

Expected Inflation and Welfare: The Role of Consumer Search

Francisca Sara-Zaror^{*‡}

November 14, 2021

[\[Click here for the latest version\]](#)

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.

*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 Adrien Auclert, Kieran Larkin, Lun Li, Simon Mongey, Harald Uhlig, and Joe Vavra, as well as participants of the Capital Theory and AMT seminars at the University of Chicago, and the Essex/RHUL/Bristol Junior SaM Workshop for helpful comments.

[†]University of Chicago. E-mail: mfransara@uchicago.edu

[‡]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, the micro-prices underlying the CPI are assigned to a product category, not to a specific good. Thus, cross-sectional comparison of prices within a category would not take product heterogeneity into account.

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 2017. 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 200 geographically dispersed markets. With these data, I can compare prices of products with identical barcodes sold across different retailers on the same date and geographic market. 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. On the other hand, 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 -20% and 30%. Thus, my approach is to exploit cross-sectional variation studying the relationship between product-level inflation and price dispersion.

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 is consistent with my findings and exhibits realistic price-setting behavior of retailers. In the model, the representative household has a worker who supplies labor and a measure-one continuum of shoppers who purchase a homogeneous good. A measure-one continuum of retailers produces and sells the good using labor. Retailers set nominal prices and face a menu cost; they adjust because real prices erode at the deterministic inflation rate and idiosyncratic shocks hit their productivity. The non-degenerate and time-invariant cross-sectional distribution of real prices results from both inflation and idiosyncratic shocks.

Every period, a continuum of shoppers enter, search, purchase, and leave the product market. They take the real-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 good from one of many nearly identical retailers. Therefore, shopping behavior determines the equilibrium demand curve. 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 to set a price low enough to attract searchers. It turns out that, since real 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 real 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. Thus, these retailers can even have incentives to increase their prices. As a result, price dispersion rises sharply.

I jointly calibrate the parameters of the model to match the average markup, the absolute size and frequency of price changes, and the Υ -shaped pattern of price dispersion in inflation. The latter is crucial to identify the parameters of the search-cost distribution. When search is relatively cheap, search intensity and price dispersion at zero inflation are relatively high. Therefore, deviations from zero inflation have a small impact on equilibrium price dispersion. Conversely, when search is relatively expensive, a positive inflation rate increases price dispersion and the returns to search considerably.

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 3% less than a household visiting only one store. Moreover, simulated data from the calibrated model closely matches this empirical result, which I did not target.

In my model, the welfare effects of 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 aggregate 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 (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 $[s, S]$.¹ In this case, the coefficient of variation is increasing and concave in the max-min price ratio, S/s .² 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 is in line with my findings. On the other hand, 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,

¹Caplin 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.

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

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.

On the other hand, 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, in contrast to the empirical evidence in my paper.

Head and Kumar (2005) 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.³

In the 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

³In 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.

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.⁴ 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.⁵ 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; 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

⁴In models with inflation and non-sequential search (Head et al., 2012; Burdett and Menzio, 2018), firms follow mixed pricing strategies, making the concept of a firm-level price spell elusive.

⁵Mostly 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 Matthey (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 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, real 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).⁶ Moreover, in the presence of a zero lower bound on nominal interest rates, the benefits from a positive level of inflation more than offset

⁶Burstein 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.

such costs (Blanco, 2021).⁷

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 the menu-cost model with endogenous consumer search that reproduces the Υ -shaped pattern. Section 4 explains the intuition behind this result and shows the model fit to the data. Section 5 presents evidence on shopping behavior and inflation that supports the mechanism. Section 6 discusses the welfare implications of inflation in the calibrated model. Section 7 concludes.

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 dataset is available only for a relatively low and stable aggregate inflation period, implying the variation for a time-series analysis is insufficient. On the other hand, the variation in market- and product-level inflation rates is substantial. Thus, I identify the relationship between inflation and price dispersion exploiting cross-sectional variation and discuss the limitations of this estimation strategy in the model section.

⁷Danziger (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.

2.1 Data

The NielsenIQ Retail Scanner Data are available for the 2006-2017 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 1,075 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).⁸ Over 35,000 stores are located across the U.S., classified into around 200 geographic markets (Designated Market Area – DMA).⁹ Each store can be associated to a retail chain. I use the Retail Scanner dataset to compute disaggregate inflation measures and retailer-level statistics on price dispersion.¹⁰

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. Therefore, we can compute price dispersion controlling for product heterogeneity. Moreover, each product is linked to a store, enabling us to control for sellers’ heterogeneity; 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.

⁸I 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.

⁹A DMA region is defined as a non-overlapping group of counties in which the commercial TV stations in the Metro/Central area achieve the largest audience share. Each DMA has, on average, 15 counties. Examples are Chicago IL, Milwaukee WI, and Columbus OH.

¹⁰In the robustness checks, I define a geographic market as a 3-digit ZIP code and distinguish between retail chains and stores when computing the relevant measures.

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 \times geographic market \times month level. Then, I aggregate product-level statistics into product modules – or categories – using total annual sales as weights. Therefore, the unit of observation in my analysis is a category \times geographic market \times month (c, m, t) triple.

We work with category- instead of product-level statistics for at least two reasons. The first is to make the dataset computationally manageable.¹¹ The second is to overcome the issue that these products are typically short lived and that we want to carry out a comparative static analysis.¹² Therefore, if we keep a panel of products (as opposed to categories) present in the data for several years, we might be selecting goods that are not representative of the rest of the economy (e.g., products with large market shares).¹³

To minimize the incidence of missing values, I only keep product \times store pairs that appear in every month of a given year, and I use monthly instead of weekly prices. To compute meaningful measures of price dispersion, I require that each product is sold by at least five stores in each market and week.

Inflation I define the price that store i in market 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 market m ,

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

¹¹The RMS dataset has information about more than 3 million UPCs, sold by 35,000 stores in over than 620 weeks; its estimated size is over 7 TB.

¹²Argente et al. (2018) show one third of all products in the U.S. are either created or destroyed in a given year, and more than 20% are less than one year old.

¹³Although a product-level analysis for top-selling brands in each category qualitatively reproduces the main category-level findings of this paper.

where $N_{km,t}$ is the number of stores.¹⁴ Category \times market-level monthly inflation corresponds to

$$\pi_{cm,t} = \sum_{k=1}^{K_{cm}} \omega_{km} \pi_{km,t},$$

where K_{cm} is the number of products in the category and market, and ω_{km} are product weights from annual sales:

$$\omega_{km} = \frac{\sum_{t=1}^{12} \sum_{i=1}^{N_{km,t}} Rev_{ikm,t}}{\sum_{k=1}^{K_{cm}} \sum_{t=1}^{12} \sum_{i=1}^{N_{km,t}} Rev_{ikm,t}}.$$

To validate this inflation measure, I compute a sales-weighted average over all food-related category \times market pairs and compare it with the official statistics. Figure A.1 shows aggregate inflation for food categories 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, market, and month¹⁵:

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

To obtain price dispersion at the category \times market level, I aggregate product-level measures weighting by total sales in the corresponding year:

$$\sigma_{cm,t} = \sum_{k=1}^{K_{cm}} \omega_{km} \sigma_{km,t}.$$

Using the same methodology, I compute alternative measures of price dispersion — the IQR, the max-min, the 90-10, the 90-50, and the 50-10 ratios of the price distribution for each (c, m, t) , as well as the skewness and kurtosis of the price distribution.

¹⁴Weighting by annual stores sales or using week-level prices does not affect the main results.

¹⁵Weighting prices using annual stores sales, using weekly instead of monthly prices, regular instead of posted prices, or removing persistent differences across sellers yield similar qualitative results.

2.3 Estimation and results

In the theoretical section, we study how price dispersion varies with inflation. These comparative statics assume inflation is constant and that 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.¹⁶ Annual inflation and price dispersion for each category \times market correspond to the average across months for each year:

$$\bar{\pi}_{cm,t} = \frac{\sum_{j=1}^{12} \pi_{cm,j}}{12}; \quad \bar{\sigma}_{cm,t} = \frac{\sum_{j=1}^{12} \sigma_{cm,j}}{12}.$$

Table A.1 shows descriptive statistics for these measures and for a sales-weighted average inflation. The final sample for estimation includes 138,485 category \times market pairs for 1,042 categories in 204 geographic markets. Each pair is present in the data for an average of 8.8 years and contains information for an average of 50 stores. Disaggregate inflation exhibits significantly larger variation than aggregate inflation, with a coefficient of variation (CV) of 4.647 and 1.267, respectively.

To understand the relationship between price dispersion and inflation, I start by constructing a binned scatterplot: I divide annual disaggregate inflation $\bar{\pi}_{cm,t}$ into 100 equally sized bins and obtain average price dispersion $\bar{\sigma}_{cm,t}$ 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 category, market, or time dimensions. In particular, I estimate

¹⁶Simulations from the calibrated model suggest the economy takes five months to adjust between two equilibria with different inflation rates. On the other hand, in the appendix, I present the results when assuming convergence to the ergodic distribution in 6, 24, or 36 months.

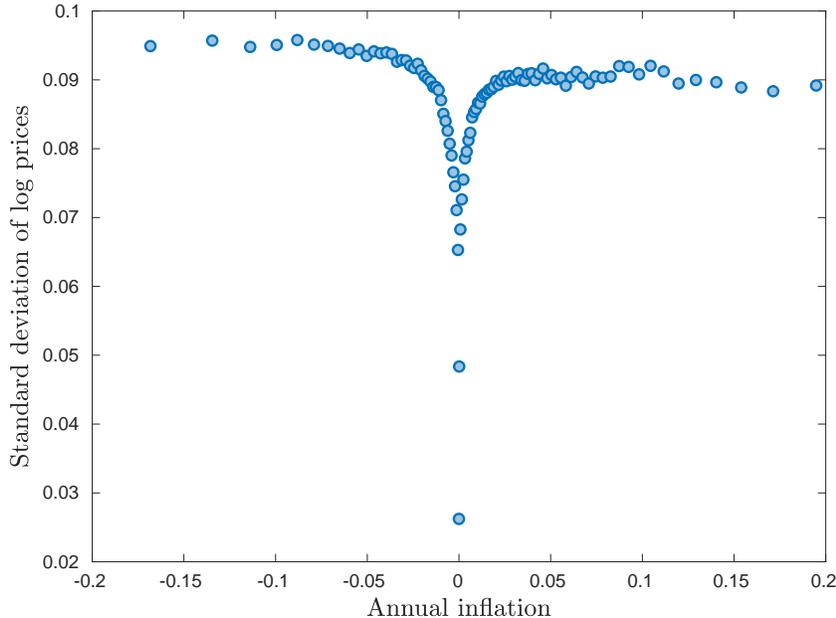


Figure 1: Price dispersion and inflation, raw data

Each dot corresponds to average price dispersion for each of 100 equally sized inflation bins. The unit of observation is a category \times market \times year, for 1,219,414 in total.

the following non-parametric regression of price dispersion on inflation:

$$\bar{\sigma}_{cm,t} = \sum_{n=1}^{100} \beta_n \mathbf{1}_{\{\bar{\pi}_{cm,t} \in B_n\}} + a_{c,t} + b_{m,t} + \alpha \log N_{cm,t} + \epsilon_{cm,t}. \quad (1)$$

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: $a_{m,t}$ and $b_{c,t}$ are market-year and category-year fixed effects, respectively, and $N_{cm,t}$ denotes the average number of stores across goods in the category \times market pair. This specification implies we identify $\{\beta_n\}_{n=1}^{100}$ by exploiting the cross-category inflation variation within each geographic market and year, taking into account unobservable category-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

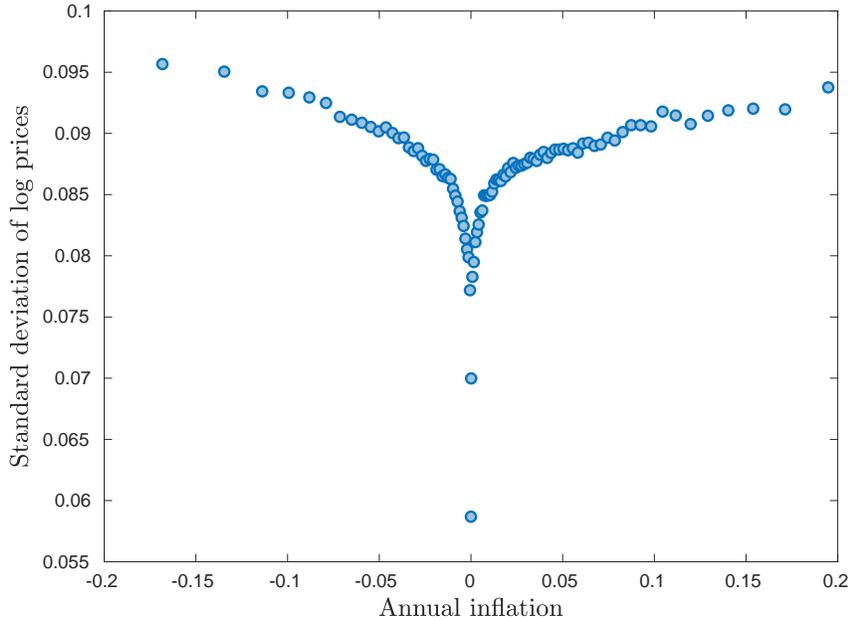


Figure 2: Price dispersion and inflation

The dots correspond to average price dispersion conditional on covariates for each of 100 equally sized inflation bins as predicted by equation (1). The unit of observation is a category \times market \times year, for 1,219,414 in total.

inflation is prevalent in the data, even when comparing different categories in a given market-year, or a category across markets for the same year. 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% (Figure A.2).

Moreover, the Υ -shaped relationship is robust to controlling linearly (Figure A.3) or non-parametrically (Figure A.4) for the number of stores; that is, binned scatterplots of $\bar{\pi}_{cm,t}$ and $\bar{\sigma}_{cm,t}$ by quartile of the number of stores display the same pattern.

On the other hand, price dispersion is approximately symmetric around zero. Figure A.5 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

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 category \times market pair to control for within-pair error correlation over time. Because the confidence band in Figure A.6 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 market and year separately. I start by assuming price dispersion is a symmetric, continuous, and differentiable function of inflation, which I identify exploiting within-market-year variation across categories¹⁷:

$$\sigma_{cm,t} = f_{m,t}(|\pi_{cm,t}|) + \alpha \log N_{cm,t} + \epsilon_{cm,t}. \quad (2)$$

Under these assumptions, I test two hypotheses for each market-year function $f_{m,t}$. The first is that the function is monotonically non-decreasing; the second is that it is concave¹⁸:

$$\begin{aligned} \inf_{\pi > 0} f'_{m,t}(\pi) &\geq 0; \\ \sup_{\pi > 0} f''_{m,t}(\pi) &\leq 0. \end{aligned}$$

The top panel in Table A.2 shows the first hypothesis is rejected for 23% and the second for 18% of the total market-year combinations. When we restrict absolute inflation to values lower than 2%, the same hypotheses are rejected for only 5% and 6%, respectively, of the market-year pairs. The bottom panel of the table shows similar results when we estimate the relationship for each category and year, exploiting cross-market variation.

