use standard errors that are two-way clustered by household and product department × market × quarter combination.

6 Parametrization and calibration

The model in Section 3 has the necessary elements to reproduce the ϒ-shaped relationship between product-level inflation and price dispersion we observe in the data. For calibration, we extend the model in two dimensions. First, on the buyers' side, we allow the marginal search cost to vary monotonically with the number of draws, which now might be finite. With this extension, we can reproduce the patterns of average prices paid when inflation is zero as discussed in the previous section.[40] Second, on the retailers' side, we include random low-cost price adjustment opportunities (CalvoPlus model in Nakamura and Steinsson, 2010). By introducing some "Calvoness", we can generate the small price changes found in the literature (more on this below) and also

[40]The generalized search problem is described and solved in Section D.

show the results in Section 4 are not limited to a "pure" menu-cost model.

We assume the marginal search cost takes the form

$$c(S) = \alpha\left[1-(1-\chi_2)\left(\frac{S-2}{\overline{S}-2}\right)\right]^{\chi_1} \quad \forall S \in \left\{2,\ldots,\overline{S}\right\},$$

where $\bar{S}$ denotes the maximum number of draws, with $\chi_1 \geq 0$ and $\chi_2 \geq 0$. If $\chi_2 < 1$ ($\chi_2 > 1$), the marginal cost decreases (increases) with search so the reservation prices also decrease (increase) with search. The parameter χ_1 controls the curvature: $\chi_1 < 1$ ($\chi_1 > 1$) generates a concave (convex) search-cost structure. We choose a flexible search cost distribution capable of producing the Υ-shaped pattern, as explained in Section 4: $\gamma \sim Beta\left(g_a, g_b\right)$, where $g_a > 1$, $g_b > 1$, and $\gamma \in [0,1]$.[41]

For simplicity, we assume log utility in consumption ($\sigma = 1$), set $\alpha = 1$, and $\theta = \sigma^{-1}$ so the product-level equilibrium is independent from economy-wide aggregates. Following Nakamura and Steinsson (2010), we assume the low menu cost is 1/40 of the high menu cost and has an arrival rate equal to the frequency of regular price changes at $\pi_k = 2\%$. As in the previous literature, we set the monthly discount rate to $\rho = 0.96^{1/12}$. Finally, we set $\overline{S} = 40$, close to the maximum number of visits to retailers in the data.

We jointly calibrate the size of the high menu cost κ_H, the persistence ρ_v and volatility σ_v of idiosyncratic shocks, the shape parameters g_a and g_b of search-cost distribution, the remaining parameters of the marginal search-cost structure $\{\chi_1, \chi_2\}$, the price elasticity of demand η of shoppers, and the disutility cost of labor τ to match the following: (i) the product-level standard deviation of log regular prices and its Υ-shaped relationship with product-level inflation at all levels of π_k; (ii) the frequency of regular price changes and the average size of regular price changes at $\pi_k = 2\%$; (iii) an average markup of 30% at $\pi_k = 2\%$; (iv) a steady-state labor supply of 1/3 at $\pi_k = 2\%$; and (v) the estimated average prices paid at zero inflation $\{\hat{\beta}_{0,S}\}_{S=1}^{10}$.[42]

To produce the model counterpart of $\{\hat{\beta}_{0,S}\}_{S=1}^{10}$, I take the equilibrium posted-price

[41] Under these assumptions, the density of shoppers at $\gamma = 0$ is zero.

[42] The regular prices for the statistics in this section result from removing V-shaped patterns where the price returns to a regular price within 8 weeks after falling (what I call "Regular price A" in Section B). Table A.2 describes the corresponding data sample; Figures A.16-A.18 plot the relationship between different pricing statistics and product-level inflation controlling non-parametrically for the total number of stores. Since in the model we assume a continuum of sellers offer the product, I target the moments when there are at least 50 stores. I explain how I compute additional pricing moments in appendix C.

Table 1: Additional moments, model fit

Moment	Model	Data	Target
Monthly frequency of regular price changes at $\pi_k = 2\%$	21.37%	22.00%	Yes
Average abs. size of regular price changes at $\pi_k = 2\%$	9.70%	9.80%	Yes
Average markup at $\pi_k = 2\%$	1.31	1.30	Yes
Steady-state labor supply at $\pi_k = 2\%$	0.329	0.333	Yes
Average fraction of regular price increases at $\pi_k = 2\%$	54.12%	58.00%	No

distribution and the optimal search policies of shoppers j for each π_k, and simulate the search stage. Using the relative prices paid – constructed as in the data – and the number of sellers visited, I estimate the equivalent of equation (12) in the model:

$$\log p_k^j = \sum_{S=1}^{10} \beta_{0,S} \mathbf{1}_{\{N_k^j = S\}} + \sum_{S=1}^{10} \beta_{1,S} \mathbf{1}_{\{N_k^j = S\}} \times |\pi_k| + c_B + \epsilon_k^j.$$

A department $\times$ market pair in equation (12) is identified exclusively by its level of product-level inflation in the model, so I replicate the empirical cross-sectional distribution of product-level inflation when I simulate search outcomes. The fixed effects c_B include 100 equally sized search cost bins to control for shoppers' heterogeneity.

Figure 9 shows the model closely fits the new empirical findings presented in this paper. At the same time, the calibration produces pricing moments and average markups consistent with the empirical evidence and the previous literature, as seen in Table 1. It is noteworthy that, although the absolute size and frequency of regular price changes is only targeted at $\pi_k = 2\%$, the model matches their relationship with inflation (Figure A.19). Moreover, as Figure 9b shows, the calibrated model replicates the untargeted $\{\hat{\beta}_{1,S}\}_{S=1}^{10}$, corresponding to the effect that search has on prices paid as product-level inflation increases. The calibrated parameters are in Table 2.

6.1 Comparison with the literature

The calibrated model shows the theory in this paper explains the Υ-shaped relationship between inflation and price dispersion in the data. At the same time, it is consistent with two additional pricing facts for low inflation levels. The first is that, on average, around half of the price changes are decreases (54% in the model, 58% in the data). The second is that the absolute size of regular price changes is significantly larger than

(a) Price dispersion

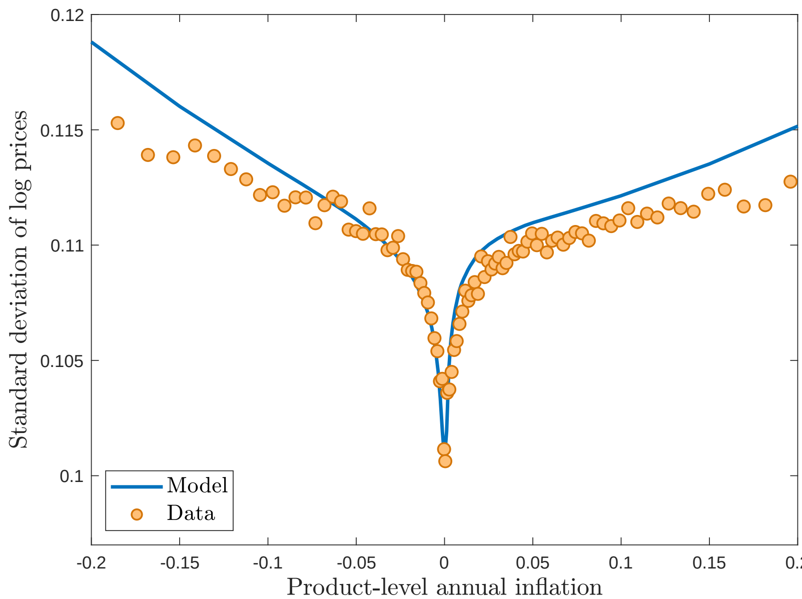

(b) Shopping behavior and prices paid

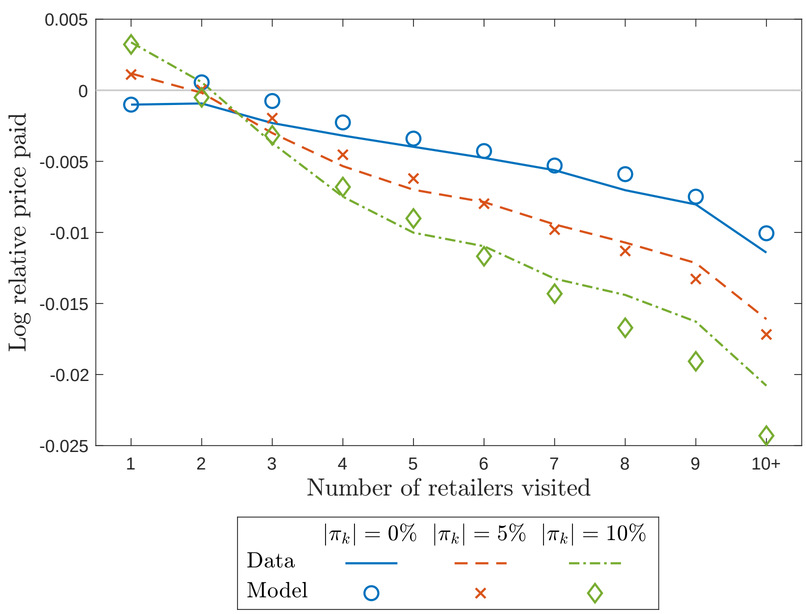

Figure 9: Price dispersion, shopping behavior, and inflation - model fit

The top figure shows the relationship between product-level inflation and price dispersion in the calibrated model (line) and in the data (dots); the bottom figure shows the empirical relationship between log relative prices paid and number of retailers visited for different levels of absolute inflation (lines) and its model counterpart (markers). Price dispersion is computed for regular prices of products sold by at least 50 stores.

Table 2: Calibrated parameter values

Retailers	
Size of the high menu cost	$\kappa_H = 0.032$
Mean-reversion rate of idiosyncratic shocks	$\rho_v = 0.870$
Volatility of idiosyncratic shocks	$\sigma_v = 0.063$
Shoppers	
Curvature of the marginal search-cost structure	$\chi_1 = 1$
Direction of the marginal search-cost structure	$\chi_2 = 0.100$
Shape of search cost distribution, $Beta\,(g_a, g_b)$	$g_a = 1.164$
Shape of search cost distribution, $Beta\,(g_a, g_b)$	$g_b = 2.000$
Price elasticity of demand of shoppers	$\eta = 3.000$
Disutility cost of labor	$\tau = 2.350$

average inflation (10% versus 2%). I argue that the alternative models in the literature cannot replicate these three pricing facts at the same time.

To develop the argument, it is convenient to note that the model in this paper nests others in the literature: (i) if $\sigma \to 0$ and $\rho \to 1$, we mute the idiosyncratic shocks; (ii) if we set the probability of the low menu cost equal to zero, we get pure menu cost adjustment; and (iii) if $g_a \to 1$ and $g_b \to \infty$ – so $\underline{\gamma} \to 1$ – we shut down search. Under assumptions (i) and (ii), the model converges to a version of Benabou (1992) without entry and with a general search cost distribution and marginal search cost structure; under (ii) and (iii), it converges to Golosov and Lucas (2007); and under (iii) and different values of the high menu cost κ_H, it gives a model with Calvo ($\kappa_H \to \infty$) or CalvoPlus ($\kappa_H < \infty$) adjustment.

Sticky-price models with idiosyncratic shocks but without search (e.g., Golosov and Lucas 2007; Nakamura and Steinsson 2010) match the additional pricing facts but miss the Υ-shaped relationship between inflation and price dispersion. Moreover, as Figure 10 shows, these models predict both price dispersion and the absolute size of the price changes are flat around zero inflation. In the model with search – as in the evidence in this paper – both moments are Υ-shaped around zero inflation.[43]

To understand the intuition behind these results, decomposing price dispersion as the sum of two components is helpful:

$$\mathbb{V}\text{ar}_k\,(p;\pi_k) = \mathbb{E}_k\left[\mathbb{V}\text{ar}_k\,(p\,|v\,;\pi_k)\right] + \mathbb{V}\text{ar}_k\left[\mathbb{E}_k\,(p\,|v\,;\pi_k)\right].$$

[43] The frequency of price changes displays a similar pattern, while the fraction of price increases is smooth around zero inflation, as Figure A.20 shows.

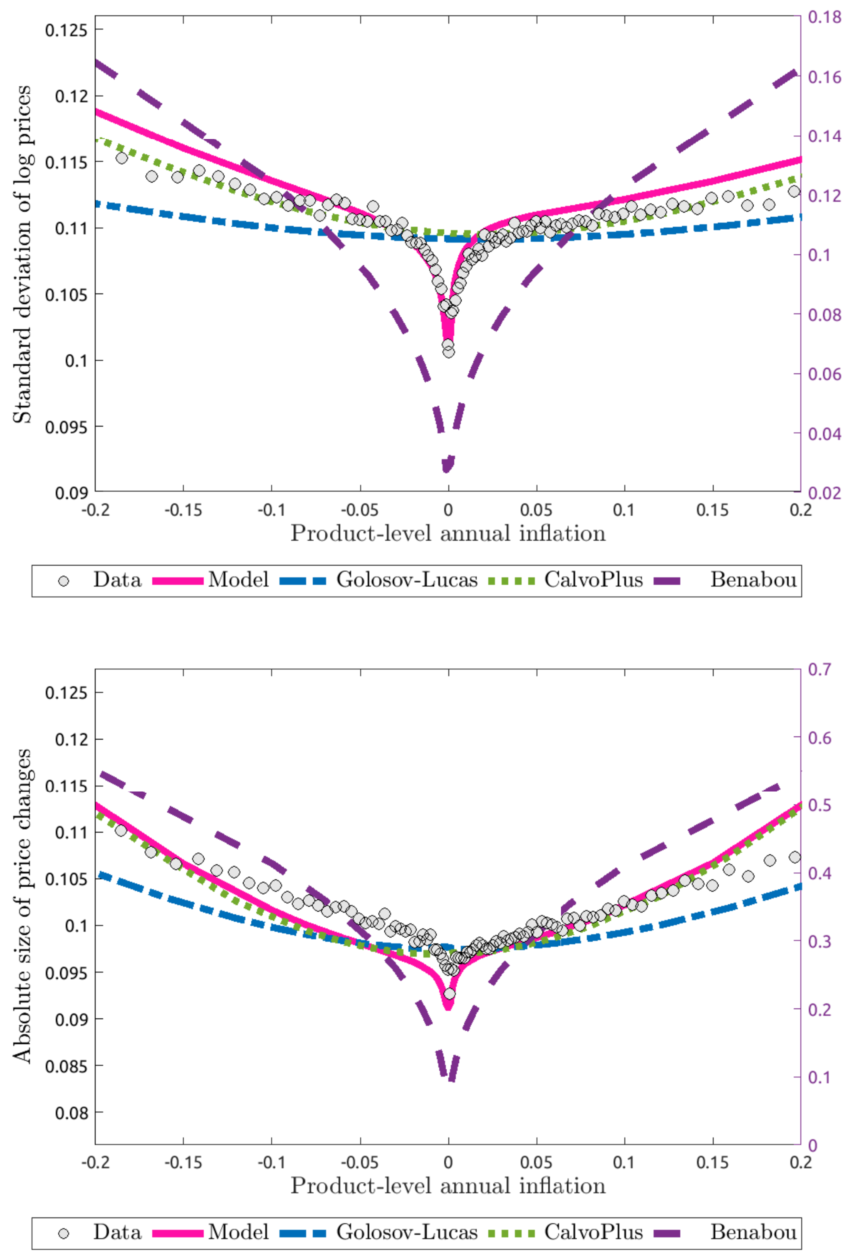

Figure 10: Pricing moments in different models

The figure plots pricing moments and inflation in the data, model, pure menu-cost model with idiosyncratic shocks but without search (Golosov and Lucas, 2007), CalvoPlus price-setting model with idiosyncratic shocks but without search, and pure menu-cost model with search but without idiosyncratic shocks (Benabou, 1992; right y-axis). The top panel shows product-level price dispersion measured by the standard deviation of log prices, and the bottom panel plots the absolute size of price changes. Pricing moments are computed for regular prices of products sold by at least 50 stores.

The first term on the right-hand side corresponds to average price dispersion conditional on productivity; the second, to variation of the average price across retailers with different levels of productivity. In menu-cost models with idiosyncratic shocks and no consumer search, the first term increases with inflation: retailers need to adjust more often as relative prices erode faster, but because adjustments are costly, they optimally choose wider inaction regions $[p_{L,k}(v), p_{U,k}(v)]$. Around zero inflation, optimal pricing behavior is mainly determined by the relatively large idiosyncratic shocks, so this term remains roughly unchanged. The second term is essentially exogenous because it approaches the cross-sectional dispersion of productivity (Alvarez et al., 2019). Thus, price dispersion is smooth at zero inflation.[44]

In the model with idiosyncratic shocks and search, *both* sources of price dispersion increase around zero inflation. On the one hand, as inflation goes from zero to a positive rate, the optimal price conditional on adjustment, $\hat{p}_k(v)$ decreases more for higher productivity draws, as shown in Figure 7. Thus, the variance of the conditional expected price increases. On the other hand, most productive retailers set a fixed relative price under $\pi_k = 0$. When inflation deviates from zero, their relative prices drift downward, so they optimally widen their adjustment bands: price dispersion conditional on productivity increases.

