Inflation Distorts Relative Prices: Theory and Evidence*

Klaus Adam (University of Mannheim & CEPR)

Andrey Alexandrov (Tor Vergata University of Rome)

Henning Weber (Deutsche Bundesbank)

May 6, 2024

Abstract

We empirically identify the effect of inflation on relative price distortions, using a novel identification approach derived from sticky price theories with time or state-dependent adjustment frictions. Our approach can be directly applied to micro price data, does not rely on estimating the gap between actual and flexible prices, and only assumes stationarity of unobserved shocks. Using the micro price data underlying the U.K. CPI, we document that suboptimally high (or low) inflation is associated with distortions in relative prices. At the level of individual products, the marginal effect of inflation on relative price distortions is highly statistically significant and aligns well with theoretical predictions. Cross-sectional price dispersion turns out to be predominantly driven by movements in the dispersion of flexible prices and thus fails to comove with inflation over time. In contrast, cross-sectional price distortions are found to increase with aggregate inflation.

JEL Class. No.: E31, E58

*We thank our discussants Hassan Afrouzi, Erwan Gautier, Lorenza Rossi, Viacheslav Sheremirov and Alex Wolman. We are also grateful to Markus Brunnermeier, Jordi Galí, Yuriy Gorodnichenko, Loretta Mester, Emi Nakamura, Krzysztof Pytka, Benjamin Schoefer, Frank Schorfheide, Jon Steinsson, Frank Smets and Carsten Trenkler, seminar and conference participants in Berkeley, Konstanz, the 2023 NBER Conference on Methods and Applications for DSGE Models in Philadelphia, the 2023 Bank of Portugal Conference on Monetary Economics, the 2nd XAmsterdam Macroeconomic Workshop, and the Warsaw 7th Annual Workshop of the ESCB Research Cluster on Monetary Economics for helpful input and discussions. Klaus Adam acknowledges support from the Deutsche Forschungsgemeinschaft through CRC-TR 224 (Project C02). The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

1 Introduction

The monetary models employed in academia and central banks assert that too high (or too low) rates of inflation give rise to distortions in relative prices. The asserted price distortions drive many of the trade-offs and policy prescriptions of monetary models, e.g., the recommendation to implement low and stable inflation rates.[1] Despite its centrality in monetary theory, there exists no structural empirical evidence validating the notion that inflation has a distorting impact on relative prices.

To fill this gap, we derive a novel theory-consistent empirical approach that allows estimating the marginal effect of inflation on relative price distortions. We apply this approach to the micro price data underlying the U.K. Consumer Price Index and find that inflation is associated - at the level of individual products - with economically significant amounts of price distortions. In the cross-section of products, price distortions comove positively with aggregate inflation over time, but inflation-induced price distortions account for at most 1% of the observed cross-sectional dispersion of prices.

Documenting the empirical relationship between inflation and relative price distortions is challenging and the present paper makes progress in two important directions. First, it is nearly impossible to recover inflation-induced distortions in relative prices from observed actual prices. To see why, let p_{jt} denote the relative price *actually* charged for product j in period t and decompose it into its *flexible* relative price p^*_{jt} and the price gap gap_{jt} that is due to nominal rigidities:[2]

$$\ln p_{jt} = \ln p^*_{jt} + gap_{jt}. \tag{1}$$

Price distortions due to nominal rigidities for product j can then be summarized by the variance (over time) of the product's price gap $Var(gap_{jt})$.

Sticky price theories postulate that inflation affects price distortions, $Var(gap_{jt})$, and that there exists a product-specific optimal inflation rate $\ln \Pi = \ln \Pi^*_j$ that minimizes these distortions for product j.[3] Higher or lower-than-optimal inflation is predicted to increase price distortions, so that $\partial^2 Var(gap_{jt}) / (\partial \ln \Pi)^2$ is strictly positive at the point $\ln \Pi = \ln \Pi^*_j$. We refer to this second derivative as the *marginal effect* of inflation on relative price distortions. Many optimal policy recommendations in monetary economics rest on the postulated positive sign of this marginal effect.

[1] See, for instance, Woodford (2003), Galí (2015), Adam and Weber (2019) or Archarya, Challe and Dogra (2023).

[2] The flexible relative price is the price that would be charged in the absence of price adjustment frictions. It may itself be distorted, e.g., due to market power. Distortions due to nominal rigidities come on top of the distortions present in the flexible price.

[3] We define $\Pi = P_t/P_{t-1}$ as the gross inflation rate, where P_t denotes the price level in period t, so that $\ln \Pi$ is the net inflation rate.

A key difficulty with empirically testing for the presence of a positive marginal effect is that the price gap in equation (1) typically cannot be measured because the flexible relative price $\ln p^*_{jt}$ is not observed.[4] In fact, we show that there exists an important identification problem: the level of price distortions $Var(gap_{jt})$ cannot be recovered from actual relative prices $\ln p_{jt}$, whenever the flexible price $\ln p^*_{jt}$ contains some stationary stochastic component.[5] This may explain why the previous literature stopped short of identifying how price gaps depend on inflation (Wulfsberg (2016) and Nakamura, Steinsson, Sun and Villar (2018)), instead highlighted the difficulties associated with empirically recovering price gaps.[6]

A main contribution of the present paper is to show that the *marginal* effect of suboptimal inflation on price distortions can be identified from actual relative prices, even though the *level* of price distortions is not identified from actual prices. Intuitively, identification of the marginal effect is feasible because the variance of the flexible relative price $\ln p^*_{jt}$ on the right-hand side of equation (1) is independent of inflation. This causes the *variance* of the actual relative price $\ln p_{jt}$ on the left-hand side of equation (1) to be informative about the variance of the price gap and hence price distortions.

A second challenge this paper addresses is that it is generally difficult to establish a *causal* relationship in the data that goes from inflation to price distortions, when exploiting variation in aggregate inflation over time: outside hyperinflationary episodes or periods with large energy price shocks, aggregate inflation tends to move only slowly over time. As a result, movements in inflation are often hard to distinguish from a slow-moving time trend. Observed time trends in price dispersion might then reflect an observed trend in inflation or other trends that operate concurrently but are unrelated to inflation, e.g., a time trend in the variety of products or secular shifts in the distribution of mark-ups and productivity.

Our empirical approach overcomes this identification challenge by exploiting cross-sectional variation in the product-specific optimal inflation rate $\ln \Pi^*_j$ during a period in which the actual inflation rate $\ln \Pi$ was relatively stable in the U.K. economy.[7] The optimal

[4] In rare cases, additional information about marginal costs and the desired mark-up is available, which identifies the flexible relative price and thereby the price gap. Eichenbaum, Jaimovich and Rebelo (2011) estimate price gaps for supermarket goods using such information, but do not analyze how inflation affects price distortions.

[5] Alvarez, Lippi and Oskolov (2022) and Baley and Blanco (2021) recover the price gap distribution under the assumption that flexible prices follow a pure random walk (and have no stationary stochastic component). In our data, this assumption is strongly rejected, see appendix D.

[6] See section IV.A in Nakamura et al. (2018).

[7] To avoid the possibility that our results are driven by energy price or other particular shocks, we exclude the Covid and post-Covid period from our analysis.

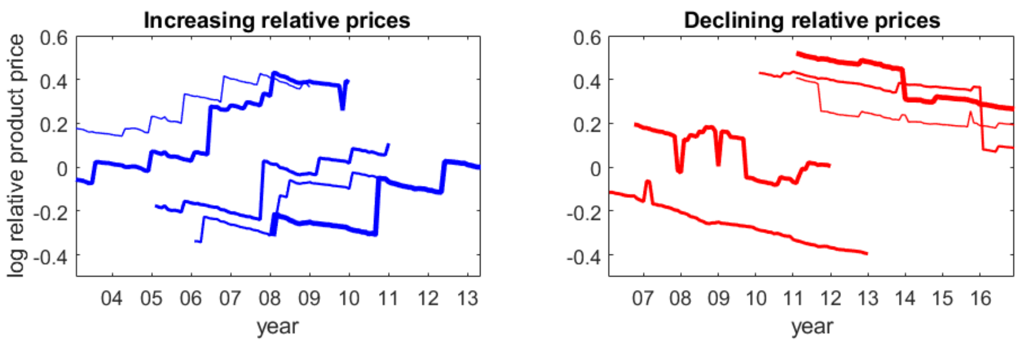

Figure 1: Different relative price trends in the expenditure category 'takeaway coffee latte'.

inflation rate $\ln \Pi_j^*$ is determined by product-specific fundamentals, such as the rate of productivity progress, input cost dynamics, or time trends in monopoly power. These fundamentals induce quasi-exogenous variation in the gap between actual and optimal inflation across products, $\ln \Pi - \ln \Pi_j^*$, that we can exploit to estimate the *causal* effect of suboptimal inflation on price distortions.

Figure 1 illustrates this approach. It depicts relative price time series for different products in the expenditure item 'takeaway coffee latte'. The left panel shows products with a positive trend in relative prices and the right panel shows products with a downward trend in relative prices.[8] Jumps in the relative price are associated with nominal price adjustments, while periods with no price adjustment display a constant reduction in the relative price due to positive and nearly constant coffee price inflation. Importantly, products in the left panel of figure 1 desire rising relative prices and thus prefer deflation in coffee prices ($\ln \Pi_j^* < 0$): an appropriate rate of deflation would cause their relative prices to increase at the desired speed without the need for nominal price adjustments, causing price adjustment frictions to be irrelevant. Conversely, products in the right panel prefer positive inflation in coffee prices ($\ln \Pi_j^* > 0$).

We show in the paper how one can exploit this kind of heterogeneity in product-specific optimal inflation rates $\ln \Pi_j^*$ to estimate the marginal effect of inflation on relative price distortions. Specifically, we show that both time and state-dependent pricing models imply a simple two-step estimation approach: in a first step, one takes out a time trend from the product's actual relative price, as depicted in figure 1, and computes the variance of residuals; in a second step, one regresses products' residual variance on the squared

[8] For illustration, we consider all products with at least 60 price observations in the expenditure item and then depict the five products with the most positive (left panel) and most negative (right panel) trends in relative prices in figure 1. Appendix A depicts all price time series with more than 60 price observations and the price index.

deviation of actual from optimal inflation $(\ln \Pi - \ln \Pi_j^*)^2$. We show that the regression coefficient in this second step identifies the *marginal* effect of suboptimal inflation on price distortions, i.e., $\partial^2 Var(gap_{jt}) / (\partial \ln \Pi)^2$ at the point $\ln \Pi = \ln \Pi_j^*$. Importantly, this structural empirical approach works without imposing restrictions on the behavior of the cross-sectional distribution of flexible relative prices over time.

Using U.K. micro price data, we show that the estimated marginal effect of inflation on relative price distortions is positive, in line with the theoretical predictions of time or state-dependent pricing models. The estimated marginal coefficient has the predicted positive sign in 98% of the expenditure items underlying the U.K. consumer price index and is statistically significant in 94% of them. Squared suboptimal inflation also has surprisingly high explanatory power for the residual variance of relative prices in the cross-section of products: for the median expenditure item, it explains 16% of residual variance. And in line with sticky price theories, we find that the marginal effect of suboptimal inflation on price distortions is stronger when prices are more rigid.

Having established that suboptimal inflation gives rise to price distortions over the lifetime of individual products, we turn consideration to the dispersion of prices in the cross-section of products. We document that cross-sectional relative price dispersion increased strongly in the U.K. over the considered sample period (1996 - 2016). Interestingly, this happened despite U.K. inflation displaying no time trend and only moderately-sized fluctuations.

We show that sticky price theory allows decomposing the cross-sectional dispersion of relative prices, at any given point in time, into the dispersion that is due to the deterministic component of flexible relative prices and a residual component. The residual component contains the effects of price distortions induced by nominal rigidities in the presence of suboptimal inflation and the effect of stochastic components of the flexible relative price. We show that the deterministic flexible price component accounts for 99% of the observed cross-sectional dispersion of actual relative prices and for nearly all of the observed increase in the cross-sectional price dispersion over time. This implies that relative price distortions due to nominal rigidities and suboptimal inflation account for at most 1% of the observed cross-sectional variance in relative prices.[9]

Sticky price theories also predict that residual price dispersion comoves positively (negatively) with inflation over time, whenever actual inflation lies above (below) the optimal inflation rate of the average product within an expenditure item. This is so because relative price distortions are predicted to comove with inflation in this way. Interestingly, we find considerable support for this prediction of sticky price theories in the data.

[9] For reasons explained in the paper, there is no covariance term.

For the average expenditure item, residual price dispersion comoves moderately positively with aggregate inflation over time: the correlation is equal to +0.47 and is statistically significant at the 5% level. This finding is consistent with the widely held notion in monetary economics that higher inflation is associated with larger relative price distortions. It should thus increase confidence in the economic relevance of key policy recommendations derived from monetary models, e.g., regarding the desirability of targeting low and stable inflation rates. It also fits with recent findings in Ascari, Bonmolo and Haque (2022), that high inflation rates are associated with a loss in the economy's output potential. Relative price distortions are one source of potential output losses associated with high inflation rates, as emphasized in the literature that infers price-induced misallocations from product-specific mark-ups (Baqaee, Farhi and Sangani (2022), Meier and Reinelt (2022)).

The paper is also related to Alvarez, Beraja, Gonzalez-Rozada and Neumeyer (2019) who estimate a nonlinear relationship between the cross-sectional dispersion of prices and inflation using data from Argentina. They find that cross-sectional price dispersion responds only weakly to inflation for inflation rates below 10%, but rises strongly for higher rates and eventually levels off. Relatedly, Sheremirov (2020) uses supermarket scanner data for the U.S. and documents how local cross-sectional price dispersion correlates with local inflation over time.[10] Instead of estimating a reduced-form relationship between the *cross-sectional* dispersion of prices and inflation over time, our structural approach calls for estimating *across-time* dispersion of prices at the level of *individual products* and relating it to a product-specific measure of suboptimal inflation in the cross-section of products.

Section 2 illustrates the empirical approach developed in this paper using the simplest possible setup. Section 3 derives the empirical approach for the full theory with time or state-dependent pricing frictions. Section 4 introduces the micro price data and section 5 presents our main empirical results. Sections 6 and 7 consider alternative estimation approaches, while section 8 analyzes the decomposition of cross-sectional relative price dispersion and the comovement of residual dispersion with inflation over time. A conclusion briefly summarizes.

2 The Approach in a Nutshell

This section illustrates with the help of a simple example how heterogeneity in the product-specific optimal inflation rate $\ln \Pi_j^*$ allows identifying the marginal effect of suboptimal inflation on relative price distortions. Suppose prices get adjusted in regular intervals

[10] Sara-Zaror (2022) extends the empirical approach of Sheremirov (2020) and documents that cross-sectional price dispersion strongly rises with the absolute deviation of inflation from zero, with the relationship becoming flatter for larger absolute inflation rates.

every $N > 1$ periods (Taylor (1979)) and the flexible relative price $p^*_{jt} = P^*_{jt}/P_t$ of product j evolves deterministically according to

$$\ln p^*_{jt} = \ln p^*_j - t \cdot \ln \Pi^*_j, \tag{2}$$

where $\ln p^*_j$ is a product-specific intercept and $\ln \Pi^*_j$ is a product-specific time trend, capturing differences in dynamics of marginal costs (or other factors) across products. Gross inflation is constant and equal to Π.

In this setting, the optimal inflation rate for product j is given by $\ln \Pi = \ln \Pi^*_j$ because its relative price then gets eroded at the desired rate $\ln \Pi^*_j$: the nominal price for product j can remain constant, so that price setting frictions do not matter for tracking the flexible relative price. When $\ln \Pi > \ln \Pi^*_j$ ($\ln \Pi < \ln \Pi^*_j$), the relative price gets eroded too quickly (slowly). As a result, adjustments of the nominal price have to be made to correct for the 'wrong' trend induced by inflation during non-adjustment periods. Due to price stickiness, these adjustments occur only occasionally, so that suboptimal inflation leads to deviations of the relative price from the flexible relative price.

Figure 2 illustrates this fact. It depicts the flexible relative price $\ln p^*_{jt}$ for three products ($j = 1, 2, 3$), for which the flexible relative price falls at rate $0 < \ln \Pi^*_1 < \ln \Pi^*_2 < \ln \Pi^*_3$.[11] Assuming that inflation $\ln \Pi$ is equal to $\ln \Pi^*_1$, the flexible relative price of product 1 coincides with the sticky relative price $\ln p_{jt}$, so that there are no relative price distortions. For product $j = 2$, inflation is such that the relative price falls insufficiently during non-adjustment periods. It then becomes optimal to choose a relative price that is lower than the flexible price in adjustment periods, to reduce the average gap between the sticky and the flexible relative price over the time period for which the price is sticky. Note that through this mechanism, suboptimally low inflation leads to deviations of the sticky relative price from the flexible relative price. These deviations are even stronger for product $j = 3$, for which the trend in the flexible relative price is even further away from the trend in relative prices induced by inflation. A larger gap between inflation and optimal inflation thus gives rise to larger deviations of the sticky relative price from the flexible relative price.

Since symmetric arguments apply when inflation is higher than optimal inflation, it is easy to verify that the variance of the gap_{jt} between the log of the sticky relative price and its time trend, i.e., the price distortion for product j, is a function of the square of suboptimal inflation:[12]

$$Var(gap_{jt}) = c \cdot (\ln \Pi - \ln \Pi^*_j)^2 \tag{3}$$

[11] The case where flexible relative prices rise over time is symmetric and thus omitted here.

[12] See appendix B for a proof.

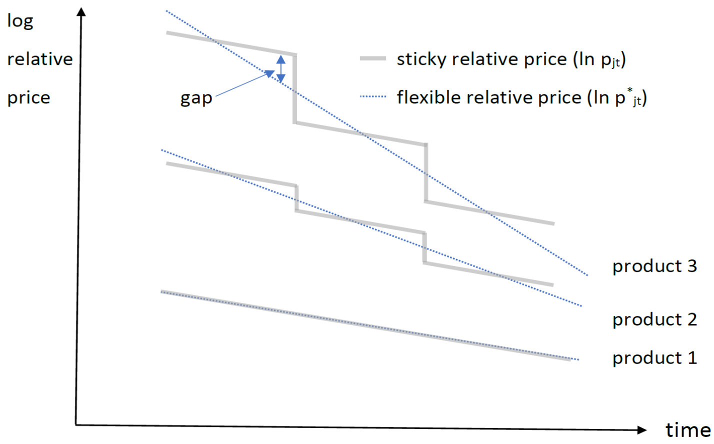

Figure 2: Relative price trends and price gaps

where $c = (N-1) \cdot (N+1)/12 > 0$ depends positively on the degree of price stickiness $N > 1$.

An important insight developed in this paper is that the relationship between suboptimal inflation and price distortions in equation (3) can actually be estimated using micro price data because (i) the product-specific optimal inflation rate Π_j^* is identified by the time trend in the sticky relative price, see figure 2, and (ii) price distortions, i.e., the gaps between the actual and the flexible price, are identified by the residuals of a regression of the sticky relative price on a time trend, see also figure 2.

Property (ii) fails to be true more generally when the flexible price also depends on unobserved idiosyncratic shocks. Yet, the next section shows that the presence of such shocks only requires adding a constant to equation (3). This holds true even when considering empirically more plausible price setting frictions, such as Calvo or menu-cost frictions.

Interestingly, considering the absolute size of price changes as a measure of relative price distortions, as in Nakamura et al. (2018), can result in misleading conclusions about the relationship between relative price distortions and suboptimal inflation. Appendix B.1 shows that the absolute size of price changes may respond to inflation in a setting where price distortions fail to do so, and that the absolute size of price changes may fail to respond to inflation in a setting where relative price distortions do indeed respond.

3 Inflation and Price Distortions: Theory

This section uses sticky price theory to derive a regression approach that allows identifying the *marginal* effect of suboptimal inflation on price distortions using micro price data. The regression approach turns out to be independent (to a second-order approximation) of whether price adjustment frictions are of a time-dependent or state-dependent nature and can be directly applied to micro price data. An attractive feature of our approach is that it does not require imposing any assumptions on the behavior of the cross-sectional distribution of flexible prices over time.

We consider the price setting problem of a firm facing a demand structure that closely matches the implicit demand structure underlying the price aggregation procedure performed by the U.K. Office of National Statistics (ONS). In particular, aggregate consumption is made up of Z different expenditure items, where an expenditure item $z \in \{1, ..., Z\}$ is a narrow product category, e.g., "Flatscreen TV, 30-inch display" or "CD-player, portable". Each expenditure item contains a large range of individual products $j \in [0, 1]$ with item-level consumption C_{zt} being given by a Dixit-Stiglitz aggregate of individual products j,

$$C_{zt} = \left(\int_0^1 C_{jzt}^{\frac{\theta-1}{\theta}} \mathrm{dj} \right)^{\frac{\theta}{\theta-1}}, \tag{4}$$

where C_{jzt} denotes the units of product j in item z consumed in period t and $\theta > 1$ the elasticity of substitution between products within the item. Product j should be interpreted as representing a physical object or service sold in a specific location over time, so that j indexes both products and selling locations. Aggregate consumption is given by

$$C_t = \prod_{z=1}^{Z} (C_{zt})^{\psi_z}, \tag{5}$$

where $\psi_z \geq 0$ denotes the (ONS) expenditure weight for item z, with $\sum_{z=1}^{Z} \psi_z = 1$. With this setup, demand for product j in item z is given by

$$C_{jzt} = \psi_z \left(\frac{P_{jzt}}{P_{zt}} \right)^{-\theta} \left(\frac{P_{zt}}{P_t} \right)^{-1} C_t, \tag{6}$$

where the item price level is defined as $P_{zt} = \left(\int_0^1 P_{jzt}^{1-\theta} \mathrm{dj} \right)^{\frac{1}{1-\theta}}$ and the aggregate price level is defined as $P_t = \prod_{z=1}^{Z} \left(\frac{P_{zt}}{\psi_z} \right)^{\psi_z}$.

Individual products are produced using a constant returns-to-scale production function

$$Y_{jzt} = \frac{A_{zt}}{G_{jzt} X_{jzt}} L_{jzt}, \tag{7}$$

where L_{jzt} denotes labor input and A_{zt} the level of productivity common to all producers of products in item z at time t.[13] $1/G_{jzt}$ is a deterministically evolving idiosyncratic productivity component, while $1/X_{jzt}$ is a stochastic idiosyncratic productivity component. Products in expenditure item z face each period an idiosyncratic exogenous exit risk $\delta_z > 0$.

[13]The setup can be generalized to include also capital in production, but this will not provide any additional insights as long as one has constant returns to scale jointly in all inputs.

Exiting products get replaced by new products that carry - for simplicity - the same product index, but that feature potentially different productivity levels $1/G_{jzt}$ and $1/X_{jzt}$. In equilibrium, the quantity of products consumed C_{jzt} is equal to the quantity produced Y_{jzt} in each period.

Firms can freely adjust inputs but face frictions for adjusting prices. Section 3.1 considers time-dependent price-setting frictions, while section 3.2 presents the case with state-dependent pricing frictions.[14]

3.1 Time-Dependent Price Setting Frictions

The price setting problem. Let $p_{jzt} \equiv P_{jzt}/P_{zt}$ denote the relative price charged for product j, where P_{jzt} denotes the products' nominal price and let p^*_{jzt} denote the flexible relative price, i.e., the price the firm would charge for product j in period t in the absence of price setting frictions.[15] Given the demand structure introduced above, appendix E.1 derives the following second-order approximation to the firm's nonlinear price setting problem with Calvo price adjustment frictions:

$$\max_{\ln p_{jzt}} -E_t \sum_{i=0}^{\infty} (\alpha_z \beta)^i \left(\ln p_{jzt} - i \ln \Pi_z - \ln p^*_{jzt+i}\right)^2, \tag{8}$$

where the parameter $\beta \in (0,1)$ denotes the firm's discount factor, $\alpha_z \in (0,1)$ the Calvo probability that the price cannot be adjusted in the period, and Π_z the gross inflation rate in item z. The firm's relative price in period $t+i$ is given by $\ln p_{jzt} - i \ln \Pi_z$, which shows that the reset price $\ln p_{jzt}$ chosen by the firm gets eroded by inflation as long as prices do not adjust. Deviations of the firm's relative price from its flexible relative price $\ln p^*_{jzt+i}$ give rise to profit losses that are quadratic in the size of the deviation.

The dynamics of the flexible price. A key object of interest in problem (8) is the flexible relative price p^*_{jzt}. This price is observed by the firm but not by the econometrician. We consider the following general stochastic process:

$$\ln p^*_{jzt} = \ln p^*_{jz} - t \cdot \ln \Pi^*_{jz} + \ln x_{jzt}. \tag{9}$$

The term $\ln p^*_{jz}$ is an unobserved product fixed-effect that is drawn at the time of product entry from some arbitrary and potentially time-varying distribution. It is a stand-in for

[14] While our approach also applies to a unified setup with time- *and* state-dependent adjustment frictions (Calvo plus), we treat the two frameworks separately. Analytic tractability in the two frameworks relies on slightly different assumptions for the shock process X_{jzt}, and pooling the two frameworks together would require imposing these assumptions simultaneously, reducing the generality of our results.