Future inflation Up to this point, we have analyzed the relationship between current inflation and price dispersion. In the model, nonetheless, price-setting behavior

¹⁷I 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.

¹⁸Both tests require differentiability of $f_{m,t}$. The estimates using pooled data suggest this function might be non-differentiable at zero. I assume $f_{m,t}$ is symmetric around zero to overcome this issue.

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.7 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.

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. The top panel of Figure A.8 shows the relationship between the max-min ratio of the price distribution and inflation, which, again, is Υ -shaped. The same pattern holds for the 90-10 ratio (bottom panel), 90-50, and 50-10 ratios of the price distribution and inflation (Figure A.9).

3 Model

In this section, I develop a one-good, one-sector monetary model that generates a Υ -shaped relationship between price dispersion and aggregate inflation as in the previous section. 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 shoppers who purchase a homogeneous good.

The consumption good is sold by a measure-one continuum of monopolistically competitive retailers. Retailers are infinitely lived, set nominal prices, and have a production technology linear in labor. A fixed cost is incurred when changing nominal prices, and retailers would like to adjust them for two reasons: (i) they face idiosyncratic transitory shocks to their productivity, and (ii) the nominal wage is increasing at a deterministic rate. These elements change retailers' production costs and therefore their desired prices.¹⁹ At the same time, both wage inflation and idiosyncratic shocks

¹⁹Equivalently, 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.

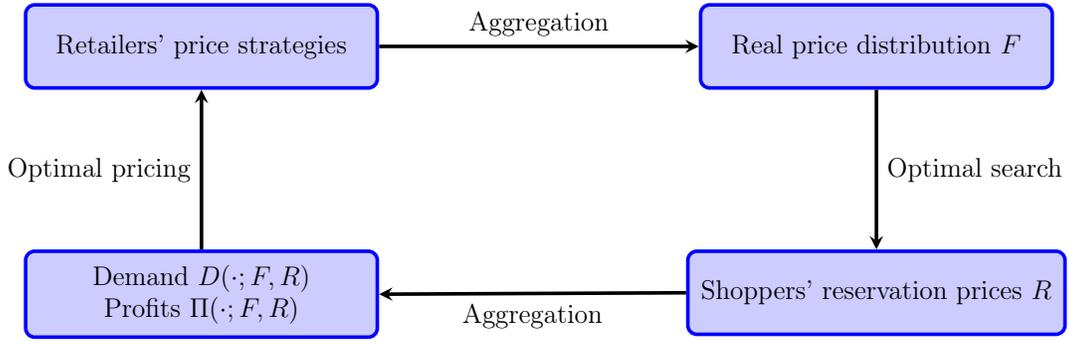


Figure 3: General model structure

are key to produce a non-degenerate and stationary cross-sectional distribution of real prices. 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 product market. Shoppers take the real-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 real reservation price at which they are indifferent between buying and searching again. Each buyer ends up purchasing the good 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-sloping demand curve for the retailer. Through 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, the price level of the good and nominal wages grow at the same rate. We assume such an inflation rate is the policy parameter that the monetary authority can control.

Relationship between model and data In the model, we study the relationship between aggregate inflation, that is, the growth rate of an economy-wide nominal production cost, and price dispersion. On the other hand, the empirical evidence 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 appendix C, I extend the baseline model to multiple goods and sectors. In such a model, sectoral variation in inflation is explained by sector-specific productivity growth rates: the mapping between the model and data is direct. Nevertheless, I show the sectoral equilibrium in this extended model is isomorphic to the aggregate equilibrium in the simple model. Thus, to focus on the mechanisms behind the Υ shape between inflation and price dispersion, I analyze the simple model for the remainder of the paper.

3.1 Retailers

A measure-one continuum of monopolistically competitive retailers is indexed by i . These retailers are infinitely lived, with discount rate ρ . Money is used as the unit of account, so retailers set their prices $P_{i,t}$ in nominal terms. The production technology of each retailer is linear in labor $l_{i,t}$, and the retailer-specific labor productivity is given by $v_{i,t}$. We assume the retailer’s productivity follows the mean-reverting process:

$$d \log v_{i,t} = -\rho_v \log v_{i,t} dt + \sigma_v dZ_{i,t}, \quad \rho_v > 0,$$

where $Z_{i,t}$ 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 all the aggregate prices, particularly the nominal wage, are growing at a constant rate π . Because real prices are defined with respect to the nominal wage, that is, $p_{i,t} \equiv P_{i,t}/W_t$, a fixed nominal price implies a real price eroding at a rate π . 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 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 $D(p)$. The instantaneous nominal profit function for a retailer charging a nominal price $P_{i,t}$, given

the nominal wage W_t and the stochastically determined labor productivity $v_{i,t}$, is

$$\Pi_t(P_{i,t}, v_{i,t}) = \left(P_{i,t} - \frac{W_t}{v_{i,t}} \right) \times D \left(\frac{P_{i,t}}{W_t} \right).$$

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

$$\Pi(p_{i,t}, v_{i,t}) = \left(p_{i,t} - \frac{1}{v_{i,t}} \right) \times D(p_{i,t}).$$

Let $\psi(p_{i,0}, v_{i,0})$ denote the present value of a retailer that begins at $t = 0$ with the relative price $p_{i,0}$ and productivity $v_{i,0}$. Given the strategies of shoppers and other retailers, this retailer chooses a time $T_i \geq 0$ to adjust and a reset price p'_i so as to solve

$$\psi(p_{i,0}, v_{i,0}) = \max_{T_i} \mathbb{E} \left\{ \int_0^{T_i} e^{-\rho t} \Pi(p_{i,t}, v_{i,t}) dt + e^{-\rho T_i} \max_{p'_i} [\psi(p'_i, v_{i,T_i}) - \kappa] \right\}.$$

This time-invariant Bellman equation is standard in the literature, and its solution is well known.²⁰ The optimal policy of a retailer with productivity level $v_{i,t}$ is to leave its nominal price unchanged if $p_{i,t}$ is between $p_L(v_{i,t})$ and $p_U(v_{i,t})$. If the real price hits any of these bounds, the retailer pays the menu cost κ and adjusts to $\hat{p}(v_{i,t})$.

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

$$F(p) = \int_{\underline{p}}^p \int_v \phi(x, v) dv dx, \quad p \in [\underline{p}, \bar{p}],$$

where the bounds $\{\underline{p}, \bar{p}\}$ are determined in equilibrium by the retailers' optimal pricing behavior.

²⁰In 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.

3.2 Representative household

A representative household has a worker and a measure-one continuum of shoppers buying a homogeneous good. The worker supplies labor and shares income equally across shoppers. Each instant t , the shoppers enter, search, purchase, and leave the product market. 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.^{21 22}

Shoppers have incomplete information about the prices: they know the posted-price distribution F , but not the price that each seller charges. Because the good is homogeneous and price dispersion exists, buyers have incentives to search for the lowest price. Shoppers j search sequentially for a retailer i . They receive a first price quote for free but need to pay γ_j labor units for each subsequent draw. We assume the search cost is heterogeneous, randomly distributed across shoppers according to G .²³ The shopper searches S_j times to find a retailer i from whom he buys q_j units of the good at the nominal price $P_{i(j)}$.

Given the search outcomes of shoppers, $P_{i(j)}$ and S_j , the household head chooses labor supply L and quantities of the consumption good to solve

$$\begin{aligned} & \max_{q_j, L} \int_0^1 \left(\frac{q_j^{1-\frac{1}{\eta}}}{1-\frac{1}{\eta}} \right) dj - L \\ \text{s.t.} \quad & \int_0^1 P_{i(j)} q_j dj + W \int_0^1 \gamma_j (S_j - 1) dj = WL + \mathcal{D}, \end{aligned}$$

where \mathcal{D} are the dividends the household gets from the retailers. We assume $\eta > 0$.

²¹Under 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.

²²The assumption of instantaneous consumers also implies search happens quickly compared with real-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.

²³Heterogeneity in γ 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.

The solution to this problem implies the indirect utility of the household is

$$\begin{aligned}\mathcal{W} &= \int_0^1 [V(p_{j(i)}) - \gamma_j(S_j - 1)] dj + \frac{\mathcal{D}}{W}; \\ V(p) &= \frac{p^{1-\eta}}{\eta - 1},\end{aligned}\tag{3}$$

where $V(p)$ is the surplus that each shopper derives from buying $q(p) = -V'(p)$ units of the good at price p .

Shoppers take as given the retailers' pricing strategies F and the utility they contribute to the household as a function of the relative price, $V(p)$. The optimal strategy of the shoppers is to search for a retailer until they find an offer below their reservation price.²⁴ To compute the optimal stopping rule, we first define the value that an offer p has for a shopper j , $B_j(p)$. This value is the maximum between accepting such an offer and continuing to search after paying the search cost γ_j :

$$B_j(p) = \max \left\{ V(p), -\gamma_j + \int_{\underline{p}}^{\bar{p}} B_j(u) dF(u) \right\}.$$

Then, we define the reservation price r_j as the relative price that makes shopper j indifferent between buying and searching again:

$$V(r_j) = \int_{\underline{p}}^{\bar{p}} [\max \{V(p), V(r_j)\}] dF(p) - \gamma_j.$$

If this equation has a solution, the optimal policy for the shopper is to accept any offer p if $p \leq r_j$, and to continue searching otherwise.²⁵ Alternatively, the reservation price equates the expected benefit and the cost of the marginal search:

$$\Gamma(r_j) \equiv \int_{\underline{p}}^{r_j} [V(p) - V(r_j)] dF(p) = \gamma_j.\tag{4}$$

A buyer with $\gamma_j < \tilde{\gamma} \equiv \lim_{r \rightarrow \infty} \Gamma(r)$ 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 $\Gamma(r)$ is increasing and

²⁴Because 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.

²⁵Lippman and McCall (1976) present a detailed discussion and proof.

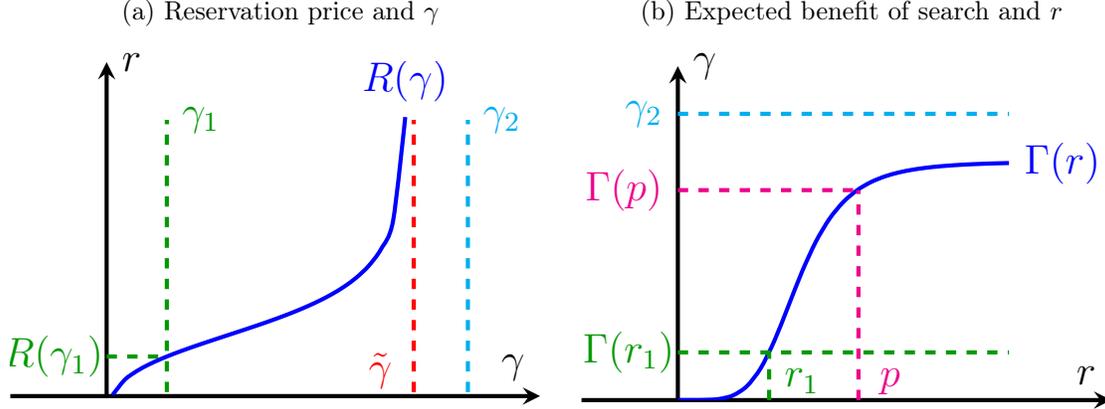


Figure 4: Optimal search strategies

continuously differentiable on $[\underline{p}, \infty)$, the reservation price is well defined and given by $r_j = R(\gamma_j) \equiv \Gamma^{-1}(\gamma_j)$, with $\underline{p} \leq r_j < \infty$.²⁶ Buyers with $\gamma_j > \tilde{\gamma}$ accept the first free offer as long as it does not exceed their maximum willingness to pay, so we let $R(\gamma_j) \rightarrow \infty$. These results are illustrated in Figure 4a.

The reservation price $R(\gamma_j)$ is an increasing function of the search cost. This finding implies all the buyer types γ_j for which $p \leq R(\gamma_j)$ 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, $\gamma_j \geq \Gamma(p)$, are the ones that accept such an offer (see Figure 4b).

3.3 Demand function

Assume $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ 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 $F(R(\gamma))$ retailers charging $p \leq R(\gamma)$. Each of these retailers retains $g(\gamma)/F(R(\gamma))$ buyers type γ , and each of them purchases $q(p)$ units of the good. From the previous section, we know the buyers with $\gamma \geq \Gamma(p)$ would accept the offer p . The equilibrium demand function for a retailer charging p results from aggregating the individual demands of all of those buyers:

$$D(p) = \int_{\Gamma(p)}^{\bar{\gamma}} q(p) \times \frac{g(\gamma)}{F(R(\gamma))} d\gamma. \quad (5)$$

²⁶By noting $V(p) = \int_p^\infty q(u) du$ and integrating by parts, it is possible to show that $\Gamma(r_j) = \int_{\underline{p}}^{r_j} q(p) F(p) dp$. Therefore, $\Gamma'(r_j) = q(r_j) F(r_j) \geq 0$.

We can define the total number of transactions at a given price p as

$$N(p) = \int_{\Gamma(p)}^{\tilde{\gamma}} \frac{g(\gamma)}{F(R(\gamma))} d\gamma.$$

Because the quantity each buyer purchases is independent of its type γ , we can express the endogenous demand curve (5) as the product of two components:

$$D(p) = q(p) \times N(p).$$

The first term on the right represents the *intensive margin* of demand, or the units that each buyer purchases at price p , $q(p) = p^{-\eta}$. The second and novel component for menu-cost models corresponds to the *extensive margin* of demand, $N(p)$. This function indicates the number of shoppers that a seller serves by setting a price p .

Using that $\Gamma'(p) = q(p) F(p)$, we can show the extensive margin is decreasing in the price:

$$N'(p) = -q(p) \times g(\Gamma(p)) \leq 0.$$

Intuitively, shoppers can take advantage of price dispersion and flee from higher prices. Nonetheless, in an equilibrium where $\gamma > \tilde{\gamma}$, 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

$$N(p) = \int_{\Gamma(p)}^{\tilde{\gamma}} \frac{g(\gamma)}{F(R(\gamma))} d\gamma + 1 - G(\tilde{\gamma}),$$

where $1 - G(\tilde{\gamma})$ 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} the minimum reservation price in the market, $R(\underline{r})$. A seller who sets $p \leq \underline{r}$ serves all shoppers and reaches the maximum number of transactions, $N(\underline{r})$.