At the same time, because random low-cost adjustment opportunities let retailers incorporate small changes in relative prices due to inflation, the average absolute size of price changes increases discretely.[45] In this way, the current theory matches price dispersion and the absolute size of price changes Υ-shaped around $\pi_k = 0$.

Menu-cost models without idiosyncratic shocks, with search (Benabou, 1988, 1992) or without (Sheshinski and Weiss, 1977), match price dispersion Υ-shaped in inflation but miss other features of realistic price-setting behavior. In particular, for positive inflation, all price changes are price increases. Additionally, the only motive for adjustment is inflation, implying the absolute size of price changes is close to zero around zero inflation and increases steeply with inflation.

[44] Everything else equal, random low-cost adjustment opportunities (CalvoPlus) keep price dispersion smooth at $\pi_k = 0$ but make it increase faster with inflation since large idiosyncratic shocks are not required for retailers to adjust their prices.

[45] At $\pi_k = 0$, without random low-cost adjustment opportunities, only retailers with an optimum far from $\underline{r}_k$ adjust (i.e., low-productivity retailers). Since the adjustment bands are narrower at lower prices, price changes when $\pi_k = 0$ are relatively large. For low positive inflation, price changes from high-productivity retailers also occur, so the average absolute size of price changes tends to decrease.

7 Welfare implications of inflation

The theory in this paper matches the ϒ-shaped relationship between inflation and price dispersion, exhibits realistic pricing behavior, and is supported by empirical evidence on shopping behavior and inflation. What are the implications of this model for the costs and benefits of inflation?

Welfare per instant of time is given by equation (4). Using $\theta = \sigma = 1$, the consumer surplus is

$$\log Q - 1 = \int_0^1 \log Q_k dk - 1,$$

where the product-level consumption aggregate Q_k is a transaction-weighted average of the relative prices:

$$Q_k = \left(\int_{\underline{p}_k}^{\overline{p}_k} (\tau p)^{1-\eta} N_k(p)\, dF_k(p) \right)^{-\frac{1}{1-\eta}}.$$

Taking into account that the first price draw is free, the total cost of searching for better prices is

$$\bar{C} = \int_0^1 \bar{C}_k dk = \int_0^1 \int_{\underline{\gamma}}^{\overline{\gamma}} \breve{C}_k(\gamma) dG(\gamma)\; dk,$$

where $\breve{C}_k(\gamma)$ is the expected search cost of shopper type γ buying product k.[46] Moreover, real dividends are aggregate real profits net of aggregate adjustment costs:

$$\begin{aligned} \frac{\mathcal{D}}{W} &= \int_0^1 \left(\bar{\Pi}_k - \mathcal{A}_k \right) dk; \\ \bar{\Pi}_k &= \int_{\underline{p}_k}^{\overline{p}_k} \int_{\underline{v}}^{\overline{v}} \Pi_k(p, v)\, \phi_k(p, v)\, dv\, dp; \\ \mathcal{A}_k &= \int_0^1 \kappa \Lambda_k dk, \end{aligned}$$

where Λ_k is the fraction of retailers that reprice product k each dt. Then, social welfare is the average contribution of each product, which is identified by its inflation π_k:

$$\begin{aligned} U &= \int_0^1 U_k dk - 1; \\ U_k &= \log Q_k + \tau \bar{\Pi}_k - \tau \mathcal{A}_k - \alpha \bar{C}_k; \\ \Rightarrow \tilde{U}(\pi_k) &= \log \tilde{Q}(\pi_k) + \tau \tilde{\bar{\Pi}}(\pi_k) - \tau \tilde{\mathcal{A}}(\pi_k) - \alpha \tilde{\bar{C}}(\pi_k). \end{aligned} \tag{13}$$

[46] See appendix D for the expression and its derivation.

The first two terms on the right correspond to the product-level aggregate gains from trade, and the last two terms to the resources spent on market frictions (i.e., search and price adjustment). In the calibrated model, inflation increases price dispersion, and thus the returns to search for low prices. As search activity increases, the resource cost of search tends to increase. Price-adjustment costs are also increasing in inflation: the higher inflation is, the faster relative prices drift away from their optimal levels, requiring more frequent price adjustments.

The gains from trade reflect the allocative role of prices through search. If, with positive inflation, consumers search for better prices, markups of more productive retailers decrease, increasing efficiency. Thus, whether inflation improves welfare depends on the size of the efficiency gains – if any – from lower markups.

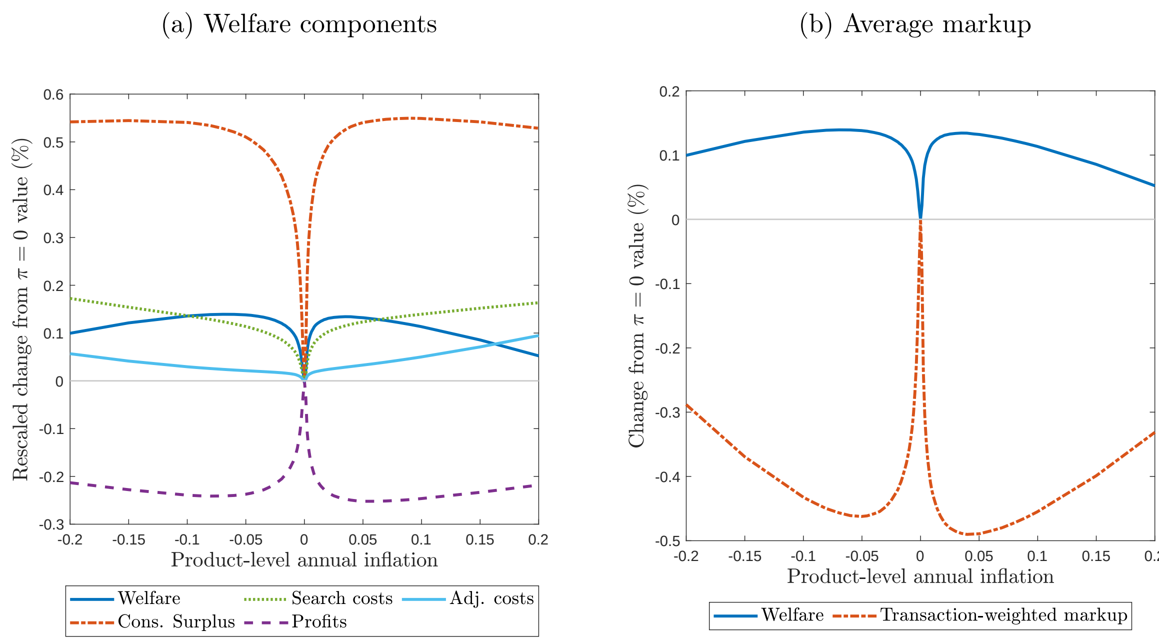

Figure 11: Product-level welfare and inflation

The figure in the left panel shows the relationship between product-level inflation and welfare, $\tilde{U}(\pi_k)$, and the contribution of each of its components in equation (13) (plotted as changes from zero inflation, times their welfare share at zero inflation). The figure in the right panel also plots the average transaction-weighted markup.

Figure 11 shows the relationship between product-level welfare and inflation, determined by equation (13), in the calibrated model. We see that, as inflation departs from zero, aggregate consumer surplus jumps. Higher consumer surplus offsets lower profits and the positive search and adjustment costs that come with inflation, generating a net welfare gain.

On the flip side, the average transaction-weighted markup,

$$\mu_k \quad = \quad \int_v \int_{\underline{p}_k}^{\overline{p}_k} pvN_k(p)\,\phi_k(p,v)\,dp\,dv,$$

decreases, reflecting the source of the efficiency gains. Furthermore, the efficiency gains from positive inflation are limited: welfare is maximized at a finite inflation rate. What determines such a rate?

When inflation is zero, production is shared relatively evenly among retailers. For a small positive inflation rate, a large part of the production shifts from less to more productive retailers: with increased search activity (Figure 12a), retailers with high productivity draws charge lower prices to attract more shoppers, and their market share increases (Figure 13a). As a result, market concentration increases – both in the model and in the data (Figure 13b) – but the average markup decreases since more productive retailers decrease their markups (Figure 12b).

Nevertheless, as inflation increases and more shoppers flee from higher prices, most of the customers of the least productive retailers become non-searchers, or captive shoppers. As Figure 12a shows, their search behavior does not change with inflation. Thus, for retailers with low productivity draws, charging higher prices to captive shoppers, raising their markups, is optimal. When the latter effect exceeds the former, efficiency gains from positive inflation dissipate, reducing welfare.

Role of decreasing marginal search cost Everything else equal, a decreasing versus a linear marginal search cost reduces the inflation rate that maximizes benefits. Going from a linear to a decreasing marginal search cost structure, the extensive margin of search remains almost constant (i.e., the cutoff of types who search, $\tilde{\gamma}_k$, is basically unchanged), but the intensive margin of search increases. Since types with $\gamma \to 0$ were already at or close to the maximum number of searches, $\overline{S}$, higher types will search proportionally more. Nevertheless, their extra search activity is relatively small in absolute terms, so their reservation prices (thus, the prices they pay) do not drop significantly. In sum, the surplus gains from lower prices paid by higher types are small relative to the higher search costs they incur.

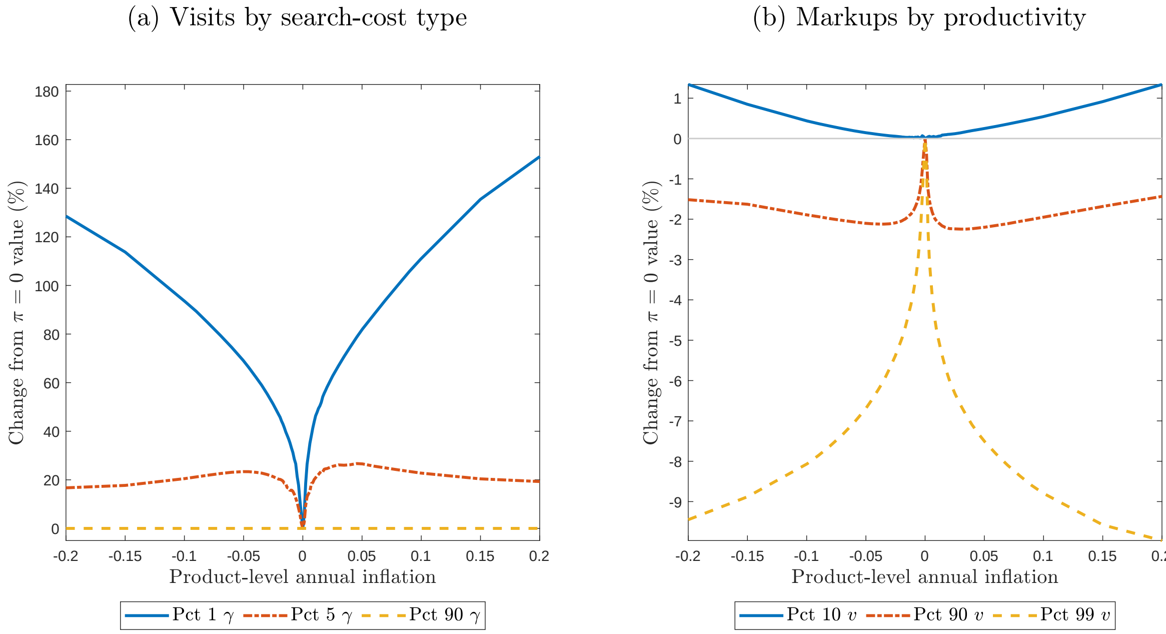

Figure 12: Heterogeneous effects of inflation

The left panel plots the number of visits as a function of inflation for shoppers in different percentiles of the search-cost distribution; the right panel plots average transaction-weighted markups as a function of inflation for retailers in different percentiles of the productivity distribution.

Aggregate inflation and welfare The product k price level is defined as the deflator of nominal product sales:

$$P_{k,t} = \frac{W_t \int_{\underline{p}_k}^{\overline{p}_k} p \left(\tau p\right)^{-\eta} N_k\left(p\right) dF_k\left(p\right)}{A_{k,t} \int_{\underline{p}_k}^{\overline{p}_k} \left(\tau p\right)^{-\eta} N_k\left(p\right) dF_k\left(p\right)}.$$

Therefore, we can verify that the product-level inflation rate is $\pi_k = \lambda_W - \lambda_{A,k}$. Assuming an economy-wide price index such as

$$P_t = \exp\left\{\int_0^1 \log P_{k,t} dk\right\}$$

yields the following expression for aggregate inflation as a function of product-level inflation rates:

$$\pi = \int_0^1 \pi_k dk = \lambda_W - \int_0^1 \lambda_{A,k} dk = \lambda_W - \bar{\lambda}_A.$$

(a) Number of transactions by productivity (b) Herfindahl–Hirschman Index

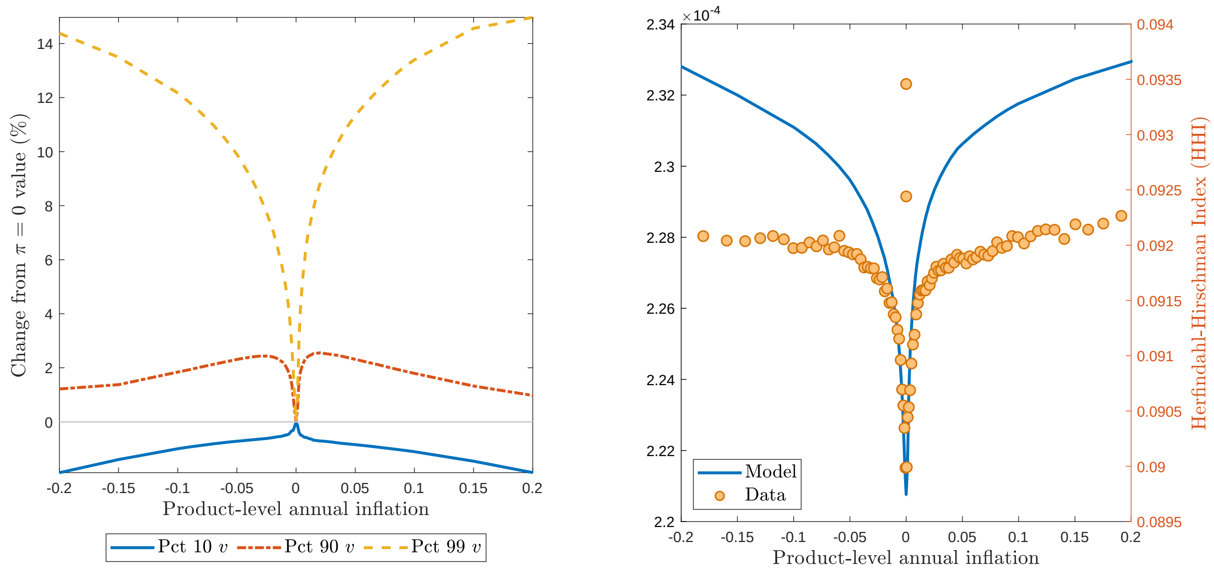

Figure 13: Product-market concentration and inflation

The left panel plots average number of transactions as a function of inflation for retailers in different percentiles of the productivity distribution; the right panel plots the Herfindahl–Hirschman Index (HHI) as a function of inflation in the model and in the data (right y-axis). The dots correspond to the average HHI for each of 100 equally sized disaggregate inflation bins as predicted by equation (1), when the dependent variable is the HHI.

Let H denote the sales-weighted distribution of cross-sectional inflation rates in the data (Figure A.2). Using that nominal sales grow 4% annually in the Retail Scanner Data, we can recover the real component of product-level inflation as $\lambda_{A,k} = \lambda_W - \pi_k$. By setting the distribution of $\lambda_{A,k}$ equal to H, we can study how social welfare changes with aggregate steady-state inflation according to the calibrated model and the associated function $\tilde{U}(\cdot)$:

$$U(\pi) = \int_{\underline{\lambda}}^{\bar{\lambda}} \tilde{U}\left(\pi + \bar{\lambda}_A - \lambda\right) dH(\lambda) - 1. \tag{14}$$

I assess the quantitative effects of a positive inflation rate in terms of consumption-equivalent welfare changes. That is, I compute the percentage change in consumption, Δ, needed to make households in economy dot ($\pi = 0$) equally well off as households in economy hat ($\pi > 0$):

$$\log\left(\dot{Q}\left(1+\Delta\right)\right) - \tau\dot{L} - \alpha\dot{\bar{C}} = \log\left(\hat{Q}\right) - \tau\hat{L} - \alpha\hat{\bar{C}}.$$

For each level of aggregate inflation, we can also compute an aggregate price dispersion measure using the equivalent of equation (14) to average product-level price dispersion.