[15] Due to the presence of product-specific monopoly mark-ups, the flexible relative price can differ from the socially efficient relative price. In the special case, where desired monopoly mark-ups are identical across products or simply absent, the frictionless relative price is equal to the efficient relative price.

unobserved location-specific effects such as differences in the level of marginal costs, wages, rents, service or quality components of the product. It also captures the presence of product and location-specific flexible-price mark-ups.

The term $\ln \Pi_{jz}^{*}$ in equation (9) captures a product-specific time trend in the relative price and also denotes the product-specific optimal inflation rate, as discussed in section 2. It is drawn at the time of product entry from an arbitrary distribution that may also depend on time. The trend in the relative price may reflect a product-specific rate of productivity progress, induced for instance by learning-by-doing effects, or product-specific marginal cost trends induced by trends in wages or rents that are specific to the particular location where the product is sold. Our empirical approach will exploit variation in $\ln \Pi_{jz}^{*}$ across products j to identify the distortionary effects of inflation in expenditure item z.[16] While we consider a linear time trend in relative prices here, our empirical analysis will also consider nonlinear time trends.[17]

Finally, there is an idiosyncratic stochastic component $\ln x_{jzt}$ in equation (9), which captures idiosyncratic fluctuations induced by changes in productivity or service components at the product level. There is no common component in these shocks because the left-hand side of equation (9) features the log *relative* price, which already absorbs common components in the nominal price (at the level of a narrowly-defined expenditure category).

The stochastic process governing the idiosyncratic components $\ln x_{jzt}$ is assumed to be stationary and Markov and to be the same for all products within a narrowly-defined expenditure item z.[18] We can thus normalize idiosyncratic shocks so that $E[\ln x_{jzt}] = 0$. We effectively rule out that idiosyncratic shocks $\ln x_{jzt}$ follow a random walk. This seems innocuous because our data strongly reject a random walk in $\ln x_{jzt}$, as shown in appendix D.[19]

Equation (9) allows the cross-sectional distribution of flexible prices in expenditure item z is to vary over time in important ways, even when abstracting from idiosyncratic shocks: (i) for a given set of products, heterogeneity in the relative price trends $\ln \Pi_{jz}^{*}$ changes the cross-sectional distribution of flexible relative prices over time; (ii) as products exit and enter the market, newly entering products may have different product-specific intercepts $\ln p_{jz}^{*}$ and time trends $\ln \Pi_{jz}^{*}$ than exiting products. Since the parameters $(\ln p_{jz}^{*}, \ln \Pi_{jz}^{*})$ of

[16] Our introductory figure 13 illustrates that relative price trends can differ considerably across products, even within relatively narrowly defined expenditure categories, and appendix I formally rejects the absence of trend heterogeneity using a bootstrapping approach.

[17] The relative price dynamics of products appear well-approximated by a linear trend, see figure A.XI in the November 2018 working paper version of Argente and Yeh (2022).

[18] We relax the assumption of identical idiosyncratic shock processes across products in section 6.

[19] This finding does not depend on assuming Calvo frictions.

newly incoming products are drawn from an arbitrarily time-varying distribution, the setup leaves the evolution of the cross-sectional distribution of flexible relative prices unrestricted over time.

The optimal reset price. Considering the limit $\beta \to 1$, the optimal reset price $\ln p_{jzt}^{opt}$ solving problem (8) is given by[20]

$$\ln p_{jzt}^{opt} = (\ln p_{jzt}^{*} - \ln x_{jzt}) + \left(\frac{\alpha_z}{1-\alpha_z}\right)(\ln \Pi_z - \ln \Pi_{jz}^{*}) + f(x_{jzt}), \tag{10}$$

where

$$f(x_{jzt}) \equiv (1-\alpha_z)E_t \sum_{i=0}^{\infty} \alpha_z^i \ln x_{jzt+i}. \tag{11}$$

The first term on the r.h.s. of equation (10), $\ln p_{jzt}^{*} - \ln x_{jzt}$, captures the deterministic component of the flexible price (9). The second term captures the effects induced by deviations of actual inflation $\ln \Pi_z$ from the product-specific optimal inflation rate $\ln \Pi_{jz}^{*}$. The last term in equation (10) captures effects due to the presence of time-varying idiosyncratic components. Equation (11) shows that it is the expected discounted value of the idiosyncratic shock over the lifetime of the price that matters for this component.

Only the second term on the r.h.s. of equation (10) depends on inflation. If inflation exceeds optimal inflation ($\ln \Pi_z > \ln \Pi_{jz}^{*}$), then the reset price gets pushed up to compensate for the suboptimally high rate of future erosion of the relative price during periods in which the price does not adjust. The opposite is true if inflation falls short of optimal inflation ($\ln \Pi_z < \ln \Pi_{jz}^{*}$).

Importantly, the optimal reset price $\ln p_{jzt}^{opt}$ is equal to the expected value of the flexible price over the expected lifetime of the reset price. Therefore, an initial period in which relative prices lie above (below) the flexible price is followed - in expectation - by a period in which the relative price falls short of (exceeds) the flexible price. This explains how - according to the theory - deviations of inflation from its optimal level induce *additional* dispersion of prices around the flexible level. This effect is stronger if prices are more sticky: for a given deviation of inflation from its optimal level, reset prices react by more, the higher is the degree of price stickiness (α_z).

The dynamics of the actual relative price. While equation (10) determines the optimal reset price in periods where prices adjust, the dynamics of the actual relative price for product j in expenditure item z are given by

$$\ln p_{jzt} = \xi_{jzt}(\ln p_{jzt-1} - \ln \Pi_z) + (1-\xi_{jzt}) \ln p_{jzt}^{opt}, \tag{12}$$

[20] See appendix E.1 for a derivation.

where $\xi_{jzt} \in \{0,1\}$ is an *iid* random variable capturing periods with price adjustment ($\xi_{jzt} = 0$ with probability $1-\alpha_z$) and no-adjustment ($\xi_{jzt} = 1$ with probability α_z). In periods in which the price does not adjust, the relative price falls with inflation.

It also follows from equation (12) that the actual relative price inherits the product-specific time trend present in the optimal price p_{jzt}^{opt}, which in turn inherits the trend from the flexible price p_{jzt}^{*}, see equation (10). We show next that the variability of the actual price $\ln p_{jzt}$ around this trend is a function of (i) the deviation of inflation from its optimal level, and (ii) the idiosyncratic shocks $\ln x_{jzt}$. This insight turns out to be key for identifying the marginal effect of suboptimal inflation on price distortions.

The first-stage regression. Equation (12) implies that the dynamics of the actual relative price can be expressed as[21]

$$\ln p_{jzt} = \ln p_{jz}^{*} - t \cdot \ln \Pi_{jz}^{*} + u_{jzt}, \tag{13}$$

with regression residuals given by

$$u_{jzt} = \xi_{jzt}(u_{jzt-1} - (\ln \Pi_z - \ln \Pi_{jz}^{*})) + (1-\xi_{jzt})(f(x_{jzt}) + \frac{\alpha_z}{1-\alpha_z}(\ln \Pi_z - \ln \Pi_{jz}^{*})), \tag{14}$$

where $f(x_{jzt})$ is defined in equation (11). Since $E[u_{jzt}|p_{jz}^{*}, \Pi_{jz}^{*}] = E[u_{jzt}] = 0$, the coefficients and residuals in equation (13) can be recovered via OLS estimation.[22]

Equation (13) shows that the actual relative price inherits - in terms of its level and time trend - the deterministic components of the flexible relative price (9). This implies that the effects of price distortions can only be contained in the residuals of regression (13). In fact, the residuals u_{jzt} are the main reason why regression (13) is of interest. We now discuss the properties of these residuals.

The level of price distortions is not identified. Due to price stickiness ($\alpha_z > 0$), the regression residuals u_{jzt} in equation (14) fail to be informative about the idiosyncratic shocks, as previously emphasized by Nakamura, Steinsson, Sun and Villar (2018). The underlying intuition is quite straightforward: in periods where prices are not adjusted, residuals reveal no new information about idiosyncratic shocks; and in periods where prices are adjusted, their adjustment depends considerably on the expected future values of the idiosyncratic shock, particularly when prices are sticky, see equation (11).

Due to this dependence on expected future shock values, the information that becomes available upon a price adjustment - the term $f(x_{jzt})$ in equation (14) - fails to identify the underlying process of idiosyncratic shocks $\ln x_{jzt}$. Appendix C proves the following result:

---

[21] See appendix E.3 for a derivation.

[22] To simplify the exposition, we abstract here from small sample issues, which are discussed in section 4 and in appendix G.

Proposition 1 *In the presence of price stickiness, observed prices* $\ln p_{jzt}$ *fail to identify the process for idiosyncratic shocks* $\ln x_{jzt}$*. Consider, for example, a stationary discrete* N*-state Markov process for* $f(x_{jzt})$*. It can be generated either by a stationary Markov process for* $\ln x_{jzt}$ *with* N *states or an infinite number of different Markov processes with* $M > N$ *states, where* M *is arbitrary and where* $M - N$ *states in the* M*-state process are not states in the* N*-state process.*

Intuitively, different fundamental processes for $\ln x_{jzt}$ give rise to identical processes for $f(x_{jzt})$, because they imply the same conditional expectation in equation (11). Since the process for $\ln x_{jzt}$ cannot be identified from observed prices, it is impossible to estimate the level of price distortions, i.e., the gap between the actual and flexible price. This may explain why the literature has to date not come up with an estimate of how price distortions respond to (suboptimal) inflation.

One way to deal with the identification problem in proposition 1 is to bring in additional information. This is the strategy pursued in Eichenbaum, Jaimovich and Rebelo (2011) who exploit information on marginal costs in supermarkets to identify price distortions (but do not analyze how they depend on inflation). Yet, information on marginal costs is only rarely observed.

An alternative approach to handle the identification problem is to impose additional identifying assumptions. This is the approach pursued in Baley and Blanco (2021) and Alvarez, Lippi and Oskolkov (2022), who show that the distribution of price distortions can be recovered from observed price changes, whenever $\ln x_{jzt}$ is a pure random walk, i.e., does not contain stationary shock components.[23] In our data, the hypothesis of a pure random walk in $\ln x_{jzt}$ is strongly rejected, see Appendix D.

We now show that it is simply not necessary to identify the *level* of price distortions to estimate the marginal effect of suboptimal inflation on price distortions.

Second-stage regression: the marginal effect of suboptimal inflation. While the level of price distortions cannot be identified from observed prices, the theory predicts that the *marginal effect* of suboptimal inflation on price dispersion can be identified. In fact, equation (10) highlights that any non-zero gap $\ln \Pi_z - \ln \Pi^*_{jz}$ generates front-loading

[23] With a random walk, we have $f(x_{jzt}) = \ln x_{jzt}$, so that the size of innovations between price reset times identifies the innovation variance of the random walk. Yet, the result in proposition 1 applies also in the case where $\ln x_{jzt}$ is non-stationary but still contains some stationary component, e.g., when $\ln x_{jzt}$ is the sum of a random walk process $\ln y_{jzt}$ and an independent stationary Markov process $\ln z_{jzt}$. We then have $f(\ln x_{jzt}) = \ln y_{jzt} + f(\ln z_{jzt})$, so that the process $\ln z_{jzt}$ and thus $\ln x_{jzt}$ can again not be identified, even if the process for $\ln y_{jzt}$ could be perfectly recovered from the data.

of prices at the time of price adjustment, as captured by the term $\frac{\alpha_z}{1-\alpha_z}(\ln \Pi_z - \ln \Pi^*_{jz})$. Likewise, during non-adjustment periods, a gap between actual and optimal inflation leads to a drift in the gap between actual and flexible relative prices. Both of these features contribute to the variance of residuals u_{jzt} in the first-stage regression (14).

Therefore, the variance of first-stage residuals satisfies the following relationship:[24]

Proposition 2 *The variance of the first-stage residual in equation (13) is given by*

$$Var(u_{jzt}) = v_z + c_z \cdot (\ln \Pi_z - \ln \Pi^*_{jz})^2, \tag{15}$$

where the intercept

$$v_z \equiv Var\left((1-\alpha_z)E_t \sum_{i=0}^{\infty} \alpha_z^i \ln x_{jzt+i}\right) \tag{16}$$

is a function of the idiosyncratic shock $\ln x_{jzt}$ *and the price stickiness parameter* α_z*, and*

$$c_z \equiv \frac{\alpha_z}{(1-\alpha_z)^2}. \tag{17}$$

The intercept term v_z in equation (16) contains both idiosyncratic flexible price components and price distortions due to price stickiness. In particular, price stickiness causes the loading on the current idiosyncratic shock to be too low relative to the flexible price case. Yet, without additional information, it is impossible to further decompose to what extent v_z reflects efficient or inefficient forces, which is precisely the feature preventing identification of the *level* of price distortions from observed actual prices. The second term on the r.h.s. of equation (15) captures the marginal effect of suboptimal inflation on price distortions. According to the theory, the coefficient c_z is positive and an increasing function of the degree of the Calvo price stickiness parameter α_z.

Equation (15) is a second-stage regression equation and a key equation we shall empirically exploit in the present paper. It uses the residual variance from the first-stage equation (13) as left-hand side variable, and the gap between the (item-level) inflation rate Π_z and the product-specific optimal inflation Π^*_{jz} as right-hand side variable, where Π^*_{jz} is also identified by the first-stage regression, see equation (13). Equation (15) implies that the marginal effect of suboptimal inflation on price distortions can be estimated using a cross-section of products for which price stickiness and the process driving idiosyncratic shocks are the same. (A more general estimation approach allowing for heterogeneous price stickiness and heterogeneous idiosyncratic shock processes is derived in sections 6 and 7).

The next section briefly shows that the results derived thus far are not specific to the case with Calvo frictions, but apply in slightly different form also in a setting with menu-cost frictions.

[24] See appendix E.4 for a derivation.

3.2 State-Dependent Price Setting Frictions

To obtain closed-form solutions in a model with state-dependent pricing, we consider a continuous-time setup and derive the continuous-time analogue to proposition 2 using only slightly different assumptions on the idiosyncratic shock process. The firm's objective (8) becomes:

$$\max_{\{\tau_{jzk},\Delta \ln p_{jzk}\}_{k=1}^{\infty}} -E\left[\int_t^{\infty} e^{-\rho(s-t)}\left(\ln p_{jzt+s}-\ln p^*_{jzt+s}\right)^2 ds+\kappa_z\sum_{k=1}^{\infty} e^{-\rho(\tau_{jzk}-t)}\right] \tag{18}$$

The parameter $\rho > 0$ is the discount rate, τ_{jzk} are the random adjustment times and κ_z is the cost paid at the times of adjustment. As with time-dependent frictions, the firm's relative price in period $\tau_{jzk}+s$ is given by $\ln p_{jz\tau_{jzk}} - s\ln \Pi_z$ until it is adjusted again, reflecting relative price erosion due to inflation.

The flexible relative price $\ln p^*_{jzt}$ follows a continuous-time analogue of (9) with an additional restriction on the idiosyncratic process $\ln x_{jzt}$, namely that it assumes values from a finite grid $\{\ln x_{1z},\ldots,\ln x_{Nz}\}$ and switches from grid point i to grid point j with Poisson intensity λ^X_{izj}.[25]

Appendix F shows that under $\rho \to 0$ and for sufficiently small adjustment cost κ_z,[26] relative price dynamics in the menu-cost model also follow equation (13), with residuals satisfying $E[u_{jzt}|p^*_{jz},\Pi^*_{jz}] = E[u_{jzt}] = 0$ and residual variance that depends on product-specific suboptimal inflation:

$$Var(u_{jzt}) = Var(\ln x_z) + c^{MC}_z\cdot(\ln \Pi_z - \ln \Pi^*_{jz})^2 + O((\ln \Pi_z/\Pi^*_{jz})^4), \tag{19}$$

where the intercept is again a function of the idiosyncratic shock process, the quadratic term depends on suboptimal inflation, and $O((\ln \Pi_z/\Pi^*_{jz})^4)$ denotes a fourth-order approximation error. The coefficient c^{MC}_z is now a function of the shock process parameters $\lambda^X_{iz} = \sum_{j\neq i}\lambda^X_{izj}$:

$$c^{MC}_z \equiv E\left[\frac{1}{(\lambda^X_{iz})^2}\right].$$

If λ^X_{iz} is constant across states, then

$$c^{MC}_z = \frac{1}{\Lambda_z^2} \tag{20}$$

where Λ_z is equal to the adjustment frequency (again up to a fourth-order approximation error $O((\ln \Pi_z/\Pi^*_{jz})^4)$) and thus can be directly estimated from the data. The coefficient

[25] The restriction is very mild because we do not impose any assumption on the switching intensities. Even though we are ruling out all processes with continuous paths, we can still approximate them well with a sufficiently fine grid.

[26] Note that we *do not* consider a limiting case $\kappa_z \to 0$, instead our result holds *for all* $\kappa_z \leq \bar{\kappa}$ for some $\bar{\kappa} > 0$.

c_z^{MC} differs slightly from the one in the discrete time setup with Calvo friction, see equation (17), for which $\Lambda_z = 1 - \alpha_z$. This is so because multiple price adjustments can happen per unit of time under continuous time modeling. Also, the coefficient c_z^{MC} does not depend on the menu cost κ_z, under the maintained assumption that menu costs are small enough. Differences in κ_z have only fourth-order effects in equation (19). This is the reason why equation (19) now holds only up to a fourth-order approximation error, while it was exact in the Calvo setup (given the quadratic approximation to the firm's objective), see equation (15).

Perhaps surprisingly, the results obtained from the state-dependent model are (to the considered order of approximation) virtually the same as obtained in proposition 2 for the time-dependent model.

4 Micro Price Data and Empirical Product Definition

In our empirical analysis, we use the micro price data underlying the official U.K. consumer price index (CPI). The advantage of using CPI micro price data is that it covers a wide range of consumer expenditures. Moreover, the UK CPI data display quite strong relative price trends and significant variation of these trends across products.[27] This is essential for our identification approach, which relies on cross-product variation in relative price trends.

We consider about 20 years of micro price data (February 1996 to December 2016), which is obtained from the Office of National Statistics (ONS). The data are monthly and classified into narrowly defined expenditure items (e.g., flat panel TV 33inch, men's shoes trainers, vegetarian main course). Given the sample selection described further below, we consider 1048 different expenditure items and 15.2 million price observations over the sample period.

A product within an item is a sequence of price observations for a particular physical object or service sold in a particular location. Otherwise identical objects or services sold in different locations will thus be treated as different products in our empirical approach. The same holds true when a product in a specific location and expenditure item gets substituted by a new product: so-called 'comparable' and 'non-comparable' substitutions will be treated as separate products. This allows us to account for location and product-specific components in the most flexible way.

Using this data, we estimate the first-stage equation (13) for every product in the sample. Thereafter, we estimate the second-stage equation (15) at the level of the expenditure

[27] See Adam and Weber (2023) who estimate the optimal aggregate inflation rate for the U.K. from relative price trends.

item $z = 1, ...1048$, considering all products j in the item:

$$\widehat{Var}(u_{jzt}) = v_z + c_z \cdot (\widehat{\ln \Pi_z/\Pi^*_{jz}})^2 + \varepsilon_{jz}, \tag{21}$$

where $\widehat{Var}(u_{jzt})$ is the estimated variance of first-stage residuals of product j in item z and $\widehat{\ln \Pi_z/\Pi^*_{jz}}$ the corresponding first-stage estimate of the gap between item-level inflation and product-specific optimal inflation. As shown in appendix G, the suboptimal inflation rate $\ln \Pi_z/\Pi^*_{jz}$ can be estimated as the time trend coefficient of the product's *nominal* price:

$$\ln P_{jzt} = \ln \widetilde{a}_{jz} + t \cdot \ln \Pi_z/\Pi^*_{jz} + \widetilde{u}_{jzt} \tag{22}$$

Intuitively, a trend in a product's nominal price indicates that the trend in the relative price induced by inflation is suboptimal and hence requires systematic adjustments in the nominal price. Appendix G shows that equations (13) and (22) jointly form a seemingly unrelated first-stage regression system and that the second-stage estimate of the coefficient c_z in equation (21) is biased towards zero as a result of first-stage estimation error. Simulations of a calibrated price setting model in appendix H verify that our procedure indeed generates at most a downward bias in the estimated coefficient. The simulations also shows that this downward bias is less pronounced when using products with a larger number of price observations. The second-stage estimate of c_z thus provides a *lower bound* of the true marginal effect of suboptimal inflation on price distortions. Since we are interested in rejecting the null hypothesis of inflation creating *no* price distortions ($H_0 : c_z = 0$), small sample bias works against our main finding.

Since our empirical approach exploits heterogeneity in suboptimal inflation rates $\ln \Pi_z/\Pi^*_{jz}$ across products j within an item z, we use a bootstrapping procedure to show that we can indeed reject the null hypothesis of no heterogeneity in suboptimal inflation, see appendix I.

The estimation of the second-stage equation (21) delivers 1048 estimates of c_z, i.e., one for each expenditure item z. These estimates identify the marginal effect of suboptimal inflation on price distortions in each item, provided our two key identifying assumptions (identical degrees of price rigidity and identical stochastic processes driving idiosyncratic shocks within expenditure items) are satisfied. These assumptions will be relaxed in sections 6 and 7, where we consider alternative estimation approaches.

The data methodology follows the one used in Adam and Weber (2023), who provide further details. Starting from the raw micro price data, we delete products with duplicate price observations in a given month[28] and also delete all price observations flagged by ONS

[28] Duplicate price quotes can arise because ONS does not disclose all available locational information underlying the data, so that in rare cases we cannot uniquely identify the product price.

Total number of price quotes used:	15.2 million			
	mean	median	min	max
Number of products per item	734	573	50	3,201
Number of price quotes per item	14,485	10,772	407	73,313

Table 1: Basic product and price statistics

as "invalid." Furthermore, we split observed price trajectories for ONS product identifiers, whenever ONS indicates a change in the underlying product, i.e., a comparable or non-comparable product substitution, and whenever price quotes are missing for two months or more. This conservative splitting approach insures that we do not lump together products that might in fact be different. It leads to a refined product definition that we use to compute relative prices by deflating nominal product prices with a quality-adjusted item price index.

We consider only expenditure items for which the item price index, computed from our micro price data, replicates the official item price index provided by ONS sufficiently well. We exclude cigarette items because their price dynamics are largely the result of tax changes. In addition, we only consider products with a minimum length of six price observations after eliminating sales prices from the sample.[29] We account for outliers by eliminating the 5% of products with the highest values for $\widehat{Var}(u_{jzt})$ and for $(\widehat{\ln \Pi_z/\Pi^*_{jz}})^2$ within each expenditure item and consider only expenditure items containing at least 50 products.[30] This leads us to the 1048 expenditure items that we use in our empirical analysis. Table 1 reports basic statistics on the number of products and price observations per item. The average number of price observations per product is equal to 21 monthly observations and the average number of price changes per products is equal to 2.

4.1 Descriptive Statistics of the Regression Inputs

This section presents key descriptive statistics about the variables entering the first and second-stage regression equations. Since we run these regressions for more than one thousand expenditure items, we report the distribution of key moments of the variables of interest in the cross-section of items.

The left column in figure 3 depicts the distribution of the mean and standard deviation of the length of product life. For most items, the mean product length ranges between 10

[29] We identify sales prices using the sales flag recorded by the ONS price collectors.

[30] We also eliminate expenditure items for which the estimated residual variances are zero for all products. The latter occurs when prices never adjust within an item, which is the case for less than a handful of items capturing administered prices.

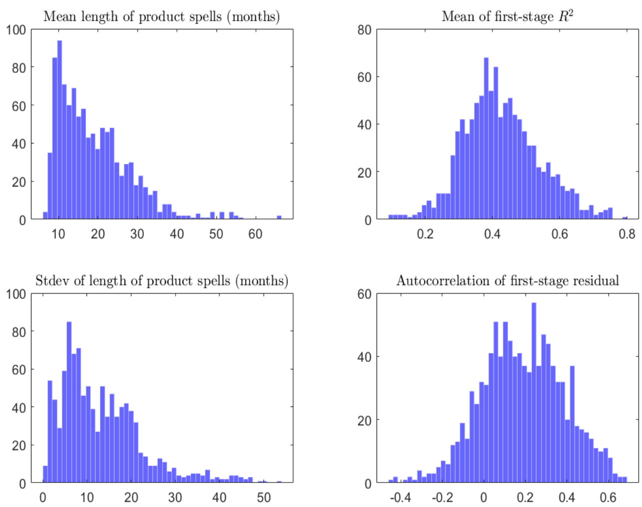

Figure 3: Descriptive statistics: first-stage regression

and 30 months, which is long enough to estimate an intercept and slope parameter in our first stage. The bottom left panel in figure 3 highlights that there is a considerable amount of variation in the length of product lives within each item. We exploit this feature below to also present estimates that are based on products whose price can be observed for at least 12 or 24 months (instead of 6 months in our baseline).