The extensive margin plays a crucial role in determining the price elasticity of demand that a retailer faces. For $p > \underline{r}$, the extensive margin tends to increase the

overall elasticity of the demand curve:

$$\epsilon_D(p) = \begin{cases} \eta & p \leq \underline{r} \\ \eta + \epsilon_N(p) & p > \underline{r} \end{cases} \quad (6)$$

$$\epsilon_N(p) = \frac{p}{N(p)} \times q(p) g(\Gamma(p)) \geq 0. \quad (7)$$

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 (4) depend on the price distribution F , which in turn depends on inflation. Thus, the demand elasticity at each price changes with inflation through shopping behavior, affecting markups and efficiency. This feature is absent in the standard Dixit-Stiglitz framework used in the literature, where monopolistic competition is attained by assuming product differentiation.

3.4 Equilibrium

In a steady-state equilibrium, the optimal strategies correspond to the following:

1. reservation prices,

$$R(\gamma; F) = \begin{cases} \Gamma^{-1}(\gamma; F) & \gamma < \tilde{\gamma} \\ \infty & \gamma > \tilde{\gamma}, \end{cases}$$

that each buyer $\gamma \in [\underline{\gamma}, \tilde{\gamma}]$ chooses given the time-invariant posted-price distribution F ; and

2. pricing strategies,

$$\Psi(v; R, F, \pi) = \{p_L(v; R, F, \pi), p_U(v; R, F, \pi), \hat{p}(v; R, F, \pi)\},$$

that solve the dynamic retailer problem $\psi(p, v; R, F)$ for each v , given the search behavior of buyers (R), the strategies of competitors (F), and the level of inflation π . The optimal policy functions $\Psi(v; R, F, \pi)$ are consistent with an invariant distribution $\phi(p, v; R, \pi)$ and a posted-price distribution $F(p; R, \pi)$, for $p \in [\underline{p}(R, \pi), \bar{p}(R, \pi)]$.²⁷ Including π as an explicit argument in the retailer's policies will be useful when conducting comparative static analyses.

²⁷Retailer pricing strategies are monotonic in productivity v , implying $\underline{p}(R, \pi) = p_L(\bar{v}; R, \pi)$ and $\bar{p}(R, \pi) = p_U(\underline{v}; R, \pi)$ in equilibrium.

The equilibrium is solved as a fixed-point problem for any inflation level. Given a guess for the posted-price distribution $F_0(p)$, we compute the optimal search strategies of each buyer γ , $R_0(\gamma; F_0)$. Then, we aggregate search rules into the retailer-level demand function, $D_0(p; R_0, F_0)$, to obtain the endogenous profit function, $\Pi_0(p; R_0, F_0)$. We solve the retailers' problem to get the policy functions, $\Psi_0(v; R_0, F_0)$, and the associated invariant distribution, $\phi_1(p, v; R_0, F_0)$. Finally, we update the posted-price distribution by taking the marginal of ϕ over v :

$$F_1(p) = \int_{\underline{p}_1}^p \int_v \phi_1(x, v; R_0, F_0) dv dx, \quad p \in [\underline{p}_1, \bar{p}_1].$$

We repeat this process until $F_1(p) = F_0(p)$ for every p and find an equilibrium if the total number of transactions equals the total number of buyers:

$$\int_{\underline{p}}^{\bar{p}} N(p; R, F) dF(p; R) = 1.$$

The posted-price distribution weighted by the number of transactions, NdF , can be interpreted as the distribution of *transaction* prices.

Defining the transaction-weighted average markup in this economy is helpful:

$$\tilde{\mu} = \int_v \int_{\underline{p}}^{\bar{p}} pv N(p) \phi(p, v) dp dv.$$

Given the distributional assumptions on the productivity shocks, we can write $\tilde{\mu}$ as the sum of transaction-weighted average markups conditional on productivity, $\tilde{\mu}(v)$:

$$\begin{aligned} \tilde{\mu} &= \int_v \tilde{\mu}(v) \check{\phi}(v) dv; \\ \tilde{\mu}(v) &= v\tilde{p}(v); \\ \tilde{p}(v) &= \int_{\underline{p}}^{\bar{p}} p N(p) \hat{\phi}(p|v) dp. \end{aligned}$$

We denote the marginal distribution of the shocks by $\check{\phi}(v)$, the distribution of prices conditional on productivity by $\hat{\phi}(p|v)$, and the transaction-weighted average price conditional on productivity by $\tilde{p}(v)$.

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 inflation. To understand the mechanism behind this result, starting to analyze the equilibrium at zero inflation as a fixed point is useful. Figure 5 describes such an equilibrium.

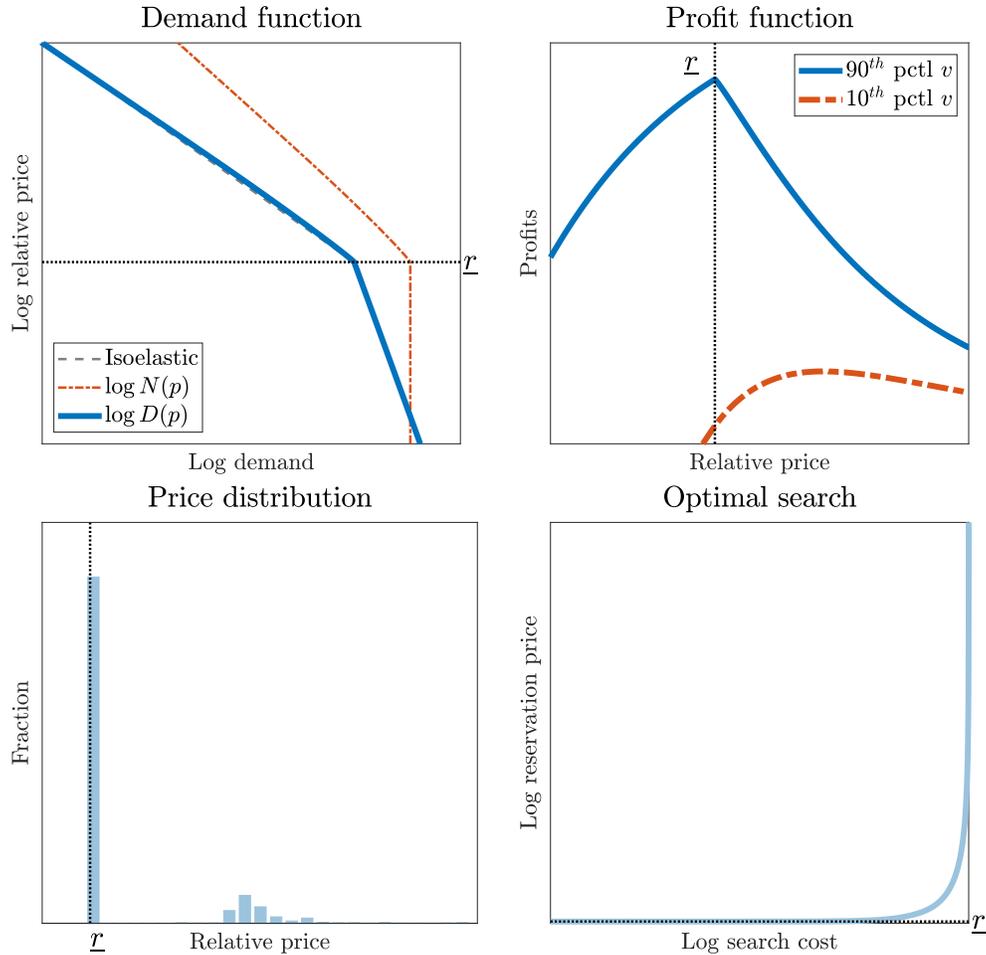


Figure 5: Equilibrium at $\pi = 0$

The figures describe an equilibrium with zero inflation. The value \underline{r} indicates the lowest reservation price among shoppers. The upper-right panel shows the profit function for retailers with 90th and 10th percentile productivity draws.

At zero inflation, retailers face a demand with a sharp kink at the lowest reservation price, \underline{r} . For retailers with higher productivity draws, \underline{r} maximizes static profits. If they charge a price slightly lower than \underline{r} , they will not attract more customers, because

they are already serving the maximum possible. On the other hand, if they charge a price slightly larger than \underline{r} , low-search-cost shoppers will reject the offer and search: the extensive margin of demand activates and demand elasticity increases discretely.

Moreover, staying at the kink is not only optimal but also feasible for these retailers: real prices are fixed, because zero inflation implies that real and nominal prices are equal. Therefore, more productive retailers bunch at \underline{r} , which generates a point mass in the price distribution. When inflation equals zero, price dispersion is relatively low: its only source is idiosyncratic productivity shocks, which manifest in dispersion at the top of the price distribution.

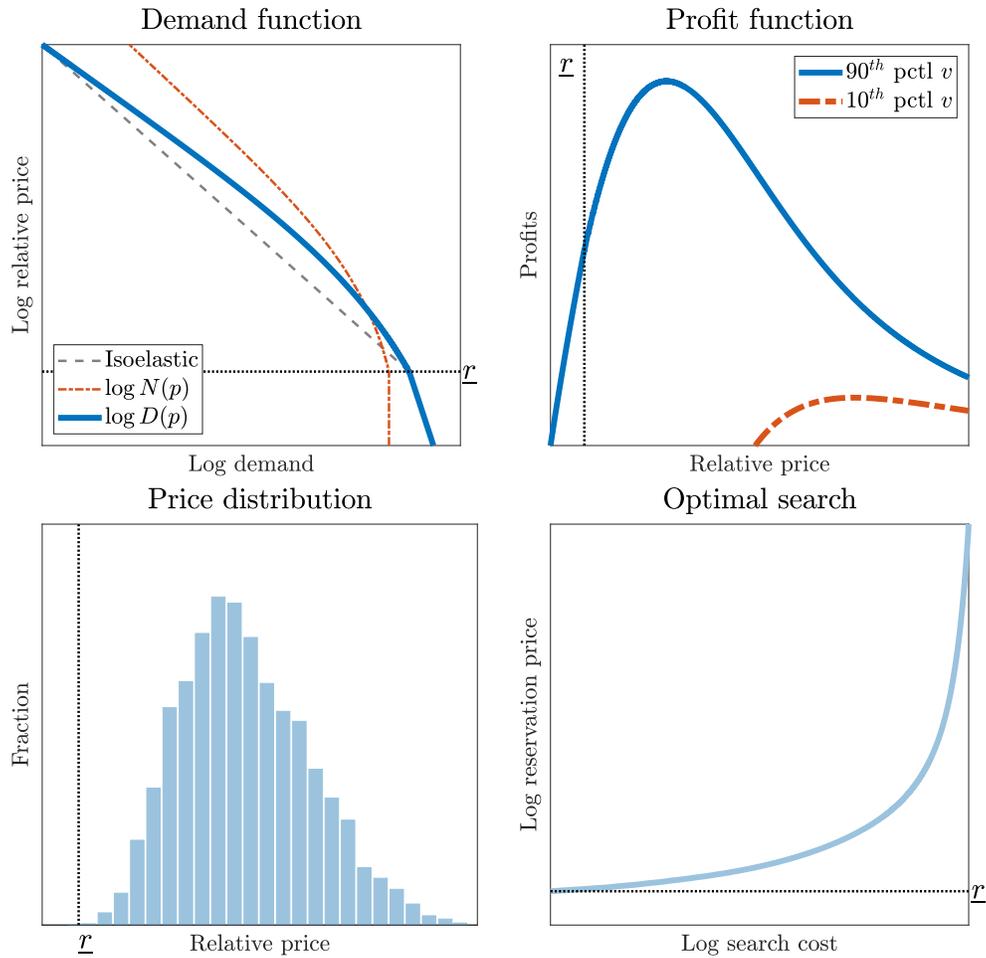


Figure 6: Equilibrium for $\pi > 0$

The figures describe an equilibrium with positive inflation. The value \underline{r} indicates the lowest reservation price among shoppers. The upper-right panel shows the profit function for retailers with 90th and 10th percentile productivity draws.

The point mass in the posted-price distribution implies shoppers who search will likely find an offer $p = \underline{r}$. Thus, only those shoppers with very high search cost will accept a price greater than \underline{r} ; those who actively search will only accept \underline{r} (lower-right panel in Figure 5). From the retailer’s perspective, a small increase in the price decreases discretely the number of shoppers who accept the offer, explaining the sharp kink in demand at \underline{r} .²⁸

For a positive level of inflation (Figure 6), real prices are continuously drifting downward.²⁹ To stay at the kink, retailers would need to pay the menu cost every period: bunching at any price is not optimal. Thus, retailers – particularly those with higher productivity draws – set wider adjustment bands, letting their real prices erode to save on menu costs. As a result, the minimum market price is lower than in the zero-inflation equilibrium, and dispersion at lower prices is higher.

Inflation-induced price dispersion increases the returns to search, so low-search-cost shoppers search more, and their reservation prices decrease.³⁰ To attract these shoppers, productive retailers charge even lower prices, and price dispersion increases further.

In sum, the Υ -shaped relationship between inflation and price dispersion is tightly related to consumer search and market power. When inflation is zero, the price dispersion generated by productivity shocks is not enough for significant search activity to occur. Thus, productive retailers set the maximum price that low-search cost shoppers will accept, and enjoy relatively high markups (Figure 7).

From zero to positive inflation, price dispersion increases discontinuously through two channels. The first is the menu-cost channel: relative prices drift downward continuously, so more productive retailers allow lower (and higher) price levels to pay the menu cost less often. Thus, price dispersion conditional on productivity increases. The second is the search channel: higher price dispersion increases the returns to search, and to attract more shoppers, more productive retailers charge even lower prices. With

²⁸The kinked demand complicates the analytical proof of uniqueness of the equilibrium at $\pi = 0$ considerably. To evaluate whether the equilibrium is unique, I solve the model for different guesses of the posted-price distribution $F_0(p)$ which I assume non-degenerate. I find convergence to the same equilibrium always exists.

²⁹Moving forward, I focus on the comparison between zero and positive inflation with monetary policy in mind.

³⁰On the other hand, more dispersion at the bottom of the distribution implies finding the minimum market price is harder: search intensity decreases as the individual search cost increases. Thus, retailers can increase their prices without a discrete loss in the number of consumers they serve. Inflation smooths out the kink in demand, making demand more inelastic at lower prices and more elastic at higher prices.