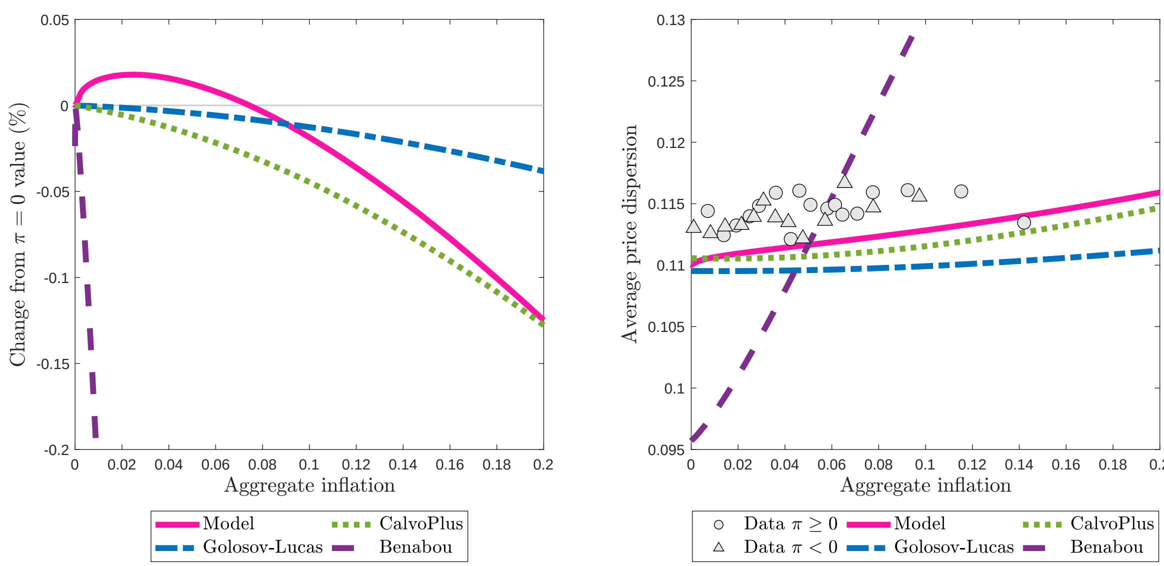

(a) Social welfare

(b) Average price dispersion

Figure 14: Aggregate inflation, price dispersion, and welfare

The figure plots average price dispersion, social welfare, and aggregate inflation in the model, pure menu-cost model with idiosyncratic shocks but without search (Golosov and Lucas, 2007), CalvoPlus price-setting model with idiosyncratic shocks but without search, and pure menu-cost model with search but without idiosyncratic shocks (Benabou, 1992). The left panel shows social welfare as average product-level welfare (equation (14)) in terms of consumption-equivalent welfare changes from $\pi = 0$. The right panel shows price dispersion as average product-level standard deviation of log prices. Both aggregates consider the sales-weighted distribution of π_k. The data for average price dispersion corresponds to 179 monthly observations divided into 30 aggregate inflation bins.

Figure 14a shows that, in the calibrated model with search, aggregate inflation of 2.53% generates the largest consumption-equivalent welfare gains (0.02%). Although the magnitude is seemingly low, the key takeaway is that this is the only one among the four models considered – that is, within this relatively simple framework – allowing a beneficial role for inflation. We would miss this feature by exclusively looking at the relationship between aggregate – instead of product-level – inflation and price dispersion, as Figure 14b suggests.

8 Conclusion

In this paper, I show the empirical relationship between product-level inflation and price dispersion is ϒ-shaped. Current sticky-price models cannot simultaneously account for this fact and other features of pricing behavior. I develop a menu-cost model with idiosyncratic shocks and endogenous consumer search that can. Furthermore, evidence on shopping behavior and inflation supports the theory.

In the model, the costs of inflation arise from two market frictions: price adjustment and search. If inflation carries benefits, as the calibrated model suggests, they stem from higher price dispersion and returns to search. As search activity increases, competition intensifies, decreasing markups. The positive welfare-maximizing inflation rate optimally trades off the efficiency gains from lower markups and the resources spent on search.

References

Alvarez, F., M. Beraja, M. Gonzalez-Rozada, and P. A. Neumeyer (2019). From hyperinflation to stable prices: Argentina's evidence on menu cost models. *The Quarterly Journal of Economics 134*(1), 451–505.

Barro, R. J. (1972). A theory of monopolistic price adjustment. *The Review of Economic Studies 39*(1), 17–26.

Beaulieu, J. and J. Mattey (1999). The effects of general inflation and idiosyncratic cost shocks on within-commodity price dispersion: evidence from microdata. *Review of Economics and Statistics 81*(2), 205–216.

Benabou, R. (1988). Search, price setting and inflation. *The Review of Economic Studies 55*(3), 353–376.

Benabou, R. (1992). Inflation and efficiency in search markets. *The Review of Economic Studies 59*(2), 299–329.

Bils, M. and P. J. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of political economy 112*(5), 947–985.

Blanco, A. (2021). Optimal inflation target in an economy with menu costs and a zero lower bound. *American Economic Journal: Macroeconomics 13*(3), 108–41.

Burdett, K. and K. L. Judd (1983). Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society*, 955–969.

Burdett, K. and G. Menzio (2017). The (q, s, s) pricing rule: A quantitative analysis. *Research in Economics 71*(4), 784–797.

Burdett, K. and G. Menzio (2018). The (q, s, s) pricing rule. *The Review of Economic Studies 85*(2), 892–928.

Burstein, A. and C. Hellwig (2008). Welfare costs of inflation in a menu cost model. *American Economic Review 98*(2), 438–43.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of monetary Economics 12*(3), 383–398.

Caplin, A. S. and D. F. Spulber (1987). Menu costs and the neutrality of money. *The Quarterly Journal of Economics 102*(4), 703–725.

Coibion, O., Y. Gorodnichenko, and G. H. Hong (2015). The cyclicality of sales, regular and effective prices: Business cycle and policy implications. *American Economic Review 105*(3), 993–1029.

Danziger, L. (1988). Costs of price adjustment and the welfare economics of inflation and disinflation. *The American Economic Review*, 633–646.

Debelle, G. and O. Lamont (1997). Relative price variability and inflation: evidence from us cities. *Journal of Political Economy 105*(1), 132–152.

DellaVigna, S. and M. Gentzkow (2019). Uniform pricing in us retail chains. *The Quarterly Journal of Economics 134*(4), 2011–2084.

Diamond, P. A. (1993). Search, sticky prices, and inflation. *The Review of Economic Studies 60*(1), 53–68.

Golosov, M. and R. E. Lucas (2007). Menu costs and phillips curves. *Journal of Political Economy 115*(2), 171–199.

Head, A. and A. Kumar (2005). Price dispersion, inflation, and welfare. *International Economic Review 46*(2), 533–572.

Head, A., L. Q. Liu, G. Menzio, and R. Wright (2012). Sticky prices: A new monetarist approach. *Journal of the European Economic Association 10*(5), 939–973.

Hitsch, G. J., A. Hortacsu, and X. Lin (2019). Prices and promotions in us retail markets: Evidence from big data. Technical report, National Bureau of Economic Research.

Kaplan, G. and G. Menzio (2015). The morphology of price dispersion. *International Economic Review 56*(4), 1165–1206.

Kehoe, P. and V. Midrigan (2015). Prices are sticky after all. *Journal of Monetary Economics 75*, 35–53.

Lach, S. and D. Tsiddon (1992). The behavior of prices and inflation: An empirical analysis of disaggregat price data. *Journal of political economy 100*(2), 349–389.

Lippman, S. A. and J. J. McCall (1976). The economics of job search: A survey. *Economic inquiry 14*(2), 155–189.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 113–126.

Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics 123*(4), 1415–1464.

Nakamura, E. and J. Steinsson (2010). Monetary non-neutrality in a multisector menu cost model. *The Quarterly journal of economics 125*(3), 961–1013.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018). The elusive costs of inflation: Price dispersion during the us great inflation. *The Quarterly Journal of Economics 133*(4), 1933–1980.

Parsley, D. C. (1996). Inflation and relative price variability in the short and long run: new evidence from the united states. *Journal of Money, Credit and Banking 28*(3), 323–341.

Reinsdorf, M. (1994). New evidence on the relation between inflation and price dispersion. *The American Economic Review 84*(3), 720–731.

Sheremirov, V. (2020). Price dispersion and inflation: New facts and theoretical implications. *Journal of Monetary Economics 114*, 59–70.

Sheshinski, E. and Y. Weiss (1977). Inflation and costs of price adjustment. *The Review of Economic Studies 44*(2), 287–303.

Stevens, L. (2020). Coarse pricing policies. *The Review of Economic Studies 87*(1), 420–453.

Stiglitz, J. E. (1987). Competition and the number of firms in a market: Are duopolies more competitive than atomistic markets? *Journal of political Economy 95*(5), 1041–1061.

Tommasi, M. (1993). *Optimal Pricing, Inflation, and the Cost of Price Adjustment*, Chapter Inflation and Relative Prices: Evidence from Argentina. Cambridge, Mass.: MIT Press.

Van Hoomissen, T. (1988). Price dispersion and inflation: Evidence from israel. *Journal of Political Economy 96*(6), 1303–1314.

A Additional figures and tables

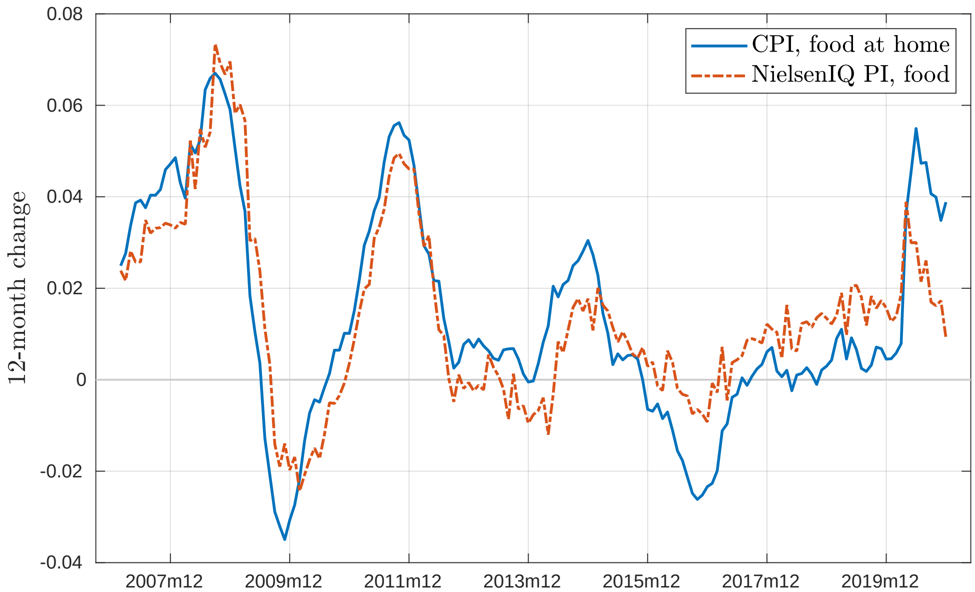

Figure A.1: Aggregate inflation, validation

The solid blue line corresponds to the Consumer Price Index for food at home reported by the Bureau of Labor Statistics. The red dashed line is the NielsenIQ Price Index for food categories, which results from aggregating product $\times$ county level inflation $\pi_{km,t}$ using annual sales as weights. I first aggregate product-level measures at the county level before doing so at the national level.

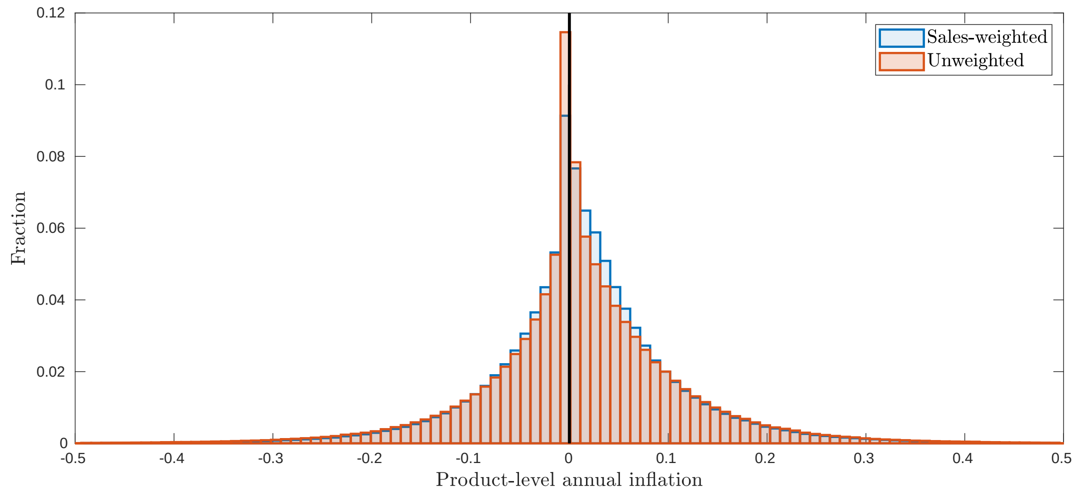

Figure A.2: Distribution of product-level annual inflation

Each bar shows the fraction of product $\times$ county $\times$ year observations within each of 100 equally spaced disaggregate inflation bins. The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. I compute inflation for each UPC and month, which I define as the annualized average size of price changes including stores that did *not* change their prices. The final sample for estimation corresponds to annual averages of the monthly variables for products sold by at least 10 stores. The total number of product $\times$ county $\times$ year observations is 40.1 million.

(a) Product-year (b) County-year

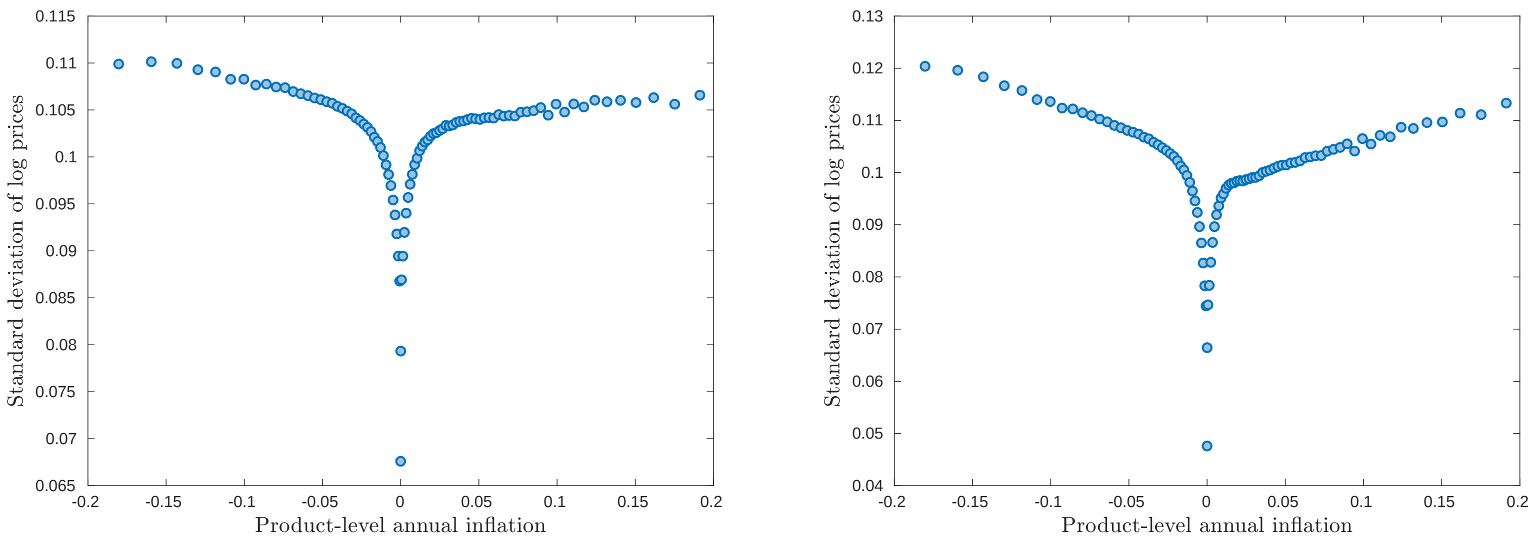

Figure A.3: Controlling for different fixed effects, product level

The dots correspond to average standard deviation of log prices for each of 100 equally sized disaggregate inflation bins as predicted by equation (1), for different combinations of the fixed effects. The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables for products sold by at least 10 stores. The unit of observation is a product $\times$ county $\times$ year, with a total of 40.1 million.