The top right panel in figure 3 reports the distribution of the mean R^2 values of the first-stage regression (13) across items. For most items, the time trend captures between 30% and 60% of the observed variation in relative prices. The remainder of the variation goes into the regression residual, the variance of which enters our second-stage regression. The bottom right panel in figure 3 depicts the distribution of the mean autocorrelation of these residuals, which is considerably smaller than one. Appendix D provides formal tests showing that residuals do not follow a random walk.[31]

The top left panel of figure 4 reports the mean standard deviation of the regression residuals across items.[32] For most items, the mean standard deviation ranges between 2% and 4%. The standard deviation of the standard deviation of residuals is shown in the bottom left panel of figure 4. It highlights that there is a considerable amount of variation in the left-hand side variable of our second-stage regression, which is desirable.

The top right panel in figure 4 depicts the distribution of item-level means of the

[31] The random walk is rejected even if one abstracts from the presence of a deterministic trend in relative prices.

[32] We report moments of the non-squared variables entering the second-stage regression to increase readability of the figures.

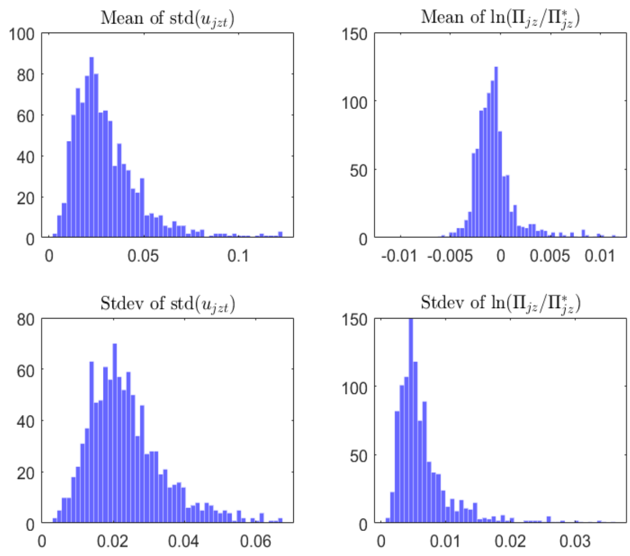

Figure 4: Descriptive statistics: second-stage regression

suboptimal inflation rate.[33] For the vast majority of items, the average suboptimal inflation rate lies between $\pm 0.5\%$ per month. The lower right panel in figure 4 shows the within-item standard deviation of suboptimal inflation. The cross-product variation is significant, with a standard deviation ranging between 1/3 and 2/3 of a percent on a monthly basis in most items. This shows that our second-stage right-hand side variable also displays a considerable amount of variation.

5 Price Distortions at the Product Level: Empirical Results

This section reports our estimates of the coefficient c_z in equation (21), which captures the marginal effect of suboptimal inflation on relative price distortions. The estimation will be carried out separately for each of the 1048 U.K. expenditure items in our sample, to maximize the chances that the key identifying assumptions used in deriving equation (21) are satisfied (identical degrees of price rigidity and identical stochastic processes for idiosyncratic disturbances for products within an item). Sections 6 and 7 will present alternative estimation approaches that relax these assumptions.

5.1 Baseline Results

Figure 5 presents our baseline estimation outcome. The left panel depicts the distribution of the estimated coefficient c_z in equation (21) across expenditure items z. The coefficient

[33]See appendix K for information on the cross-sectional distribution of the product-specific optimal inflation rate Π^*_{jz}.

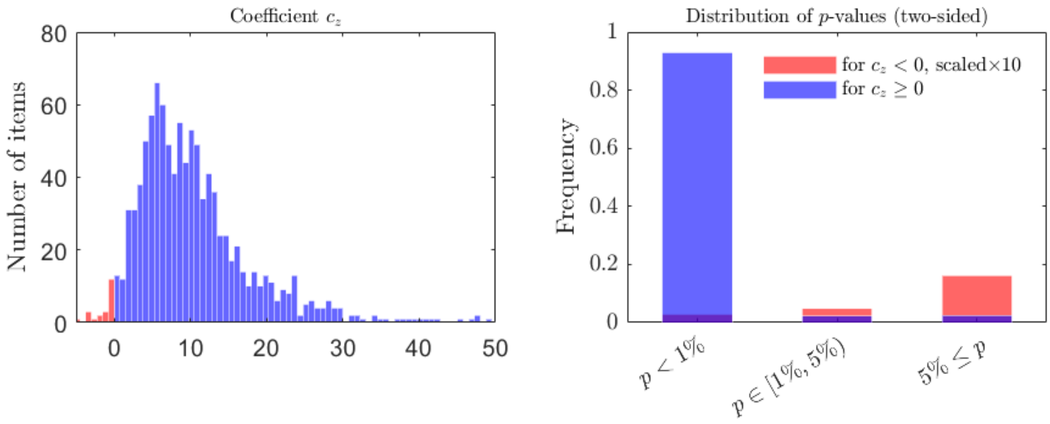

Figure 5: Baseline results from estimating equation (21), bootstrapped p-values

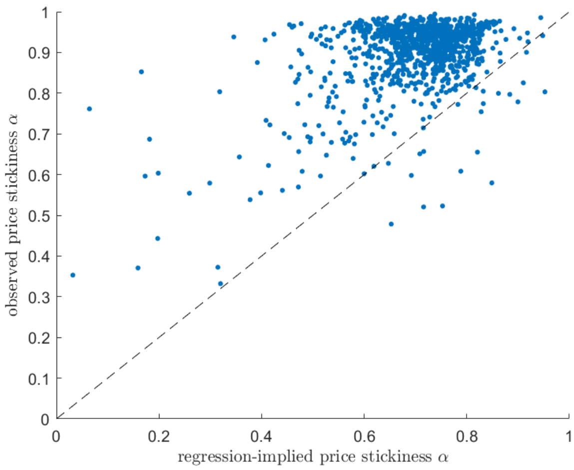

Figure 6: Observed versus regression-implied price stickiness (α_z)

c_z captures the marginal effect of suboptimal inflation on relative price distortions and we find that 98% of the estimated coefficients are positive (shown in blue), in line with the predictions of sticky price theories, while only 2% of the estimated coefficients are negative (shown in red). The right panel in figure 5 depicts the distribution of bootstrapped p-values. 94% of the estimates are significantly positive at the 5 % level and 92% at the 1% level. Only 1% of the coefficients are negative and significant at the 5% level. Figure 5 thus provides overwhelming support for the notion that suboptimal inflation gives rise to relative price distortions at the product level.

Row 1 in table 5.2 reports further details of the regression outcome. Interestingly, the median adjusted R^2 value of the second-stage regression (21) is 16%. The square of suboptimal inflation is thus not only statistically significant but also explains a sizable part of the cross-product variance of first-stage residuals within each item. This is the case despite the fact that first-stage estimation error likely contributes to unexplained variance on the left-hand side of the second-stage regression (21).

The point estimates for c_z are not only positive and statistically significant, but also quantitatively large: the average point estimate is close to 12. It implies that a monthly inflation rate that lies 1% above (or below) its optimal level[34] increases the standard deviation of first-stage residuals by 3.5 percentage points.[35]

Since first-stage estimation error causes the second-stage estimates of c_z to be biased towards zero, we refrain here from a further quantitative interpretation of the point estimates. Section 8 will assess the quantitative importance of relative price distortions using the unbiased first-stage residuals.

Sticky price theories suggest that the coefficient c_z is determined by the adjustment rate for prices, see equations (17) and (20). We can thus compute the price adjustment rate *implied* by any given coefficient estimate and see how it covaries (in the cross-section of items) with the *actual* price adjustment rate measured directly from price data. Inverting equation (17) to solve for the regression-implied share of non-adjusting products α_z, we obtain:[36]

$$\alpha_z = \frac{1 + 2c_z - \sqrt{1 + 4c_z}}{2c_z}$$

Figure 6 presents a scatter plot with the regression-implied α_z (x-axis) and the share of non-adjusters α_z measured directly from the data (y-axis). The two measures display a positive correlation equal to +0.45, which is significant at the 1% level. This shows that relative price distortions due to suboptimal inflation are larger for items featuring lower price-adjustment rates, as predicted by sticky price theory. However, the vast majority of items lie above the 45-degree line depicted in figure 6, while theory predicts the two measures to align. The downward bias in the regression-implied value for α_z likely arise due to the downward bias in our estimated coefficients c_z associated with first-stage estimation error, see appendix G.

Overall, our baseline results provide strong support for the notion that suboptimal inflation distorts relative prices.

[34] The 1% number corresponds roughly to a 2 standard deviation variation for the typical item, as the standard deviation of suboptimal inflation ranges between 1/3% and 2/3% per month for most items, see the lower right panel in figure 4.

[35] The predicted increase in the variance is $0.12\% = 12 \cdot (0.01)^2$ from which we obtain $\sqrt{0.12\%} = 3.5\%$.

[36] We perform the inversion only for items with strictly positive estimated c_z, which is true for 98% of items. The other root of the polynomial is larger than one and can be ruled out. In the discrete-time setup, the adjustment rate is equal to $1 - \alpha_z$. For the continuous-time setup, we can recover the adjustment rate as $\Lambda_z = -\ln \alpha_z$ and obtain very similar results.

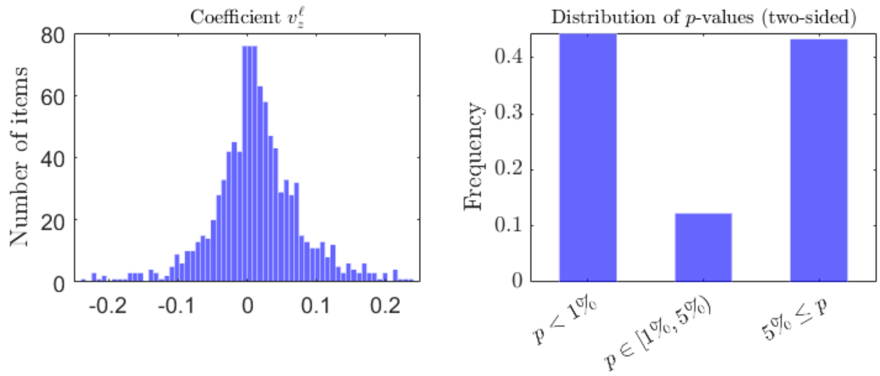

Figure 7: Robustness to adding a linear term (equation (23), estimates for c_z are reported in table 5.2)

5.2 Robustness of Baseline Approach

We now explore the robustness of our baseline results. The outcomes of all robustness exercises are summarized in table 5.2, which reports also the baseline outcome for reference.

Adding Linear Terms. Sticky price theories predict that only the squared deviation of inflation from its optimal level explains the variance of first-stage regression residuals to second order. In particular, the linear gap between inflation and optimal inflation should have a zero coefficient when added to the right-hand side of equation (21). One can test this overidentifying restriction by running an alternative second-stage regression of the form

$$\widehat{Var}(u_{jzt}) = v_z + v_z^l \cdot \ln \widehat{\Pi_z / \Pi_{jz}^*} + c_z \cdot (\ln \widehat{\Pi_z / \Pi_{jz}^*})^2 + \varepsilon_{jz}. \tag{23}$$

We then check whether the coefficients v_z^l are indeed approximately equal to zero and whether the estimates for c_z remain unaffected by the presence of the linear term.

Figure 7 reports the distribution of the estimated v_z^l (left panel) and the associated distribution of bootstrapped p-values (right panel). In line with sticky price theory, the estimates for v_z^l are tightly centered around zero and often statistically insignificant. Row 2 in table 5.2 shows that estimates for c_z are hardly affected by the presence of a linear term, in line with sticky price theory. Also, the R^2 increases only marginally when including the linear term.

Inflation versus Suboptimal Inflation as Regressor. It turns out to be important for our empirical results that the right-hand side of equation (21) features the squared value of *suboptimal* inflation rather than simply the squared value of inflation itself. To illustrate this point, let $\ln \widehat{\Pi_z(j)}$ denote the average item-level inflation rate prevailing *over the lifetime of product* j and consider the following alternative formulation of the second-stage regression

$$\widehat{Var}(u_{jzt}) = v_z + c_z \cdot (\ln \widehat{\Pi_z(j)})^2 + \varepsilon_{jz}, \tag{24}$$

which counterfactually assumes that the optimal inflation rate equals zero for all products.

Row 3 in table 5.2 shows that outcomes differ radically from the baseline: (i) one obtains about as many positive as negative estimates for c_z; (ii) 61% of the coefficients are statistically insignificant; and (iii) the R^2 value of the regression drops to zero.

Based on regression (24), which - in line with textbook sticky price models - assumes product-specific optimal inflation to be equal to zero, one would wrongly conclude that inflation is *not* associated with relative price distortions. This highlights that our baseline findings emerge predominantly due to differences in the *optimal inflation rate* Π^*_{jz} across products j within an item z.

Positive versus Negative Deviations Inflation from Optimal Inflation. We now explore whether the direction of the deviation from optimal inflation makes a difference for the magnitude of relative price distortions. In particular, when $\ln \Pi_z/\Pi^*_{jz} < 0$, then nominal prices have to *fall* to keep relative prices at their desired level, while nominal prices have to *increase* when $\ln \Pi_z/\Pi^*_{jz} > 0$. If the degree of price rigidity depends on the direction of the price adjustment, then positive versus negative deviations from optimal inflation generate price distortions of different magnitudes. We can study this by estimating the baseline equation (21) using coefficients that depend on the sign of the deviation:

$$\widehat{Var}(u_{jzt}) = v_z + \left(c_z + c_z^- \cdot I_{\{\ln \widehat{\Pi_z/\Pi^*_{jz}} < 0\}}\right) \cdot (\ln \widehat{\Pi_z/\Pi^*_{jz}})^2 \; + \varepsilon_{jz},$$

where $I_{\{x\}}$ is an indicator function that is equal to 1 if x is true and zero otherwise. Row 4 in table 5.2 shows that the estimated value for c_z^- is positive in about three quarters of the cases and also often significantly so, while it is rarely significantly negative. Price distortions thus tend to be larger when firms have to decrease prices to counteract the effects of suboptimal inflation, compared to the case where suboptimal inflation requires an increase in prices. This suggests that downward rigidity of prices is more pronounced than upward rigidity.

Row	Specification	Share of positive point estimates (in%) ($c_z > 0$)	Share of estimates (in%) with $c_z < 0$ & p-value		$c_z > 0$ & p-value		Median c_z	Mean c_z	Median adj. R^2 (in %)	No. of price obs. (millions)
			<0.01	<0.05	< 0.05	<0.01				
1	**Baseline**	98	0	1	94	92	8.94	12.08	16	15.2
	Alternative 2nd Stage									
2	Add linear term to r.h.s.	96	1	1	90	85	8.49	11.86	17	15.2
3	Only squared inflation on r.h.s.	53	12	18	23	18	0.23	-3.27	0	15.2
4	Sign-dependent coefficients:									
	squared term (c_z)	97	1	1	87	81	8.10	10.40	15	14.7
	negative branch (c_z^-)	74	5	8	43	30	2.77	5.72	15	14.7
	Alternative Selection									
	Using only products with									
5	$\geq$ 12 price observations	98	1	1	92	88	12.9	16.34	16	12.4
6	$\geq$ 24 price observations	96	0	1	84	77	23.68	28.61	15	8.5
7	$\geq$ 2 nominal price changes	91	2	3	70	59	6.84	8.95	5	8.2
8	$\geq$ 4 nominal price changes	85	3	5	52	40	7.03	11.12	3	4.8
9	Include sales prices	98	0	0	92	90	8.93	13.26	12	16.5
	Alternative 1st Stage									
10	Nonlinear 1st stage	97	0	1	93	88	6.52	9.4	12	15.2
	Trend stability test									
11	- with p-value $\geq$ 10%	97	0	0	88	81	7.24	11.31	13	6.1
12	- with p-value $\geq$ 20%	97	0	1	86	77	7.37	11.62	13	4.8
	Within Product Regression									
13	All products	99	0	0	98	96	3.57	4.17	23	14.9
	Only products with									
14	$\geq$ 1 price changes *per* half life	98	1	1	91	82	4.25	5.24	12	6.1
15	$\geq$ 2 price changes *per* half life	93	1	2	76	63	5.98	7.1	7	3.3

Table 2: Baseline estimates & robustness exercises with bootstrapped p-values

Reducing First-Stage Estimation Error. One possible concern with the baseline estimation approach is that first-stage estimation errors are large and might lead to substantial attenuation in the second stage. We address these concerns by selecting products for which estimation errors are likely smaller. We do so in two ways.

First, we select products with a higher minimum number of price observations, i.e., 12 or 24 monthly price observations instead of 6 observations in the baseline approach. This allows for a more reliable estimation of the optimal inflation trend Π^*_{jz} and the rate of suboptimal inflation. The regression outcomes are reported in rows 5 and 6 in table 5.2. While results barely change in terms of the share of positive coefficients c_z and their statistical significance, the magnitude of the mean and median estimate increases considerably relative to the baseline. This suggests that first-stage estimation error indeed causes a considerable downward bias in the second-stage estimates for c_z, which is in line with the simulation evidence presented in appendix H.

Second, we use the number of nominal price changes as a selection criterion for including products in the regressions. Excluding products with only few price changes avoids that the variation of residuals in the cross-section of products is purely driven by whether or not a price change is observed over the product lifetime. To this end, we perform the second-stage regression using only products with 2 or more (4 or more) price changes. Rows 7 and 8 in table 5.2 show that one obtains again a very large number of positive point estimates and high statistical significance levels, albeit somewhat lower values than in the baseline. Also, the R^2 value of regression and the magnitudes of coefficient estimates decline. Despite this, support for the notion that suboptimal inflation distorts relative prices remains overall strong.

Including Sales Prices. Our baseline estimation removes all sales prices from the sample, mainly because the underlying sticky price theories typically do not model sales. Row 9 in table 5.2 shows that our results are robust to including sales prices into the estimation.

Nonlinear Time Trends/Testing for Breaks in Time Trends. Our baseline approach allows for a linear time trend in relative prices in the first-stage regression equation (13). Since the presence of nonlinear time trends may be a source of concern, we recompute the first-stage residuals allowing also for a quadratic time trend. We then use the resulting residuals in our second-stage regression (21). Row 10 in table 5.2 reports the regression outcomes. Again, we obtain a very high number of positive point estimates and very high levels of statistical significance.

An alternative approach to deal with potential non-linearities in relative price trends is to test for trend stability. To this end, we run a Chow test for trend stability in the first-stage regression, using the first and second half of product life. We exclude in the second stage all products with p-value for the null hypothesis of no trend break below 10% or 20%.[37] The estimation outcomes are reported in rows 11 and 12 of table 5.2 and hardly change compared to the baseline.

6 Exploiting Within-Product Variation

This section pursues an alternative estimation strategy that allows relaxing key identifying assumptions of the baseline estimation approach in the previous section. It exploits *within-product* variation and thus can deal with settings in which idiosyncratic shock processes and price rigidities *both differ* across products within the *same* expenditure item. While this strategy addresses key concerns one might have with the baseline approach, it comes at the cost of increased second-stage attenuation bias.

The key idea consists of splitting the sample life of every product into two equally long subsamples and to exploit variation in the inflation rate across the two subsamples. Specifically, let $\ln \Pi_{jz1} - \ln \Pi^*_{jz}$ denote the suboptimal inflation rate of product j in item z in the first half of product life, $\ln \Pi_{jz2} - \ln \Pi^*_{jz}$ the suboptimal rate in the second half.[38] We consider the case with Calvo frictions below, but appendix J.1 shows that similar arguments apply for the case with menu cost frictions.

Equation (15) holds separately in the first and the second half of product life. Taking differences across the product half lives, we obtain

$$\begin{aligned} &Var_1(u_{jzt}) - Var_2(u_{jzt}) \\ &\quad = c_z \cdot ((\ln \Pi_{jz1} - \ln \Pi^*_{jz})^2 - (\ln \Pi_{jz2} - \ln \Pi^*_{jz})^2), \end{aligned} \tag{25}$$

where $Var_1(u_{jzt})$ and $Var_2(u_{jzt})$ denote the residual variances in the first and second half of the product lifetime, respectively.[39] The key feature of equation (25) is that it eliminates the constant present in the baseline regression specification (15). This allows testing whether the coefficient $c_z = \alpha_z/(1-\alpha_z)^2$ in equation (25) is positive *without* requiring that idiosyncratic shock processes are identical across products. If the Calvo adjustment frequencies $\alpha_{jz} \in [0,1]$ also vary across products within the *same* expenditure item, then the OLS estimate $\widehat{c}_z$ of the coefficient c_z in equation (25) will recover the average

[37] As is well-known, the Chow test is oversized in small samples (Candelon and Lütkepohl (2001)), i.e., it rejects the null hypothesis of no-trend-break too often in small samples when the null hypothesis is true. In our application, this only increases the strictness of selecting products featuring a constant trend.

[38] The suboptimal rates can be estimated using equation (61) in appendix G, separately for the first and second half of product lifetime.

[39] These are estimated using the same regression as in the baseline approach.

price distortion across products, i.e.,

$$E\left[\widehat{c}_z\right] = E\left[\frac{\alpha_{zj}}{\left(1-\alpha_{zj}\right)^2}\right]$$

provided the product-specific coefficients $\alpha_{jz}/\left(1-\alpha_{jz}\right)^2$ display conditional mean independence from the regressor in (25).[40] Under this condition, one can allow for product-specific idiosyncratic shock processes *and* product-specific price stickiness, but still test whether (on average across products within an item) suboptimal inflation distorts relative prices.

While this within-product approach generalizes our baseline approach, the second-stage (25) likely features larger right-hand side measurement error, as now one has to estimate (in the first stage) how suboptimal inflation *changes* over the product life, rather than just the average level of suboptimal inflation over the product life. This likely results in increased second-stage attenuation bias for the estimates of the coefficient c_z.

Row 13 in table 5.2 reports the outcomes from estimating equation (25). Results are even stronger than in the baseline case: 99% of the estimated coefficients are now positive and the share of significantly positive coefficients is also higher. Yet, the point estimates are now considerably smaller, which is likely due to increased attenuation bias.

To document that these results do not emerge because there is a price change in one product half-life but not in the other half life, rows 14 and 15 in table 5.2 repeat the within-estimation approach using only products that have at least 1 or at least 2 nominal price changes *per half life*. In these cases, one still obtains very strong support for the notion that suboptimal inflation is associated with relative price distortions.

7 Exploiting Across-Item Variation

This section estimates the effects of suboptimal inflation on price distortions exploiting only variation in suboptimal inflation across items (rather than variation across products within an item considered thus far). This is motivated by the fact that average relative price trends, and thus the optimal inflation rates, vary significantly in the cross-section of items (Adam and Weber (2023), Adam, Gautier, Santoro and Weber (2022)).

To this end, we estimate the first-stage equation (13) jointly for all products within an item, imposing the restriction $\Pi^*_{jz} = \Pi^*_z$, which delivers a vector of first-stage residuals u_z for all products in the item. We then estimate suboptimal inflation at the item level, using equation (22) and imposing a common time trend ($\ln \Pi_z/\Pi^*_{jz} = \ln \Pi_z/\Pi^*_z$). Finally, we estimate the second-stage equation using the cross-section of items:

$$Var(u_z) = v_0 + c_0 \cdot (\ln \Pi_z - \ln \Pi^*_z)^2 + \varepsilon_z. \tag{26}$$

[40] See appendix J for a proof and further details.