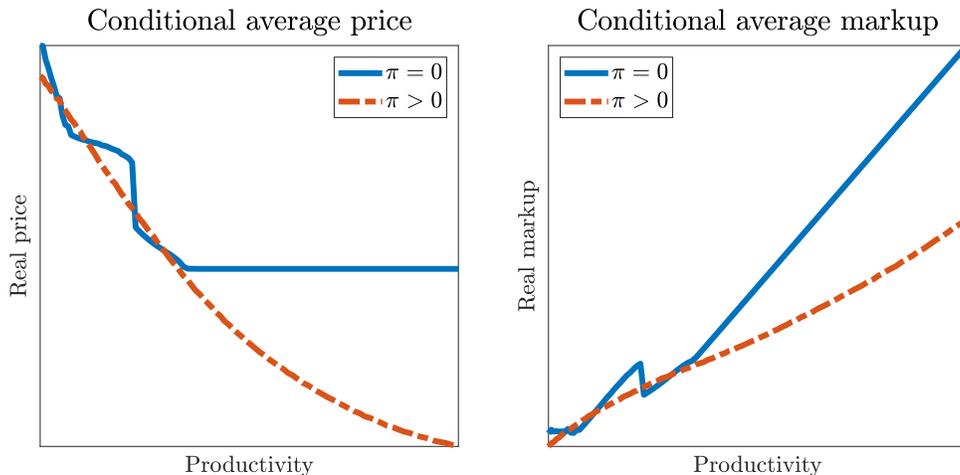


Figure 7: Optimal pricing and productivity

The left panel shows transaction-weighted average prices conditional on productivity, $\tilde{p}(v)$; the right panel, transaction-weighted average markups conditional on productivity, $\tilde{\mu}(v)$.

more competition, their markups decrease.

In the next subsection, we assess the model fit to the data by matching the Υ -shaped pattern in Figure 2.

4.1 Parametrization and calibration

We choose a flexible search cost distribution: $\gamma \sim \text{Beta}(\alpha, \beta)$, where $\alpha > 1$, $\beta > 1$, and $\gamma \in [0, 1]$.³¹ We jointly calibrate the size of menu cost κ , the persistence ρ_v , and volatility σ_v of idiosyncratic shocks, the shape parameters α and β of search-cost distribution, and the price elasticity of demand η of shoppers to match the following: (i) the standard deviation of log prices at $\pi = 0$, $\pi = \pm 2\%$, and $\pi = \pm 20\%$; (ii) the frequency of price changes and the average size of price changes at $\pi = 2\%$; and (iii) an average markup of 30% at $\pi = 2\%$.³² Following the literature, we set the monthly discount rate to $\rho = 0.96^{1/12}$.

The Υ -shaped relationship between inflation and price dispersion is key to calibrate the search-cost distribution. When this distribution has a small mean and variance (low α or high β), search activity is relatively cheap; search intensity and price dispersion

³¹Under these assumptions, the density of shoppers at $\underline{\gamma} = 0$ is zero, and the density smoothly increases from $\underline{\gamma} = 0$. Thus, the presence of the kink is unrelated to the specific distribution of the search cost.

³²In appendix B I explain how I compute additional pricing moments using the same scanner data.

Table 1: Additional moments, model fit

Moment	Model	Data	Targeted
Monthly frequency of adjustment at $\pi = 2\%$	14.41%	15.00%	Yes
Average abs. size of price changes at $\pi = 2\%$	17.86%	18.00%	Yes
Average markup at $\pi = 2\%$	1.30	1.30	Yes
Average fraction of price increases at $\pi = 2\%$	53.82%	51.00%	No

at $\pi = 0$ are relatively high. Thus, going from zero to positive inflation increases price dispersion by a relatively small amount. The opposite happens when search is relatively expensive (high α or low β): inflation generates a steep increase in price dispersion and in the returns to search.

Figure 8 shows the model matches the Υ -shaped relationship between inflation and the standard deviation of log prices in the data. At the same time, the calibration produces pricing moments and average markups consistent with the empirical evidence, as seen in Table 1. It is noteworthy that, although the absolute size of price changes is only targeted at $\pi = 2\%$, the model matches its relationship with inflation (bottom panel in Figure 8). The calibrated parameters are in Table 2.

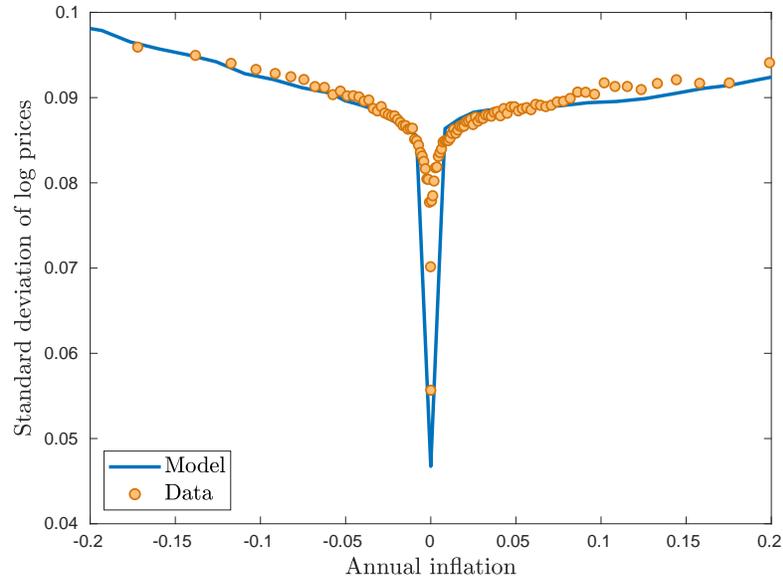
Table 2: Calibrated parameters

Parameter	Description	Value
κ	Menu cost	0.061
ρ_v	Mean-reversion rate of shocks	0.553
σ_v	Volatility of shocks	0.123
α	Shape of search cost distribution	1.056
β	Shape of search cost distribution	4.087
η	Intensive margin elasticity of demand	1.947

4.2 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 positive and low inflation. The first is that, on average, around half of the price changes are decreases. The second is that the absolute size of price changes is significantly larger than average inflation (18% versus 2%), and as absolute inflation increases, it does not increase as steeply as price dispersion.

(a) Price dispersion



(b) Absolute size of price changes

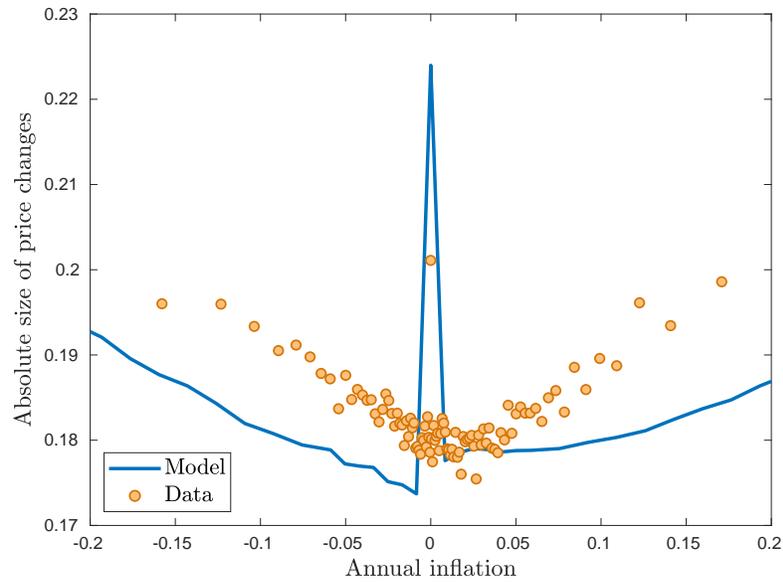


Figure 8: Pricing moments and inflation, model fit

The top figure shows the relationship between inflation and price dispersion in the calibrated model (dashed line) and in the data (dots); the bottom figure shows the equivalent for the absolute size of price changes.

Menu-cost models with idiosyncratic shocks but without search (Goloso and Lucas, 2007; Alvarez et al., 2019; Nakamura et al., 2018) match the additional pricing facts but miss the Υ -shaped relationship between inflation and price dispersion. Moreover, as Figure 9 shows, these models predict both price dispersion and the absolute size of the price changes are flat in inflation. In the model with search, the direct relationship between those pricing moments breaks.

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

$$\bar{\sigma}^2(p; \pi) = \mathbb{E} [\mathbb{V}(p|v; \pi)] + \mathbb{V} [\mathbb{E}(p|v; \pi)].$$

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 real prices erode faster, but because adjustments are costly, they optimally choose wider inaction regions $[p_L(v), p_U(v)]$. 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.

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 transaction-weighted average price $\tilde{p}(v)$ increases 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 real price under $\pi = 0$. When inflation deviates from zero, their real prices drift downward, so they optimally widen their adjustment bands: price dispersion conditional on productivity increases.

At the same time, those retailers increase nominal prices by less than the retailers that were adjusting under $\pi = 0$: the average absolute size of price changes decreases discretely. Nonetheless, this pricing moment increases with inflation and remains relatively flat around zero because the size of the adjustment bands stabilizes for low and positive inflation rates. In this way, the current theory matches price dispersion Υ -shaped in inflation and the absolute size of price changes flat in inflation (except at

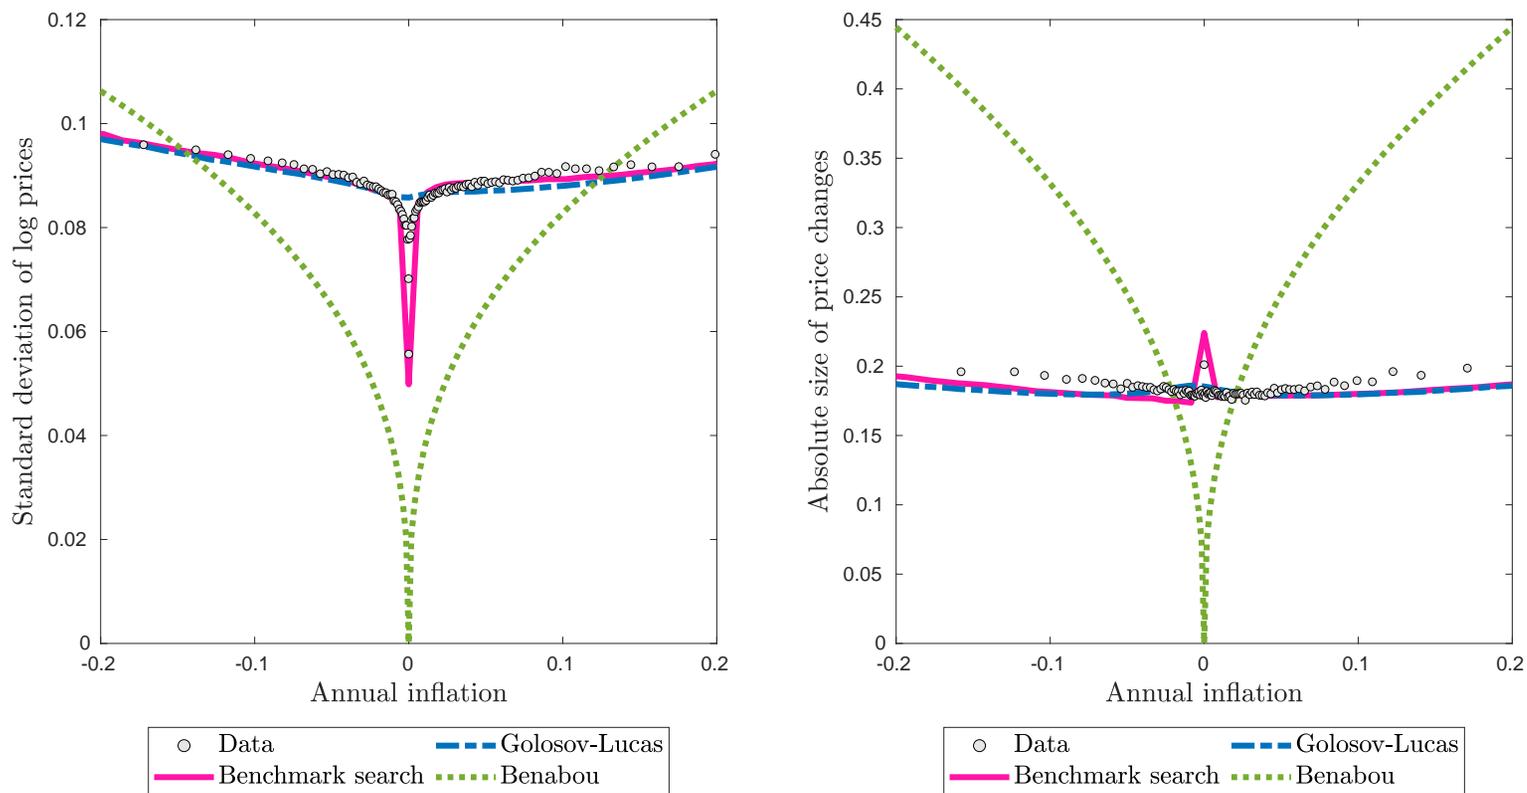


Figure 9: Pricing moments in different models

The figure plots pricing moments and inflation in Benabou (1992), Golosov and Lucas (2007), the data, and my model. The left panel shows price dispersion measured by the standard deviation of log prices, and the right panel plots the absolute size of price changes.

$\pi = 0$).

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.

In the next section, we see the menu-cost model with idiosyncratic shocks and consumer search not only matches key price-setting statistics, but is also supported by empirical evidence on households' shopping behavior.

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. In the model, the latter variable is inversely correlated with the search cost of shoppers.

Households are sampled from 54 geographically dispersed Scantrack markets (each roughly corresponding to an MSA), and detailed demographic information about them is available.³³ The product categories are the same as in the NielsenIQ 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

³³NielsenIQ provides projection factors to make the sample demographically representative of the Scantrack market population.

for a given good while controlling for several sources of heterogeneity.