(a) Raw data (b) Including fixed effects

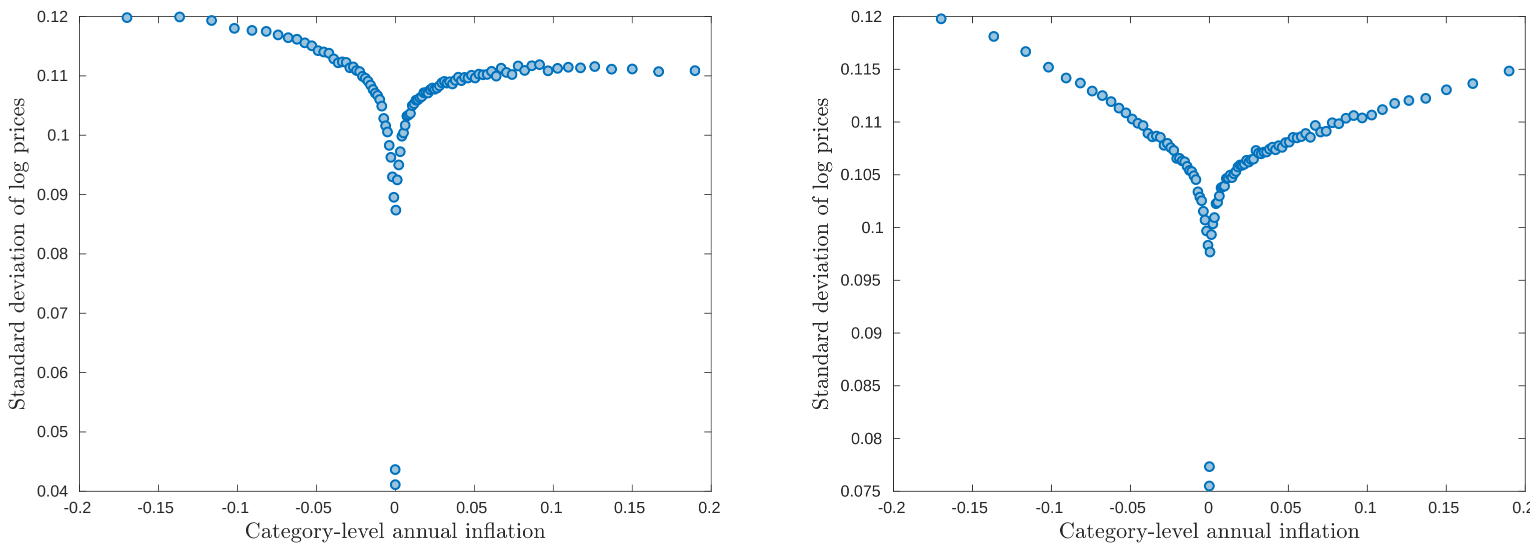

Figure A.4: Category-level price dispersion and inflation

The dots correspond to average standard deviation of log prices for each of 100 equally sized category-level inflation bins. The right figure conditions on controls as in equation (1) but using category-level, instead of product-level, statistics. The underlying microdata correspond to monthly prices for a given store and product (UPC). Monthly prices are defined as total revenues over quantity sold in a given month. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and the standard deviation of log prices among stores. I average products in the same category using annual sales as weights to produce category-level statistics. The final sample for estimation corresponds to annual averages of the monthly variables. The unit of observation is a category $\times$ county $\times$ year, with a total of 5,234,083.

(a) Product level (b) Category level

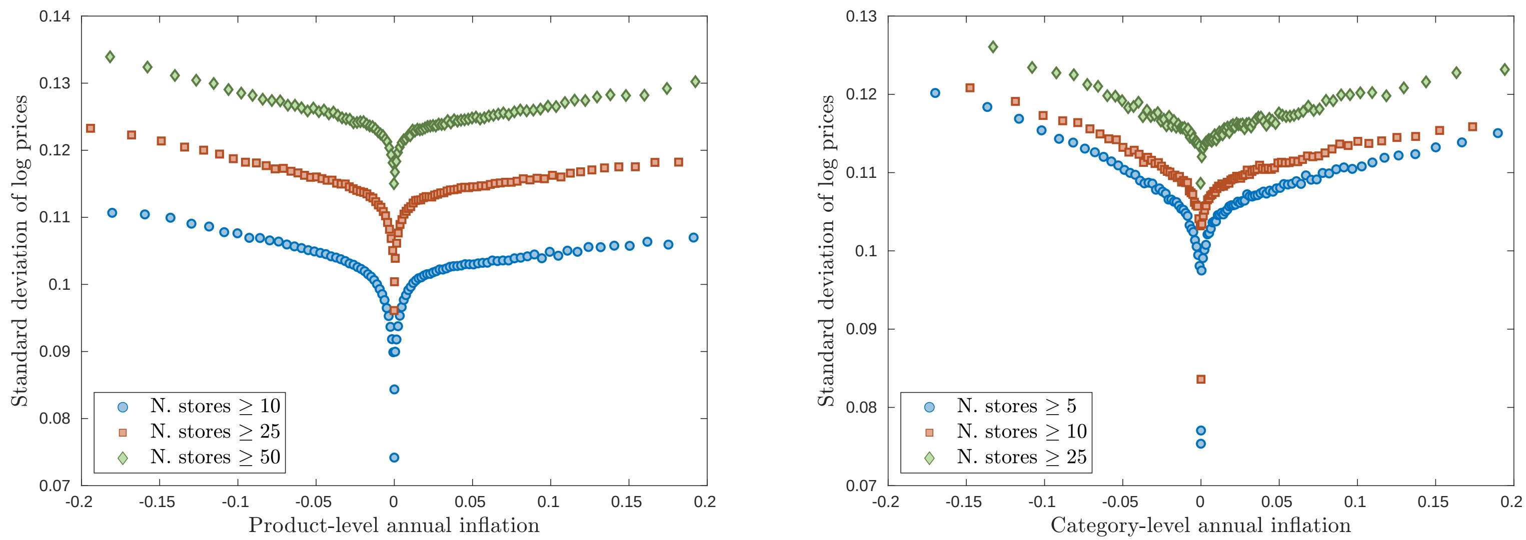

Figure A.5: Dividing sample by number of stores

The dots correspond to average standard deviation of log prices for each of 100 equally sized disaggregate inflation bins conditional on controls. The labels on the left (right) figure show results when we estimate equation (1) keeping products (categories) sold by at least 10 (5), 25 (10), and 50 (25) stores, which are roughly the minimum, median, and 90th percentile of the distribution. The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The right figure also averages among products in the same category (using annual sales as weights). For the right figure, the unit of observation is a product $\times$ county $\times$ year, with a total of 40.1 million when at least 10 stores sell the product. For the left figure, the unit of observation is a category $\times$ county $\times$ year, with a total of 5,234,083 when at least an average of 5 stores sell the products in the category.

(a) Product level (b) Category level

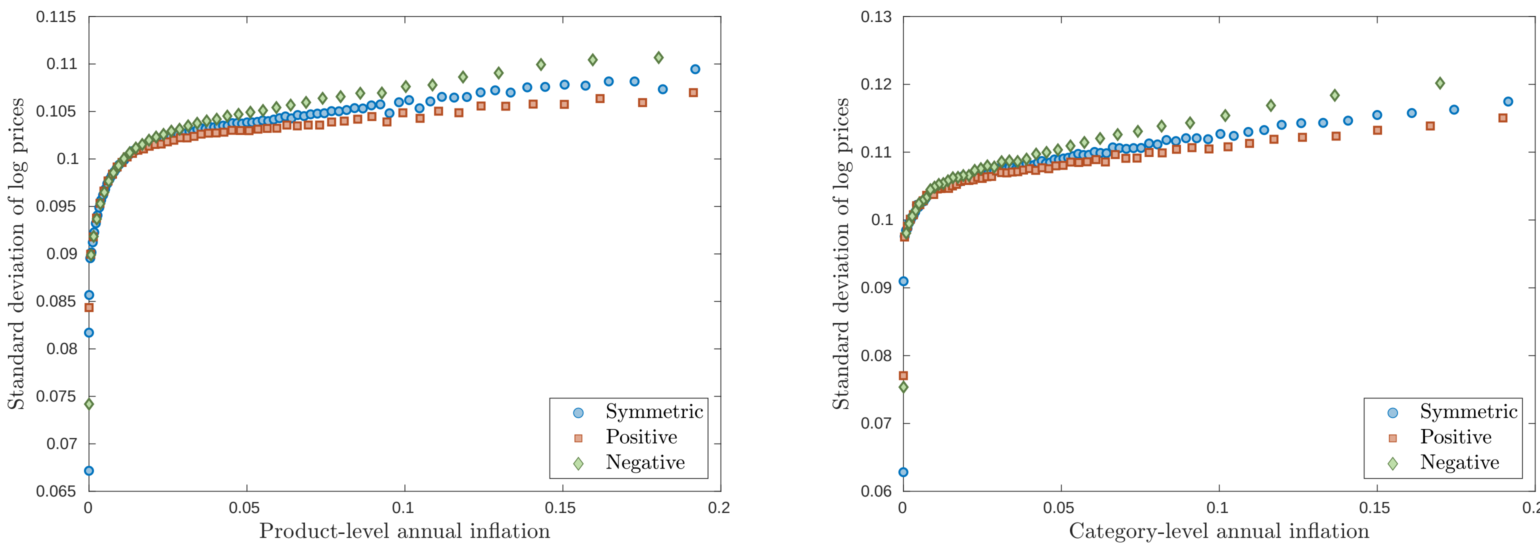

Figure A.6: Price dispersion and absolute inflation

The dots correspond to average standard deviation of log prices for each of 100 equally sized disaggregate inflation bins conditional on controls. The label “Symmetric” shows results when we estimate equation (1) for absolute inflation. The labels “Positive” and “Negative” indicate estimates in Figure 2 (left figure) and Figure A.4b (right figure) with absolute inflation on the x-axis. The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product $\times$ county $\times$ year, with a total of 40.1 million. For the right figure, the unit of observation is a category $\times$ county $\times$ year, with a total of 5,234,083.

(a) Product level

(b) Category level

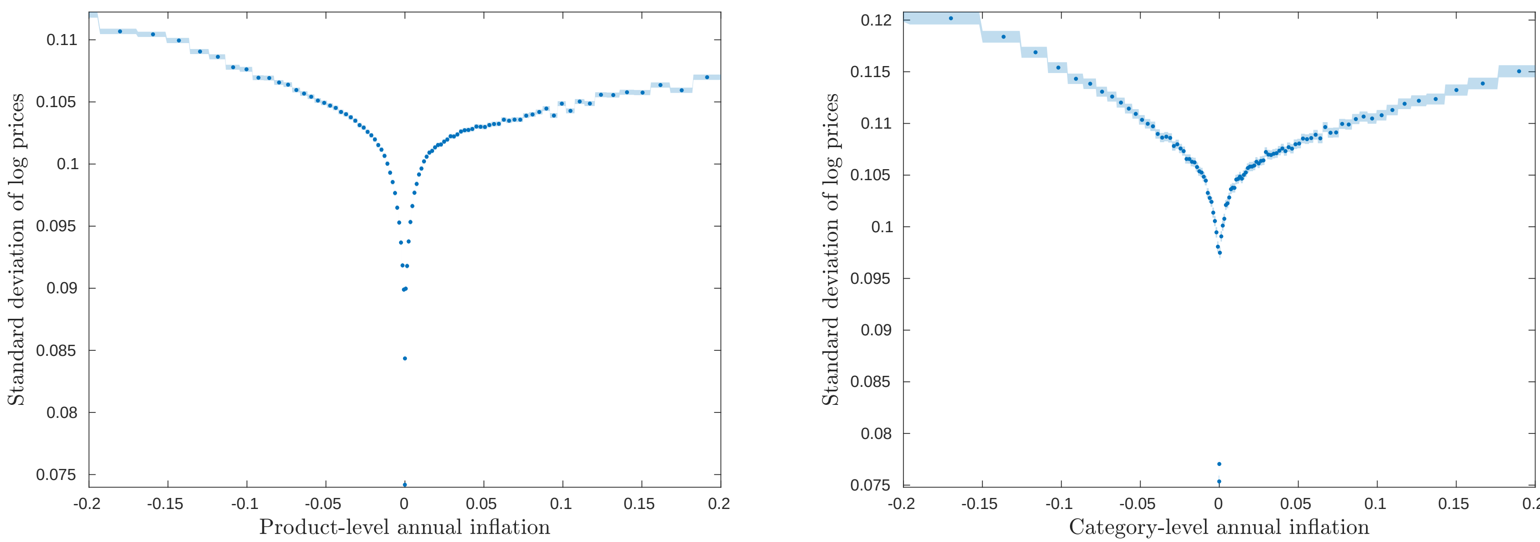

Figure A.7: Statistical significance

The dots correspond to average standard deviation of log prices for each of 100 equally sized disaggregate inflation bins as predicted by equation (1). The band covers the entire function with probability 0.95. Standard errors are clustered by product × county pair for the figure on the left, and by category × county pair for the figure on the right. The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product × county × year, with a total of 40.1 million. For the right figure, the unit of observation is a category × county × year, with a total of 5,234,083.

(a) Product level (b) Category level

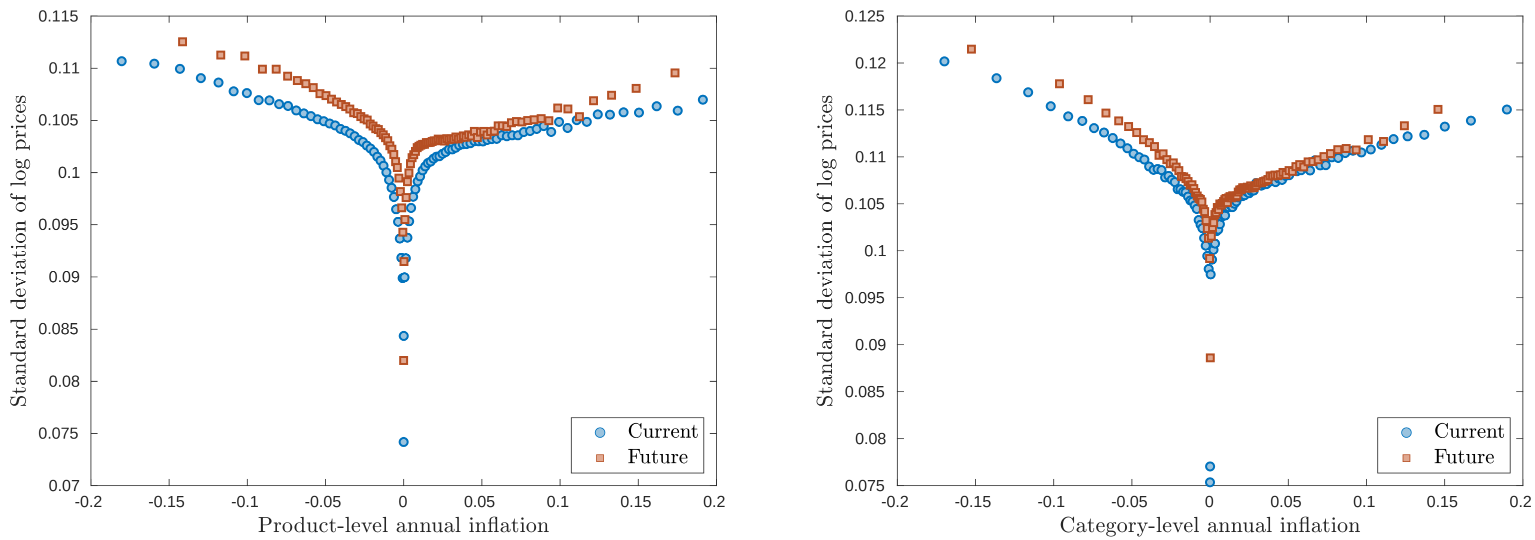

Figure A.8: Current and future inflation

The dots correspond to average standard deviation of log prices for each of 100 equally sized disaggregate inflation bins conditional on controls. The label “Future” shows results when we estimate equation (1) for next year’s observed inflation, when available. The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product $\times$ county $\times$ year, with a total of 40.1 million for “Current” and 29.9 million for “Future”. For the right figure, the unit of observation is a category $\times$ county $\times$ year, with a total of 5,234,083 for “Current” and 4,480,422 for “Future”.