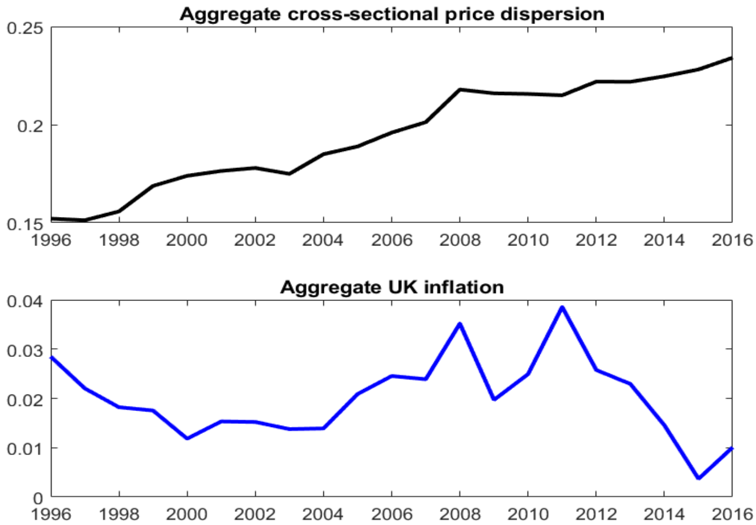

Figure 8: Aggregate cross-sectional price dispersion and inflation

Proposition 12 in appendix J.2 provides conditions under which the coefficient c_0 identifies the average relative price distortion generated by suboptimal inflation in the cross-section of items. Importantly, these conditions allow for heterogeneity in price rigidities and thus accommodate the well-known fact that price adjustment frequencies differ significantly across expenditure items. Estimating equation (26), we find $\widehat{c}_0 = 18.1$ with a t-statistic equal to 9.98. The estimate is thus highly significant and is larger than the average item-level estimate for c_z obtained in the baseline estimation approach in table 5.2. Estimating the right-hand side variable in equation (26) using many products likely decreases measurement error and the associated attenuation bias.

8 Understanding Cross-Sectional Price Dispersion

The empirical analysis up to now has focused on price distortions over time at the level of individual products or items. We now shift focus and consider the contribution of price distortions to cross-sectional dispersion of prices at any given point in time.

The top panel in figure 8 depicts an aggregate measure of cross-sectional price dispersion. It is constructed by computing for every year and every expenditure item the cross-sectional variance of relative prices, $Var^j(\ln p_{jzt})$. We then aggregate these variances across items using expenditure weights. The figure shows that cross-sectional price dispersion has increased by more then 50% over the sample period. Importantly, price dispersion increased despite aggregate inflation not displaying any time trend over the sample period, see the bottom panel in figure 8.

This section shows that the aggregate cross-sectional *dispersion* of prices in the United

Kingdom is to 99% the result of flexible-price dispersion. The same holds true for the large increase in cross-sectional price dispersion over time. The level and time trend of cross-sectional price dispersion is thus largely unrelated to price *distortions*. Interestingly, we also show that price distortions due to inflation comove over time with inflation in line with the predictions of sticky price theory.

Section 8.1 shows how we decompose aggregate cross-sectional price dispersion into a component capturing identifiable components of the flexible price dispersion and a remainder component that contains the effects of price distortions and idiosyncratic shocks. Section 8.2 analyzes the comovement of this remainder component with inflation over time.

8.1 Decomposing Cross-Sectional Price Dispersion

The sticky price theories analyzed in section 3 imply that the relative price of product j in expenditure item z evolves *over time* according to[41]

$$\ln p_{jzt} = \ln p^*_{jz} - \ln \Pi^*_{jz} \cdot t + u_{jzt}, \tag{27}$$

where the residuals u_{jzt} satisfy $E[u_{jzt}|p^*_{jz}, \Pi^*_{jz}] = 0$, are independent across j and z, and have variance *over time* equal to

$$Var(u_{jzt}) = v_z + c_z \cdot (\ln \Pi_z - \ln \Pi^*_{jz})^2. \tag{28}$$

We now decompose the *cross-sectional* variance of prices for products j in item z at time t, which we denote by $Var^j(p_{jzt})$. To this end, we consider a setting with a unit mass of products j in item z, where each month a share $\delta_z > 0$ of products randomly exits the sample and gets replaced by newly sampled products. Newly sampled products may have different characteristics than the products that leave the sample, so that the distribution of product characteristics $\{p^*_{jz}, \Pi^*_{jz}\}$ may change over time. We allow the new characteristics to be drawn from arbitrarily time-varying distributions.[42] Once products are in the sample, their relative price evolves according to equations (27)-(28). We then have the following cross-sectional decomposition result:[43]

Proposition 3 *Let $Var^j(\cdot)$ denote the variance in the cross-section of products j. The cross-sectional variance of relative prices in expenditure category z at time t is then given*

[41] See equation (52) in appendix E.3 for the case with Calvo frictions and equation (53) in appendix F for the case with menu costs.

[42] We assume that upon the time of product entry, the initial residual u_{jzt} is drawn from the stationary residual distribution for products with characteristics (p^*_{jz}, Π^*_{jz}). This is justified by the fact that newly sampled products in our data typically do not represent truly new products, but rather relatively mature products that get newly sampled by the ONS.

[43] See appendix L for the proof.

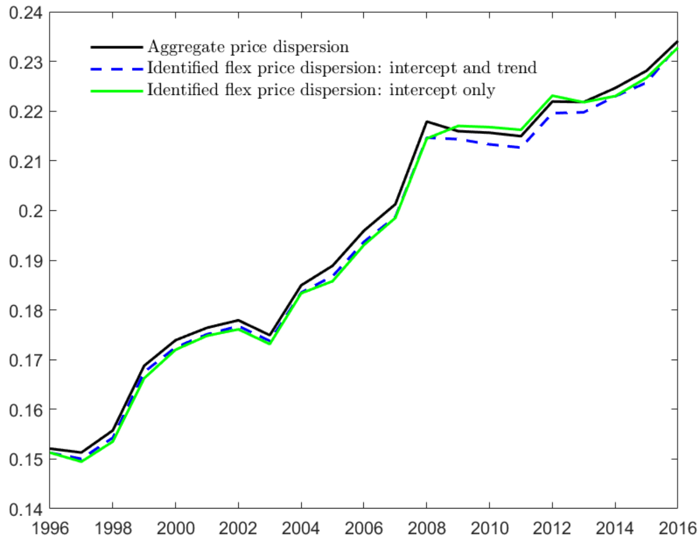

Figure 9: Aggregate cross-sectional log price dispersion and its flexible price components (various identified parts)

by

$$Var^j(\ln p_{jzt}) = Var^j(\ln p^*_{jz} - \ln \Pi^*_{jz} \cdot t) + Var^j(u_{jzt}), \tag{29}$$

where

$$Var^j(u_{jzt}) = v_z + c_z \cdot E^j[(\ln \Pi_z - \ln \Pi^*_{jz})^2]. \tag{30}$$

Equation (29) decomposes the cross-sectional price dispersion into two components. The first component, $Var^j(\ln p^*_{jz} - \ln \Pi^*_{jz} \cdot t)$, captures dispersion from the deterministic parts of products' flexible price dynamics. The second component on the right-hand side of equation (29) captures price distortions from suboptimal inflation, as captured by the term $c_z \cdot E^j[(\ln \Pi_z - \ln \Pi^*_{jz})^2]$ in equation (30), and a constant variance term $v_z \geq 0$ that is due to price dispersion resulting from idiosyncratic shocks.[44]

The decomposition in proposition 3 holds at each point in time and its most interesting feature is the absence of a covariance term on the right-hand side of equation (29). Since the deterministic parts of the flexible price $(\ln p^*_{jz}, \ln \Pi^*_{jz})$ can be estimated using the first-stage regression (13), proposition 3 allows decomposing cross-sectional price dispersion into a flexible price component and a remainder term $(Var^j(u_{jzt}))$. The remainder provides an *upper bound* for the contribution of price distortions to overall price dispersion, as $v_z \geq 0$ in equation (30).

[44]The constant v_z is defined in equation (16) for the case with Calvo frictions and in equation (20) for the case with menu costs.

Figure 9 depicts the aggregate cross-sectional price dispersion (black line), as previously shown in the top panel of figure 8. It also depicts the dispersion associated with the identifiable deterministic parts of flexible prices (blue dashed line).[45] This identifiable part of the flexible-price dispersion accounts for the bulk of aggregate price dispersion and also closely tracks its increase over time. Since time variation in the distribution of optimal inflation rates ($\ln \Pi^*_{jz}$) is quite limited, see appendix K, movements in indentifiable flexible price dispersion over time is mostly due to the changing dispersion of the intercept terms $\ln p^*_{jz}$. The green line in figure 9 illustrates this fact and depicts the cross-sectional dispersion explained by the intercept term only.[46]

This shows that cross-sectional price dispersion is to a large extent driven by the dispersion in flexible prices, which strongly increased over time. This increase may reflect a number of economic forces, for instance a widening cross-sectional dispersion of mark-ups, productivities, and (unmeasured) product qualities. Analyzing the forces underlying the widening dispersion of flexible prices is interesting but beyond the scope of the paper. The key takeaway here is that the large increase in flexible-price dispersion explains why aggregate inflation fails to covary with observed price dispersion over time in figure 8.

8.2 Inflation and Cross-Sectional Price Distortions over Time

We now analyze the time series properties of the residual dispersion $Var^j(u_{jzt})$ in proposition 3. According to equation (30), the residual dispersion is determined by a constant term and by price distortions. Price distortions depend on the item-level inflation rate Π_z and on the cross-sectional distribution of optimal inflation rates $\{\Pi^*_{jz}\}$. In the data, the cross-sectional distribution of optimal inflation rates $\{\Pi^*_{jz}\}$ is nearly constant over time, see appendix K, which allows considering a setting with a constant cross-sectional distribution of optimal inflation rates. Variation in residual dispersion $Var^j(u_{jzt})$ over time then exclusively reflects variation in price distortions induced by changes in inflation, as we show next.

Consider a setting with constant cross-sectional distribution of optimal inflation rates $\{\Pi^*_{jz}\}$.[47] Furthermore, let inflation in year t in expenditure item z be denoted by Π_{zt} and assume that inflation changes from year to year according to a random walk. Price setters

[45] As before, we aggregate at any given point in time across expenditure items using expenditure weights.

[46] Due to negative covariance between the intercept and slope terms, the dispersion explained by intercepts alone can exceed the overall observed dispersion. Also, to make comparisons meaningful over time, figure 9 reports the dispersion coming from intercepts using the normalized intercepts $\ln p^*_{jz} - \Pi^*_{jz} \cdot t^0_{jz}$, where t^0_{jz} is the time period in which the product first enters the sample.

[47] We impose no restrictions on the distribution of intercept terms $\{p^*_{jt}\}$, thus allow for the large increase in intercept dispersion documented in the previous section.

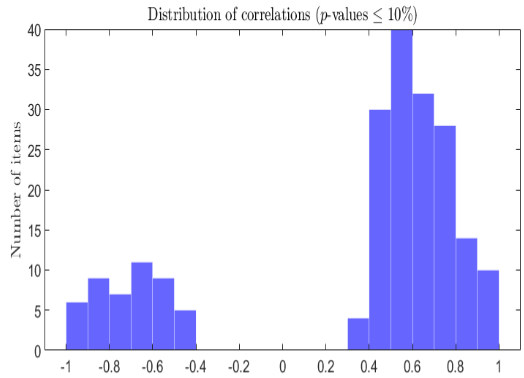

Figure 10: Correlation between inflation and price distortions at the item level

who adjust prices during year t will expect future inflation to be equal to the current inflation rate Π_{zt}, so that they set relative prices *as if* the steady-state inflation rate was equal to Π_{zt}.[48] With the vast majority of prices adjusting over the course of a year, the cross-sectional dispersion of relative prices at the end of each year will depend only on the inflation rate Π_{zt} that prevailed during year t and will thus be given by equation (30) with $\Pi_z = \Pi_{zt}$.

Equation (30) thus provides a theory-implied relationship linking (yearly) inflation rates Π_{zt} to the cross-sectional dispersion of first-stage residuals $Var^j(u_{jzt})$ at the end of each year. It predicts that a marginal increase in the inflation rate Π_{zt} from one year to the next increases (decreases) residual dispersion, whenever the average optimal inflation rate in the item, $E^j[\Pi^*_{jz}]$, lies below (above) actual inflation. Residual dispersion is thus predicted to comove positively (negatively) with inflation whenever actual inflation is above (below) average optimal inflation.

This prediction can be tested in the data. It is of economic interest because equation (30) implies that time-variation in residual dispersion exclusively captures time-variation in price distortions. By analyzing how residual dispersion commoves with inflation, one effectively analyzes how price distortions comove with inflation. Section 8.2.1 analyzes this comovement pattern at the item level and section 8.2.2 at the aggregate level.

8.2.1 Item Level Results

To test whether price distortions correlate positively/negatively with inflation as predicted by sticky price theory, we compute for each expenditure item z the correlation between Π_{zt} and $Var^j(\widehat{u}_{jzt})$ over time.[49] Figure 10 depicts the resulting distribution of correlations across items, using all correlations for which p-value are below 10%. In the data, there

[48] This is so because certainty equivalence applies for the considered order of approximation.

[49] We consider the 680 expenditure items with at least three years of data.

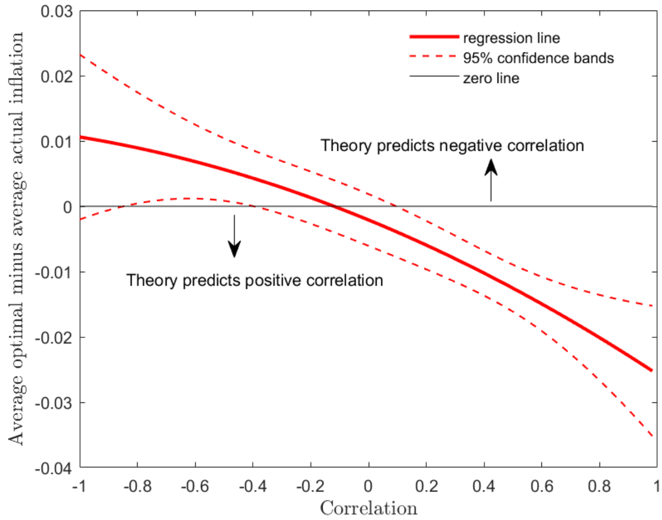

Figure 11: The relationship between average optimal minus actual inflation (y-axis) and the correlation between inflation and cross-sectional price distortions (x-axis)

are items with significantly positive and significantly negative correlations, even if positive ones dominate.[50]

According to the arguments in the previous section, positive (negative) correlations should emerge whenever average optimal inflation $E^j[\Pi^*_{jz}]$ lies above (below) actual inflation in the item. Figure 11 confirms this prediction: it depicts the outcome of a regression of the gap between optimal and actual inflation on the correlation and its square. The regression line behaves fully in line with the theoretical predictions:[51] the correlation between price distortions and inflation is positive, if actual item inflation is above average optimal inflation, and it is negative otherwise. This is particularly true for the statistically significant parts of the regression line. This shows that price distortions at the item level comove with inflation over time as predicted by sticky price theory.

8.2.2 Aggregate Results

We now consider an economy-wide measure of price distortions by aggregating the item-level variances $Var^j(\widehat{u}_{jzt})$ using expenditure weights. Under the assumptions maintained in this section, time variation in this measure again reflects time variation in cross-sectional price distortions.

[50] This result also emerges if we consider smaller p-values or consider all correlations independent of p-values.

[51] This continues to be true when restricting consideration to a linear regression or when including a third order term into the regression. The coefficient on the third order term is not statistically signficant.

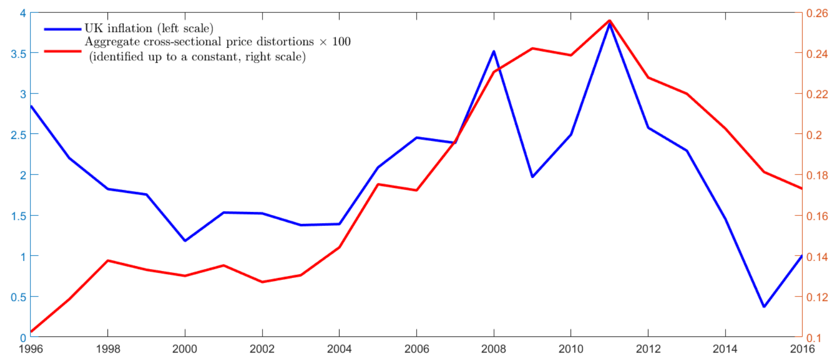

Figure 12: Aggregate inflation and aggregate cross-sectional price distortions

Figure 12 depicts the resulting aggregate distortion measure together with aggregate inflation.[52] Aggregate price *distortions* do not display a time trend, unlike aggregate *dispersion* in figure 8. Aggregate price distortions also covary positively with aggregate inflation: the correlation is equal to +0.46 and significant at the 5% level. Higher aggregate inflation is thus associated with larger amounts of cross-sectional price distortions in the data. This reflects the fact that in expenditure-weighted terms, there are more items for which actual inflation exceeds the average optimal inflation rate of products in the item. As we have seen in the previous section, these items display a positive comovement between inflation and price distortions.

Proposition 3 implies that the variance of first-stage residuals represents an upper bound on the amount of price distortions that is due to inflation.[53] The upper bound of the variance reached in figure 12 is approximately $2.5 \cdot 10^{-3}$. Therefore, inflation-induced price distortions alone give rise to a standard deviation of prices of at most $\sqrt{2.5 \cdot 10^{-3}} = 5\%$ in the absence of flexible-price dispersion. A lower bound on the maximum contribution of inflation to price distortions over the sample period is given by the min-max range of the variance of first-stage residuals, as the time-varying component is - according to theory - solely due to inflation. This range is approximately equal to $1.5 \cdot 10^{-3}$ in figure 12 and implies that inflation-induced distortions of relative prices gives rise to a standard deviation of relative prices of at least $\sqrt{1.5 \cdot 10^{-3}} = 3.87\%$ over the sample period (again in the absence of flexible-price dispersion).

[52] Note that aggregate inflation is also an expenditure-weighted average of item-level inflation rates. Figure 12 displays annual dispersion and annual inflation to remove within-year seasonalities in price dispersion and inflation. Both measures are computed as a 12 month average of monthly dispersion and monthly year-over-year inflation rate.

[53] This is so because the constant v_z in equation (30) is positive.

9 Conclusions

Using a structural empirical approach, the present paper derives four empirical insights: (i) at the product-level, relative price distortions robustly increase with the (squared) deviation of inflation from the (product-specific) optimal level and with nominal rigidities; (ii) in the cross-section of products, price dispersion and its evolution over time predominantly reflect the dispersion of flexible prices and its movements: at most 1% of aggregate price dispersion is due to price distortions associated with suboptimal inflation; (iii) cross-sectional price distortions at the item level comove positively or negatively with inflation over time, with sticky price theory correctly predicting the sign of this comovement; (iv) aggregate price distortions comove positively with aggregate inflation over the sample period.

Collectively, these findings offer substantial empirical support for the theoretical foundations of sticky price models and the monetary policy implications they engender, but they also raise new important questions: what is behind the large increase in aggregate cross-sectional price dispersion? Are the price distortions we identify associated with corresponding demand distortions and how big are these? Exploring these questions appears to be an important task for future research.

References

Acharya, S., E. Challe, and K. Dogra (2023): "Optimal Monetary Policy According to HANK," *American Economic Review (forthcoming).*

Adam, K., E. Gautier, S. Santoro, and H. Weber (2022): "The Case for a Positive Euro Area Inflation Target: Evidence from France, Germany and Italy," *Journal of Monetary Economics (forthcoming).*

Adam, K., and H. Weber (2019): "Optimal Trend Inflation," *American Economic Review*, 109(2), 702–737.

——— (2023): "Estimating the Optimal Inflation Target from Trends in Relative Prices," *American Economic Journal: Macroeconomics*, 15(3), 1–42.

Alvarez, F., M. Beraja, M. Gonzalez-Rozada, and P. A. Neumeyer (2019): "From Hyperinflation to Stable Prices: ArgentinaŠs Evidence on Menu Cost Models," *The Quarterly Journal of Economics*, 134(1), 451–505.

Alvarez, F., F. Lippi, and A. Oskolov (2022): "The Macroeconomics of Sticky Prices and Generalized Hazard Functions," *Quarterly Journal of Eonomics*, 137, 989–1038.

Argente, D., and C. Yeh (2022): "Product Life Cycle, Learning, and Nominal Shocks," *The Review of Economic Studies*, 89(6), 2992Ű3054.

Ascari, G., P. Bonmolo, and Q. Haque (2022): "The Long-Run Phillips Curve is ... a Curve," Dutch Central Bank mimeo.

Baley, I., and A. Blanco (2021): “Aggregate Dynamics in Lumpy Economies,” *Econometrica*, 89, 1235–1264.

Baqaee, D., E. Farhi, and K. Sangani (2022): “Supply Side Effects of Monetary Policy,” NBER Working Paper 28345.

Candelon, B., and H. Lütkepohl (2001): “On the Reliability of Chow-Type Test for Parameter Constancy in Multivariate Dynamic Models,” *Economics Letters*.

Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011): “Reference prices, costs, and nominal rigidities,” *American Economic Review*, 101(1), 234–62.

Galí, J. (2015): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New KeynesianFramework and Its Applications*. Second Edition, Princeton University Press.

Meier, M., and T. Reinelt (2022): “Monetary Policy, Markup Dispersion, and Aggregate TFP,” *Review of Economics and Statistics (forthcoming)*.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018): “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation,” *The Quarterly Journal of Economics*, 133(4), 1933–1980.

Sara-Zaror, F. (2022): “Expected Inflation and Welfare: The Role of Consumer Search,” Federal Reserve Bank of San Francisco mimeo.

Sheremirov, V. (2020): “Price Dispersion and Inflation: New Facts and Theoretical Implications,” *Journal of Monetary Economics*, 114, 59–70.

Taylor, J. B. (1979): “Staggered wage setting in a macro model,” *American Economic Review, Papers and Proceedings*, 69(2), 108Ű13.

Woodford, M. (2003): *Interest and Prices*. Princeton University Press, Princeton.

Wulfsberg, F. (2016): “Inflation and Price Adjustment: Micro Evidence from Norwegian Consumer Prices 1975-2004,” *American Economic Journal: Macroeconomics*, 8, 175–193.

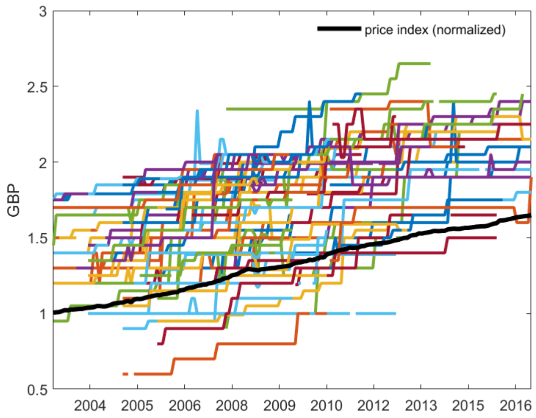

Figure 13: Nominal prices and price index (normalized to one at the start) for "takeaway coffee latte", all products with ≥60 price observations.

A Nominal Prices and Average Price in the Expenditure Item "Takeaway coffee latte"

Figure 13 depicts nominal price time series observations in the expenditure item "takeaway coffee latte" for all products with at least 60 price observations. It also depicts the price index, which is computed using all products in the item, including those with less than 60 price observations. To increase readability of the chart, the price index is normalized to one at the start of the sample. Figure 13 shows that the price index steadily increases over time, while coffee prices display stepwise price increases at different trend rates.

B Details of the Introductory Model with Taylor Frictions

Consider the Taylor (1979) model as outlined in Section 2. The firm's objective is as follows:

$$\max_{\ln p_{jt}} - \sum_{i=0}^{N-1} \left(\ln p_{jt+i} - \ln p^*_{jt+i}\right)^2 = \max_{\ln p_{jt}} - \sum_{i=0}^{N-1} \left(\ln p_{jt} - \ln p^*_{jt} - i\ln(\Pi/\Pi^*_j)\right)^2$$

The first order condition yields:

$$\ln p^{opt}_{jt} = \ln p^*_{jt} + \frac{N-1}{2}\ln(\Pi/\Pi^*_j)$$

If an adjustment happens in period t, then for all $0 \leq i < N$:

$$\ln p_{jt+i} = \ln p^{opt}_{jt} - i\ln\Pi = \ln p^*_{jt} + \frac{N-1}{2}\ln(\Pi/\Pi^*_j) - i\ln\Pi$$

Since the flexible price is given by $\ln p^*_{jt} - i\ln\Pi^*_j$, relative price distortions are:

$$u_{jt+i} = \left(\frac{N-1}{2} - i\right)\ln(\Pi/\Pi^*_j)$$

Summing squared distortions over all $0 \leq i < N$:

$$Var(u_j) = \frac{1}{N}\sum_{i=0}^{N-1} u_{jt+i}^2 = \frac{1}{N}(\ln \Pi - \ln \Pi_j^*)^2 \sum_{i=0}^{N-1}\left(\frac{N-1}{2} - i\right)^2$$
$$= \frac{(N-1)(N+1)}{12}(\ln \Pi - \ln \Pi_j^*)^2$$

Note that adjustment size is given by:

$$\ln P_{jt}^{opt} - \ln P_{jt-1} = \ln p_{jt}^{opt} - \ln p_{jt-1} + \ln \Pi$$
$$= \ln p_{jt}^{opt} - \ln p_{jt-N}^{opt} + (N-1)\ln \Pi + \ln \Pi$$
$$= N(\ln \Pi - \ln \Pi_j^*)$$

B.1 Absolute Price Changes May Miss Price Distortions

We first consider an example in which the absolute size of price changes may respond to inflation despite price distortions failing to do so. Thereafter, we consider a setting where the absolute size of price changes fails to respond to inflation even though relative price distortions do respond.