5.1 Variable construction

The unit of analysis is a household \times department \times 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 \times module level, for instance, shoppers visit, on average, 1.1 retailers each quarter (CV of 0.20), whereas at the household \times department level, they visit three retailers (CV of 0.71). 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

$$\bar{P}_{km,t} = \frac{\sum_{h=1}^{H_{km,t}} \sum_{l=1}^{L_{km,t}^h} P_{km,t}^{h,l} q_{km,t}^{h,l}}{\sum_{h=1}^{H_{km,t}} \sum_{l=1}^{L_{km,t}^h} q_{km,t}^{h,l}},$$

where $P_{km,t}^{h,l}$ is the price household h paid for $q_{km,t}^{h,l}$ 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 $L_{km,t}^h$ 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:

$$p_{dm,t}^h = \frac{X_{dm,t}^h}{\bar{X}_{dm,t}^h} = \frac{\sum_{k=1}^{D_{dm,t}} \sum_{l=1}^{L_{km,t}^h} P_{km,t}^{h,l} q_{km,t}^{h,l}}{\sum_{k=1}^{D_{dm,t}} \sum_{l=1}^{L_{km,t}^h} \bar{P}_{km,t} q_{km,t}^{h,l}},$$

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

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

$$X_{dm,t}^h = \sum_{i=1}^{N_{dm,t}^h} \sum_{l=1}^{L_{km,t}^{h,i}} X_{dm,t}^{h,l},$$

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

Table A.3 shows descriptive statistics for the sample. The total number of household \times department \times quarter observations is 19,376,570, with 130,262 households, 10 departments, 44 quarters, and 53 geographic markets.

5.2 Estimation and results

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

$$\log p_{dm,t}^h = \sum_{s=1}^{10} \alpha_s \mathbf{1}_{\{N_{dm,t}^h=s\}} + \sum_{s=1}^{10} \theta_s \mathbf{1}_{\{N_{dm,t}^h=s\}} \times |\bar{\pi}_{dm,t}| + \mu' X_t^h + a_d + b_m + c_t + \epsilon_{dm,t}^h. \quad (8)$$

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 $\{\alpha_s\}_{s=1}^{10}$ correspond to average relative prices paid by households that visit s retailers; more

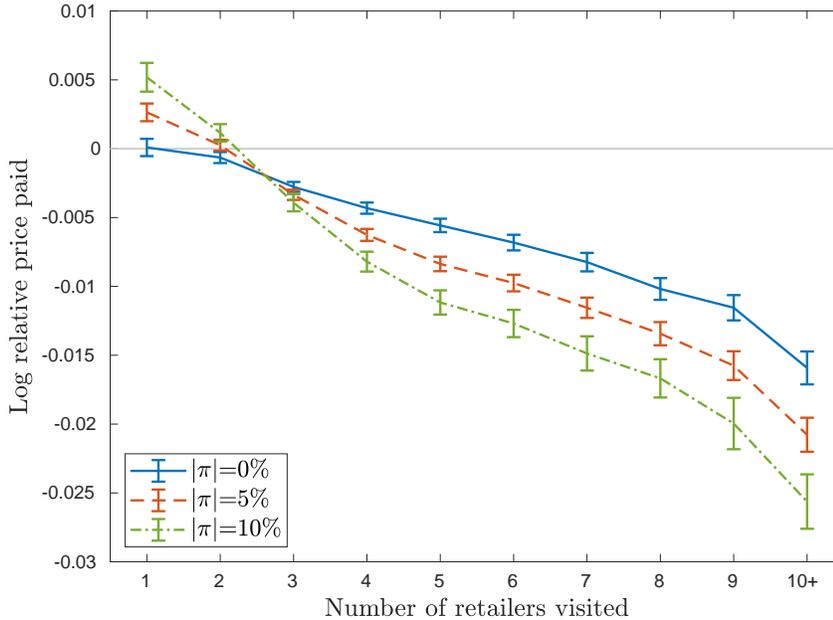


Figure 10: Shopping behavior and absolute inflation

The figure shows the relationship between log relative prices paid, number of retailers visited, and absolute inflation as predicted by equation (8). The unit of observation is a household \times product department \times quarter, and the total number of observations is 19,376,570. Observations are weighted using NielsenIQ sampling weights. The confidence intervals (bars) use standard errors that are two-way clustered by household and product department \times market \times quarter combination.

importantly, the coefficients $\{\theta_s\}_{s=1}^{10}$ indicate how and by how much absolute inflation affects the relative prices paid by households according to their shopping behavior. The theory predicts $\theta_{10} < 0$ and that the sequence $\{\theta_s\}_{s=1}^{10}$ is increasing: when absolute inflation is high, shoppers that search more find the lowest prices.

The vector X_t^h contains a set of demographic controls (household size, household income bin, household head age and education, employment status, and presence of children) and shopping behavior controls (number of shopping trips and fraction of transactions paid with coupons). In addition, I include department, market, and quarter fixed effects. 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 10 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% less than when inflation is zero, and 3% less than a household that visits only one store.

In the theory, inflation is relevant for the shopping behavior of households only through price dispersion. Therefore, I provide additional evidence by estimating equation (8) using inflation-induced price dispersion rather than absolute inflation. To obtain a measure of price dispersion that is only explained by inflation, I estimate equation (1) at the quarter level and predict price dispersion for each category \times market pair:

$$\hat{\sigma}_{cm,t} = \sum_{n=1}^{100} \hat{\beta}_n \mathbf{1}_{\{\bar{\pi}_{cm,t} \in B_n\}}.$$

I compute inflation-induced price dispersion for each department \times market \times quarter, $\hat{\sigma}_{dm,t}$, by aggregating category-level measures using annual sales as weights.

Figure 11 shows the effects of shopping behavior on prices paid depending on the amount of inflation-induced price dispersion in the market. A household that visits 10 stores when price dispersion is high (at the 95th percentile of the distribution) pays 1% less than when price dispersion is low (at the 5th percentile of the distribution), and 3% less than a household that visits only one store. As with absolute inflation, the evidence is consistent with the model.³⁴

Additionally, I test the fit of the calibrated model to these results. For each level of inflation, I take the equilibrium posted-price distribution and the optimal search policies of shoppers j , 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 (8) in the model:

$$\log p_{dm}^j = \sum_{s=1}^{10} \alpha_s \mathbf{1}_{\{N_{dm}^j = s\}} + \sum_{s=1}^{10} \theta_s \mathbf{1}_{\{N_{dm}^j = s\}} \times |\bar{\pi}_{dm,t}| + \mu' X^j + \epsilon_{dm}^j.$$

A department \times market pair is identified exclusively by its level of inflation, so I choose department \times market pairs to replicate the cross-sectional distribution of inflation in the data. The shopper-level controls include deciles of the search cost – the only source of heterogeneity across buyers. Figure 12 shows the calibrated model closely replicates the data, especially when we repeat the exercise for inflation-induced price dispersion.

³⁴In Figure A.11, I show the qualitative results hold if we use raw – rather than predicted – price dispersion.

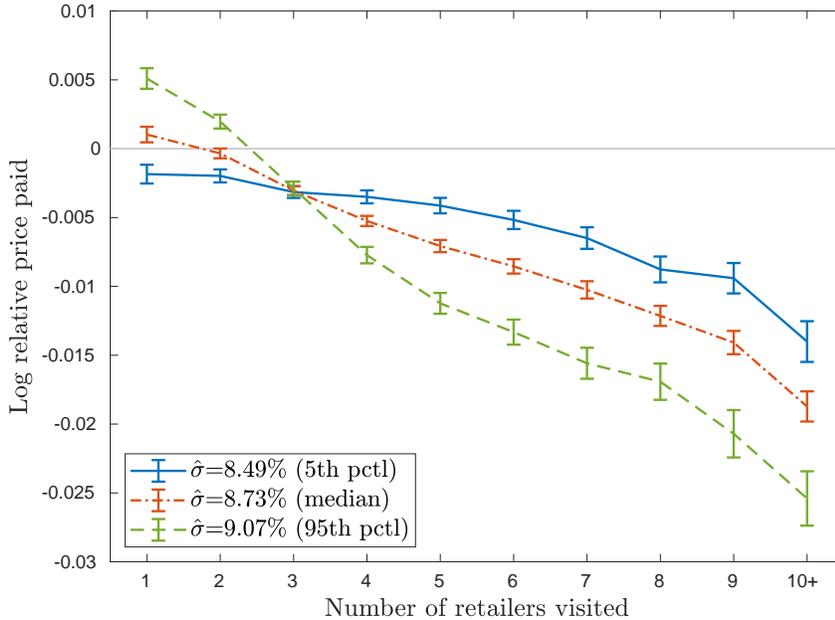


Figure 11: Shopping behavior and price dispersion

The figure shows the relationship between log relative prices paid, number of retailers visited, and inflation-induced price dispersion as predicted by equation (8). The unit of observation is a household \times product department \times quarter, and the total number of observations is 19,376,570. Observations are weighted using NielsenIQ sampling weights. The confidence intervals (bars) use standard errors that are two-way clustered by household and product department \times market \times quarter combination.

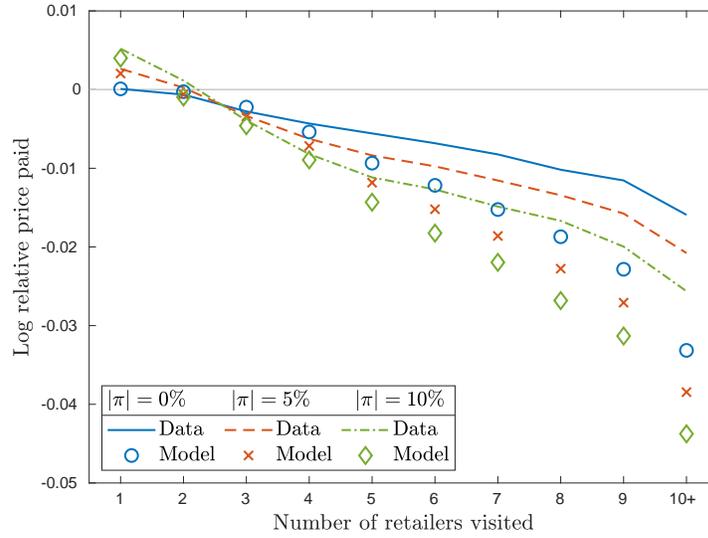
6 Welfare implications of inflation

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 dt is given by equation (3). Expressing its components in terms of the posted-price and the search-cost distributions is useful. The household gross surplus from consumption \bar{V} is the transaction-weighted average of the shopper's surplus in terms of real prices:

$$\bar{V} = \int_0^1 V(p_{j(i)}) dj = \int_{\underline{p}}^{\bar{p}} V(p) N(p) dF(p).$$

(a) Absolute inflation



(b) Price dispersion

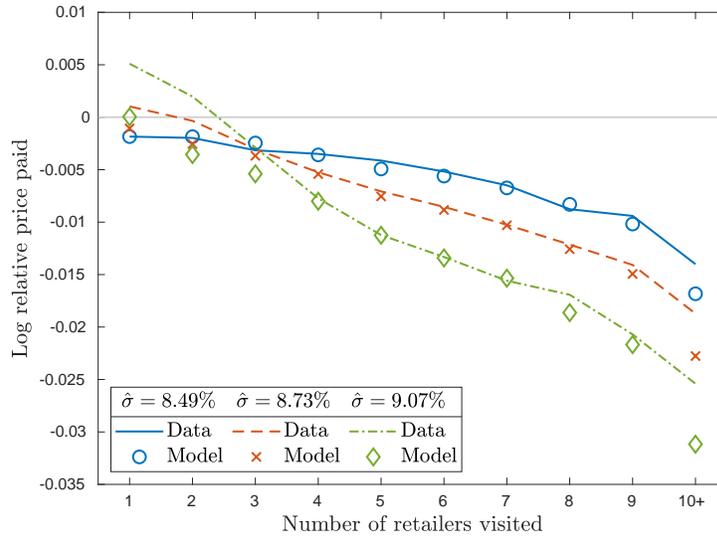


Figure 12: Shopping behavior and prices paid, data and model

The lines show the empirical relationship between log relative prices paid and number of retailers visited for different levels of absolute inflation (upper panel) and inflation-induced price dispersion (lower panel), as predicted by equation (8). The dots correspond to simulated data using the calibrated model.

To derive the total cost of search, note a shopper type γ accepts an offer $p \leq R(\gamma)$ with probability $F(R(\gamma))$ and continues searching with probability $1 - F(R(\gamma))$. Therefore, the probability of receiving a successful offer after s searches is

$$\Pr[K(\gamma) = s] = F(R(\gamma)) \times [1 - F(R(\gamma))]^{s-1}.$$

Because the number of searches they can conduct is not limited, a shopper with cost γ searches, on average, $S(\gamma)$ times, where

$$S(\gamma) = \sum_{s=1}^{\infty} s \times \Pr[K(\gamma) = s] = \frac{1}{F(R(\gamma))}.$$

Taking into account that the first price draw is free, the total resources spent on the search for better prices are

$$C = \int_0^1 \gamma_j (S_j - 1) dj = \int_{\underline{\gamma}}^{\bar{\gamma}} \gamma \left[\frac{1}{F(R(\gamma))} - 1 \right] dG(\gamma).$$

On the other hand, each dt , real dividends are aggregate real profits net of aggregate adjustment costs. Gross aggregate real profits are

$$\bar{\Pi} = \int_0^1 \Pi(p_{i,t}, v_{i,t}) di = \int_{\underline{p}}^{\bar{p}} \int_v \Pi(p, v) \phi(p, v) dv dp.$$

If Λ is the number of retailers that reprice each dt , total adjustment costs are $\kappa\Lambda$.

Then, social welfare is aggregate consumer surplus net of search costs plus total real profits net of adjustment costs. Each one of these terms depends on inflation:

$$\mathcal{W}(\pi) = \bar{V}(\pi) + \bar{\Pi}(\pi) - C(\pi) - \kappa\Lambda(\pi).$$

The first two terms on the right correspond to the 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 real 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 at higher prices decrease, increasing efficiency. Thus, whether inflation improves welfare depends on the size of the efficiency gains from lower markups.

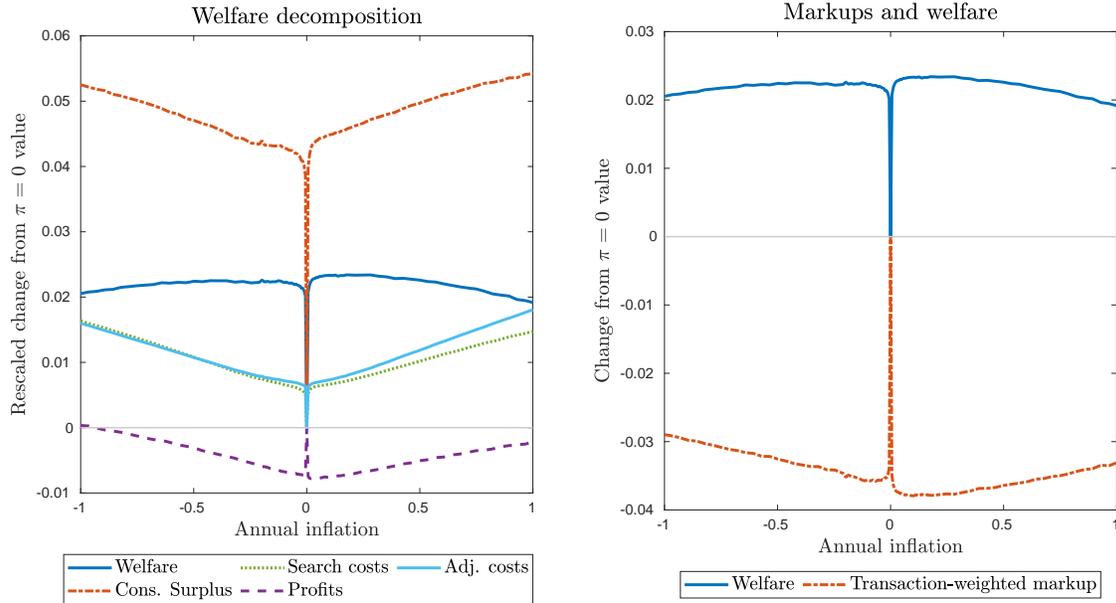


Figure 13: Welfare and inflation

The figure in the left panel shows the relationship between inflation and welfare and the contribution of each of its components. The one in the right panel also plots the average transaction-weighted markup.