(a) Product level (b) Category level

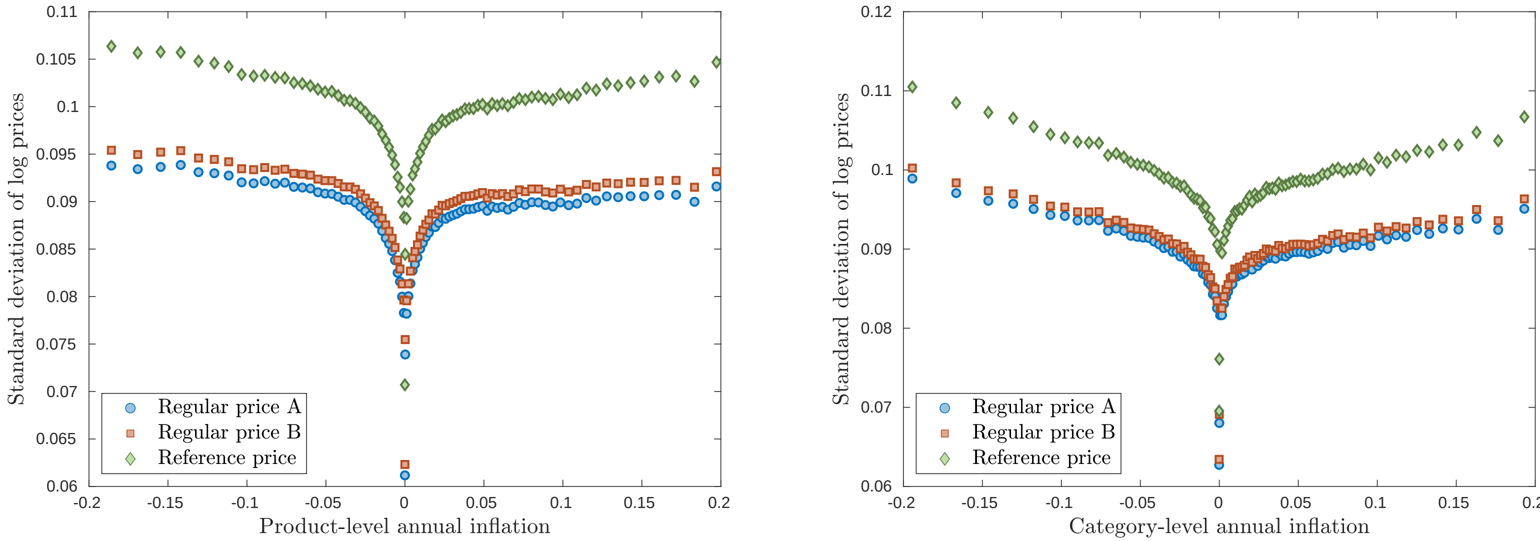

Figure A.9: Dispersion of regular prices and inflation, various filters

The dots correspond to average standard deviation of log prices for each of 100 equally sized disaggregate inflation bins conditional on controls as predicted by equation (1). The underlying microdata are weekly prices for a given store and product (UPC), defined as total revenues over quantity sold each week. I use three different filters to remove temporary price changes from weekly posted prices. The first two identify and remove V-shaped patterns, as in Nakamura and Steinsson (2008): "Regular price A" requires that the price returns to a regular price within 8 weeks after falling, while "Regular price B" is analogous but considers 4 weeks. The third filter follows Kehoe and Midrigan's algorithm to construct a regular price ("Reference price" in the plots). I define the reference price as the modal price within a window of 4 weeks around a given price. When the posted price equals the reference price, I set the regular price equal to the reference price; otherwise, I set it equal the previous' period regular price. To construct the monthly-level statistics, I keep the regular price in the third week of the month. For each UPC and month, I compute inflation (defined as the unconditional annualized average size of *posted* price changes) and the standard deviation of log prices among stores. The final sample corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product × county × year, for 11.7 million in total. For the right figure, the unit of observation is a category × county × year, for 2,584,702 in total.

(a) Product level

(b) Category level

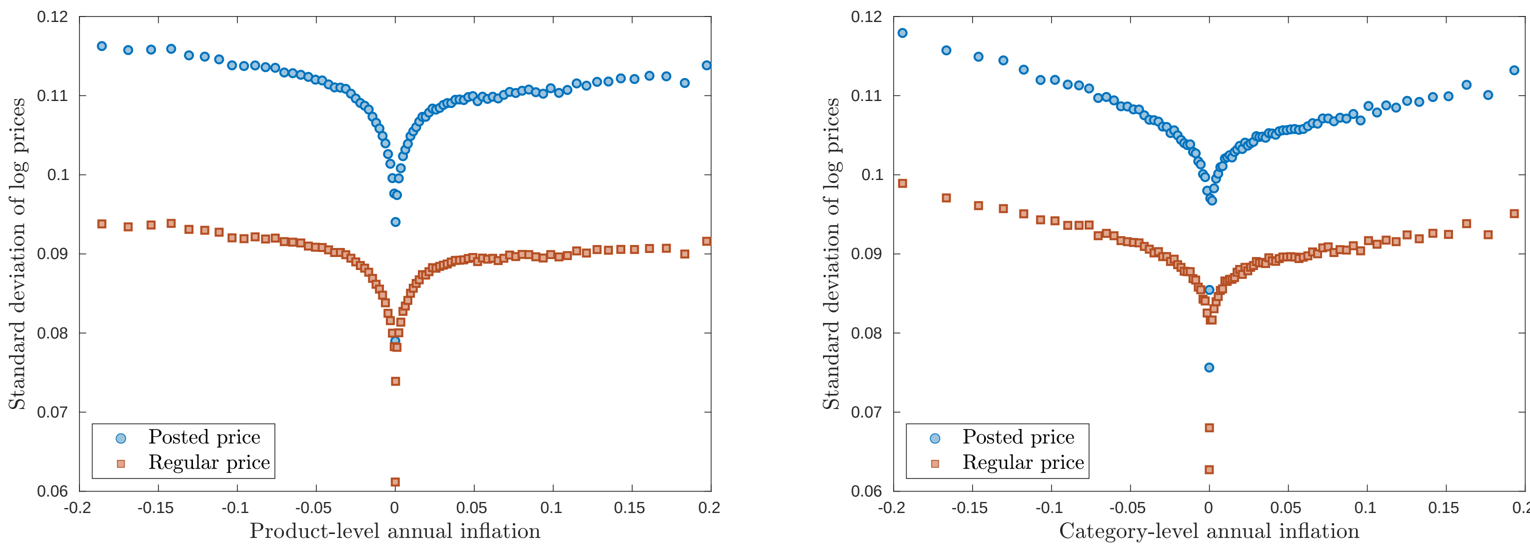

Figure A.10: Price dispersion and inflation, posted and regular prices

The dots correspond to average standard deviation of log prices for each of 100 equally sized disaggregate inflation bins conditional on controls as predicted by equation (1). The underlying microdata are weekly prices for a given store and product (UPC), defined as total revenues over quantity sold each week. "Posted prices" is the raw price series, while "regular prices" is the series after removing temporary sales using a filter as in Nakamura and Steinsson (2008). To construct the monthly-level statistics, I keep the price in the third week of the month. For each UPC and month, I compute inflation (defined as the annualized average size of *posted* price changes including stores that did *not* change their prices) and the standard deviation of both posted and log regular prices among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product $\times$ county $\times$ year, for 11.7 million in total. For the right figure, the unit of observation is a category $\times$ county $\times$ year, for 2,584,702 in total.

(a) Product level (b) Category level

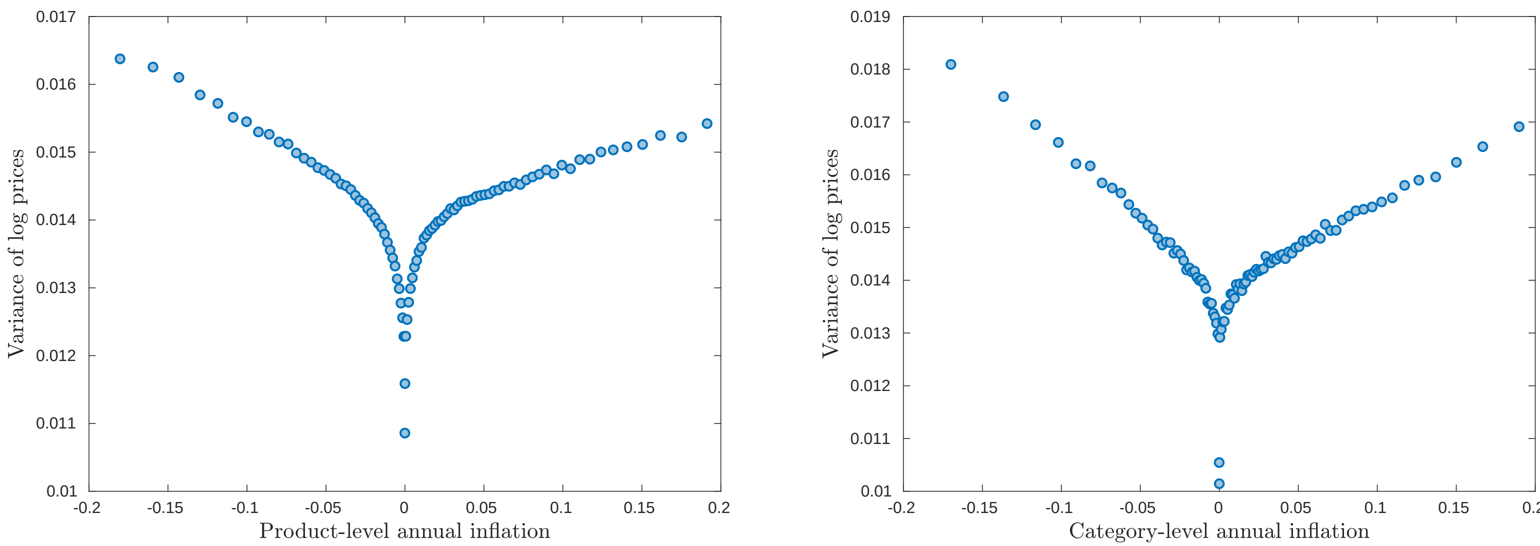

Figure A.11: Variance of log prices and inflation

The dots correspond to average variance of log prices for each of 100 equally sized disaggregate inflation bins as predicted by equation (1). The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product $\times$ county $\times$ year, with a total of 40.1 million. For the right figure, the unit of observation is a category $\times$ county $\times$ year, with a total of 5,234,083.

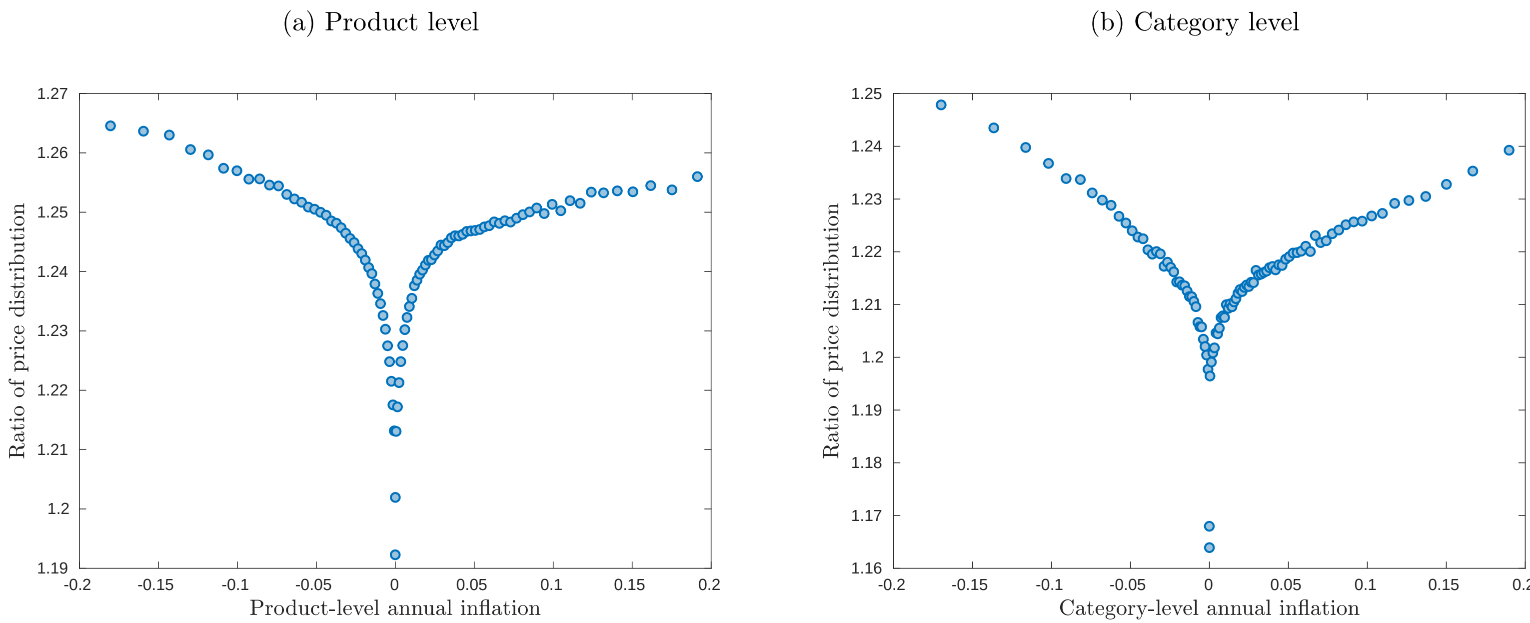

Figure A.12: 90-10 ratio of the price distribution and inflation

The dots correspond to the average 90-10 ratio of the price distribution for each of 100 equally sized disaggregate inflation bins as predicted by equation (1). The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product $\times$ county $\times$ year, with a total of 40.1 million. For the right figure, the unit of observation is a category $\times$ county $\times$ year, with a total of 5,234,083.

(a) Product level (b) Category level

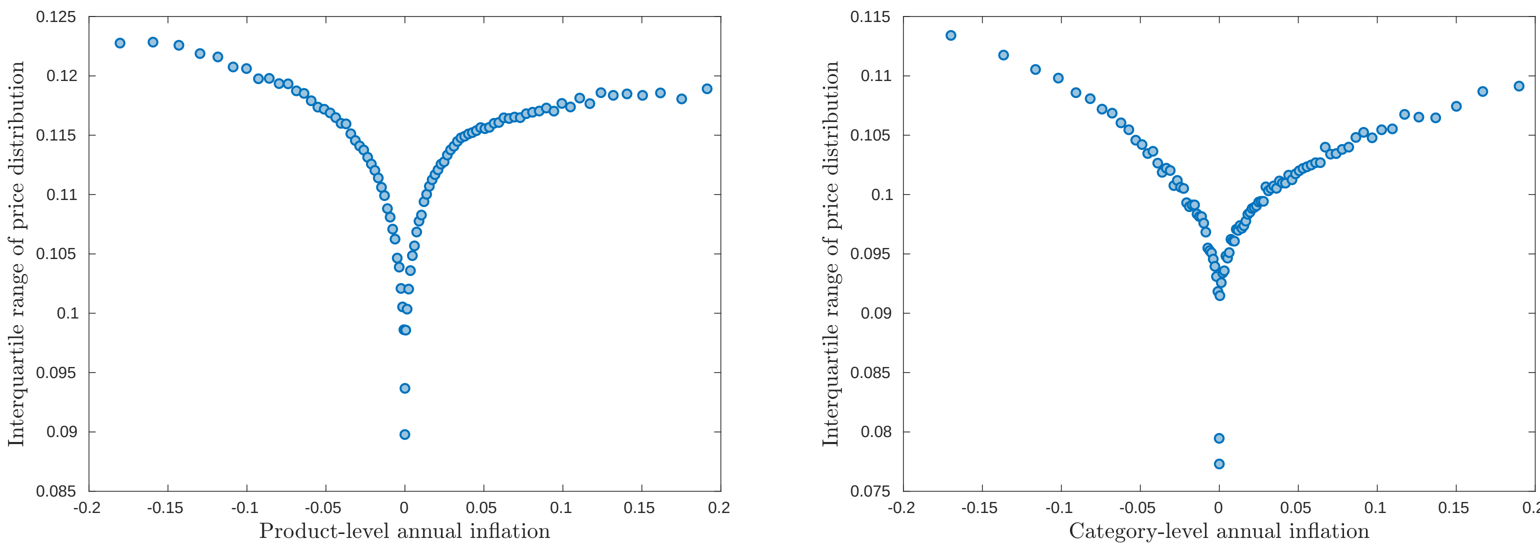

Figure A.13: Interquartile range of the price distribution and inflation

The dots correspond to the average interquartile range of the price distribution for each of 100 equally sized disaggregate inflation bins as predicted by equation (1). The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product $\times$ county $\times$ year, with a total of 40.1 million. For the right figure, the unit of observation is a category $\times$ county $\times$ year, with a total of 5,234,083.