The first point is simple. Consider the example discussed in section 2. The absolute size of log nominal price changes is simply a function of price stickiness and suboptimal inflation and equal to $N \cdot \left|\Pi - \Pi_j^*\right|$. In the limit where prices become fully flexible ($N \to 1$), the absolute size of nominal price changes is given by $\left|\Pi - \Pi_j^*\right|$ and varies one-to-one with the gap between actual and optimal inflation. The absolute size of price changes thus suggests a relationship between suboptimal inflation and relative price distortions, even in a setting where prices are fully flexible and price distortions are absent.[54]

This contrasts with the detrended residuals gap_{jt} proposed in figure 2: in the limit with flexible prices, relative prices follow the dotted lines in the figure, so that the residuals gap_{jt} are all equal to zero. Their variance will thus not covary with suboptimal inflation in the cross-section of products. In fact, for the limit $N \to 1$, the coefficient c in equation (3) converges to zero: one arrives at the correct conclusion that suboptimal inflation does *not* lead to relative price distortions.

Next consider the case with sticky prices. We show below - using a setting with a stochastic component in the flexible price - that the absolute size of price changes may fail to respond to changes in suboptimal inflation, even in a setting where price distortions do change with suboptimal inflation.

[54] This argument holds not only in the cross-section of goods, but equally applies in the time dimension when considering the effects of a change in the steady-state inflation rate Π for the price distortions present at the level of some product with optimal rate Π_j^*.

Suppose that the frictionless price $\ln p_{jt}^*$ has an additional idiosyncratic component x_{jt} that follows a two-state Markov chain ($x_{jt} \in \{-\bar{x}, \bar{x}\}, \bar{x} > 0$) and switches states with probability one at the times of price adjustment and with probability zero otherwise:

$$\ln p_{jt}^* = \ln p_j^* - t \ln \Pi_j^* + x_{jt}$$

Since the value of x_{jt} does not change during a price spell, it is straightforward to verify that, as before:

$$\begin{aligned}
\ln p_{jt}^{opt} &= \ln p_{jt}^* + \frac{N-1}{2} \ln(\Pi/\Pi_j^*) \\
u_{jt+i} &= \left(\frac{N-1}{2} - i\right) \ln(\Pi/\Pi_j^*) \\
Var(u_j) &= \frac{(N-1)(N+1)}{12} (\ln \Pi - \ln \Pi_j^*)^2
\end{aligned}$$

Conditional on x_{jt}, the size of adjustment becomes:

$$\begin{aligned}
\ln P_{jt}^{opt} - \ln P_{jt-1} &= \ln p_{jt}^{opt} - \ln p_{jt-1} + \ln \Pi \\
&= \ln p_{jt}^{opt} - \ln p_{jt-N}^{opt} + (N-1) \ln \Pi + \ln \Pi \\
&= N(\ln \Pi - \ln \Pi_j^*) + 2x_{jt}
\end{aligned}$$

The average absolute adjustment size is then:

$$\mathbb{E}\left[|\ln P_{jt}^{opt} - \ln P_{jt-1}|\right] = \frac{1}{2}\left[|N \ln\left(\Pi/\Pi_j^*\right) + 2\bar{x}| + |N \ln\left(\Pi/\Pi_j^*\right) - 2\bar{x}|\right]$$

Suppose that $N \ln\left(\Pi/\Pi_j^*\right) \in (-2\bar{x}, 2\bar{x})$. Then:

$$\begin{aligned}
\mathbb{E}\left[|\ln P_{jt}^{opt} - \ln P_{jt-1}|\right] &= \frac{1}{2}\left[\left(N \ln\left(\Pi/\Pi_j^*\right) + 2\bar{x}\right) - \left(N \ln\left(\Pi/\Pi_j^*\right) - 2\bar{x}\right)\right] \\
&= 2\bar{x}
\end{aligned}$$

Therefore, as long as $N \ln\left(\Pi/\Pi_j^*\right) \in (-2\bar{x}, 2\bar{x})$, suboptimal inflation has no effect on the average absolute size of adjustments, while still affecting price distortions.

C Proof of Proposition 1

In this section we prove that it is impossible to recover the price gap distribution if shocks are stationary. To lighten notation in this appendix, we drop the z subscript referring to the expenditure category. Suppose an econometrician observes the infinite path of actual prices $\ln p_{jt}$ and it is known that this path is generated under the time-dependent friction and stationary shocks $\ln x_{jt}$. The econometrician can recover the N values of the vector $\mathbf{f} \equiv [f_1, \dots, f_N]'$ of $f(x_{jt})$ as defined in (11):

$$f(x_{jt}) \equiv (1-\alpha) E_t \sum_{i=0}^{\infty} (\alpha)^i \ln x_{jt+i}.$$

In addition, the econometrician can recover the $N \times N$ transition matrix Λ^f:

$$\Lambda^f = \begin{bmatrix} \lambda^f_{11} & \cdots & \lambda^f_{1N} \\ \vdots & \ddots & \vdots \\ \lambda^f_{N1} & \cdots & \lambda^f_{NN} \end{bmatrix},$$

where λ^f_{ij} is the probability of observing f_j in the subsequent period, conditional on observing f_i in the previous period.[55] From the definition of $f(x_{jt})$ it follows that:

$$\mathbf{f} = (1-\alpha)\ln \mathbf{x} + \alpha \Lambda^x \mathbf{f}$$

where $\ln \mathbf{x}$ is the state vector of the process $\ln x_{jt}$ and Λ^x is its transition matrix. Setting $\Lambda^x = \Lambda^f$ and solving the above equation for $\ln \mathbf{x} \equiv [\ln x_1, \ldots, \ln x_N]$ provides a candidate for the process $\ln x_{jt}$ that leads to the observed process $f(x_{jt})$. However, as we show below, this candidate solution is not unique and the observed N-state process of $f(x_{jt})$ can be equally supported by an $(N+1)$-state process $\ln \tilde{x}_{jt}$, defined on the grid $\ln \tilde{\mathbf{x}} \equiv [\ln \tilde{x}_1, \ldots, \ln \tilde{x}_N, \ln \tilde{x}_{N+1}]$ with $(N+1)\times(N+1)$ transition matrix $\tilde{\Lambda}^x$. Such a process would lead to an $(N+1)$-state process of $\tilde{f}(x_{jt})$, with $\tilde{f}_i = f_i$ for all $i < N$ and $\tilde{f}_N = \tilde{f}_{N+1} = f_N$, making $\tilde{f}(x_{jt})$ and $f(x_{jt})$ observationally equivalent, provided the transition probabilities of $\tilde{\Lambda}^x$ imply Λ^f. To construct such a process, set $\ln \tilde{x}_i = \ln x_i$ for all $i < N$, $\ln \tilde{x}_N = \ln x_N - \varepsilon$ and $\ln \tilde{x}_{N+1} = \ln x_N + \varepsilon$ for a sufficiently small $\varepsilon > 0$.[56] We now construct the transition matrix $\tilde{\Lambda}^x$ in the following way:

$$\tilde{\Lambda}^x = \left[\begin{array}{cccc|cc} \lambda^x_{11} & \lambda^x_{12} & \cdots & \lambda^x_{1(N-1)} & \lambda^x_{1N}/2 & \lambda^x_{1N}/2 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \\ \lambda^x_{(N-1)1} & \lambda^x_{(N-1)2} & \cdots & \lambda^x_{(N-1)(N-1)} & \lambda^x_{(N-1)N}/2 & \lambda^x_{(N-1)N}/2 \\ \hline \tilde{\lambda}^x_{N1} & \lambda^x_{N2} & \cdots & \lambda^x_{N(N-1)} & \tilde{\lambda}^x_{N} & \tilde{\lambda}^x_{N} \\ \tilde{\lambda}^x_{(N+1)1} & \lambda^x_{N2} & \cdots & \lambda^x_{(N+1)(N-1)} & \tilde{\lambda}^x_{N+1} & \tilde{\lambda}^x_{N+1} \end{array}\right]$$

All elements in black are borrowed directly from the Λ^x matrix, whereas elements in red are to be solved for.[57] The first $(N-1)$ rows of $\tilde{\Lambda}^x$ ensure that for all $i < N$:

$$\begin{aligned} \tilde{f}_i &= (1-\alpha)\ln \tilde{x}_i + \alpha \sum_{j=1}^{N+1} \tilde{\lambda}^x_{ij} \tilde{f}_j \\ &= (1-\alpha)\ln x_i + \alpha \sum_{j=1}^{N-1} \lambda^x_{ij} f_j + \left(\frac{\lambda^x_{iN}}{2} + \frac{\lambda^x_{iN}}{2} \right) f_N = f_i \end{aligned}$$

[55] This can be achieved by conditioning on price spells of length one.

[56] One requirement for ε is that $\ln \tilde{x}_N$ and $\ln \tilde{x}_{N+1}$ do not coincide with existing values of $\ln x_i$. A stricter condition on the size of ε is introduced below.

[57] We order states such that $\lambda^x_{N1} > 0$ and $\lambda^x_{NN} > 0$. This is without loss of generality since $\ln x_{jt}$ is a stochastic process, implying that there exists a state i such that for at least two states j_1 and j_2, $\lambda^x_{ij_1} > 0$ and $\lambda^x_{ij_2} > 0$.

We now have to set the elements in red ($\tilde{\lambda}^x_{N1}$, $\tilde{\lambda}^x_N$, $\tilde{\lambda}^x_{(N+1)1}$, $\tilde{\lambda}^x_{N+1}$) such that $\tilde{f}_N = \tilde{f}_{N+1} = f_N$. For $i = N$ it requires:

$$\begin{aligned}\tilde{f}_N &= (1-\alpha)(\ln x_N - \varepsilon) + \alpha\tilde{\lambda}^x_{N1} f_1 + \alpha \sum_{j=2}^{N-1} \lambda^x_{Nj} f_j + 2\tilde{\lambda}^x_N f_N \\ &= f_N - (1-\alpha)\varepsilon + \alpha(\tilde{\lambda}^x_{N1} - \lambda^x_{N1}) f_1 + \alpha(2\tilde{\lambda}^x_N - \lambda^x_{NN}) f_N \overset{!}{=} f_N\end{aligned}$$

Denote $\sum_{j=2}^{N-1} \lambda^x_{Nj} \equiv \lambda$, then it must be the case that $\tilde{\lambda}^x_{N1} + \lambda + 2\tilde{\lambda}^x_N = 1$ to ensure that $\tilde{\Lambda}^x$ is a proper transition matrix. The same applies to the elements of Λ^x: $\lambda^x_{N1} + \lambda + \lambda^x_{NN} = 1$. Substituting $\tilde{\lambda}^x_N$ and λ^x_{NN} in the above equation and rearranging terms yields:

$$\tilde{\lambda}^x_{N1} = \lambda^x_{N1} + \frac{1-\alpha}{\alpha} \frac{\varepsilon}{f_1 - f_N}$$

For $i = N+1$, a similar line of arguments leads to:

$$\tilde{\lambda}^x_{(N+1)1} = \lambda^x_{N1} - \frac{1-\alpha}{\alpha} \frac{\varepsilon}{f_1 - f_N}$$

and the remaining elements $\tilde{\lambda}^x_N$ and $\tilde{\lambda}^x_{N+1}$ can then be recovered using the fact that all rows of $\tilde{\Lambda}^x$ sum up to one. ε must be small enough to ensure that $\tilde{\lambda}^x_{N1}$, $\tilde{\lambda}^x_N$, $\tilde{\lambda}^x_{(N+1)1}$ and $\tilde{\lambda}^x_{N+1}$ are all $\in [0,1]$. Such ε always exists since we have ordered the states to ensure $\lambda^x_{N1} > 0$ and $\lambda^x_{NN} > 0$ and there are infinitely many of them. It remains to show that transition probabilities in $\tilde{\Lambda}^x$ imply Λ^f. This holds trivially for all transitions between states f_i and f_j such that $i, j < N$. It is also true for transitions from f_i to f_N when $i < N$ since the probability of transiting from f_i to f_N is then equal to $\frac{\lambda^x_{iN}}{2} + \frac{\lambda^x_{iN}}{2} = \lambda^x_{iN}$. Finally, note that states $\ln x_N$ and $\ln x_{N+1}$ have the same unconditional probability,[58] and therefore the probability of moving from f_N to f_i is equal to $\frac{1}{2}\left(\tilde{\lambda}^x_{Ni} + \tilde{\lambda}^x_{(N+1)i}\right) = \lambda^x_{Ni}$ for all $i < N$. This implies that the probability of staying in f_N is also the same as in the original process (λ^x_{NN}).

Therefore, we have constructed an $N+1$-state process $\ln \tilde{x}_{jt}$ that leads to the same process $f(x_{jt})$ as the N-state process $\ln x_{jt}$. By induction this step can be repeated arbitrary many times.

D Testing for a Random Walk in Idiosyncratic Shocks

This appendix shows that our data strongly rejects the presence of a pure random walk in $\ln x_{jzt}$. One can test for a random walk in $\ln x_{jzt}$ by exploiting the fact that the optimal reset price upon price adjustment involves a constant gap relative to the flexible price,

[58] The unconditional probability satisfies $p = (\tilde{\Lambda}^x)'p$, and the last two columns of $\tilde{\Lambda}^x$ are identical, implying identical values of p_N and p_{N+1}.

whenever $\ln x_{jzt}$ is a random walk. This holds true with Calvo frictions, see equation (10), but also for the case with menu cost frictions.

Consider the times t_n ($n = 1, 2, ...N_{jz}$) during which the price of some product j in expenditure item z adjusts. Given the constant gap property, we have

$$\ln p^{opt}_{jzt_{n+1}} - \ln p^{opt}_{jzt_n} = -\ln \Pi^*_{jz} \cdot (t_{n+1} - t_n) + \ln e_{jzn+1} \tag{31}$$

where

$$\ln e_{jzn+1} \equiv \ln x_{jzt_{n+1}} - \ln x_{jzt_n}.$$

With a random walk in $\ln x$, the residuals $\ln e$ are uncorrelated over time, which can be tested. To do so, we re-scale residuals according to $(\ln e_{jzn+1}) / \sqrt{t_{n+1} - t_n}$ to make them homoskedastic under the null hypothesis of a random walk. We then compute the autocorrelations $\widehat{Corr}_z = \widehat{Cov}_z / \widehat{Var}_z$ of these re-scaled residuals within each item z, using the variance and covariance estimates for all products with $N_{jz} > 3$:

$$\widehat{Var}_z = \sum\nolimits_j \left(\frac{N_{jz} - 2}{\sum_k (N_{kz} - 2)} \sum\nolimits_{n=2}^{N_{jz}} \frac{\left(\frac{\ln e_{jzn}}{\sqrt{t_n - t_{n-1}}} \right)^2}{N_{jz} - 2} \right)$$
$$\widehat{Cov}_z = \sum\nolimits_j \left(\frac{N_{jz} - 3}{\sum_k (N_{kz} - 3)} \sum\nolimits_{n=2}^{N_{jz}-1} \frac{\frac{\ln e_{jzn}}{\sqrt{t_n - t_{n-1}}} \frac{\ln e_{jzn+1}}{\sqrt{t_{n+1} - t_n}}}{N_{jz} - 3} \right)$$

The top left panel in figure 14 depicts the estimated autocorrelations across items. Almost all of the estimates are negative, and most of them sizably so, which is inconsistent with $\ln x_{jzt}$ following a random walk. The right panel in the figure reports the bootstrapped p-values for the autocorrelation being weakly larger than zero, as implied by the random walk, and shows that these values are very low.

We then repeat the analysis when exogenously imposing $\Pi^*_{jz} = 0$ for all products in the first-stage regression. This is motivated by the possibility that the estimated time trends Π^*_{jz} could be purely spurious in the presence of a random walk in $\ln x_{jzt}$. The auto-correlations of the resulting residuals are then even more negative, see the lower left panel in figure 14. The bootstrapped p-values of the auto-correlations remain again very low (lower right panel).

Based on these findings, which relies exclusively on prices that are not sales prices, we can conclude that unobserved shocks in our data do *not* follow a pure random walk.

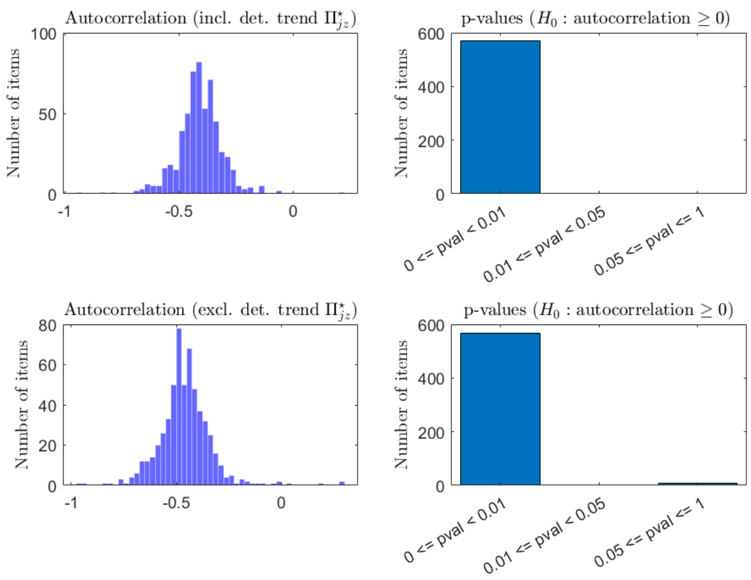

Figure 14: Autocorrelation of residuals (left panel) and bootstrapped p-values (right panel): random walk implies autocorrelation of zero

E Details of the Calvo Model

E.1 Firm Problem

The price-setting problem of firm j in item z in price-adjustment period t consists of choosing a nominal price P_{jzt} that maximizes the expected discounted sum of profits,

$$\max_{P_{jzt}} E_t \sum_{i=0}^{\infty} \alpha_z^i \frac{\Omega_{t,t+i}}{P_{t+i}} \left[(1+\tau)P_{jzt} - \frac{W_{t+i}}{A_{zt+i}} G_{jzt+i} X_{jzt+i} \right] Y_{jzt+i} \tag{32}$$

$$s.t. \quad Y_{jzt+i} = \psi_z \left(\frac{P_{jzt}}{P_{zt+i}} \right)^{-\theta} \left(\frac{P_{zt+i}}{P_{t+i}} \right)^{-1} Y_{t+i}, \tag{33}$$

where $\Omega_{t,t+i}$ denotes the stochastic discount factor of the representative household, Y_{jzt} output of product j in item z, and $W_{t+i}G_{jzt+i}X_{jzt+i}/A_{zt+i}$ the firm's nominal marginal costs, with firm productivity given by $A_{zt+i}/(G_{jzt+i}X_{jzt+i})$, as in equation (7), and the nominal wage given by W_{t+i}. The parameter τ is a sales subsidy (tax if negative). Maximization is subject to equation (33), which is derived from the cost-minimizing household demand function (6) using market clearing conditions.

E.1.1 Balanced Growth Path

We approximate the profit maximization problem (32) around a deterministic balanced growth path of the economy, in which aggregate and item-level output and consumption grow at constant rates, aggregate and item-level inflation rates are constant, and in which the amount of labor L^e_{zt} allocated to production in item z is also constant over time.

All idiosyncratic shocks continue to operate, i.e., there is product entry and exit and idiosyncratic shocks move the product's optimal relative price over time. Without loss of generality, we consider the efficient deterministic balanced growth path.

Within each item z, the efficient allocation of labor across products j maximizes the item-level output in equation (4) subject to the production function (7) and the feasibility constraint that $L^e_{zt} = \int L^e_{jzt}$ dj. This implies that the efficient level of output in item z is

$$Y^e_{zt} = \frac{A_{zt}}{\Delta^e_{zt}} L^e_{zt}, \tag{34}$$

where the productivity parameter $1/\Delta^e_{zt}$ in the efficient allocation is given by

$$1/\Delta^e_{zt} \equiv \left(\int_0^1 (1/(G_{jzt}X_{jzt}))^{\theta-1} \,\mathrm{dj} \right)^{\frac{1}{\theta-1}}. \tag{35}$$

We consider a balanced growth path in which $1/\Delta^e_{zt} = 1/\Delta^e_z$, so that equation (34) implies that item-level productivity is given by[59]

$$\Gamma^e_{zt} \equiv A_{zt}/\Delta^e_z. \tag{36}$$

Using equation (5), aggregate productivity Γ^e_t of the economy is given by[60]

$$\Gamma^e_t \equiv \prod_{z=1}^{Z} (\Gamma^e_{zt})^{\psi_z}. \tag{37}$$

Equation (36) and the previous equations imply that the steady-state growth rate of aggregate output and consumption along the balanced growth path, $\gamma^e \equiv \Gamma^e_t/\Gamma^e_{t-1}$, is given by

$$\gamma^e = \prod_{z=1}^{Z} a_z^{\psi_z} \tag{38}$$

where a_z denotes the steady-state growth rate of item-level productivity A_{zt}. From equation (36), we also obtain that the steady-state growth rate of item-level output and consumption, $\gamma^e_z \equiv \Gamma^e_{zt}/\Gamma^e_{zt-1}$, is given by

$$\gamma^e_z = a_z.$$

[59] It is straightforward to accommodate also a trend in $1/\Delta^e_{zt}$ in the balanced growth path, but this does not yield any additional insights.

[60] To see why, substitute equilibrium output for equilibrium consumption in equation (5) and detrend all output variables in the resulting equation by their growth trends. This yields

$$\frac{Y^e_t}{\Gamma^e_t} = \left[\frac{\prod_{z=1}^{Z} (\Gamma^e_{zt})^{\psi_z}}{\Gamma^e_t} \right] \prod_{z=1}^{Z} \left(\frac{Y^e_{zt}}{\Gamma^e_{zt}} \right)^{\psi_z},$$

so that the aggregate growth trend is given by equation (37).