Figure 13 shows the relationship between welfare and inflation 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, average transaction-weighted markups decrease, 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, a large part of the production shifts from less to more productive retailers: with increased search activity (right panel in Figure 14), retailers with high productivity draw charge lower prices to attract more shoppers, and their markups decrease (left panel in Figure 14).

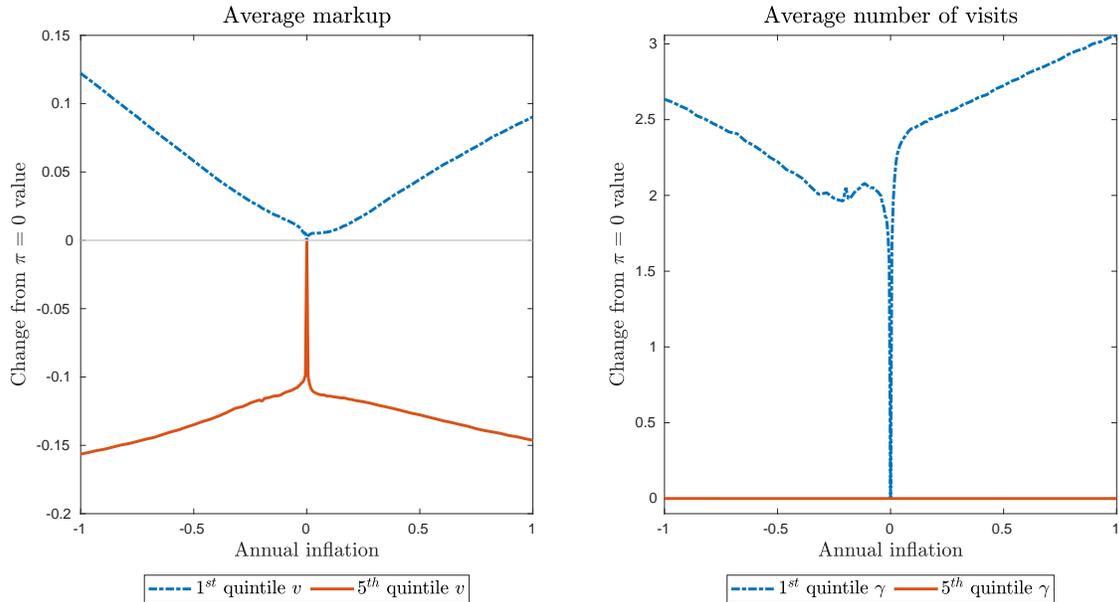


Figure 14: Heterogeneous effects of inflation

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

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 14 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.

For large positive inflation, we see from Figure 13 that aggregate consumer surplus is still increasing in inflation: shoppers who search are paying even less for the good as price dispersion increases with inflation. On the other hand, aggregate gross profits start to increase: the least productive retailers increase their markups by extracting the surplus of captive shoppers. If the latter effect exceeds the former, efficiency gains from positive inflation dissipate, reducing welfare.

7 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.
- Argente, D., M. Lee, and S. Moreira (2018). Innovation and product reallocation in the great recession. *Journal of Monetary Economics* 93, 1–20.
- Beaulieu, J. and J. Matthey (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 (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.
- 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.
- Klenow, P. J. and J. L. Willis (2016). Real rigidities and nominal price changes. *Economica* 83(331), 443–472.
- Lach, S. and D. Tsiddon (1992). The behavior of prices and inflation: An empirical analysis of disaggregated 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., 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.
- 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

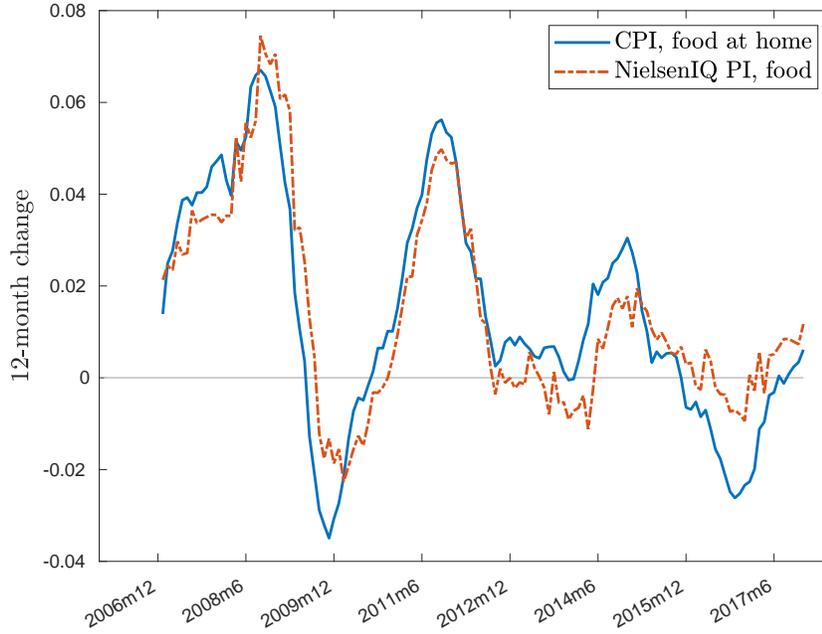


Figure A.1: Aggregate inflation, validation

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 category \times market level inflation $\pi_{cm,t}$ using annual sales as weights.

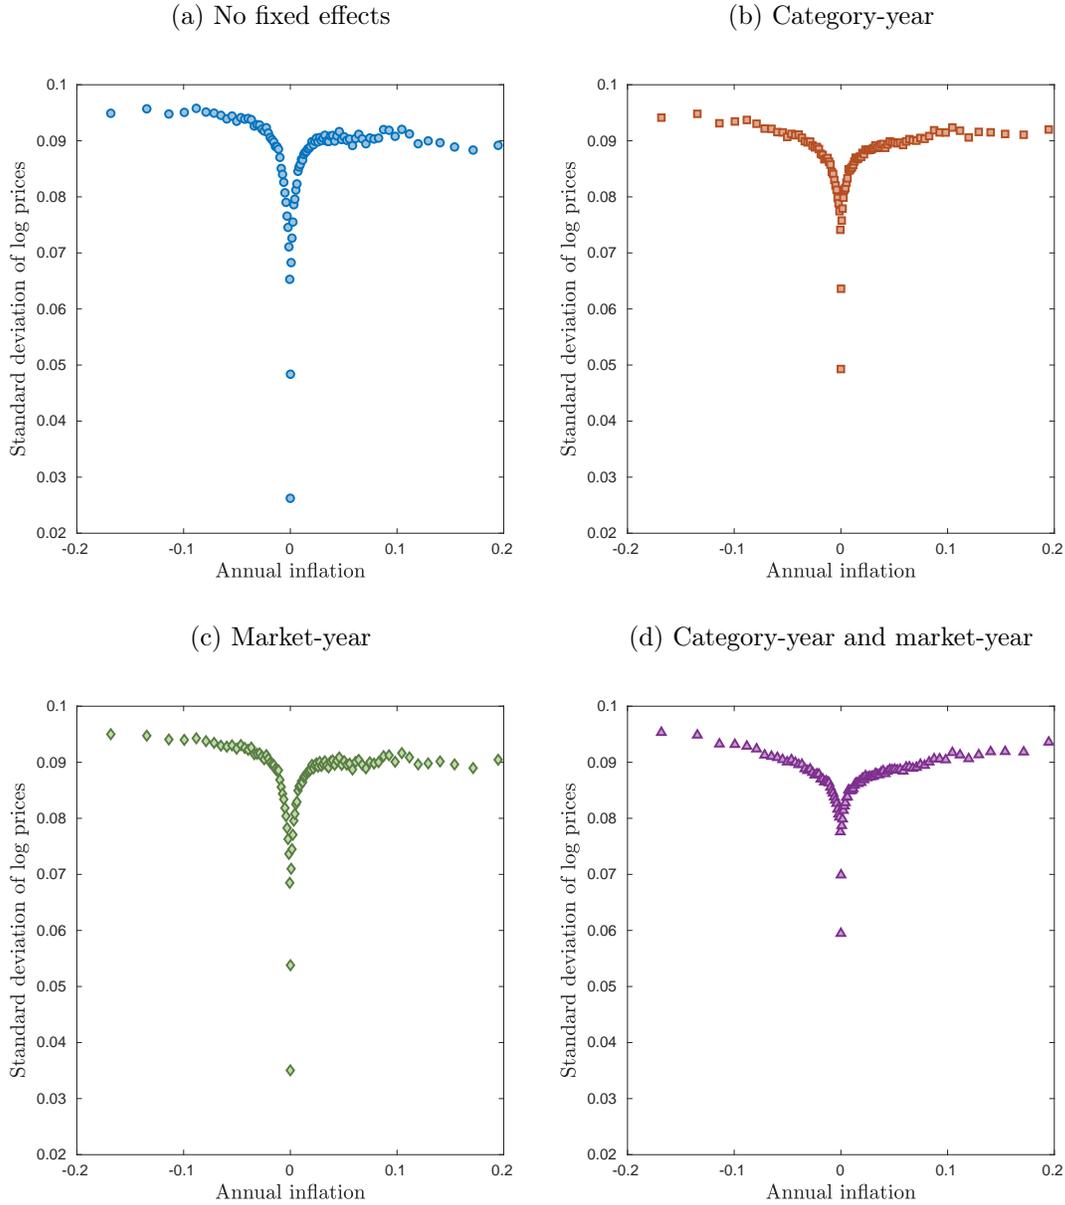


Figure A.2: Controlling for different fixed effects

The dots correspond to average price dispersion for each of 100 equally sized inflation bins as predicted by equation (1), for different combinations of the fixed effects and excluding the number of stores. The unit of observation is a category \times market \times year, for 1,219,414 in total.

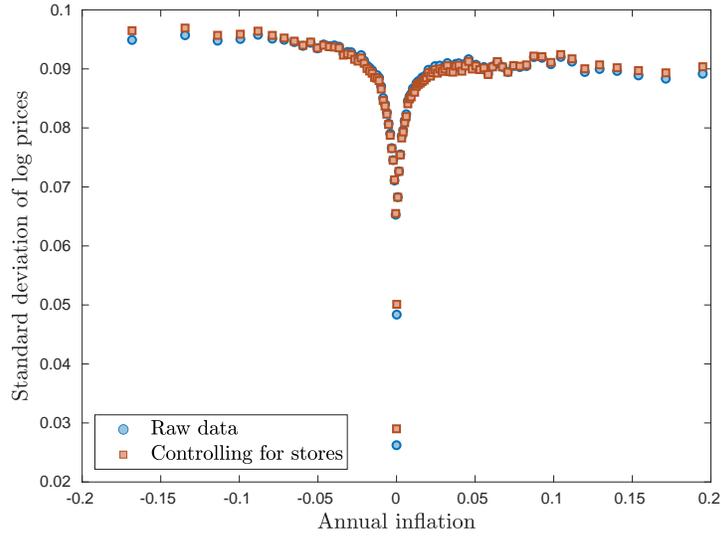


Figure A.3: Controlling for number of stores

The dots correspond to average price dispersion for each of 100 equally sized inflation bins as predicted by equation (1), without fixed effects. The unit of observation is a category \times market \times year, for 1,219,414 in total.

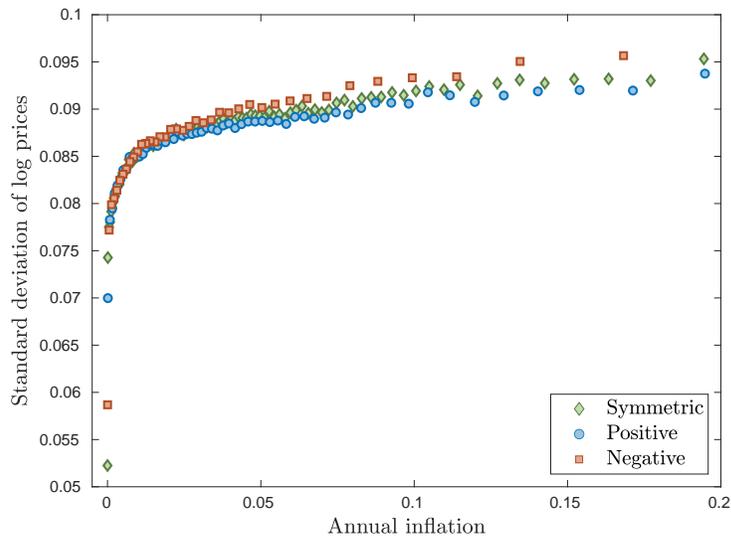


Figure A.5: Price dispersion and absolute inflation

The dots correspond to average price dispersion for each inflation bin 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 with absolute inflation on the x-axis. The unit of observation is a category \times market \times year, for 1,219,414 in total.