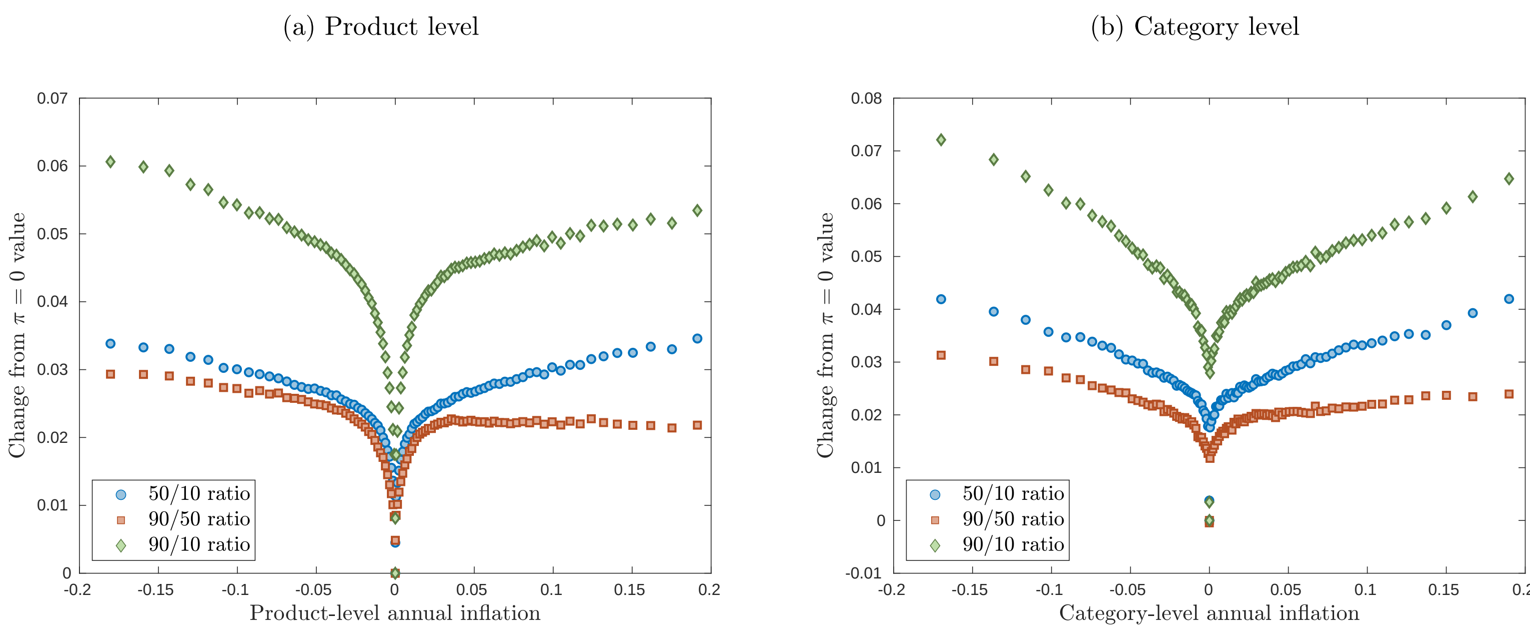

Figure A.14: Ratios of the price distribution and inflation

The dots correspond to average price dispersion for each of 100 equally sized disaggregate inflation bins as predicted by equation (1). Price dispersion is measured using the 90-10, 90-50 or 50-10 ratio of the price distribution. The y-axis shows deviations from predicted price dispersion values when $\pi = 0$. The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The left figure restricts the sample to products sold by at least 10 stores, while the right one averages among products in the same category (using annual sales as weights). For the left figure, the unit of observation is a product $\times$ county $\times$ year, with a total of 40.1 million. For the right figure, the unit of observation is a category $\times$ county $\times$ year, with a total of 5,234,083.

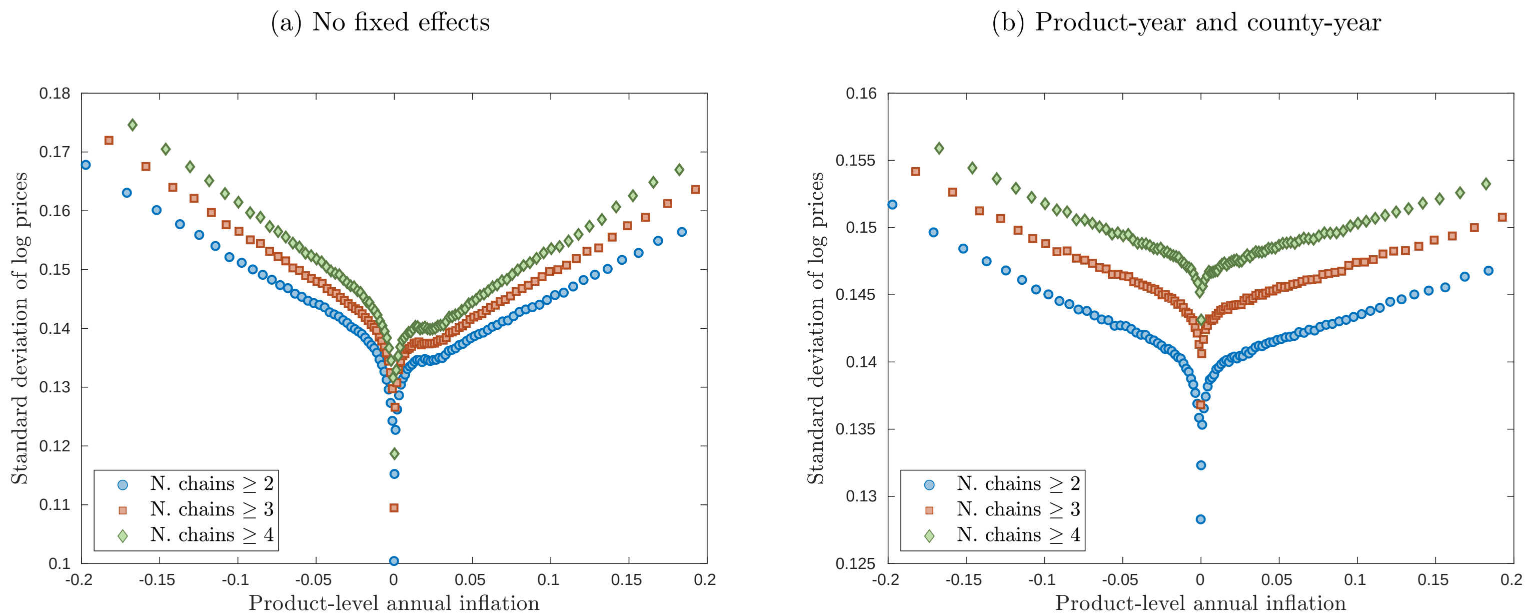

Figure A.15: Price dispersion at the chain level

The dots correspond to average standard deviation of log prices for each of 100 equally sized disaggregate inflation bins as predicted by equation (1), when the price is defined at the retail chain/parent company level. The figure on the left excludes fixed effects and the log-number of stores. The labels show results when keeping products sold by at least 2, 3, and 4 chains/parent companies, which are roughly the 25th, median, and 75th percentile of the distribution. The underlying microdata correspond, for a given retail chain/parent company, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation and price dispersion statistics among chains/parent companies. The final sample for estimation corresponds to annual averages of the monthly variables for products sold by at least 10 stores. The unit of observation is a product $\times$ county $\times$ year, with a total of 33.3 million when there are at least 2 chains, 24.7 million when there are at least 3 chains, and 15.6 million when there are at least 4 chains.

(a) Levels

(b) Deviation from $\pi = 0$

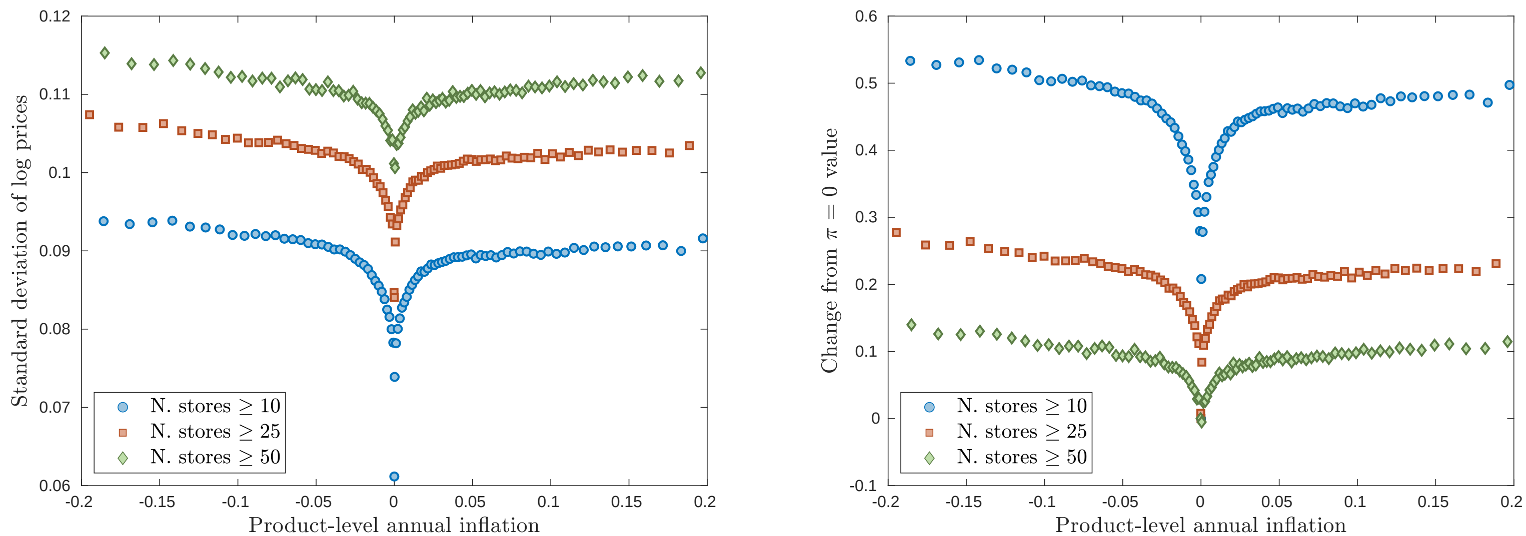

Figure A.16: Dispersion of regular prices and inflation by number of stores

The dots correspond to average standard deviation of log regular prices for each of 100 equally sized disaggregate inflation bins conditional on controls as predicted by equation (1). The plots show estimates keeping products sold by at least 10, 25, and 50 stores, which are roughly the minimum, median, and 90th percentile of the distribution. The left figure is in levels; the right one in deviations from $\pi = 0$ values. The underlying microdata are weekly prices for a given store and product (UPC), defined as total revenues over quantity sold each week. Regular prices are raw weekly prices after removing temporary sales using a filter as in Nakamura and Steinsson (2008) ("Regular price A" in the main text). To construct the monthly-level statistics, I keep the price in the third week of the month. For each UPC and month, I compute inflation (defined as the annualized average size of *posted* price changes including stores that did *not* change their prices) and the standard deviation of log regular prices among stores. The final sample for estimation corresponds to annual averages of the monthly variables. The unit of observation is a product $\times$ county $\times$ year, with a total of 11.7 million, 3,007,475, and 829,541 when at least 10, 25, and 50 stores, respectively, sell the product.

(a) Levels

(b) Deviation from $\pi = 0$

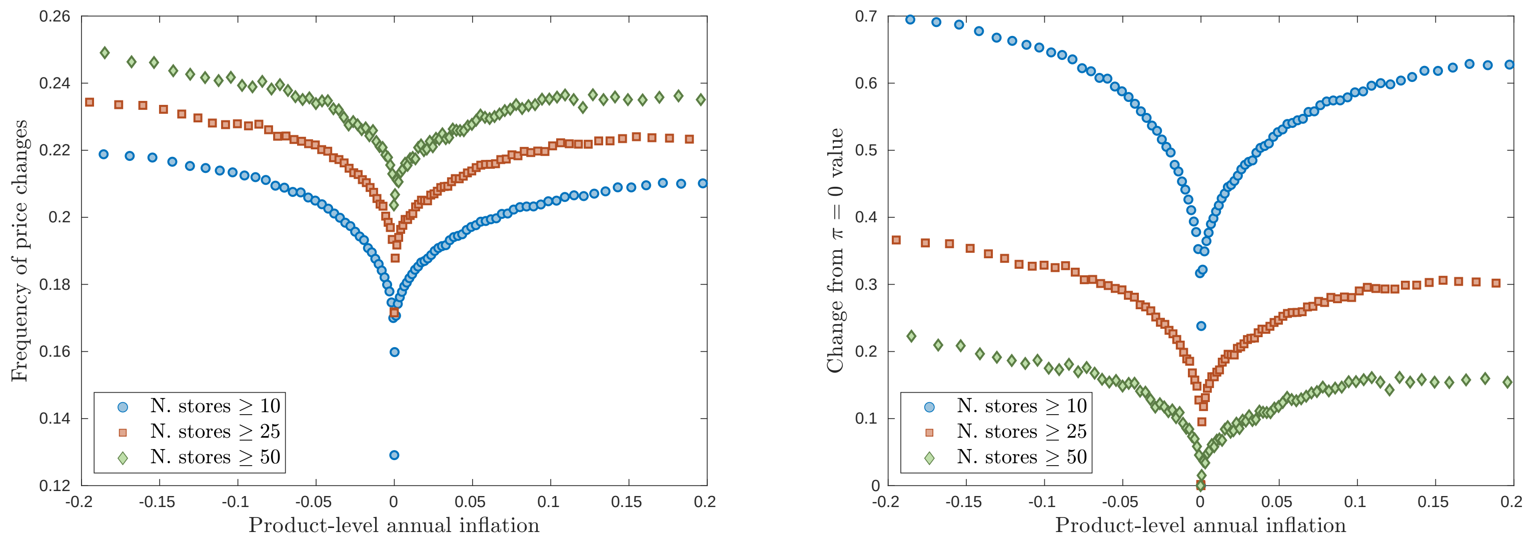

Figure A.17: Frequency of regular price changes and inflation by number of stores

The dots correspond to average frequency of regular price changes for each of 100 equally sized disaggregate inflation bins conditional on controls. The estimation is analogous to equation (1) when the dependent variable is the average frequency of regular price changes. The plots show estimates keeping products sold by at least 10, 25, and 50 stores, which are roughly the minimum, median, and 90th percentile of the distribution. The left figure is in levels; the right one in deviations from $\pi = 0$ values. The underlying microdata are weekly prices for a given store and product (UPC), defined as total revenues over quantity sold each week. Regular prices are raw weekly prices after removing temporary sales using a filter as in Nakamura and Steinsson (2008) ("Regular price A" in the main text). To construct the monthly-level statistics, I keep the price in the third week of the month. For each UPC and month, I compute inflation (defined as the annualized average size of *posted* price changes including stores that did *not* change their prices) and the frequency of regular price changes (or fraction of stores that adjusted regular prices in a given month). The final sample for estimation corresponds to annual averages of the monthly variables. The unit of observation is a product $\times$ county $\times$ year, with a total of 11.7 million, 3,007,475, and 829,541 when at least 10, 25, and 50 stores, respectively, sell the product.

(a) Levels | (b) Deviation from $\pi = 0$

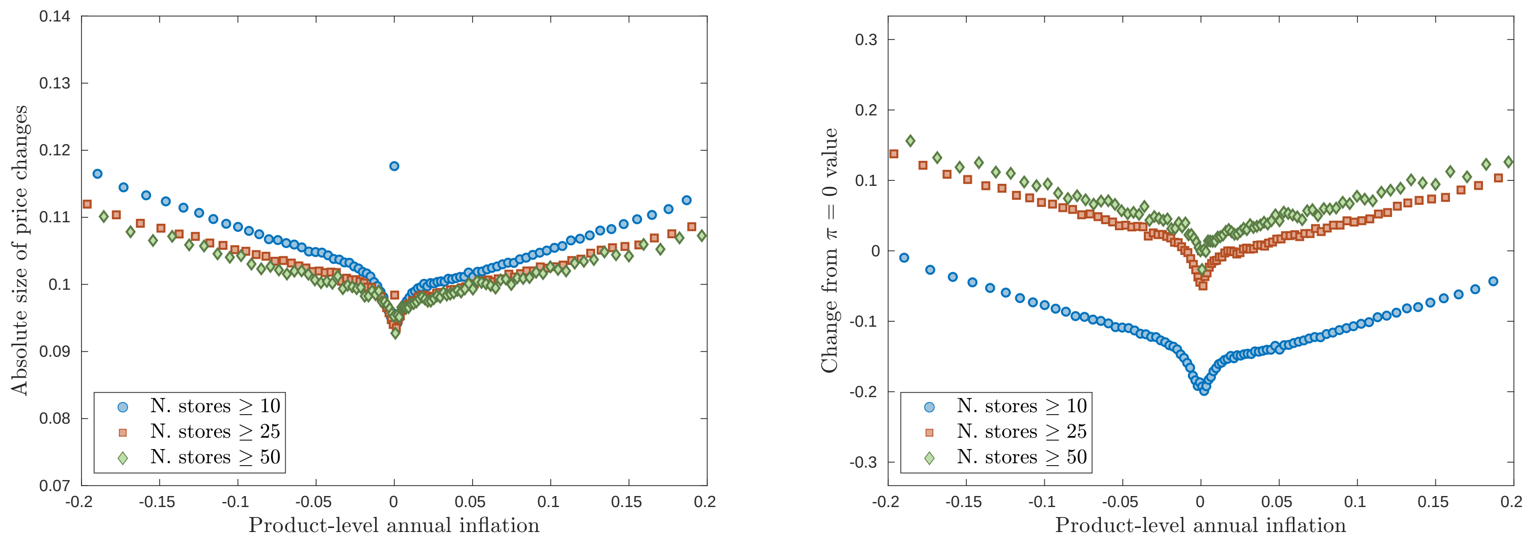

Figure A.18: Absolute size of regular price changes and inflation by number of stores

The dots correspond to average absolute size of regular price changes for each of 100 equally sized disaggregate inflation bins conditional on controls. The estimation is analogous to equation (1) when the dependent variable is the average absolute size of regular price changes. The plots show estimates keeping products sold by at least 10, 25, and 50 stores, which are roughly the minimum, median, and 90th percentile of the distribution. The left figure is in levels; the right one in deviations from $\pi = 0$ values. The underlying microdata are weekly prices for a given store and product (UPC), defined as total revenues over quantity sold each week. Regular prices are raw weekly prices after removing temporary sales using a filter as in Nakamura and Steinsson (2008) ("Regular price A" in the main text). To construct the monthly-level statistics, I keep the price in the third week of the month. For each UPC and month, I compute inflation (defined as the annualized average size of *posted* price changes including stores that did *not* change their prices) and the absolute size of regular price changes (conditional on adjustment). The final sample for estimation corresponds to annual averages of the monthly variables. The unit of observation is a product $\times$ county $\times$ year, with a total of 11.7 million, 3,007,475, and 829,541 when at least 10, 25, and 50 stores, respectively, sell the product.