E.1.2 Detrended Firm Problem

With growth-consistent preferences that exhibit constant relative risk aversion, the one-period household discount factor is given by $\Omega = \omega\left(\gamma^e\right)^{-\sigma} < 1$, where σ denotes relative risk aversion and ω is the rate of time preference. Using this expression, the firm problem (32)-(33) along the balance growth path can be written as

$$\Gamma_t^e E_t \sum_{i=0}^{\infty} \left(\alpha_z \omega (\gamma^e)^{1-\sigma}\right)^i \left[(1+\tau)\frac{P_{jzt}}{P_{t+i}} - \frac{W_{t+i} G_{jzt+i} X_{jzt+i}}{P_{t+i} A_{zt+i}}\right] \psi_z \left(\frac{P_{jzt}}{P_{zt+i}}\right)^{-\theta} \left(\frac{P_{zt+i}}{P_{t+i}}\right)^{-1} y,$$

where $y = Y_{t+i}/\Gamma_{t+i}^e$ denotes detrended output. Furthermore, using equation (36) to substitute for A_{zt+i} in the previous equation, augmenting the wage rate by the aggregate growth trend Γ_{t+i}^e and denoting the detrended real wage by $w = \frac{W_t}{P_t \Gamma_t^e}$, we obtain

$$\Gamma_t^e E_t \sum_{i=0}^{\infty} \left(\alpha_z \omega (\gamma^e)^{1-\sigma}\right)^i \left[(1+\tau)\frac{P_{jzt}}{P_{t+i}} - w \frac{G_{jzt+i} X_{jzt+i}}{\Delta_z^e} \frac{\Gamma_{t+i}^e}{\Gamma_{zt+i}^e}\right] \psi_z \left(\frac{P_{jzt}}{P_{zt+i}}\right)^{-\theta} \left(\frac{P_{zt+i}}{P_{t+i}}\right)^{-1} y.$$

Augmenting the relative product price P_{jzt}/P_{t+i} in the previous equation by the item price level and rearranging yields

$$\Gamma_t^e E_t \sum_{i=0}^{\infty} \left(\alpha_z \omega (\gamma^e)^{1-\sigma}\right)^i \left[\psi_z (1+\tau)\frac{P_{jzt}}{P_{zt+i}} - w \frac{G_{jzt+i} X_{jzt+i}}{\Delta_z^e} \left\{\psi_z \left(\frac{P_{zt+i}\Gamma_{zt+i}^e}{P_{t+i}\Gamma_{t+i}^e}\right)^{-1}\right\}\right] \left(\frac{P_{jzt}}{P_{zt+i}}\right)^{-\theta} y.$$

To show that the term in curly brackets in the previous equation is constant along the balanced growth path, we divide each output variable in the demand for item-level output, $Y_{zt} = \psi_z (P_{zt}/P_t)^{-1} Y_t$, by its respective growth trend. This yields

$$\frac{y_z}{y} = \psi_z \left(\frac{P_{zt}\Gamma_{zt}^e}{P_t \Gamma_t^e}\right)^{-1}. \tag{39}$$

Shifting this equation forward and substituting it into the firm objective yields

$$\Gamma_t^e E_t \sum_{i=0}^{\infty} \left(\alpha_z \omega (\gamma^e)^{1-\sigma}\right)^i \left[\psi_z (1+\tau)\frac{P_{jzt}}{P_{zt}} \Pi_z^{-i} - \frac{w}{\Delta_z^e} \frac{y_z}{y} G_{jzt+i} X_{jzt+i}\right] \left(\frac{P_{jzt}}{P_{zt}} \Pi_z^{-i}\right)^{-\theta} y, \tag{40}$$

where we denote the steady-state inflation rate in item z by

$$\Pi_z = P_{zt}/P_{zt-1}.$$

To rewrite the firm objective (40) in terms of the relative prices and marginal costs, we define the relative reset price $p_{jzt} \equiv P_{jzt}/P_{zt}$, which is the nominal price of product j in period t over the item price level in the same period, and the relative price $\widetilde{p}_{jzt+i} \equiv p_{jzt}\Pi_z^{-i}$, which is the nominal reset price in t over the item price level in $t+i$. We also define real marginal costs in units of the good produced in item z according to

$$mc_{jzt} \equiv \frac{W_t}{P_{zt}} \frac{G_{jzt} X_{jzt}}{A_{zt}}.$$

Augmenting this definition by $P_t\Gamma_t^e$ and using equation (36) to substitute for A_{zt} yields

$$mc_{jzt} = \frac{W_t}{P_t\Gamma_t^e}\frac{G_{jzt}X_{jzt}}{\Delta_z^e}\left(\frac{P_{zt}\Gamma_{zt}^e}{P_t\Gamma_t^e}\right)^{-1},$$

and using equation (39) to substitute for the last term on the right hand side in the previous equation shows that marginal costs can be expressed as

$$\psi_z mc_{jzt} = \frac{w}{\Delta_z^e}\frac{y_z}{y}G_{jzt}X_{jzt}. \tag{41}$$

Substituting the previous equation and the definition of the relative price $\widetilde{p}_{jzt+i}$ into the firm objective in equation (40) yields, after dropping the pre-multiplying constant $\psi_z\Gamma_t^e$:

$$E_t\sum_{i=0}^{\infty}\left(\alpha_z\omega(\gamma^e)^{1-\sigma}\right)^i\left[(1+\tau)\widetilde{p}_{jzt+i} - mc_{jzt+i}\right]\left(\widetilde{p}_{jzt+i}\right)^{-\theta}y. \tag{42}$$

E.2 Quadratic Approximation of the Firm Objective

To simplify notation, we drop the item-level subscript z in the remainder of the appendix. The firm objective (42), that we seek to quadratically approximate, can then be written as

$$E_t\sum_{i=0}^{\infty}(\alpha\omega\left(\gamma^e\right)^{1-\sigma})^i\left[(1+\tau)\widetilde{p}_{jt+i} - mc_{jt+i}\right]\left(\widetilde{p}_{jt+i}\right)^{-\theta}y \tag{43}$$

where it is understood that $\alpha, \widetilde{p}_{jt+i}$ and mc_{jt+i} are item specific objects. From equation (41) follows that

$$\ln mc_{jt} = \ln mc_j - (\ln \Pi_j^*)\cdot t + \ln x_{jt}. \tag{44}$$

where $mc_j = \frac{1}{\psi_z}\frac{w}{\Delta_z^e}\frac{y_z}{y}G_{jzt_0}$, with G_{jzt_0} denoting the inverse product-specific productivity level at the time of product entry t_0; $\ln\Pi_j^* = -\ln G_{jzt}/G_{jzt-1}$ is the deterministic constant growth rate of product-specific productivity and $\ln x_{jt} = \ln X_{jzt}$ denotes the stationary stochastic idiosyncratic component of productivity. The values for mc_j and Π_j^* are drawn at the time of product entry from potentially time-varying distributions.

By equation (43), the objective for period $t+i$ is given by

$$D_{jt+i} = \left[(1+\tau)e^{\ln\widetilde{p}_{jt+i}} - e^{\ln mc_{jt+i}}\right]\left(e^{\ln\widetilde{p}_{jt+i}}\right)^{-\theta}y. \tag{45}$$

We approximate this objective to second order in the variables $\ln\widetilde{p}_{jt+i}$ and $\ln mc_{jt+i}$ around the deterministic paths of the flexible price and marginal costs, respectively. The deterministic path of the flexible price is equal to

$$\vartheta mc_{jt+i}^{\det}$$

where $mc_{jt}^{\det}$ denotes the deterministic path of marginal costs which is equal to the value of marginal costs mc_{jt} imposing $x_{jt}=1$, and $\vartheta = \frac{\theta}{\theta-1}\frac{1}{1+\tau}$ denotes the flexible-price markup.

The second-order Taylor approximation of equation (45) yields

$$
\begin{aligned}
D_{jt+i} &= \left(y\vartheta^{-\theta}\right) e^{(1-\theta)\ln mc^{\text{det}}_{jt+i}}\big[-\theta(\ln \widetilde{p}_{jt+i} - \ln(\vartheta mc^{\text{det}}_{jt+i}))^2 \\
&\quad + 2\theta(\ln \widetilde{p}_{jt+i} - \ln(\vartheta mc^{\text{det}}_{jt+i}))(\ln mc_{jt+i} - \ln mc^{\text{det}}_{t+i})\big] + O(3) \\
&= \left(-\theta y\vartheta^{-\theta}\right)\left(mc^{\text{det}}_{jt+i}\right)^{1-\theta}\left[\ln \widetilde{p}_{jt+i} - \ln(\vartheta mc^{\text{det}}_{jt+i}) - (\ln mc_{jt+i} - \ln mc^{\text{det}}_{t+i})\right]^2 + \text{t.i.p.} + O(3) \\
&= \left(-\theta y\vartheta^{-\theta}\right)\left(mc^{\text{det}}_{jt+i}\right)^{1-\theta}\left[\ln \widetilde{p}_{jt+i} - \ln(\vartheta mc_{jt+i})\right]^2 + \text{t.i.p.} + O(3), \qquad (46)
\end{aligned}
$$

where t.i.p. collects terms independent of policy and it follows from equation (44) that $mc^{\text{det}}_{jt+i} = mc_j e^{-(\ln \Pi^*_j)(t+i)}$. Thus, we rewrite the Taylor approximation coefficient in the previous equation according to

$$
-\theta y\vartheta^{-\theta}\left(mc_j e^{-(\ln \Pi^*_j)(t+i)}\right)^{1-\theta} = -\theta y\vartheta^{-\theta} mc_j^{1-\theta}(\Pi^*_j)^{(\theta-1)(t+i)}.
$$

We can now express the expected discounted sum of period profits in equation (43) accurate to second order according to

$$
-\theta y\vartheta^{-\theta} mc_j^{1-\theta}(\Pi^*_j)^{(\theta-1)t} E_t \sum_{i=0}^{\infty}(\alpha\omega\left(\gamma^e\right)^{1-\sigma}(\Pi^*_j)^{\theta-1})^i\big[\ln \widetilde{p}_{jt+i} - \ln(\vartheta mc_{jt+i})\big]^2 + \text{t.i.p.} + O(3)
$$

which is proportional to

$$
-E_t\sum_{i=0}^{\infty}(\alpha\beta_j)^i\big[\ln p_{jt} - i\ln \Pi - \ln(p^*_{jt+i})\big]^2 + \text{t.i.p.} + O(3) \qquad (47)
$$

after substituting $\widetilde{p}_{jt+i} = p_{jt}\Pi^{-i}$ and denoting the firm discount factor by $\beta_j = \omega\left(\gamma^e\right)^{1-\sigma}(\Pi^*_j)^{\theta-1}$ and defining

$$
p^*_{jt+i} = \vartheta mc_{jt+i}
$$

which implies using equation (44)

$$
p^*_{jt} = p^*_j e^{-(\ln \Pi^*_j)t} x_{jt},
$$

which is equal to (9) for $p^*_j = \vartheta mc_j$. While p^*_{jt} denotes the firm's flexible price, the ratio of two firms' flexible prices is equal to the efficient relative price for these firms, whenever price mark-ups are constant across firms and time. In this special case, p^*_{jt} denotes also the efficient relative price.

We can then express the flexible price in period $t+i$ as

$$
p^*_{jt+i} = p^*_{jt} e^{-(\ln \Pi^*_j)i} x_{jt+i} x_{jt}^{-1}.
$$

and substitute into equation (47), which delivers

$$
\max_{\ln p_{jt}} -E_t\sum_{i=0}^{\infty}(\alpha\beta_j)^i\left(\ln p_{jt} - i\ln(\Pi/\Pi^*_j) - \ln p^*_{jt} - \ln x_{jt+i} + \ln x_{jt}\right)^2. \qquad (48)
$$

The first-order condition is given by

$$0 = -2E_t \sum_{i=0}^{\infty} (\alpha\beta_j)^i \left(\ln p_{jt}^{opt} - i \ln(\Pi/\Pi_j^*) - \ln p_{jt}^* - \ln x_{jt+i} + \ln x_{jt}\right),$$

which implies that the optimal price is given by

$$\ln p_{jt}^{opt} = \ln p_{jt}^* - \ln x_{jt} + \left(\frac{\alpha\beta_j}{1-\alpha\beta_j}\right) \ln(\Pi/\Pi_j^*) + E_t(1-\alpha\beta_j) \sum_{i=0}^{\infty} (\alpha\beta_j)^i \ln x_{jt+i} \tag{49}$$

since $\sum_{i=0}^{\infty} (\alpha\beta_j)^i i = \sum_{i=1}^{\infty} (\alpha\beta_j)^i i = \frac{\alpha\beta_j}{(1-\alpha\beta_j)^2}$ with $\alpha\beta_j < 1$. For the limit $\beta_j \to 1$, this reduces to equation (10).

E.3 The First-Stage Regression

To simplify notation, we drop the item-level subscript z in the remainder of this appendix. Starting with equation (12), we substitute $\ln p_{jt}^{opt}$ using equation (10) and also use (9) to obtain

$$\ln p_{jt} = \xi_{jt}(\ln p_{jt-1} - \ln \Pi) + (1-\xi_{jt}) \left(\ln p_j^* - t \ln \Pi_j^* + \frac{\alpha}{1-\alpha} \ln(\Pi/\Pi_j^*) + f(x_{jt})\right), \tag{50}$$

where $f(x_{jt})$ is defined in equation (11).

To derive the OLS estimates of the parameters in equation (13), we rearrange equation (50) to

$$\begin{aligned} \ln p_{jt} + t \ln \Pi_j^* = {} & \xi_{jt}(\ln p_{jt-1} + (t-1) \ln \Pi_j^* - \ln(\Pi/\Pi_j^*)) \\ & + (1-\xi_{jt}) \left(\ln p_j^* + \frac{\alpha}{1-\alpha} \ln(\Pi/\Pi_j^*) + f(x_{jt})\right). \end{aligned} \tag{51}$$

Computing the unconditional expectation yields

$$\begin{aligned} E[\ln p_{jt} + t \ln \Pi_j^*] = {} & \alpha E[\ln p_{jt-1} + (t-1) \ln \Pi_j^*] - \alpha \ln(\Pi/\Pi_j^*) \\ & + (1-\alpha) \left(\ln p_j^* + \frac{\alpha}{1-\alpha} \ln(\Pi/\Pi_j^*)\right), \end{aligned}$$

using independence of ξ_{jt} and $E[f(x_{jt})] = 0$. Given stationarity of the detrended relative price $\ln p_{jt} + t \ln \Pi_j^*$, the previous equation yields

$$E[\ln p_{jt} + t \ln \Pi_j^*] = \ln p_j^*,$$

or

$$\ln p_{jt} = \ln p_j^* - t \ln \Pi_j^* + u_{jt}, \tag{52}$$

where u_{jt} denotes an expectation error. Substituting equation (52) into equation (51), shows that residuals u_{jt} are given by equation (14). They satisfy $E[u_{jt}|p_j^*, \Pi_j^*] = E[u_{jt}] = 0$, which implies that the OLS estimates of regression (13) converge to the true value as the product length increases without bound. For small product lengths, OLS estimates are unbiased but contaminated with sampling error. The effects of sampling error are discussed in appendix G.

E.4 Proof of Proposition 2 (Second-Stage Regression)

To simplify notation, we drop the item-level subscript z in the remainder of this appendix. Squaring equation (14), taking unconditional expectations, and using independence of ξ_{jt} yields

$$E[u_{jt}^2] = E[\xi_{jt}^2]E[(u_{jt-1} - \ln(\Pi/\Pi_j^*))^2] + E[(1-\xi_{jt})^2]E\big[\big(f(x_{jt}) + \frac{\alpha}{1-\alpha}\ln(\Pi/\Pi_j^*)\big)^2\big],$$

where we also used $E[(1-\xi_{jt})\xi_{jt}] = 0$. We can rewrite the previous equation using $E[\xi_t^2] = \alpha$ and $E[(1-\xi_t)^2] = 1-\alpha$, completing the squares to obtain

$$\begin{aligned} E[u_{jt}^2] &= \alpha E[u_{jt-1}^2 + \ln(\Pi/\Pi_j^*)^2 - 2u_{jt-1}\ln(\Pi/\Pi_j^*)] \\ &\quad + (1-\alpha)E\big[f(x_{jt})^2 + \left(\frac{\alpha}{1-\alpha}\ln(\Pi/\Pi_j^*)\right)^2 + 2f(x_{jt})\frac{\alpha}{1-\alpha}\ln(\Pi/\Pi_j^*)\big]. \end{aligned}$$

Recognizing that the expectation of the cross terms in the previous equation are zero because $E[u_{jt}] = 0$ and $E[f(x_{jt})] = 0$ yields

$$E[u_{jt}^2] = \alpha E[u_{jt-1}^2] + \alpha\ln(\Pi/\Pi_j^*)^2 + (1-\alpha)E[f(x_{jt})^2] + (1-\alpha)\left(\frac{\alpha}{1-\alpha}\ln(\Pi/\Pi_j^*)\right)^2.$$

Using $E[u_{jt}^2] = E[u_{jt-1}^2]$ and simplifying terms yields

$$E[u_{jt}^2] = E[f(x_{jt})^2] + \frac{\alpha}{(1-\alpha)^2}(\ln\Pi - \ln\Pi_j^*)^2.$$

Recognizing that $Var[u_{jt}] = E[u_{jt}^2]$, as $E[u_{jt}] = 0$, and $Var[f(x_{jt})] = E[f(x_{jt})^2]$, as $E[f(x_{jt})] = 0$, delivers equation (15).

F Details of the State-Dependent Model

To simplify notation, we drop the item-level subscript z in the remainder of the appendix.

F.1 Setup and OLS regression

Let $z_{jt} = \ln p_{jt} - \ln p_{jt}^*$ be the deviation of the current relative price of product j from the flexible price optimum. Then in between adjustments z_{jt} follows:

$$dz_{jt} = d\ln p_{jt} - d\ln p_{jt}^* = -\underbrace{(\ln\Pi - \ln\Pi_j^*)}_{\mu_j}dt - d\ln x_{jt}$$

$$d\ln x_{jt} = \sum_{i=1}^{N}(\ln x_i - \ln x_{jt})dJ_t^i(\ln x_{jt})$$

where $dJ_t^i(\ln x_{jt})$ is a Poisson jump process with intensity dependent on the current state $\ln x_{jt}$. Since $\ln p_{jt} = \ln p_{jt}^* + z_{jt}$, it follows that:

$$\ln p_{jt} = \ln p_j^* + \ln x_{jt} - t\ln\Pi_j^* + z_{jt} \tag{53}$$

$$E\left[\ln p_{jt} + t\ln\Pi_j^*\right] = \ln p_j^* + \underbrace{E[\ln x_{jt}]}_{=0} + E[z_{jt}]$$

And thus the estimates of OLS regression (13) converge to their true values if $E[z_{jt}|p_j^*, \Pi_j^*] = E[z_{jt}] = 0$, which is true in the limiting case as $\rho \to 0$, as shown below.[61] Furthermore, residuals and their variance can be written as:

$$
\begin{aligned}
u_{jt} &= \ln p_{jt} - \ln p_j^* + t \ln \Pi_j^* = z_{jt} + \ln x_{jt} \\
Var(u_{jt}) &= E[z_{jt}^2] + 2E[z_{jt} \ln x_{jt}] + Var(\ln x_{jt})
\end{aligned} \tag{54}
$$

F.2 Solution

The firm's objective is to maximize its value from equation (18), given by:

$$
V(z, x_i) = \max_{\{\tau_k, \Delta z_{\tau_k}\}_{k=1}^{\infty}} -E\left[\int_0^{\infty} e^{-\rho t} z_t^2 dt + \kappa \sum_{k=1}^{\infty} e^{-\rho \tau_k} \,\middle|\, z_0 = z, x_0 = x_i\right]
$$

The firm's policy consists of a collection of inaction region boundaries $\{\underline{z}(x_i), \overline{z}(x_i)\}$ and reset price gaps $\hat{z}(x_i)$, for all $i \in N$. The HJB equation for the inaction region is given by:

$$
\begin{aligned}
\rho V(z, x_i) = & -z^2 - \mu \partial_z V(z, x_i) \\
& + \sum_{j \neq i}^{N} \lambda_{ij}^X \big(V(z - (\ln x_j - \ln x_i), x_j) - V(z, x_i) \big)
\end{aligned}
$$

The optimal policy satisfies the usual smooth pasting and optimality conditions: $\partial_z V(\hat{z}(x_i), x_i) = \partial_z V(\underline{z}(x_i), x_i) = \partial_z V(\overline{z}(x_i), x_i) = 0$ and $V(\underline{z}(x_i), x_i) = V(\overline{z}(x_i), x_i) = V(\hat{z}(x_i), x_i) - \kappa$. Define $v(z, x_i) = V(z, x_i) - V(\hat{z}(x_1), x_1)$. Then:

$$
\begin{aligned}
\rho v(z, x_i) = & -z^2 - \mu \partial_z v(z, x_i) \\
& + \sum_{j \neq i}^{N} \lambda_{ij}^X \big(v(z - (\ln x_j - \ln x_i), x_j) - v(z, x_i) \big) - \rho V(\hat{z}(x_1), x_1)
\end{aligned}
$$

with $\partial_z v(\hat{z}(x_i), x_i) = \partial_z v(\underline{z}(x_i), x_i) = \partial_z v(\overline{z}(x_i), x_i) = 0$ and $v(\underline{z}(x_i), x_i) = v(\overline{z}(x_i), x_i) = v(\hat{z}(x_i), x_i) - \kappa$. We now take the limit as $\rho \to 0$.

Proposition 4 *As $\rho \to 0$, the scaled value function $\rho V(z, x)$ at any state $\{z, x\}$ converges to a constant: $\lim_{\rho \to 0} \rho V(z, x) = A \in \mathbb{R} \ \forall z, x$.*

[61] While this result is shown formally under the assumption of sufficiently small κ, it holds more generally. As $\rho \to 0$, the firms' value until adjustment becomes the negative expected squared deviation of price gaps from zero, maximizing which requires setting the expected price gap to zero.

All proofs are provided in section F.3. By Proposition 4, $\lim_{\rho \to 0} \rho v(z, x_i) = 0$ and $\lim_{\rho \to 0} \rho V(\hat{z}(x_1), x_1) = A$, so that:

$$\begin{aligned} \lambda_i^X v(z, x_i) &= -z^2 - \mu \partial_z v(z, x_i) \\ &\quad + \sum_{j \neq i}^{N} \lambda_{ij}^X v(z - (\ln x_j - \ln x_i), x_j) - A \end{aligned}$$

where $\lambda_i^X = \sum_{j \neq i}^N \lambda_{ij}^X = -\lambda_{ii}^X$ is the intensity with which $\ln x_t$ is exiting state i. Evaluate the above expression at $z = \hat{z}(x_1), x_i = x_1$ to obtain:

$$A = -\left(\hat{z}(x_1)\right)^2 + \sum_{j \neq 1}^{N} \lambda_{1j}^X v(\hat{z}(x_1) - (\ln x_j - \ln x_1), x_j)$$

Lemma 5 *There exists $\overline{\kappa} > 0$ such that firms find it optimal to adjust after every change in x for all $\kappa < \overline{\kappa}$.*

Suppose that κ is small enough in the sense of Lemma 5. Then firms find it optimal to adjust whenever idiosyncratic state x changes its value. The HJB equation becomes:

$$\begin{aligned} \lambda_i^X v(z, x_i) &= -z^2 - \mu \partial_z v(z, x_i) \\ &\quad + \sum_{j \neq i}^{N} \lambda_{ij}^X v(\hat{z}(x_j), x_j) - \lambda_i^X \kappa - A \end{aligned}$$

with

$$A = -\left(\hat{z}(x_1)\right)^2 + \sum_{j \neq i}^{N} \lambda_{1j}^X v(\hat{z}(x_j), x_j)$$

and value function satisfies:

$$\begin{aligned} v(z, x_i) &= C_i^v e^{-\alpha_i z} - \frac{z^2}{\lambda_i^X} + \frac{2z}{\alpha_i \lambda_i^X} - \frac{2}{\alpha_i^2 \lambda_i^X} + \frac{C_i}{\lambda_i^X} \\ C_i &= \sum_{j \neq i}^{N} \lambda_{ij}^X v(\hat{z}(x_j), x_j) - \lambda_i^X \kappa - A \\ \partial_z v(\hat{z}(x_i), x_i) &= \partial_z v(\underline{z}(x_i), x_i) = \partial_z v(\overline{z}(x_i), x_i) = 0 \\ v(\hat{z}(x_i), x_i) - \kappa &= v(\underline{z}(x_i), x_i) = v(\overline{z}(x_i), x_i) \end{aligned}$$

with $\alpha_i = \frac{\lambda_i^X}{\mu}$. As long as state x remains unchanged, price gaps evolve deterministically with drift $-\mu$. It thus suffices to solve for the reset price gap and only one boundary of the inaction region. From now on, we consider $\mu > 0$ and solve for $\hat{z}(x_i)$ and $\underline{z}(x_i)$ since the upper boundary of the inaction region is irrelevant. Because of symmetry properties of the model, it is straightforward to then recover the solution and all statistics for $\mu < 0$. To ease notation, let $\hat{z}(x_i) = \hat{z}_i$ and $\underline{z}(x_i) = \underline{z}_i$.

Lemma 6 *Suppose $\mu > 0$. Then for each state x_i, optimal policy is determined by the following two conditions:*

$$\underline{z}_i^2 - \hat{z}_i^2 = \lambda_i^X \kappa \tag{55}$$

$$e^{\alpha_i \hat{z}_i}(1 - \alpha_i \hat{z}_i) = e^{\alpha_i \underline{z}_i}(1 - \alpha_i \underline{z}_i) \tag{56}$$

where $\alpha_i = \frac{\lambda_i^X}{\mu}$.

Conditional on state x_i, the price gap distribution satisfies:

$$\lambda_i^X f_i(z) = \mu \partial_z f_i(z)$$

$$\int_{\underline{z}_i}^{\hat{z}_i} f_i(z) dz = 1$$

and is thus given by:

$$f_i(z) = \frac{\alpha_i e^{\alpha_i z}}{e^{\alpha_i \hat{z}_i} - e^{\alpha_i \underline{z}_i}}$$

It follows that:

$$E[z|x_i] = \int_{\underline{z}_i}^{\hat{z}_i} z f_i(z) dz = 0$$

$$E[z] = 0$$

$$E[z^2|x_i] = \int_{\underline{z}_i}^{\hat{z}_i} z^2 f_i(z) dz = \frac{\hat{z}_i + \underline{z}_i}{\alpha_i} - \hat{z}_i \underline{z}_i \tag{57}$$

$$E[z^2] = E_x \left[\frac{\hat{z}_i + \underline{z}_i}{\alpha_i} - \hat{z}_i \underline{z}_i \right] \tag{58}$$

where $E_x[\cdot]$ is the expectation with respect to stationary distribution of x.