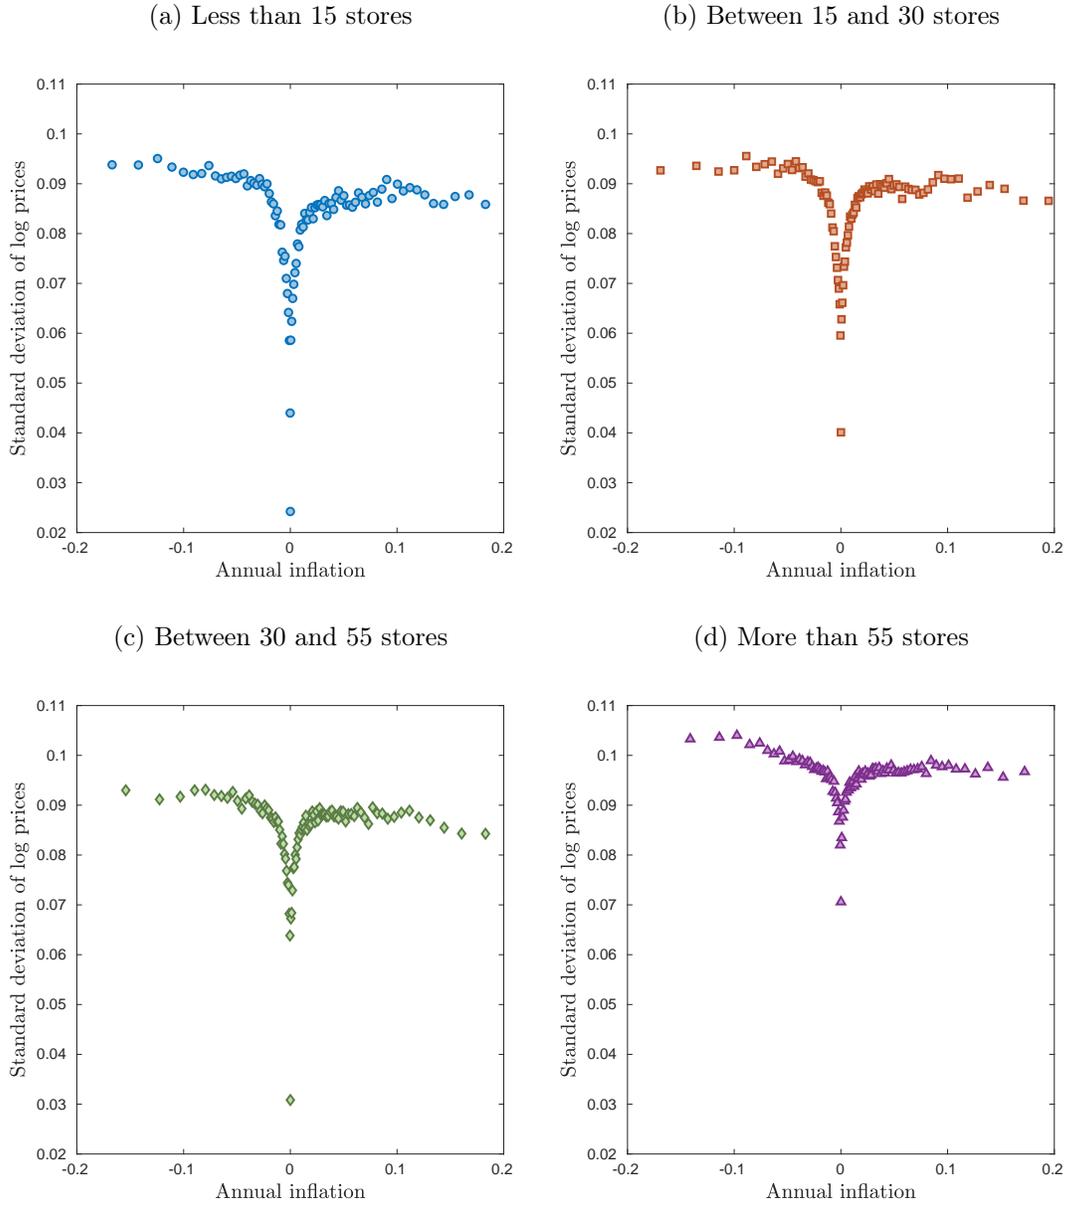


Figure A.4: Dividing sample by number of stores

The dots correspond to average price dispersion for each of 100 equally sized inflation bins when observations are divided according to quartiles of the average number of stores. The unit of observation is a category \times market \times year, for 1,219,414 in total.

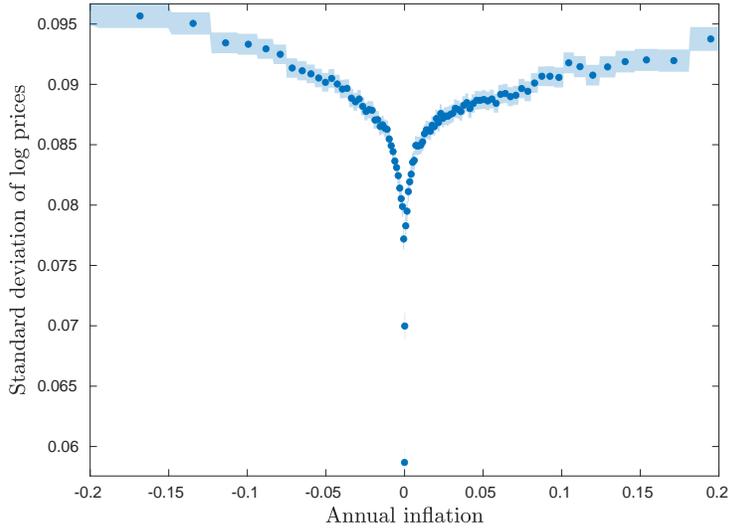


Figure A.6: Statistical significance

The dots correspond to average price dispersion for each of 100 equally sized inflation bins as predicted by equation (1). The band covers the entire function with probability 0.95. The unit of observation is a category \times market \times year, for 1,219,414 in total. Standard errors are clustered by category \times market pair.

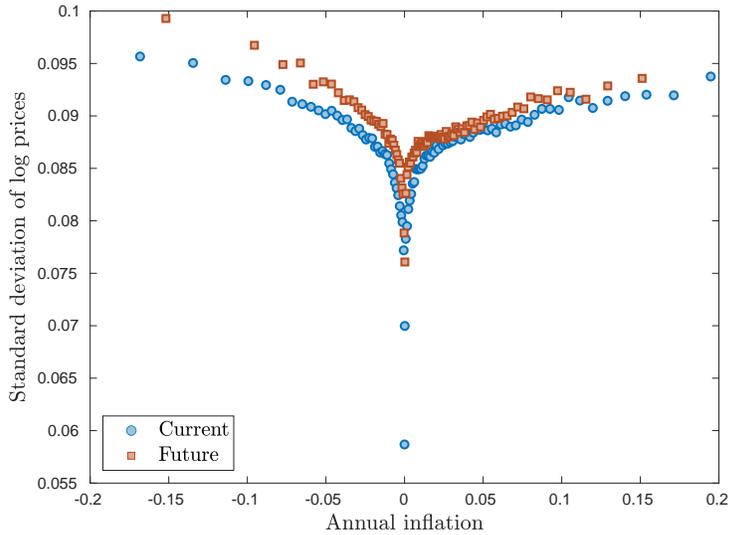
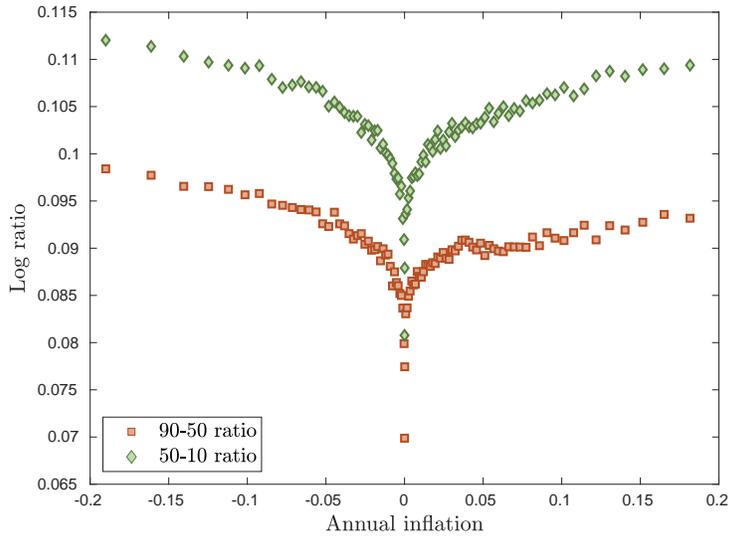


Figure A.7: Current and future inflation

The dots correspond to average price dispersion for each inflation bin as predicted by equation (1), using future instead of current inflation. The unit of observation is a category \times market \times year.

(a) Levels



(b) Decomposition

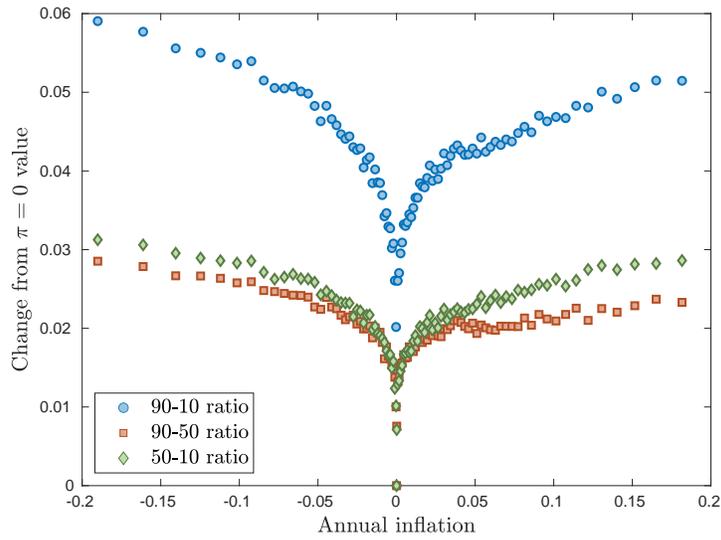
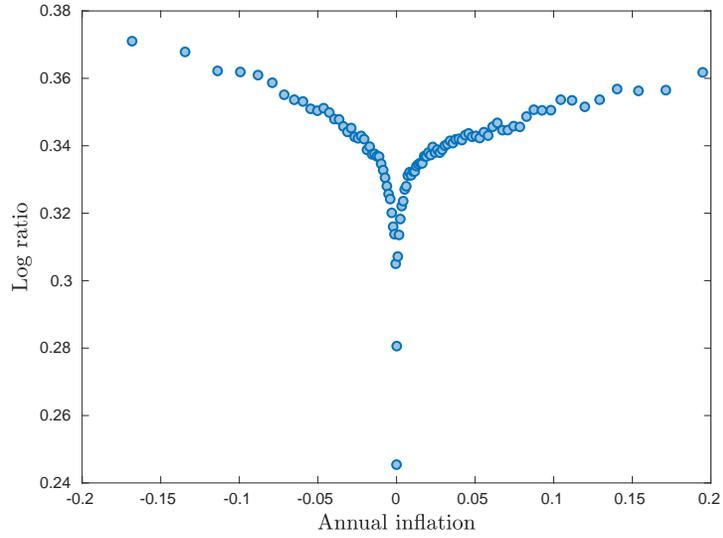


Figure A.9: Price dispersion and inflation by half of the distribution

The dots correspond to average price dispersion for each inflation bin conditional on controls. Price dispersion is measured using the log 90-10, 90-50 or 50-10 ratio of the price distribution. The top panel shows levels of the 90-50 and 50-10 ratios; the bottom panel, deviations from their $\pi = 0$ values. The unit of observation is a category \times market \times year, for 1,219,414 in total.

(a) Max-min ratio of price distribution



(b) 90-10 ratio of price distribution

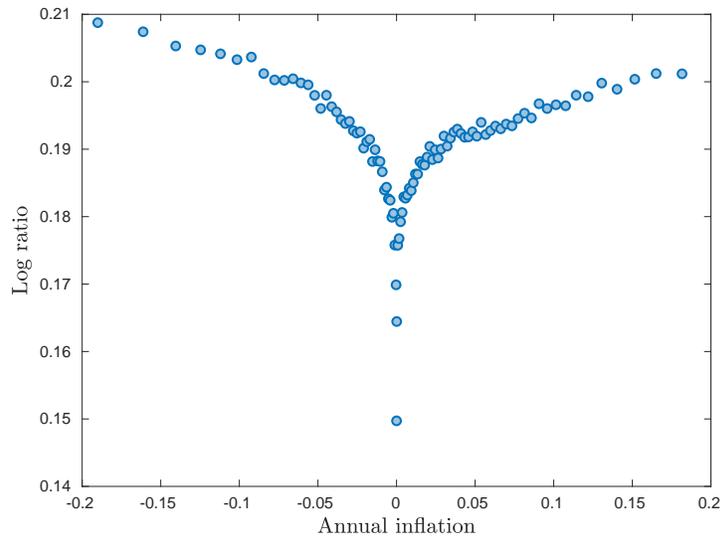


Figure A.8: Alternative price dispersion measures and inflation

The dots correspond to average price dispersion for each inflation bin as predicted by equation (1). The top and bottom panel show estimates using the log max-min and the log 90-10 ratios of the price distribution, respectively, as measures of price dispersion. The unit of observation is a category \times market \times year, for 1,219,414 in total.