(a) Absolute size of regular price changes

(b) Frequency of regular price changes

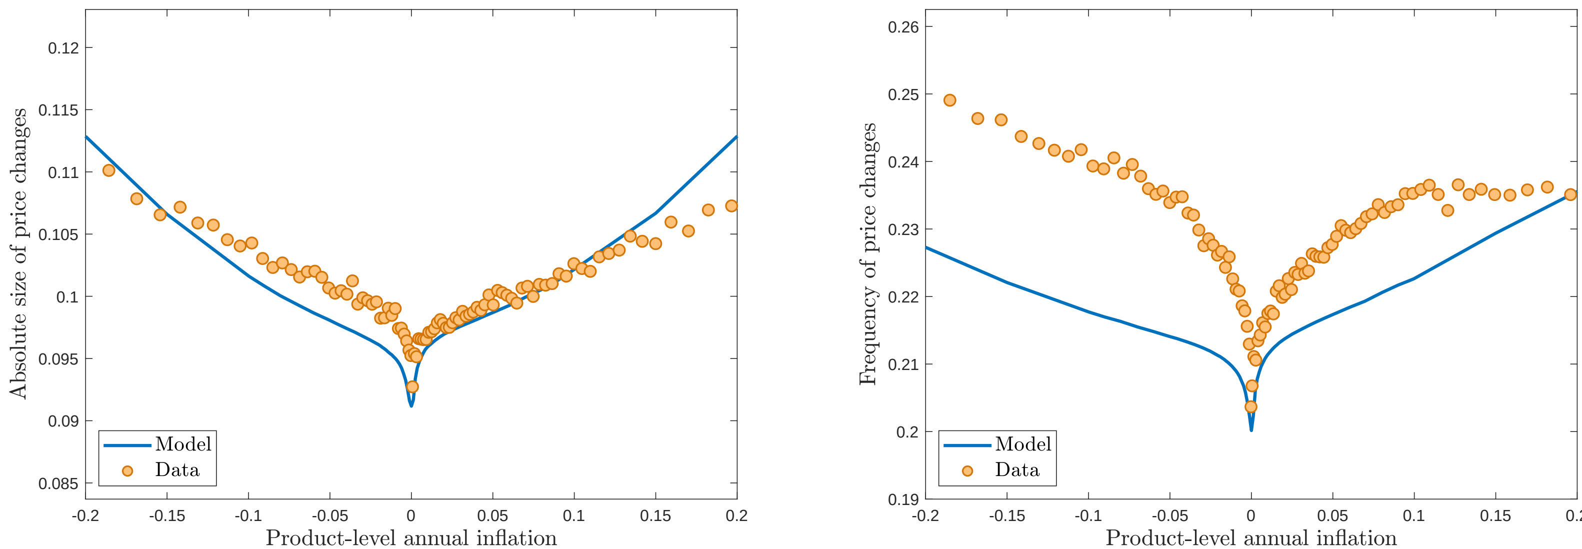

Figure A.19: Additional pricing moments and inflation, model fit

The figures show the relationship between product-level inflation and additional pricing moments in the calibrated model (line) and in the data (dots). The dots in the left (right) figure correspond to average absolute size (frequency) of regular price changes for each of 100 equally sized disaggregate inflation bins conditional on controls. The estimation is analogous to equation (1) when the dependent variable is the average absolute size or frequency of regular price changes, keeping products sold by at least 50 stores.

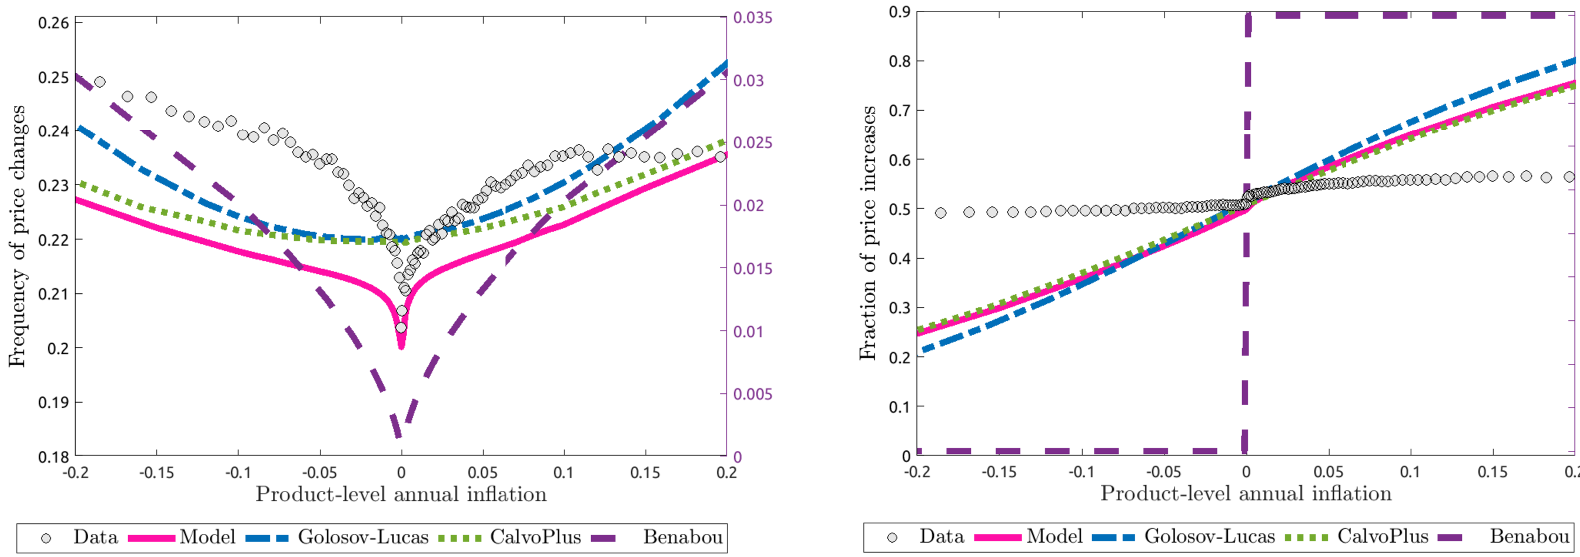

Figure A.20: Pricing moments in different models

The figure plots pricing moments and inflation in the data, model, pure menu-cost model with idiosyncratic shocks but without search (Golosov and Lucas, 2007), CalvoPlus price-setting model with idiosyncratic shocks but without search, and pure menu-cost model with search but without idiosyncratic shocks (Benabou, 1992; right y-axis). The left panel shows product-level frequency of price changes, and the right panel plots the fraction of price increases. Pricing moments are computed for regular prices of products sold by at least 50 stores.

Table A.1: Descriptive statistics, pricing behavior

	Mean	Std.Dev.	Percentile						
			1^{st}	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	99^{th}
Aggregate inflation	0.015	0.017	-0.012	-0.005	0.001	0.013	0.024	0.039	0.067
A. Product level									
Inflation	0.013	0.106	-0.292	-0.096	-0.030	0.005	0.058	0.128	0.322
Absolute inflation	0.070	0.081	0.000	0.004	0.015	0.045	0.097	0.168	0.376
Price dispersion									
Std. dev. of log prices	0.102	0.063	0.000	0.025	0.056	0.096	0.139	0.183	0.283
Max-min price ratio	1.415	0.369	1.000	1.070	1.182	1.351	1.566	1.816	2.509
90-10 percentile ratio	1.243	0.207	1.000	1.027	1.093	1.204	1.343	1.499	1.912
90-50 percentile ratio	1.109	0.119	1.000	1.005	1.026	1.075	1.152	1.249	1.543
50-10 percentile ratio	1.121	0.128	1.000	1.008	1.033	1.088	1.169	1.269	1.564
Number of stores	26.2	29.8	10.0	10.0	12.0	18.0	28.0	47.0	148.0
Number of chains	3.3	1.8	1	1	2	3	4	6	9
B. Category level									
Inflation	0.016	0.078	-0.196	-0.060	-0.018	0.009	0.047	0.100	0.255
Absolute inflation	0.052	0.060	0.000	0.004	0.013	0.033	0.069	0.123	0.283
Price dispersion									
Std. dev. of log prices	0.107	0.056	0.003	0.038	0.070	0.104	0.138	0.175	0.274
Max-min price ratio	1.335	0.248	1.002	1.077	1.175	1.302	1.448	1.612	2.082
90-10 percentile ratio	1.218	0.167	1.001	1.043	1.107	1.196	1.295	1.404	1.731
90-50 percentile ratio	1.095	0.089	1.000	1.010	1.034	1.078	1.131	1.193	1.407
50-10 percentile ratio	1.111	0.108	1.000	1.015	1.044	1.093	1.151	1.220	1.451
Number of stores	14.6	18.9	5.0	5.4	6.6	9.0	15.6	27.4	91.4
Number of UPCs	18.6	36.2	1.0	1.0	3.0	7.0	18.0	45.0	185.0

Notes. The table shows unweighted descriptive statistics for the sample used to estimate equation (1). The underlying microdata correspond, for a given store, product (UPC), and month, to total revenues over quantity sold – or monthly price. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices) and price dispersion statistics among stores. The final sample for estimation corresponds to annual averages of the monthly variables. Panel A shows product-level statistics for UPCs sold by at least 10 stores, while panel B also averages among products in the same module/category (using annual sales as weights). For panel A, the total number of observations is 40.1 million for 9,235,551 UPC × county pairs between 2007 and 2020, with 268,149 unique UPCs and 904 counties. For panel B, the total number of observations is 5,234,083 for 556,452 category × county pairs between 2007 and 2020, with 1,078 unique categories and 1,460 counties. The fraction of observations with deflation is 0.427 at the UPC and 0.397 at the category level. Aggregate inflation corresponds to an average over all product × county pairs using annual sales as weights.

Table A.2: Detailed descriptive statistics, product-level pricing behavior (weekly)

			Percentile						
	Mean	Std.Dev.	1^{st}	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	99^{th}
Inflation	0.014	0.130	-0.367	-0.125	-0.037	0.007	0.071	0.156	0.387
Absolute inflation	0.087	0.098	0.000	0.004	0.018	0.055	0.121	0.211	0.451
Aggregate inflation	0.015	0.017	-0.012	-0.005	0.001	0.013	0.024	0.039	0.067
Number of stores	23.6	25.1	10	10	12	16	25	41	128
A. Posted prices									
Price dispersion									
Std. dev. of log prices	0.108	0.066	0.000	0.026	0.059	0.103	0.150	0.196	0.285
Max-min price ratio	1.419	0.413	1.000	1.060	1.178	1.349	1.574	1.835	2.531
90-10 percentile ratio	1.245	0.215	1.000	1.014	1.080	1.202	1.357	1.525	1.911
90-50 percentile ratio	1.107	0.120	1.000	1.000	1.017	1.071	1.157	1.257	1.531
50-10 percentile ratio	1.125	0.134	1.000	1.002	1.026	1.089	1.185	1.295	1.562
Price flexibility									
Frequency of price changes	0.476	0.230	0.000	0.149	0.307	0.493	0.654	0.772	0.917
Abs. size of price changes	0.163	0.095	0.020	0.055	0.092	0.148	0.216	0.287	0.451
B. Regular prices									
Price dispersion									
Std. dev. of log prices	0.088	0.061	0.000	0.013	0.042	0.080	0.124	0.168	0.264
Max-min price ratio	1.328	0.295	1.000	1.024	1.119	1.269	1.467	1.688	2.240
90-10 percentile ratio	1.195	0.195	1.000	1.001	1.047	1.147	1.288	1.447	1.833
90-50 percentile ratio	1.086	0.118	1.000	1.000	1.004	1.044	1.121	1.229	1.534
50-10 percentile ratio	1.101	0.135	1.000	1.000	1.006	1.054	1.148	1.270	1.579
Price flexibility									
Frequency of price changes	0.196	0.142	0.000	0.023	0.084	0.175	0.282	0.392	0.605
Abs. size of price changes	0.106	0.067	0.016	0.040	0.060	0.092	0.135	0.187	0.334

Notes. The table shows unweighted descriptive statistics for the sample used to estimate equation (1). The underlying microdata correspond to monthly prices for a given store and product (UPC), defined as total revenues over quantity sold in the third week of the month. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices), price dispersion statistics among stores, the percent of stores adjusting their prices (frequency), and the average absolute size of these changes conditional on adjusting. The final sample corresponds to annual averages of monthly statistics for UPCs sold by at least 10 stores. Panel B shows the dispersion and flexibility statistics after removing temporary sales from the underlying price series using a filter ("Regular price A"), as in Nakamura and Steinsson (2008). The total number of observations is 11.7 million for 2,813,263 UPC $\times$ county pairs between 2007 and 2020, with 129,490 unique UPCs and 897 counties. The fraction of observations with deflation is 0.415. Aggregate inflation corresponds to an average over all product $\times$ county pairs using annual sales as weights. Frequency corresponds to percent per month.

Table A.3: Detailed descriptive statistics, category-level pricing behavior (weekly)

	Mean	Std.Dev.	Percentile 1^{st}	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	99^{th}
Inflation	0.017	0.104	-0.276	-0.086	-0.026	0.011	0.061	0.129	0.328
Absolute inflation	0.069	0.079	0.000	0.005	0.017	0.044	0.093	0.164	0.373
Aggregate inflation	0.015	0.017	-0.012	-0.005	0.001	0.013	0.024	0.039	0.067
A. Posted prices									
Price dispersion									
Std. dev. of log prices	0.106	0.059	0.000	0.032	0.065	0.103	0.142	0.179	0.269
Max-min price ratio	1.308	0.292	1.000	1.050	1.140	1.268	1.425	1.597	2.048
90-10 percentile ratio	1.200	0.198	1.000	1.025	1.081	1.172	1.284	1.402	1.708
90-50 percentile ratio	1.085	0.090	1.000	1.004	1.021	1.064	1.123	1.188	1.387
50-10 percentile ratio	1.106	0.140	1.000	1.008	1.033	1.083	1.151	1.228	1.451
Price flexibility									
Frequency of price changes	0.424	0.194	0.000	0.157	0.286	0.431	0.564	0.672	0.835
Abs. size of price changes	0.155	0.072	0.029	0.069	0.103	0.147	0.196	0.249	0.366
B. Regular prices									
Price dispersion									
Std. dev. of log prices	0.089	0.057	0.000	0.020	0.050	0.084	0.121	0.160	0.258
Max-min price ratio	1.254	0.270	1.000	1.026	1.102	1.214	1.356	1.515	1.941
90-10 percentile ratio	1.167	0.191	1.000	1.013	1.058	1.136	1.237	1.351	1.674
90-50 percentile ratio	1.072	0.089	1.000	1.000	1.012	1.047	1.101	1.170	1.390
50-10 percentile ratio	1.089	0.145	1.000	1.001	1.017	1.060	1.126	1.210	1.465
Price flexibility									
Frequency of price changes	0.173	0.112	0.000	0.038	0.090	0.160	0.239	0.322	0.500
Abs. size of price changes	0.099	0.042	0.026	0.051	0.069	0.093	0.121	0.152	0.227

Notes. The table shows unweighted descriptive statistics for the sample used to estimate equation (1). The underlying microdata correspond to monthly prices for a given store and product (UPC), defined as total revenues over quantity sold in the third week of the month. For each UPC and month, I compute inflation (defined as the annualized average size of price changes including stores that did *not* change their prices), price dispersion statistics among stores, the percent of stores adjusting their prices (frequency), and the average absolute size of these changes conditional on adjusting. The final sample corresponds to annual averages of monthly statistics for products in the same category, using annual sales as weights. Panel B shows the dispersion and flexibility statistics after removing temporary sales from the underlying price series using a filter ("Regular price A"), as in Nakamura and Steinsson (2008). The total number of observations is 2,584,702 for 298,987 category × county pairs between 2007 and 2020, with 1,018 unique categories and 1,456 counties. The fraction of observations with deflation is 0.399. Aggregate inflation corresponds to an average over all category × county pairs using annual sales as weights.