Proposition 7 *For μ close to zero, $E[z^2] = E\left[\frac{1}{(\lambda_i^X)^2}\right]\mu^2 + O(4)$.*

Finally, note that $E[zx] = E\big[x_i E[z|x_i]\big] = 0$ and the main object of interest – the variance of residuals from the OLS regression (13) – is given by:

$$\begin{aligned} Var(u_{jt}) &= Var(\ln x_{jt}) + E\left[\frac{1}{(\lambda_i^X)^2}\right]\mu_j^2 + O(4) \\ &= Var(\ln x_{jt}) + E\left[\frac{1}{(\lambda_i^X)^2}\right](\ln \Pi - \ln \Pi_j^*)^2 + O(4) \end{aligned}$$

F.3 Proofs

Proof of Proposition 4. The proof here extends Lemma 3 in Online Appendix of Alvarez et al. (2019) to a setting with two state variables. Let $V(z, x, \rho)$ be the value function in state $\{z, x\}$ under discount rate ρ. We can write $\rho V(z, x, \rho)$ as follows:

$$\rho V(z,x,\rho) = -E\left[\rho \int_0^{\tau_N} e^{-\rho t} z_t^2 dt\right] - \kappa E\left[\rho \sum_{k=1}^{N} e^{-\rho \tau_k}\right]$$
$$\underbrace{-\rho E\left[\int_0^{\infty} e^{-\rho(\tau_N+t)} z_{\tau_N+t}^2 dt + \kappa \sum_{k=1}^{\infty} e^{-\rho \tau_{N+k}}\right]}_{\rho E\left[e^{-\rho\tau_N} V(z_{\tau_N}, x_{\tau_N}, \rho)\right]}$$

where τ_N is the N-th adjustment and all expectation operators are conditional on $z_0 = z, x_0 = x$. Subtract $\rho E\left[e^{-\rho\tau_N} V(z,x,\rho)\right]$ from both sides and divide by $(1 - E\left[e^{-\rho\tau_N}\right])$ to obtain:

$$\rho V(z,x,\rho) = -\frac{\rho}{1-E\left[e^{-\rho\tau_N}\right]} E\left[\int_0^{\tau_N} e^{-\rho t} z_t^2 dt\right] - \frac{\rho\kappa}{1-E\left[e^{-\rho\tau_N}\right]} E\left[\sum_{k=1}^{N} e^{-\rho\tau_k}\right]$$
$$\frac{\rho}{1-E\left[e^{-\rho\tau_N}\right]} E\left[e^{-\rho\tau_N}\left(V(z_{\tau_N}, x_{\tau_N}, \rho) - V(z,x,\rho)\right)\right]$$

Take the limit as $\rho \to 0$. Note that $\frac{\rho}{1-E\left[e^{-\rho\tau_N}\right]} \to \frac{1}{E[\tau_N]}$ and thus:

$$\lim_{\rho\to 0} \rho V(z,x,\rho) = -\frac{1}{E\left[\tau_N\right]} E\left[\int_0^{\tau_N} z_t^2 dt\right] - \frac{\kappa N}{E\left[\tau_N\right]}$$
$$\frac{1}{E\left[\tau_N\right]} \lim_{\rho\to 0} E\left[V(z_{\tau_N}, x_{\tau_N}, \rho) - V(z,x,\rho)\right]$$

By Lemma 8, $|V(z_{\tau_N}, x_{\tau_N}, \rho) - V(z,x,\rho)| \leq C \in \mathbb{R}$ for all $\rho > 0$ and thus this also holds in the limit as $\rho \to 0$. As we take the limit with $N \to \infty$, the first term converges to the unconditional expected squared gap $E[z^2]$, the second term converges to adjustment frequency λ_a times adjustment cost κ, and the third term vanishes as $E\left[\tau_N\right] \to \infty$. Thus $\lim_{\rho\to 0} \rho V(z,x,\rho) = -E[z^2] - \kappa\lambda_a \equiv A$ for all z, x. ■

Lemma 8 *There exists $C \in \mathbb{R}$ such that for any $\rho > 0$ and any z, x, z', x', $|V(z,x) - V(z',x')| \leq C$.*

Proof. First, we show that $\rho V(z, x_i)$ is bounded from below. To see that, recall that $V(z, x_i)$ is achieved under the optimal adjustment policy, meaning that the value of any feasible policy is weakly lower. Consider the following policy: the firm adjusts its price gap whenever it is hit by a Poisson x shock. In addition, it also adjusts at random times with Poisson intensity λ_i, which is specific to each state x_i. These intensities satisfy the following condition: $\lambda_i^X + \lambda_i = \max_i \lambda_i^X \equiv \lambda$, such that in every state x_i firms adjust with equal intensity λ. Since adjustments occur exogenously, firms only choose the reset price gap $\hat{z}_i$ to maximize expected profits until the next adjustment:

$$\max_{\hat{z}_i} E\left[-\int_0^{\tau} e^{-\rho t} z_t^2 \,\middle|\, z_0 = \hat{z}_i\right] = \max_{\hat{z}_i} E\left[-\int_0^{\infty} e^{-(\rho+\lambda)t} z_t^2 \,\middle|\, z_0 = \hat{z}_i\right]$$

Because in between adjustments price gaps drift deterministically ($z_t = \hat{z}_i - \mu t$) and adjustment intensities are equalized across states, optimal reset price gap does not depend on x and satisfies FOC:

$$\int_0^\infty e^{-(\rho+\lambda)t}(\hat{z} - \mu t) = 0 \Longrightarrow \hat{z} = \frac{\mu}{\rho + \lambda}$$

Denote by $\tilde{V}(z, x)$ the value function under this policy. Since $\partial_z \tilde{V}(\hat{z}, x_i) = 0$, evaluating the HJB equation at $\hat{z}$ yields:

$$\begin{aligned}
\rho\tilde{V}(\hat{z}, x_i) &= -\hat{z}^2 + \lambda_i\left(\tilde{V}(\hat{z}, x_i) - \kappa - \tilde{V}(\hat{z}, x_i)\right) + \sum_{j\neq i}^{N} \lambda_{ij}^X \left(\tilde{V}(\hat{z}, x_j) - \kappa - \tilde{V}(\hat{z}, x_i)\right) \\
&= -\hat{z}^2 + \sum_{j\neq i}^{N} \lambda_{ij}^X \left(\tilde{V}(\hat{z}, x_j) - \tilde{V}(\hat{z}, x_i)\right) - \kappa \underbrace{\left(\lambda_i + \sum_{j\neq i}^{N} \lambda_{ij}^X\right)}_{=\lambda}
\end{aligned}$$

It is straightforward to show that $\tilde{V}(\hat{z}, x_i) = \tilde{V}(\hat{z}, x_j)$ for all i and j. Assume the opposite and let $\overline{v} = \max_i \tilde{V}(\hat{z}, x_i)$ and $\underline{v} = \min_i \tilde{V}(\hat{z}, x_i)$. Then:

$$\begin{aligned}
\rho\overline{v} &= -\hat{z}^2 + \sum_{j\neq i(\overline{v})}^{N} \lambda_{i(\overline{v})j}^X \underbrace{\left(\tilde{V}(\hat{z}, x_j) - \overline{v}\right)}_{\leq 0} - \lambda\kappa \\
&\leq -\hat{z}^2 + \sum_{j\neq i(\underline{v})}^{N} \lambda_{i(\underline{v})j}^X \underbrace{\left(\tilde{V}(\hat{z}, x_j) - \underline{v}\right)}_{\geq 0} - \lambda\kappa = \rho\underline{v}
\end{aligned}$$

Meaning $\underline{v} = \overline{v}$. As a result, $\rho\tilde{V}(\hat{z}, x_i) = -\hat{z}^2 - \lambda\kappa = -\frac{\mu^2}{(\rho+\lambda)^2} - \lambda\kappa \geq -\frac{\mu^2}{\lambda^2} - \lambda\kappa$ for any $\rho > 0$. Thus for the true value function evaluated at the true optimal reset price gap $\hat{z}(x_i)$ it holds that $\rho V(\hat{z}(x_i), x_i) \geq \rho\tilde{V}(\hat{z}, x_i) \geq -\frac{\mu^2}{\lambda^2} - \lambda\kappa$ for all $\rho > 0$.

Consider now the true value function $V(z, x_i)$ and pick i such that $V(\hat{z}(x_i), x_i) = \max_j V(\hat{z}(x_j), x_j)$. The HJB equation for this value function satisfies:

$$\begin{aligned}
-\frac{\mu^2}{\lambda^2} - \lambda\kappa \leq \rho V(\hat{z}(x_i), x_i) &= \underbrace{-\left(\hat{z}(x_i)\right)^2}_{\leq 0} - \mu \underbrace{\partial_z V(\hat{z}(x_i), x_i)}_{=0} \\
&\quad + \sum_{j\neq i}^{N} \lambda_{ij}^X \left(\underbrace{V(\hat{z}(x_i) - (\ln x_j - \ln x_i), x_j)}_{\leq V(\hat{z}(x_j), x_j)} - V(\hat{z}(x_i), x_i)\right) \\
&\leq \sum_{j\neq i}^{N} \lambda_{ij}^X \underbrace{\left(V(\hat{z}(x_j), x_j) - V(\hat{z}(x_i), x_i)\right)}_{\leq 0} \leq 0
\end{aligned}$$

It follows that whenever $\lambda_{ij}^X > 0$:

$$\left(-\frac{\mu^2}{\lambda^2} - \lambda\kappa\right) / \lambda_{ij}^X \leq V(\hat{z}(x_j), x_j) - V(\hat{z}(x_i), x_i) \leq 0$$

For the states j where $\lambda_{ij}^X = 0$ we can bound the difference $V(\hat{z}(x_j), x_j) - V(\hat{z}(x_i), x_i)$ iteratively because the network of x_i is connected (every two states are connected by some path). In addition, for any z, x_i:

$$V(\hat{z}(x_i), x_i) - \kappa \leq V(z, x_i) \leq V(\hat{z}(x_i), x_i)$$

Therefore there exists $C \in \mathbb{R}$ such that for all $\rho > 0$, $|V(z,x) - V(z', x')| \leq C$ for all z, x, z', x'. ■

Proof of Lemma 5. Consider a model M in which firms are forced to adjust after every change in x, but can also adjust at other times and choose the boundaries of inaction regions and reset price gaps. Suppose we now allow the firms to adjust whenever they find it to be optimal. They will adjust their policies $\{\underline{z}(x_i), \hat{z}(x_i), \overline{z}(x_i)\}_{i=1}^N$ only if changes in x keep price gaps within the bounds of inaction regions. Otherwise the optimal policy in model M and in the model of interest coincide, meaning that firms find it optimal to adjust after every change in x. To see that, compare the HJB equations in the original model (first line) and model M (second line):

$$\begin{aligned}
\lambda_i^X v(z, x_i) &= -z^2 - \mu \partial_z v(z, x_i) \\
&\quad + \sum_{j \neq i}^{N} \lambda_{ij}^X v(z - (\ln x_j - \ln x_i), x_j) - A \\
\lambda_i^X v(z, x_i) &= -z^2 - \mu \partial_z v(z, x_i) \\
&\quad + \sum_{j \neq i}^{N} \lambda_{ij}^X (v(\hat{z}(x_j), x_j) - \kappa) - A
\end{aligned}$$

If upon the change in x, $z - (\ln x_j - \ln x_i) \notin [\underline{z}(x_j), \overline{z}(x_j)]$, then $v(z - (\ln x_j - \ln x_i), x_j) = v(\hat{z}(x_j), x_j) - \kappa$ and the value functions in the two models coincide. Therefore, $\overline{\kappa}$ is such that $\min_{ij} |\ln x_i - \ln x_j| = \max_i \overline{z}(x_i) - \min_i \underline{z}(x_i)$ in model M. Such $\overline{\kappa} > 0$ always exists since for all i $\lim_{\kappa \to 0} \overline{z}(x_i) = \lim_{\kappa \to 0} \underline{z}(x_i) = 0$. ■

Proof of Lemma 6. From $\partial_z v(\hat{z}_i, x_i) = 0$ and $\partial_z v(\underline{z}_i, x_i) = 0$ it follows:

$$\begin{gathered}
-\alpha_i C_i^v e^{-\alpha_i \hat{z}_i} - \frac{2\hat{z}_i}{\lambda_i^X} + \frac{2}{\alpha_i \lambda_i^X} = 0 = -\alpha_i C_i^v e^{-\alpha_i \underline{z}_i} - \frac{2\underline{z}_i}{\lambda_i^X} + \frac{2}{\alpha_i \lambda_i^X} \\
-\alpha_i C_i^v - \frac{2\hat{z}_i e^{\alpha_i \hat{z}_i}}{\lambda_i^X} + \frac{2e^{\alpha_i \hat{z}_i}}{\alpha_i \lambda_i^X} = 0 = -\alpha_i C_i^v - \frac{2\underline{z}_i e^{\alpha_i \underline{z}_i}}{\lambda_i^X} + \frac{2e^{\alpha_i \underline{z}_i}}{\alpha_i \lambda_i^X} \\
e^{\alpha_i \hat{z}_i}(1 - \alpha_i \hat{z}_i) = e^{\alpha_i \underline{z}_i}(1 - \alpha_i \underline{z}_i)
\end{gathered}$$

Similarly:

$$
\begin{aligned}
-\alpha_i C_i^v e^{-\alpha_i \hat{z}_i} - \frac{2\hat{z}_i}{\lambda_i^X} + \frac{2}{\alpha_i \lambda_i^X} &= -\alpha_i C_i^v e^{-\alpha_i \underline{z}_i} - \frac{2\underline{z}_i}{\lambda_i^X} + \frac{2}{\alpha_i \lambda_i^X} \\
C_i^v e^{-\alpha_i \hat{z}_i} + \frac{2\hat{z}_i}{\alpha_i \lambda_i^X} &= C_i^v e^{-\alpha_i \underline{z}_i} + \frac{2\underline{z}_i}{\alpha_i \lambda_i^X} \\
C_i^v e^{\alpha_i(\underline{z}_i - \hat{z}_i)} &= C_i^v + e^{\alpha_i \underline{z}_i} \frac{2(\underline{z}_i - \hat{z}_i)}{\alpha_i \lambda_i^X}
\end{aligned} \tag{59}
$$

From $v(\hat{z}_i, x_i) - \kappa = v(\underline{z}_i, x_i)$ it follows:

$$
\begin{aligned}
C_i^v e^{-\alpha_i \hat{z}_i} - \frac{\hat{z}_i^2}{\lambda_i^X} + \frac{2\hat{z}_i}{\alpha_i \lambda_i^X} - \kappa &= C_i^v e^{-\alpha_i \underline{z}_i} - \frac{\underline{z}_i^2}{\lambda_i^X} + \frac{2\underline{z}_i}{\alpha_i \lambda_i^X} \\
C_i^v e^{\alpha_i(\underline{z}_i - \hat{z}_i)} + e^{\alpha_i \underline{z}_i} \left[\frac{2(\hat{z}_i - \underline{z}_i)}{\alpha_i \lambda_i^X} + \frac{\underline{z}_i^2 - \hat{z}_i^2}{\lambda_i^X} - \kappa \right] &= C_i^v \\
\underline{z}_i^2 - \hat{z}_i^2 &= \lambda_i^X \kappa
\end{aligned}
$$

where the last line follows from (59). ■

Lemma 9 *For every state x_i, $\hat{z}_i \frac{\partial \hat{z}_i}{\partial \mu} = \underline{z}_i \frac{\partial \underline{z}_i}{\partial \mu} = \frac{E[z^2|x_i]}{\mu}$.*

Proof. The first equality follows directly from the first order derivative of equilibrium condition (55) with respect to μ. For the second equality, differentiate equilibrium condition (56) and collect terms:

$$
\begin{aligned}
e^{\alpha_i \hat{z}_i} \left[\frac{\partial \hat{z}_i}{\partial \mu} - \frac{\hat{z}_i}{\mu} \right] \hat{z}_i &= e^{\alpha_i \underline{z}_i} \left[\frac{\partial \underline{z}_i}{\partial \mu} - \frac{\underline{z}_i}{\mu} \right] \underline{z}_i \\
(1 - \alpha_i \underline{z}_i) \left[\frac{\partial \hat{z}_i}{\partial \mu} - \frac{\hat{z}_i}{\mu} \right] \hat{z}_i &= (1 - \alpha_i \hat{z}_i) \left[\frac{\partial \underline{z}_i}{\partial \mu} - \frac{\underline{z}_i}{\mu} \right] \underline{z}_i \\
\underline{z}_i \frac{\partial \underline{z}_i}{\partial \mu} (\alpha_i \hat{z}_i - \alpha_i \underline{z}_i) &= \frac{\hat{z}_i^2 (1 - \alpha_i \underline{z}_i) - \underline{z}_i^2 (1 - \alpha_i \hat{z}_i)}{\mu} \\
\underline{z}_i \frac{\partial \underline{z}_i}{\partial \mu} &= \frac{1}{\mu} \frac{\hat{z}_i^2 - \underline{z}_i^2 - \alpha_i \hat{z}_i \underline{z}_i (\hat{z}_i - \underline{z}_i)}{\alpha_i (\hat{z}_i - \underline{z}_i)} \\
&= \frac{1}{\mu} \left[\frac{\hat{z}_i + \underline{z}_i}{\alpha_i} - \hat{z}_i \underline{z}_i \right] = \frac{E[z^2|x_i]}{\mu}
\end{aligned}
$$

where the second line uses (56) and the third line uses $\hat{z}_i \frac{\partial \hat{z}_i}{\partial \mu} = \underline{z}_i \frac{\partial \underline{z}_i}{\partial \mu}$. ■

Lemma 10 *As $\mu \to 0$, $\hat{z}_i \to 0$, $\underline{z}_i \to -\sqrt{\lambda_i^X \kappa}$ and $E[z^2] \to 0$.*

Proof. Combine equilibrium conditions (55) and (56) to obtain:

$$
\underbrace{\left(\mu + \lambda_i^X \sqrt{\lambda_i^X \kappa + \hat{z}_i^2} \right)}_{>0} = \left(\mu - \lambda_i^X \hat{z}_i \right) \underbrace{e^{\frac{\lambda_i^X}{\mu} \left(\hat{z}_i + \sqrt{\lambda_i^X \kappa + \hat{z}_i^2} \right)}}_{>0}
$$

Since the LHS is always positive, and so is the exponent on the RHS, $\lim_{\mu\to 0} \hat{z}_i = 0$. It then follows from (55) that $\lim_{\mu\to 0} \underline{z}_i = -\sqrt{\lambda_i^X \kappa}$ and from (58) that $\lim_{\mu\to 0} E[z^2] = 0$. ■

Proof of Proposition 7. From Lemmas 9 and 10, and equation (57) it follows that:

$$\underline{z}_i' \equiv \frac{\partial \underline{z}_i}{\partial \mu} = \frac{1}{\lambda_i^X} + \frac{\mu \hat{z}_i - \lambda_i^X \hat{z}_i \underline{z}_i}{\mu \lambda_i^X \underline{z}_i}$$

$$\lim_{\mu\to 0} \underline{z}_i' = \frac{1}{\lambda_i^X} - \lim_{\mu\to 0} \frac{\hat{z}_i}{\mu} = \frac{1}{\lambda_i^X} - \lim_{\mu\to 0} \hat{z}_i'$$

At the same time, by Lemma 9: $\hat{z}_i = \frac{z_i \underline{z}_i'}{\hat{z}_i'}$, and by Lemma 10: $\lim_{\mu\to 0} \frac{\underline{z}_i'}{\hat{z}_i'} = 0$. It then follows that:

$$O(1) = \frac{\underline{z}_i'}{\hat{z}_i'} = \frac{\frac{1}{\lambda_i^X} - \hat{z}_i' + O(1)}{\hat{z}_i'} = \frac{1 + O(1)}{\lambda_i^X \hat{z}_i'} - 1$$

And therefore $\lim_{\mu\to 0} \hat{z}_i' = \frac{1}{\lambda_i^X}$. From (55) it follows that $\lim_{\mu\to 0} \underline{z}_i' = 0$ and from (58) that $\lim_{\mu\to 0} \frac{\partial E[z^2]}{\partial \mu} = 0$. If $\hat{z}_i$ is twice differentiable at $\mu = 0$, then due to anti-symmetry ($\hat{z}_i(\mu) = -\hat{z}_i(-\mu)$), $\hat{z}_i''(0) = 0$. It follows that $\hat{z}_i' = \frac{1}{\lambda_i^X} + O(2)$ and $\hat{z}_i = \frac{\mu}{\lambda_i^X} + O(3)$. Using Lemma 9 we obtain that:

$$E[z^2] = E\left[\frac{1}{(\lambda_i^X)^2}\right] \mu^2 + O(4)$$

■

Lemma 11 *Suppose $\lambda_i^X = \Lambda$ for all i. Then, as $\mu \to 0$, adjustment frequency $\Lambda_a = \Lambda + O(4)$.*

Proof. Since $\lambda_i^X = \Lambda$, we can omit the i index. The expected stopping time $\tau(z)$ solves the following ODE: $\Lambda\tau(z) = 1 - \mu\partial_z\tau(z)$, together with boundary condition $\tau(\underline{z}) = 0$, and is given by $\tau(z) = \frac{1}{\Lambda}\left(1 - e^{\alpha(\underline{z} - \hat{z})}\right)$. It follows from Lemma 6 and equation (58) that:

$$\Lambda_a \equiv \frac{1}{\tau(\hat{z})} = \frac{1}{\kappa}\left(\underline{z}^2 - E[z^2]\right)$$

Lemma 10 implies that as $\mu \to 0$, $\Lambda_a \to \Lambda$. Furthermore:

$$\begin{aligned} \frac{\partial \Lambda_a}{\partial \mu} &= \frac{1}{\kappa}\left(2\underline{z}\frac{\partial \underline{z}}{\partial \mu} - \frac{\partial E[z^2]}{\partial \mu}\right) \\ &= \frac{1}{\kappa}\left(2\frac{E[z^2]}{\mu} - 2\frac{\mu}{\Lambda^2} + O(3)\right) = O(3) \end{aligned}$$

where the last line follows from Lemma 9 and Proposition 7. Therefore, $\Lambda_a = \Lambda + O(4)$. ■

G Details of the Regression Approach

This section discusses econometric details associated with estimating our key equation (15), which relates price distortions to suboptimal inflation at the product level. In our baseline empirical approach, we estimate equation (15) at the level of finely disaggregated expenditure items, exploiting variation across products within the item. Our sample contains more than 1000 expenditure items, so that we obtain a large number of estimates of the coefficient of interest c_z in equation (15).

We use a two-step estimation approach, because neither the left-hand side variable nor the right hand-side variables in equation (15) can be directly observed. This section presents this approach and discusses how first-stage estimation errors affect second-stage regression outcomes. In particular, it shows that first-stage error biases the estimates of the coefficient c_z towards zero, i.e., towards finding no marginal effect of suboptimal inflation on price distortions.

Our first-stage estimation consists of a seemingly unrelated regression (SUR) system that contains two equations. The left-hand side variable in equation (15) can be estimated using the residuals of relative-price regressions of the form

$$\ln p_{jzt} = \ln p^*_{jz} - \left(\ln \Pi^*_{jz}\right) \cdot t + u_{jzt} \tag{60}$$

where j denotes the product, $z \in \{1, ...Z\}$ the expenditure item under consideration, and Z the total number of expenditure items in our sample. Theory implies that $E[u_{jzt}|p^*_{jz}, \Pi^*_{jz}] = E[u_{jzt}] = 0$ and that the OLS estimates in (60) satisfy $E[\widehat{\ln p^*_{jz}}] = \ln p^*_{jz}$ and $E[\widehat{\ln \Pi^*_{jzt}}] = \ln \Pi^*_{jz}$. An unbiased estimate for the residual variance is given by $\widehat{Var}(u_{jzt}) = \frac{1}{T_{jz}-2} \sum_t \left(\widehat{u}_{jzt}\right)^2$, where T_{jz} denotes the number of price observations for product j in item z.

Estimation of the right-hand side variables in equation (15) requires estimating the average inflation rate, $\ln \Pi_z$, and the product specific optimal inflation rate, $\ln \Pi^*_{jz}$. However, having two first-stage estimates on the right-hand side of equation (15) is unattractive on econometric grounds.[62] A preferred way to proceed is to estimate directly the gap between the item-level inflation rate and the product-specific optimal inflation rate ($\ln \Pi_z - \ln \Pi^*_{jz}$). This can be achieved by adding the price level equation

$$\ln P_{zt} = \ln P_{z0} + \ln \Pi_z \cdot t$$

to equation (9), which delivers for every product *another* first-stage regression of the form

$$\ln P_{jzt} = \ln \widetilde{a}_{jz} + \left(\ln \Pi_z/\Pi^*_{jz}\right) \cdot t + \widetilde{u}_{jzt} \tag{61}$$

where P_{jzt} denotes the *nominal* product price and $\ln \widetilde{a}_{jz} = \ln p^*_{jz} + \ln P_{z0}$. Again, theory implies that $E[\widetilde{u}_{jzt}|p^*_{jz}, \Pi^*_{jz}] = E[\widetilde{u}_{jzt}] = 0$. Equation (61) reveals that the time trend in the nominal price of the product directly identifies the gap between item-level inflation

[62] It requires discussing, amongst other things, the covariance in the estimation errors of these two right-hand side variables.

and the product-specific optimal inflation rate. Equations (60) and (61) jointly make up our first-stage SUR system.