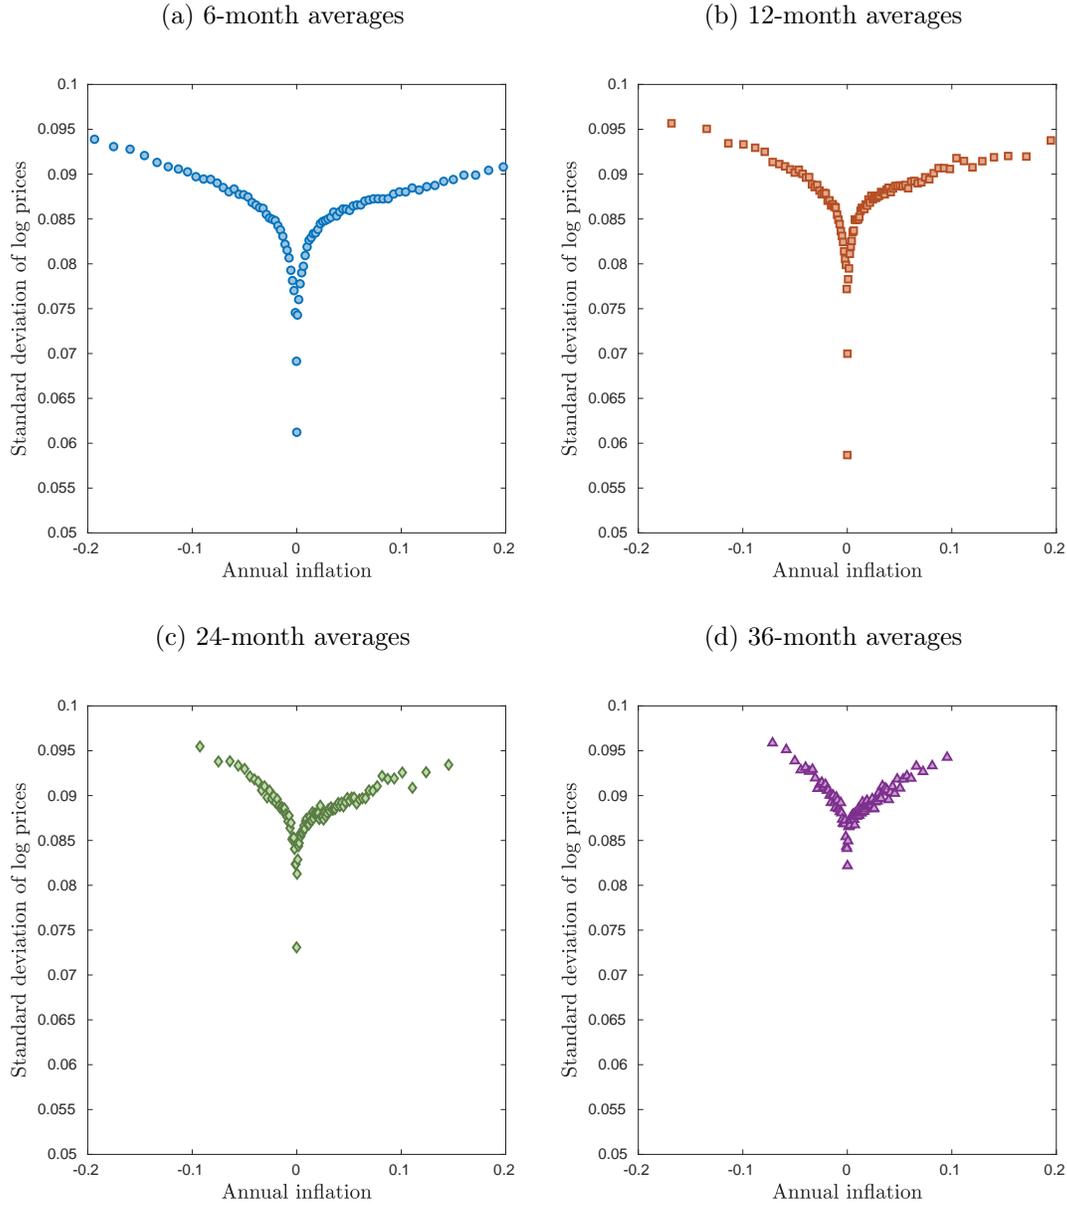


Figure A.10: Inflation for different period lengths

The dots correspond to average price dispersion conditional on covariates for each of 100 equally sized inflation bins assuming convergence to an ergodic distribution in 6, 12, 24, and 36 months. That is, we estimate equation (1) averaging monthly inflation and price dispersion over non-overlapping time periods with different lengths. The unit of observation is a category \times market \times period.

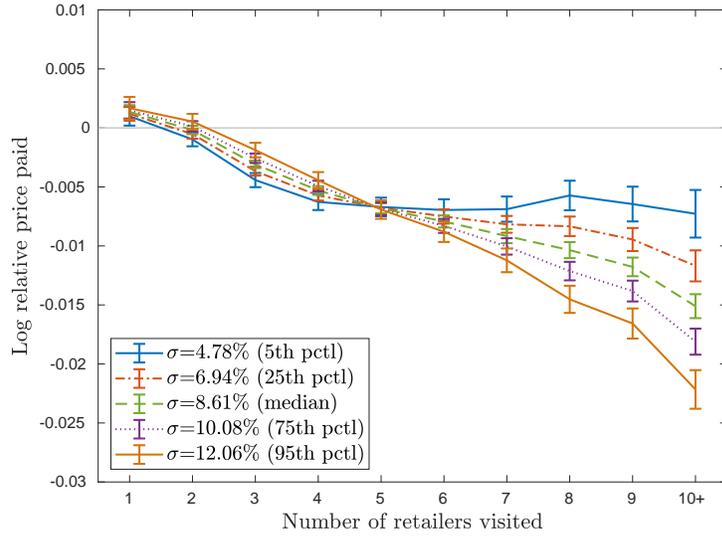


Figure A.11: Shopping behavior and raw price dispersion

The figure shows log relative prices paid by number of retailers visited when we estimate equation (8) using raw price dispersion rather than absolute inflation. The unit of observation is a household \times department \times quarter, for 19,376,570 in total. Observations are weighted using NielsenIQ sampling weights. The confidence intervals (bars) use standard errors that are two-way clustered by household and department \times market \times quarter combination.

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
Inflation	0.017	0.079	-0.194	-0.057	-0.016	0.009	0.047	0.101	0.263
Absolute inflation	0.052	0.061	0.000	0.004	0.012	0.032	0.069	0.124	0.288
Aggregate inflation	0.015	0.019	-0.011	-0.005	0.000	0.009	0.029	0.042	0.068
Price dispersion									
Std. dev. of log prices	0.087	0.049	0.001	0.027	0.054	0.084	0.115	0.147	0.231
Max-min price ratio	1.431	0.305	1.006	1.112	1.235	1.391	1.570	1.774	2.352
90-10 percentile ratio	1.217	0.172	1.000	1.028	1.095	1.193	1.305	1.425	1.740
90-50 percentile ratio	1.096	0.088	1.000	1.006	1.030	1.079	1.138	1.200	1.389
50-10 percentile ratio	1.110	0.106	1.000	1.009	1.038	1.090	1.156	1.230	1.449
Number of stores	49.548	62.102	6.507	11.275	16.655	28.208	54.397	111.977	306.507

Notes. The table shows descriptive statistics for the sample used to estimate equation (1). All variables are in annual terms. The total number of observations is 1,219,414 for 138,485 category \times market pairs between 2007 and 2017, with 1,042 unique categories and 204 geographic markets. The fraction of observations with deflation is 0.397. Aggregate inflation corresponds to an average over all category \times market pairs using annual sales as weights.

Table A.2: Numerical tests, year-by-year estimates

	Full $ \pi $ range	$ \pi < 2\%$
Market-year estimates, $f_{m,t}$		
% reject null: Increasing	22.68	4.91
% reject null: Concave	17.63	5.85
Average # categories per market-year	573.65	226.53
Total # market-year observations	2,121	1,916
Category-year estimates, $f_{c,t}$		
% reject null: Increasing	33.53	10.27
% reject null: Concave	24.13	10.75
Average # markets per category-year	139.63	96.69
Total # category-year observations	7,985	1,898

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

Table A.3: 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	-0.003	0.110	-0.348	-0.091	-0.029	0	0.035	0.085	0.264
Number of retailers visited	2.984	2.120	1	1	1	2	4	6	10
Number of shopping trips	8.775	8.694	1	2	3	6	11	18	43
Fraction of trans. with coupons	0.079	0.172	0	0	0	0	0.071	0.286	0.862
Department \times market \times quarter									
Inflation (annualized)	0.015	0.033	-0.075	-0.017	-0.002	0.011	0.029	0.054	0.112
Absolute inflation (annualized)	0.025	0.026	0.000	0.003	0.008	0.016	0.034	0.059	0.118
Std. dev. of log prices	0.082	0.023	0.029	0.052	0.066	0.083	0.098	0.110	0.132
Std. dev. of log prices, predicted	0.087	0.002	0.083	0.085	0.086	0.087	0.089	0.090	0.093

Notes. The table shows descriptive statistics for the sample used to estimate equation (8). Household-related measures are computed using the Consumer Panel Data; inflation and price dispersion, using the Retail Scanner Data. The total number of observations for estimation are 19,376,570 household \times department \times quarter triples, with 130,262 households, 10 departments, 44 quarters, and 53 Scantrack markets. At the department \times market \times quarter level, 23,296 distinct observations for inflation and price dispersion are available.

B 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.³⁵ 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.

Temporary price reductions are not flagged in the RMS data. To identify sales, I use the weekly version of Nakamura and Steinsson’s filter, as in Coibion et al. (2015). In particular, I classify a price as a sale when the price decreases (for up to three weeks) and then returns to its previous level. At the same time, I classify a price change as regular when neither period has a sale episode.

For each product \times market \times month, I obtain the number of regular price changes across weeks and stores. The frequency of adjustment is given by this number over the total price observations. 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 level, it is the unweighted average across weeks and stores in a given month. To aggregate across products within a category \times market \times month, I average the product-level statistics using annual sales as weights.

³⁵That is, I do not impute missing values using the preceding or subsequent price.

C Multi-sector model

In the data, the measure of inflation is sectoral: it is the rate at which sector-level prices grow. The baseline model contains a unique sector in the economy, where the nominal price level increases because nominal wages increase. In the multi-sector model, sectoral inflation rates have two sources: a sector-specific productivity drift and a common drift that is influenced by the monetary authority. I assume the only source of heterogeneity across sectors is the rate at which their prices grow.

Assume the existence of a continuum of sectors m in the economy. Following Klenow and Willis (2016), I let sectoral prices exhibit differential trends; I continue assuming idiosyncratic productivity shocks are stationary because relative prices within a product category exhibit mean reversion. The wage rate grows at a rate λ_W , $W_t = W_0 e^{\lambda_W t}$, and sectoral productivity grows at a rate $\lambda_{A,m}$, so $A_{m,t} = A_{m,0} e^{\lambda_{A,m} t}$.

Households The economy contains a measure-one continuum of households h . Each household consists of a continuum measure one of shoppers and a worker. Each shopper j visits every sector m , populated by a measure-one continuum of retailers i_m . Dropping time subscripts, the problem of the household is

$$\begin{aligned} & \max_{q_{j,m}^h, L^h} \int_0^1 \int_0^1 \left(\frac{Z_m^{\frac{1}{\eta}} q_{j,m}^h}{1 - \frac{1}{\eta}} \right)^{1-\frac{1}{\eta}} dj dm - L^h \\ s.t. \quad & \int_0^1 \int_0^1 P_{j(i_m),m}^h q_{j,m}^h dj dm + W \int_0^1 \int_0^1 \gamma_{j,m}^h (S_{j,m}^h - 1) dj dm = W L^h + \mathcal{D}, \end{aligned}$$

where the search cost $\gamma_{j,m}^h$ is in labor units and it is randomly and identically distributed among shoppers, sectors, and households. A deterministic demand shifter Z_m exists for each sector. For convenience, we assume $Z_m = A_m^{1-\eta}$. Household h 's head decides optimal labor supply and quantities given the prices:

$$q_{j,m}^h = \left(\frac{P_{j(i_m),m}^h Z_m^{-\frac{1}{\eta}}}{W} \right)^{-\eta} = A_m \left(\frac{P_{j(i_m),m}^h A_m}{W} \right)^{-\eta}.$$

The associated indirect utility of the household is

$$\mathcal{W}^h = \int_0^1 \int_0^1 [V(p_{j(i_m),m}^h) - \gamma_{j,m}^h (S_{j,m}^h - 1)] dj dm + \frac{\mathcal{D}}{W},$$

where the relative prices are in effective labor units required to purchase one unit of the sector m good:

$$p_{j(i_m),m}^h = \frac{P_{j(i_m),m}^h A_m}{W}.$$

Shoppers in each sector follow similar search rules as in the baseline model. The main difference is that because the posted-price distribution $F_m(p)$ is sector specific, the rules they follow are also sector specific: $R_m(\gamma)$.

Demand function Under the assumption that the search cost is randomly and identically distributed among shoppers within and across households, each sector is populated by a continuum of shoppers with the same search-cost distribution, g . The aggregation of individual search rules of shoppers yields a sector-specific demand function:

$$\begin{aligned} D_{m,t}(p) &= A_{m,t} p^{-\eta} \times N_m(p); \\ N_m(p) &= \int_{\Gamma_m(p)}^{\bar{\gamma}} \frac{g(\gamma)}{F_m(R_m(\gamma))} d\gamma. \end{aligned}$$

The intensive and extensive margins of demand are nearly identical to those in the baseline model. In this case, the demand function is time specific because the intensive margin is affected by the sector-specific trend.

Retailers Nominal profits for retailers in sector m are

$$\Pi_{m,t}(P_{i_m,t}, v_{i_m,t}) = \left(P_{i_m,t} - \frac{W_t}{A_{m,t} v_{i_m,t}} \right) \times D_{m,t} \left(\frac{P_{i_m,t} A_{m,t}}{W_t} \right).$$

The time-invariant profit function in labor units is given by

$$\Pi_m(p_{i_m,t}, v_{i_m,t}) = \left(p_{i_m,t} - \frac{1}{v_{i_m,t}} \right) \times p_{i_m,t}^{-\eta} N_m(p_{i_m,t}).$$

Equilibrium Under this simple specification of preferences, the sector-level equilibrium can be computed as the one-sector model equilibrium. Therefore, analyzing the sectoral equilibrium of this multi-sector model is equivalent to analyzing that of a one-sector model for different steady-state levels of inflation.

Aggregate inflation and welfare The sectoral price level is defined as the deflator of nominal output in sector m :

$$P_{m,t} = \frac{W_t \int p^{1-\eta} N_m(p) dF_m(p)}{A_{m,t} \int p^{-\eta} N_m(p) dF_m(p)}.$$

Therefore, the sectoral inflation rate is $\pi_m = \lambda_W - \lambda_{A,m}$. Assuming an economy-wide price index such as

$$P_t = \exp \left\{ \int \log P_{m,t} dm \right\}$$

yields the following expression for aggregate inflation as a function of cross-sectional inflation rates:

$$\pi = \lambda_W - \int \lambda_{A,m} dm.$$

The assumptions on the search-cost distribution imply $\mathcal{W}^h = \bar{\mathcal{W}}$ for every household. Moreover, we can express aggregate welfare as

$$\begin{aligned} \bar{\mathcal{W}} &= \int_0^1 \mathcal{W}_m dm; \\ \mathcal{W}_m &= \bar{V}_m - C_m + \bar{\Pi}_m - \kappa \Lambda_m. \end{aligned}$$

Alternatively, using that sectors only differ with respect to their inflation rate,

$$\begin{aligned} \bar{\mathcal{W}} &= \int_0^1 \mathcal{W}(\pi_m) dm \\ \Rightarrow \bar{\mathcal{W}}(\pi) &= \int_0^1 \mathcal{W} \left(\pi + \int \lambda_{A,m} dm - \lambda_{A,m} \right) dm. \end{aligned}$$

The calibrated baseline model provides a function $\mathcal{W}(\cdot)$ that matches the Υ -shaped pattern between sectoral inflation π_m and sectoral price dispersion. From the microdata, and using aggregate estimates of λ_W , we can recover the real component of sectoral price-level growth: $\lambda_{A,m} = \lambda_W - \pi_m$. By setting $\lambda_{A,m}$ equal to the data and using the calibrated function $\mathcal{W}(\cdot)$, we can study how aggregate welfare changes with aggregate steady-state inflation according to the model: $d\bar{\mathcal{W}}(\pi)/d\pi$.