Table A.4: Numerical tests, year-by-year product-level estimates

	Full $\|\pi\|$ range	$\|\pi\| < 2\%$
County-year estimates, $f_{m,y}$		
% reject null: Increasing	25.19	6.23
% reject null: Concave	35.31	8.66
Average # products per county-year	4,270.52	1,524.21
Total # county-year observations	9,372	7,838
Product-year estimates, $f_{k,y}$		
% reject null: Increasing	40.92	10.57
% reject null: Concave	29.86	11.23
Average # counties per product-year	142.41	90.34
Total # product-year observations	216,227	59,107

Notes. The functions $f_{m,y}$ and $f_{k,y}$ denote the relationship between price dispersion and absolute inflation for each county-year and product-year combination, respectively. The top panel shows a summary of the results for $f_{m,y}$: the percent of county-year estimates that reject the null hypothesis of monotonicity and concavity; the average number of products used to estimate each function $f_{m,y}$ (where I imposed a minimum of 50); and the total number of estimated functions $f_{m,y}$. The first column shows estimates of $f_{m,y}$ for the full absolute inflation range; the second, for absolute inflation lower than 2%. The bottom panel shows a summary of the results for $f_{k,y}$, and the description is analogous.

Table A.5: Numerical tests, year-by-year category-level estimates

	Full $\|\pi\|$ range	$\|\pi\| < 2\%$
County-year estimates, $f_{m,y}$		
% reject null: Increasing	18.65	3.62
% reject null: Concave	16.63	6.06
Average # categories per county-year	428.51	60.23
Total # county-year observations	12,017	1,106
Category-year estimates, $f_{k,y}$		
% reject null: Increasing	37.27	5.60
% reject null: Concave	31.86	6.89
Average # counties per category-year	404.21	77.16
Total # category-year observations	12,717	929

Notes. The functions $f_{m,y}$ and $f_{k,y}$ denote the relationship between price dispersion and absolute inflation for each county-year and category-year combination, respectively. The top panel shows a summary of the results for $f_{m,y}$: the percent of county-year estimates that reject the null hypothesis of monotonicity and concavity; the average number of categories used to estimate each function $f_{m,y}$ (where I imposed a minimum of 50); and the total number of estimated functions $f_{m,y}$. The first column shows estimates of $f_{m,y}$ for the full absolute inflation range; the second, for absolute inflation lower than 2%. The bottom panel shows a summary of the results for $f_{k,y}$, and the description is analogous.

Table A.6: Descriptive statistics, shopping behavior

	Mean	Std.Dev.	Percentile 1^{st}	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}	99^{th}
Household $\times$ department $\times$ quarter									
Log relative price paid	-0.004	0.105	-0.332	-0.088	-0.029	0	0.032	0.079	0.247
Number of retailers visited	3.033	2.116	1	1	2	2	4	6	10
Number of shopping trips	9.042	8.684	1	2	3	7	12	19	43
Number of transactions	25.494	42.254	1	2	5	11	26	60	217
Fraction of trans. with coupons	0.078	0.169	0	0	0	0	0.074	0.278	0.833
Department $\times$ market $\times$ quarter									
Inflation (annualized)	0.014	0.034	-0.079	-0.017	-0.012	0.011	0.027	0.051	0.115
Absolute inflation (annualized)	0.025	0.027	0.000	0.003	0.008	0.017	0.032	0.057	0.125

Notes. The table shows descriptive statistics for the sample used to estimate equation (12). Household-related measures are computed using the Consumer Panel Data; inflation, using the Retail Scanner Data. The total number of observations for estimation are 24,468,527 household $\times$ department $\times$ quarter triples, with 149,324 households, 10 departments, 56 quarters, and 53 Scantrack markets. At the department $\times$ market $\times$ quarter level, 29,572 distinct observations for inflation are available.

Table A.7: Parameters used in Section 4

Retailers	
Size of the menu cost	$\kappa = 0.02$
Mean-reversion rate of idiosyncratic shocks	$\rho_v = 0.85$
Volatility of idiosyncratic shocks	$\sigma_v = 0.07$
Shoppers	
Shape of search cost distribution, $Beta\,(g_a, g_b)$	$g_a = 1.05$
Shape of search cost distribution, $Beta\,(g_a, g_b)$	$g_b = 4.00$
Price elasticity of demand of shoppers	$\eta = 3.00$
Disutility cost of labor	$\tau = 2.50$

*Note*s. The table shows the parameters used to describe the intuition of the model in Section 4. We assume $\gamma \sim Beta\,(g_a, g_b)$, where $g_a > 1$, $g_b > 1$, and $\gamma \in [0, 1]$. In addition, we set $\theta = \sigma^{-1} = 1$ and $\rho = 0.96^{1/12}$.

B Temporary price reductions

To identify temporary price reductions, I define prices as total revenues over quantity sold each week. Provided intra-week price changes are less frequent than intra-month changes, defining prices at the weekly level would allow for a more precise measurement of the exact price the store charges and, thus, a better identification of price changes and sales using filters.

I use three different filters to remove temporary price changes from the series of posted prices and obtain a series of *regular* prices. The first two filters remove V-shaped patterns in the posted price series, as in Nakamura and Steinsson (2008). Both filters are designed to remove patterns in which a price returns to the same or to a different regular price within J weeks, after an initial drop. The parameters K and L in their algorithm – also in number of weeks – determine the transition to a new regular price. Stevens (2020) finds that the parameter J is the most important in determining the frequency of regular price changes. Therefore, I fix $K = L = 6$, and set $J = 8$ for what I call "Regular price A" and $J = 4$ for "Regular price B".

The third filter follows Kehoe and Midrigan's algorithm to construct a regular price. Consider the modal price within a window of 4 weeks before and after the current price. I call this price a reference price if there are at least 5 out of the 9 weeks of data available, and if at least a third of the observations equal the modal price. In rough terms, the regular price equals the reference price if the posted equals the reference price. If not, the regular price corresponds to the regular price of the previous period. Deviations from the regular price series are considered temporary.

To have enough data points for the filters to work, I only keep product-store pairs with at least two consecutive observations each month of the year. I do not impute missing price values. Also before applying the filters, I round prices to two decimal places and remove spurious price changes by identifying price changes in middle of the week.

C Additional pricing-behavior measures

Two relevant statistics of price stickiness are the frequency and size of regular price changes (i.e., excluding price changes due to temporary sales). To construct these measures, I start by defining a price spell as an uninterrupted sequence of weekly prices for a given product $\times$ store pair.[47] Following Coibion et al. (2015), I classify the difference between two consecutive prices within a spell as price change when it is larger than one cent or 1% in absolute value (or more than 0.5% for prices larger than \$5). The purpose of this restriction is to remove small price changes that could arise from rounding errors, given that the weekly price is constructed as revenues over units sold.

For each product $\times$ county $\times$ month, I obtain the number of regular price changes across stores in the third week of the month. The monthly frequency of adjustment is given by this number over the total price observations in the third week of the month. The size of a regular price change is the log difference between the price in the period identified as having a regular price change and the price in the preceding period. At the product $\times$ county $\times$ month level, it is the unweighted average across stores in the third week of the month.

[47] That is, I do not impute missing values using the preceding or subsequent price.

D Generalized search problem

Stiglitz (1987) argued that search costs might decrease with the number of searches if, for example, stores are clustered around a shopping area so getting there entails most of the cost. Conversely, search costs might increase with the number of searches if stores are geographically dispersed. In addition, the scarcity of time imposes a natural cap on the number of stores that shoppers can visit. In this section, I let the marginal search cost change with the number of searches and limit the number of draws. These generalizations serve two purposes: first, to match the empirical findings on shopping behavior, and second, to shed further light on the interaction between inflation, search, and welfare.

For clarity, we drop the subscript k. Assume there is a maximum of $\overline{S}$ draws and the first one is free so all shoppers draw at least one offer. Recall of previous offers is allowed.[48] The value that offer p has for γ after $S \in \left\{2, \ldots, \overline{S}\right\}$ draws is given by $B_S(p, \gamma)$, where

$$
\begin{aligned}
B_1 &= \max\left\{0, \mathbb{E}\left[B_2(p,\gamma)\right]\right\} = \mathbb{E}\left[B_2(p,\gamma)\right] > 0; \\
B_S(p,\gamma) &= \max\left\{V(p), \int_{\underline{p}}^{p} B_{S+1}(p',\gamma)\, dF(p') + B_{S+1}(p,\gamma)\left[1 - F(p)\right] - \gamma \times c(S)\right\}; \\
B_{\overline{S}+1}(p,\gamma) &= V(p).
\end{aligned}
$$

The solution to this problem is a sequence of reservation prices for each type γ,

$$
r_S = R_S(\gamma) \equiv \Gamma_S^{-1}(\gamma) \geq p \quad \forall S \in \left\{2, \ldots, \overline{S}\right\},
$$

where the search cost thresholds $\Gamma_S(p)$ indicate the types that accept each p.

To obtain the retailer-level demand function, we only need to solve for the sequence $\left\{\Gamma_S(p)\right\}_{S=2}^{\overline{S}}$. The solution will depend on the slope of the monotonic marginal cost function $c(S)$. If the marginal cost of search is constant, that is $c(S) = \alpha$, reservation prices are also constant and the search cost thresholds satisfy equation (6). If the marginal search cost increases (decreases) with S, given that $V(p)$ is strictly decreasing,

[48] If the number of draws is limited and recall of previous offers is not allowed, a constant marginal search cost leads to reservation prices increasing with S – opposite to what the empirical evidence suggests. Everything else equal, allowing recall produces reservation prices constant in S (see Lippman and McCall, 1976 for proof). Moreover, as discussed below, decreasing marginal search costs always generate decreasing reservation prices, facilitating the calibration of the model.

it is straightforward to prove that reservation prices increase (decrease) with S. Then, we can solve for the search cost thresholds recursively using that $B_S(p, \Gamma_S(p)) = V(p)\ \forall S \in \{2, \ldots, \overline{S}\}$, and noting that $\Gamma_{S-1}(p) \geq \Gamma_S(p)\ \ \forall S$ if reservation prices increase and $\Gamma_{S-1}(p) \leq \Gamma_S(p)\ \ \forall S$ if they decrease with S. For $c(S)$ increasing, the solution satisfies

$$\Gamma_S(p) = c(S)^{-1} \int_{\underline{p}}^{p} [V(u) - V(p)]\, dF(u) \quad \forall S \in \{2, \ldots, \overline{S}\},$$

and it is easy to see that, as S increases, $\Gamma_S(p)$ becomes flatter, meaning that a shopper of any type γ picks higher reservation prices. For $c(S)$ decreasing, the solution satisfies

$$\begin{aligned} \Gamma_S(p) &= \Gamma_{S+1}(u_S) \quad \forall S \in \{2, \ldots, \overline{S} - 1\}; \\ \Gamma_{\overline{S}}(p) &= c(\overline{S})^{-1} \int_{\underline{p}}^{p} [V(u) - V(p)]\, dF(u), \end{aligned}$$

where the sequence $\{u_S\}_{S=2}^{\overline{S}-1}$ solves

$$\begin{aligned} J_S(p) &= K_S(u_S); \\ J_S(p) &\equiv \sum_{s=1}^{\overline{S}-S} \int_{\underline{p}}^{p} [\Gamma_{\overline{S}}(p) - \Gamma_{\overline{S}}(u)]\, d(1 - F(u))^s + (\overline{S} - S + 1) \times \Gamma_{\overline{S}}(p); \\ K_S(u_S) &\equiv J_S(u_S) + \Gamma_{S+1}(u_S) \frac{c(S)}{c(\overline{S})} - \Gamma_{\overline{S}}(u_S). \end{aligned}$$

Retailer-level demand function Consider a retailer charging relative price p for the good. Since the first draw is free, $g(\gamma)$ shoppers type γ search at least once and draw p. Their probability of success is $F(R_2(\gamma))$; those with type $\gamma \geq \Gamma_2(p)$ accept p and those with $\gamma < \Gamma_2(p)$ continue searching. Therefore, $[1 - F(R_2(\gamma))]\, g(\gamma)$ shoppers obtain p on their second draw, and their probability of success is $F(R_3(\gamma))$; those with $\gamma \geq \Gamma_3(p)$ accept it while those with $\gamma < \Gamma_3(p)$ continue searching. On their third draw, $[1 - F(R_3(\gamma))]\, [1 - F(R_2(\gamma))]\, g(\gamma)$ shoppers get p; etc. On their final draw, $\bar{S}$, $\prod_{j=2}^{\bar{S}} [1 - F(R_j(\gamma))]\, g(\gamma)$ shoppers get the offer p and accept it if it is the minimum among all draws $\{1, \ldots, \bar{S}\}$. Summing up over those who accept the offer, the extensive margin of demand for a retailer charging p is given by:

$$\begin{aligned} N(p) &= \int_{\Gamma_2(p)}^{\bar{\gamma}} g(\gamma)\, d\gamma + \sum_{S=2}^{\overline{S}-1} \int_{\Gamma_{S+1}(p)}^{\bar{\gamma}} \prod_{j=2}^{S} \left[1 - F(R_j(\gamma))\right] g(\gamma)\, d\gamma \\ &\quad + \int_{\underline{\gamma}}^{\bar{\gamma}} \prod_{j=2}^{\overline{S}} \left[1 - F(R_j(\gamma))\right] g(\gamma)\, d\gamma. \end{aligned}$$

For computational purposes, it is convenient to express $N(p)$ in terms of $\Gamma_2(p)$:

$$\begin{aligned} N(p) &= \int_{\Gamma_2(\infty)}^{\bar{\gamma}} g(\gamma)\, d\gamma + \int_{p}^{\infty} g(\Gamma_2(r))\, d\Gamma_2(r) \\ &\quad + \sum_{i=2}^{\bar{S}-1} \int_{\Gamma_2^{-1}[\Gamma_{i+1}(p)]}^{\bar{p}} \prod_{j=3}^{i} \left[1 - F(R_j(\Gamma_2(r)))\right] \left[1 - F(r)\right] g(\Gamma_2(r))\, d\Gamma_2(r) \\ &\quad + \int_{\Gamma_2^{-1}(\underline{\gamma})}^{\bar{p}} \prod_{j=3}^{\bar{S}} \left[1 - F(R_j(\Gamma_2(r)))\right] \left[1 - F(r)\right] g(\Gamma_2(r))\, d\Gamma_2(r). \end{aligned}$$

Expected search outcomes The probability of γ drawing S offers is $F(R_2(\gamma))$ for $S = 1$,

$$F(R_{S+1}(\gamma)) \prod_{j=2}^{S} \left[1 - F(R_j(\gamma))\right]$$

for $1 < S < \overline{S}$, and $\prod_{j=2}^{\overline{S}} \left[1 - F(R_j(\gamma))\right]$ for $S = \overline{S}$. For any maximum number of draws $\overline{S} = \{1, 2, 3, \ldots\}$, monotonic marginal search cost function $c(S)$, relative posted-price distribution F, search-cost type γ, and associated search rules $R_S(\gamma)$ $\forall\, S = \left\{2, 3, \ldots, \overline{S}\right\}$, we can compute

- the expected search costs

$$\begin{aligned} \breve{C}(\gamma) \equiv \gamma \mathbb{E}\left[C(S) \,|\gamma\right] &= \gamma \times \sum_{S=2}^{\overline{S}-1} \left\{ C(S) \times F(R_{S+1}(\gamma)) \prod_{j=2}^{S} \left[1 - F(R_j(\gamma))\right] \right\} \\ &\quad + \gamma \times C\left(\overline{S}\right) \times \prod_{j=2}^{\overline{S}} \left[1 - F(R_j(\gamma))\right], \end{aligned}$$

  where $C(S) = \alpha^{-1} \sum_{j=2}^{S} c(j)$;

- the expected number of draws:

$$\begin{aligned}\breve{S}(\gamma) \equiv \mathbb{E}[S|\gamma] &= 1 \times F(R_2(\gamma)) + \sum_{S=2}^{\overline{S}-1}\left\{S \times F(R_{S+1}(\gamma)) \prod_{j=2}^{S}[1-F(R_j(\gamma))]\right\} \\ &\quad + \overline{S} \times \prod_{j=2}^{\overline{S}}[1-F(R_j(\gamma))];\end{aligned}$$

- the expected surplus:

$$\begin{aligned}\breve{V}(\gamma) \equiv \mathbb{E}[V(p)|\gamma] &= \int_{\underline{p}}^{R_2(\gamma)} V(p)\,dF(p) \\ &\quad + \sum_{S=2}^{\overline{S}-1}\left\{\int_{\underline{p}}^{R_{S+1}(\gamma)} V(p)\,dF(p) \prod_{j=2}^{S}[1-F(R_j(\gamma))]\right\} \\ &\quad + \int_{\underline{p}}^{\overline{p}} V(p)\,d\left\{1-[1-F(p)]^{\overline{S}}\right\} \times \prod_{j=2}^{\overline{S}}[1-F(R_j(\gamma))];\end{aligned}$$

- and the expected price paid:

$$\begin{aligned}\breve{p}(\gamma) \equiv \mathbb{E}[p|\gamma] &= \int_{\underline{p}}^{R_2(\gamma)} p\,dF(p) \\ &\quad + \sum_{S=2}^{\overline{S}-1}\left\{\int_{\underline{p}}^{R_{S+1}(\gamma)} p\,dF(p) \prod_{j=2}^{S}[1-F(R_j(\gamma))]\right\} \\ &\quad + \int_{\underline{p}}^{\overline{p}} p\,d\left\{1-[1-F(p)]^{\overline{S}}\right\} \times \prod_{j=2}^{\overline{S}}[1-F(R_j(\gamma))].\end{aligned}$$