Since the SUR system (60)-(61) does not feature exclusion restrictions, OLS estimation is identical to GLS estimation, despite the presence of correlated residuals. OLS estimation of (60) delivers an unbiased estimate for the residual variance of interest, $Var(u_{jzt})$, and OLS estimation of equation (61) an unbiased estimate of the gap $\ln \Pi_z/\Pi^*_{jz}$.

The first-stage estimates for each product j within expenditure item z can then be used to estimate the second-stage equation

$$\widehat{Var}(u_{jzt}) = v_z + c_z \cdot (\ln \widehat{\Pi_z/\Pi^*_{jz}})^2 + \varepsilon_{jz} \tag{62}$$

using OLS estimation. This delivers an estimate of c_z for each expenditure item $z = 1, \ldots, Z$. The error term ε_{jz} in equation (62) absorbs measurement error of the left-hand side variable, as discussed below, as well as the higher-order approximation errors implied by menu-cost models, see equation (19).

While the first-stage estimates $\widehat{Var}(u_{jzt})$ and $\ln \widehat{\Pi_z/\Pi^*_{jz}}$ are unbiased, they are contaminated by sampling error. Sampling error is an important concern because the product price time series underlying the first-stage system can be relatively short. Fortunately, the effect of the first-stage sampling error consists solely of biasing the estimate of c_z towards zero, as we show next.

To illustrate this point, we assume that the first-stage residuals $(u_{jzt}, \widetilde{u}_{jzt})$ are normally distributed. (The more general case with non-normal errors is discussed in appendix G.1 below.) When estimating the SUR system (60)-(61), the estimation error in $\ln \widehat{\Pi_z/\Pi^*_{jz}}$ is orthogonal to the estimation error in the residuals $\{\widehat{u}_{jzt}\}$, by construction of the OLS estimate. With normality, both estimation errors are also independent of each other. Therefore, the estimation error in $\widehat{Var}(u_{jzt})$ on the l.h.s. of equation (62) is independent of the estimation error in $(\ln \widehat{\Pi_z/\Pi^*_{jz}})^2$ on the r.h.s. of the equation, because both variables are nonlinear transformations of independent random variables.

First-stage estimation error on the l.h.s. of equation (62) thus takes the form of classical measurement error: it does not generate any bias in the second-stage estimates of c_z, instead gets absorbed by the regression residual ε_{jz}. However, first-stage estimation error in $(\ln \widehat{\Pi_z/\Pi^*_{jz}})^2$ biases the second-stage estimate of c_z towards zero. This is so because measurement error in $(\ln \widehat{\Pi_z/\Pi^*_{jz}})^2$ generates a classic attenuation effect. In addition, estimation error in $\ln \widehat{\Pi_z/\Pi^*_{jz}}$ raises the expected value of $(\ln \widehat{\Pi_z/\Pi^*_{jz}})^2$, which generates a further bias towards zero.

Our second-stage estimates for c_z thus provides a *lower bound* of the true marginal effect of suboptimal inflation on price distortions. Since we are interested in rejecting the

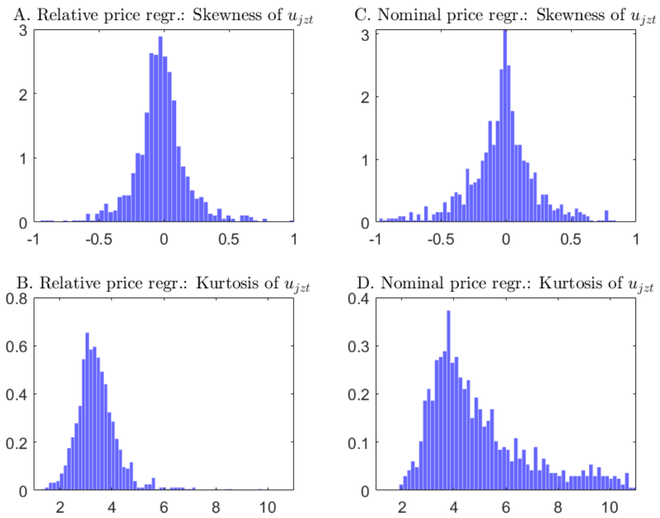

Figure 15: Skewness and kurtosis of the first-stage regression residuals

null hypothesis of inflation *not* creating price distortions, $H_0 : c_z = 0$, the bias is working against our main finding.

Finally, to insure that our results are not driven by outliers, e.g., associated with errors in price collection, we eliminate within each expenditure item all products falling into the top 5% of the distribution of residual variances $\widehat{Var}(u_{jzt})$ and the top 5% of estimated inflation gaps $(\ln \widehat{\Pi_z/\Pi^*_{jz}})^2$ when running our second-stage regression. Results are robust to choosing different thresholds.

G.1 General Case with Non-Normal First-Stage Residuals

Figure 15 reports the skewness and kurtosis of the first-stage regression residuals of equation (60) (left-hand side panels) and equation (61) (right-hand side panels) across the considered expenditure items.[63] The top panels show that skewness is centered around zero and relatively tightly so, in line with the zero skewness of the normal distribution. For kurtosis, shown in the lower panels of figure 15, the situation looks different. Kurtosis values often lie above and below the value of 3 implied by a normal distribution.

We now show that quite similar arguments apply to our second-stage estimates of c_z when first-stage residuals fail to be normal. In fact, to insure that there is at most a downward bias in the second-stage estimate of c_z, it is sufficient that the estimation error in the l.h.s. variable $\widehat{Var}(u_{jzt})$ in equation (62) is orthogonal to (rather than independent of) the estimation error in the r.h.s. regressor $(\ln \widehat{\Pi_z/\Pi^*_{jz}})^2$.

[63] The measures use outlier trimmed residuals by considering the 2.5%-97.5% quantile of the residual distribution.

Recall that the errors in $(\widehat{\ln \Pi_z/\Pi_{jz}^*})$ and $\{\widehat{u}_{jzt}\}$ are orthogonal by construction. A violation of orthogonality between $(\widehat{\ln \Pi_z/\Pi_{jz}^*})^2$ and $\widehat{Var}(u_{jzt})$ can thus only arise because these variables are nonlinear rather than linear functions of $\widehat{\ln \Pi_z/\Pi_{jz}^*}$ and $\{\widehat{u}_{jzt}\}$, respectively. This illustrates that violations of orthogonality conditions are somewhat unlikely to emerge on a priori grounds, even in the absence of normality.

We show below that orthogonality of the estimation errors in $(\widehat{\ln \Pi_z/\Pi_{jz}^*})^2$ and $\widehat{Var}(u_{jzt})$ holds whenever the residuals satisfy

$$Cov[\left((0,1)\left(X'X\right)^{-1}X'u(0,1)'\right)^2, (1,0)'u'Mu(1,0)|X] = 0, \tag{63}$$

where

$$X' \equiv \begin{pmatrix} 1\,1\,1\ldots \\ 0\,1\,2\ldots \end{pmatrix} \tag{64}$$

is the matrix of first-stage regressors and M the matrix defined in (65) below. Condition (63) is a condition on the true residuals u, which is satisfied in the special case with normal errors. Condition (63) holds by construction when replacing the true residuals u by the estimated OLS or GLS residuals $\widehat{u}$ and thus cannot be tested empirically using the regression residuals.[64]

To understand why condition (63) insures that the same outcome is obtained as with normality, consider our first-stage regression system, which takes the form of a seemingly unrelated regression (SUR) system:

$$\underbrace{Y}_{T\times 2} = \underbrace{X}_{T\times 2}\underbrace{\beta}_{2\times 2} + \underbrace{u}_{T\times 2},$$

where X denotes the (deterministic) regressors defined in (64) and Y the stacked vector of the left-hand side variables (p_{jzt}, P_{jzt}) in equations (60) and (61). Letting u_t denote the residuals at date t and u the stacked residual vector, we have $E[u_t] = 0$ and

$$Var(u_t) = \begin{pmatrix} v_{11}^2 \; v_{12} \\ v_{12} \; v_{22}^2 \end{pmatrix}.$$

Since the SUR system does not feature exclusions restrictions, OLS estimation is identical to GLS estimation. In particular, the OLS/GLS estimate $\widehat{\beta}$ of β is given by

$$\widehat{\beta} \equiv \left(X'X\right)^{-1} X'Y$$

[64]Using the notation introduced below, this follows from the fact that

$$\begin{aligned} &\left(X'V^{-1}X\right)^{-1} X'V^{-1}\widehat{u} \\ &= \left(X'V^{-1}X\right)^{-1} X'V(I - X\left(X'V^{-1}X\right)^{-1} X'V^{-1})Y \\ &= 0. \end{aligned}$$

and the regression residuals by

$$\underbrace{\widehat{u}}_{T\times 2} = MY = Mu \text{ where } M \equiv \left(I - X\left(X'X\right)^{-1}X'\right) \tag{65}$$

We have

$$\begin{aligned}
E[\widehat{u}'\widehat{u}|X] &= E[\underbrace{u'}_{2\times T}\underbrace{M'M}_{T\times T}\underbrace{u}_{T\times 2}|X] \\
&= E[\underbrace{u'}_{2\times T}\underbrace{M}_{T\times T}\underbrace{u}_{T\times 2}|X] \\
&= tr(M)E[u'u|X] \\
&= \frac{1}{T-2}\begin{pmatrix} v_{11}^2 \; v_{12} \\ v_{12} \; v_{22}^2 \end{pmatrix},
\end{aligned}$$

An unbiased estimate of the residual variance v_{11}^2 is thus given by

$$\widehat{v_{11}^2} \equiv \frac{(1,0)'\widehat{u}'\widehat{u}(1,0)}{T-2}. \tag{66}$$

The estimation errors in the second-stage regression variables $((0,1)\left(\widehat{\beta}-\beta\right)(0,1)')^2$ and $(\widehat{v_{11}^2} - v_{11}^2)$, are orthogonal if and only if

$$\begin{aligned}
& E[\left(\left((0,1)\left(\widehat{\beta}-\beta\right)(0,1)'\right)^2\right)\left(\widehat{v_{11}^2} - v_{11}^2\right)|X] \stackrel{!}{=} 0 \\
\Leftrightarrow\; & E[\left((0,1)\left(X'X\right)^{-1}X'u(0,1)'\right)^2\left(\frac{(1,0)'\widehat{u}'\widehat{u}(1,0)}{T-2} - v_{11}^2\right)|X] \stackrel{!}{=} 0 \\
\Leftrightarrow\; & E[\left((0,1)\left(X'X\right)^{-1}X'u(0,1)'\right)^2\left(\frac{(1,0)'u'u(1,0)}{T-2} - v_{11}^2\right)|X] \stackrel{!}{=} 0
\end{aligned}$$

The last equality holds if and only if

$$\begin{aligned}
& E[\left((0,1)\left(X'X\right)^{-1}X'u(0,1)'\right)^2\frac{(1,0)'u'Mu(1,0)}{T-2}|X] \\
&= E[\left((0,1)\left(X'X\right)^{-1}X'u(0,1)'\right)^2 v_{11}^2|X],
\end{aligned}$$

which is the case if and only if condition (63) holds, as $E[\frac{(1,0)'uM'Mu(1,0)}{tr(M'M)}] = v_{11}^2$.

H Simulation Evidence for the Econometric Approach

This appendix shows that our two-stage estimation approach recovers the second-stage coefficient of interest in simulated data. In particular, taking into account the observed price adjustment frequency, the distribution of estimated product-specific trends, and the short sample features of the data, we obtain at most a downward bias in the estimated second-stage coefficient.

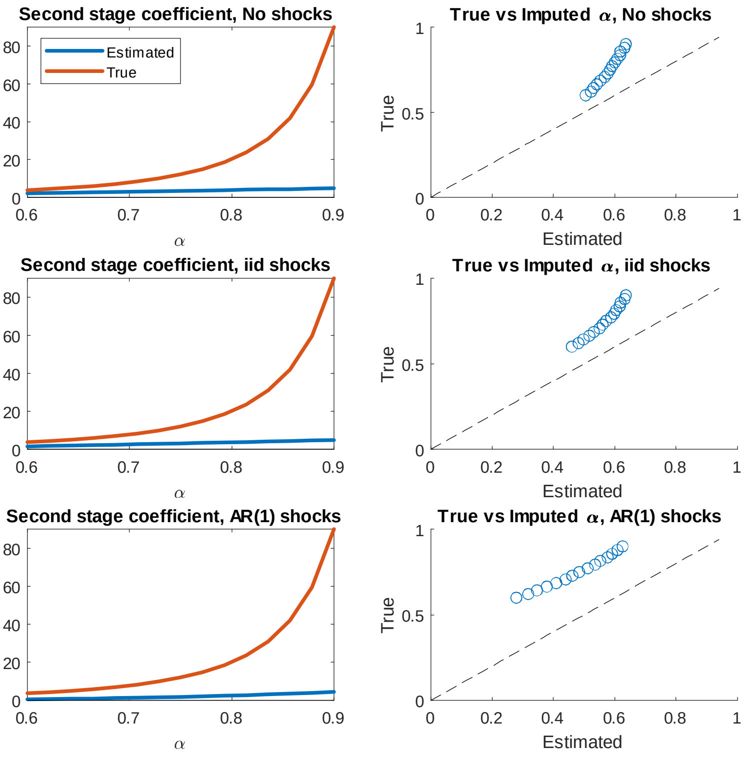

Figure 16: Simulation results, baseline

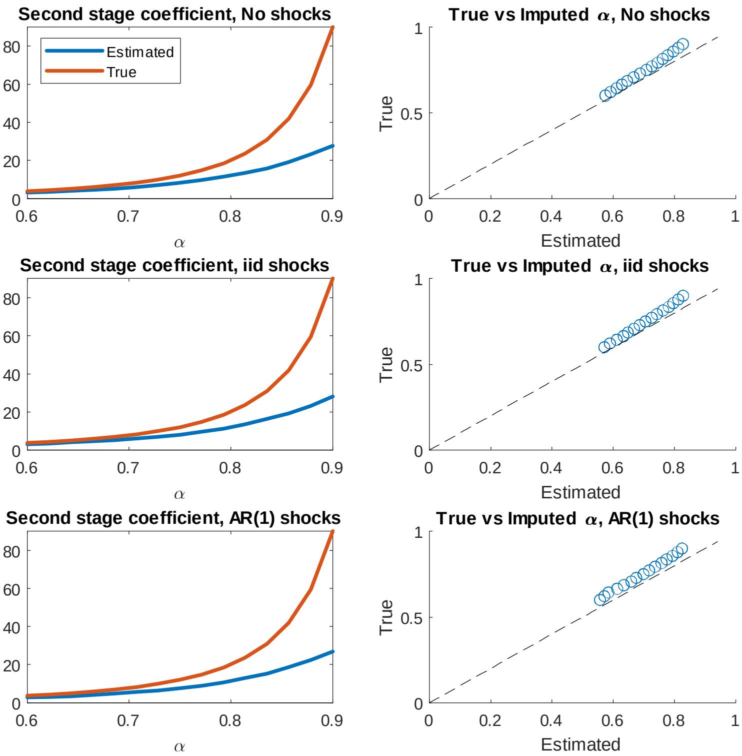

Figure 17: Simulation results, at least 24 observations per product

We simulate a Calvo model and set the non-adjustment rate for prices α equal to 0.89, which is the average (across-item) value in the data. Each simulated product draws a flexible-price trend from the normal distribution. The standard deviation of this distribution is calibrated so that the standard deviation of the estimated trends in simulated price time series matches the average (across-item) standard deviation of estimated trends in the data. Simulated products are sampled at random times for a random number of periods drawn from the average (across-item) distribution of observed product lengths. In a first step, we set idiosyncratic shocks to zero, then we vary the idiosyncratic shock process and the Calvo parameter to see how these affect the second-stage coefficient estimates.

The left column in figure 16 shows the mean of estimated second stage coefficients across simulations in blue and the true (theory-implied) coefficients in red. The right column shows the scatter plots of rates of price non-adjustment, imputed from average estimated coefficients (y-axis) and the true values (x-axis) – the simulation analogues of Figure 6 in the main text. The top panel considers a setting without idiosyncratic shocks, the middle panel one with iid normal shocks, and the bottom panel one with AR(1) shocks with normal iid innovations.[65] All graphs show substantial downward bias in estimated coefficients. The bias increases as we add idiosyncratic shocks and make them persistent. In addition, the bias is increasing in the degree of price rigidity. Importantly, none of the simulations suggests a possibility of upward bias in our estimates, and the relation between imputed and true α-s resembles the one we obtain in the data, see figure 6.

Figure 17 repeats the simulation analysis in figure 16 using only products with at least 24 observations. This is a robustness check that we also perform in the data, see table 5.2. Focusing on these longer products dramatically reduces the downward bias, independent of the assumed idiosyncratic shock process. This is in line with our empirical findings, where we also obtain larger estimates of the second stage coefficient as we increase the threshold for the minimal number of observations per product.

Overall, the simulations show that our econometric approach recovers the second-stage coefficient of interest, albeit possibly with a substantial downward bias when including all products including those with shorter lengths.

I Testing for Heterogeneity in Relative Price Trends

This appendix uses a bootstrapping procedure to show that there is significant evidence for heterogeneity in suboptimal inflation ($\ln \Pi_z/\Pi^*_{jz}$) across products j within expenditure

[65] The standard deviation of these innovations is 10 times larger than the standard deviation of the distribution from which products draw flexible-price trends. We set the AR(1) coefficient for idiosyncratic shock to 0.8.

items z. The main challenge for bootstrapping is the fact that residuals from the first-stage regression (22) are peculiar: between price adjustment periods residuals drift at a constant rate and they provide new information only in price adjustment periods. Simply drawing from the set of residuals would ignore this feature, destroy the sticky nature of observed prices, and thereby strongly confound results. We propose below a bootstrapping procedure that takes infrequent price adjustment into account and that reproduces the main features of the data under the null hypothesis of no trend heterogeneity. We then show that the data contains strong evidence against this null hypothesis.

I.1 Estimation

We start by estimating the first-stage regression under the null of no trend heterogeneity. We then impute the data generating process for residuals taking into account the stickiness of prices and potential autocorrelation of idiosyncratic shocks. The estimation is performed item-by-item and all estimated objects are item-specific. We drop the item index z below to simplify notation. The estimation uses all products that have at least 3 observations, exit in the sample, have at least one price change, and satisfy our second-stage truncation criteria. We then perform the following steps for all products j in a given item:

1. First, we estimate the common nominal price trend by pooling all products together and estimating:

$$\ln P_{jt} = \ln a_j + \ln b \cdot s_{jt} + u_{jt} \tag{67}$$

where $\ln P_{jt}$ is the (log) nominal price, s_{jt} is product's age in the sample, $\ln a_j$ is the product-specific intercept and $\ln b$ is the item-specific suboptimal inflation. We work with nominal prices since the slope coefficient in (67) directly identifies the common (item-level) suboptimal inflation, under the null of no trend heterogeneity.

2. We consider residuals u_{jt} at adjustment times $t \in \{\tau_1, \tau_2, \dots, \tau_{N_j}\}$. These residuals depend on the realization of idiosyncratic shocks and a constant frontloading component from suboptimal inflation. Any heterogeneity in these residuals we attribute to heterogeneity in realized shocks because the frontloading component is common across products under the null. We collect residuals $U = \{u_1, u_2, \dots, u_M\}$, pooling all adjustment-time residuals across all products. We create $\tilde{M}$ bins for this collection $\tilde{U} = \{\tilde{u}_1, \tilde{u}_2, \dots, \tilde{u}_{\tilde{M}}\}$, with $\tilde{M} \leq M$.

3. For each $\tilde{u} \in \tilde{U}$ we estimate the adjustment distribution function $g(\tilde{u}, \tilde{u}', t) : \tilde{U} \times (\tilde{U} \cup E) \times \mathbb{N} \to [0, 1]$ that assigns probabilities of adjusting to bin $\tilde{u}'$ after t periods conditional

on starting in bin $\tilde{u}$. The new bin $\tilde{u}'$ is one of the bins in $\tilde{U}$ or the product 'exit' bin E.

4. For the first price adjustment, we use residuals u_{jt} at the time of product entry and similarly construct bins $\tilde{U}^0$ and adjustment distribution function $g^0(\tilde{u}, \tilde{u}', t) : \tilde{U}^0 \times \tilde{U} \times \mathbb{N} \to [0, 1]$.[66] This additional step is required since products do not necessarily enter the sample at adjustment times and therefore may have a different distribution of residuals at the time of entry.

I.2 Simulation

The next step simulates nominal price time series for products (item-by-item) under the null hypothesis of a common relative price trend, bootstrapping the residuals and price adjustment times using the item-specific $g^0(\cdot)$ and $g(\cdot)$ functions. For each item we perform 5000 bootstrap repetitions. In each bootstrap repetition, we simulate the same number of products as we use for estimation. The maximum simulated product length is capped at the maximum product length observed for a given item in the data (denoted here by L).

1. For the initial prices we draw u_{j0} from the empirical distribution of initial residuals observed for the time of product entry. Without loss of generality, we assign zero intercepts for all products, as intercepts do not affect the slope estimates.

2. For each simulated product, we draw the lengths of the first price spell and the first adjustment bins from $g^0(\cdot)$, and subsequently from $g(\cdot)$, until either the product exits by drawing $\tilde{u}' = E$, or its lifespan exceeds L. If we draw bin $\tilde{u}'$, the residual assigned in the simulation is a randomly drawn residual from that bin. Together with the common slope $\ln b$ from (67) this gives us the sequence of reset prices for the product. Between adjustment periods we then assign the last reset price.

I.3 Comparison of Simulated and Actual Moments

An accurate bootstrap procedure should reproduce the key moments affecting the estimation of slopes. We consider the mean product lengths, the standard deviation of product length, the mean price adjustment size and the mean price adjustment frequency in each considered item. Specifically, for each bootstrapped sample of an item, we compute the ratio of simulated mean product length over the actual mean product length in the item. We compute corresponding ratios for the standard deviation of product length, the stan-

[66] Exiting before adjusting the price is ruled out since we only consider products with at least one price adjustment.

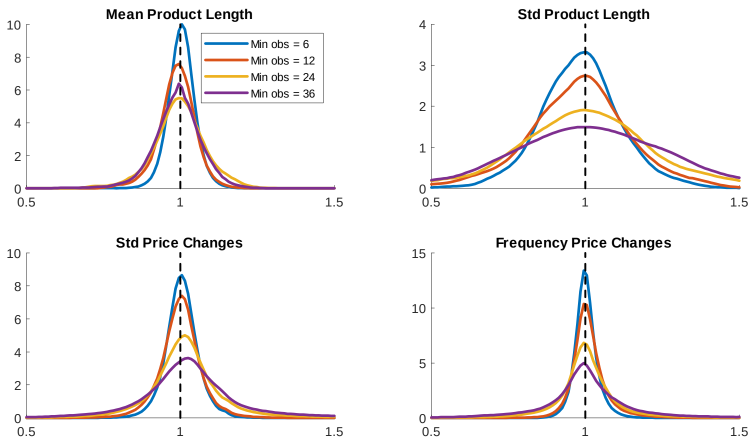

Figure 18: Bootstrapped moments relative to data moments

dard deviation of price changes and the mean frequency of price changes. Figure 18 shows the distributions of these ratios across all bootstraps and items, using different thresholds for the minimal number of price observations per product. The bootstrapping procedure matches the data moments well, even for products with longer horizons, despite the fact that neither the estimation nor the simulation procedures conditioned on product age or the number of observations.

I.4 Bootstrapped Critical Values

Finally, we run our first-stage regression on bootstrapped data, estimate product-specific slope coefficients, and compute t-statistics for null hypothesis of a common slope. From the distributions of bootstrapped t-statistics we obtain item-specific critical values for the t-statistic under the null of no trend heterogeneity. We then compute the share of products in the actual data with t-statistics falling outside the critical values. Figure 19 shows the distributions of these shares across items for 5%/95% and 10%/90% critical values, for different thresholds of minimal number of price observations per product.[67] The dotted vertical lines indicate the corresponding shares one should expect under the null of no trend heterogeneity – the level of significance (0.1 and 0.2, respectively).

All distributions are visibly shifted to the right of the level of statistical significance, providing strong evidence for the presence of trend heterogeneity in the data. Furthermore, the shift is stronger for products with more observations. For instance, the average share of products outside the 5%/95% confidence interval increases from 0.18 to 0.29 as the minimum number of price observations goes up from 6 to 36. Considering long products

[67]The bootstrapped critical values are computed separately for each considered minimum number of price observations.

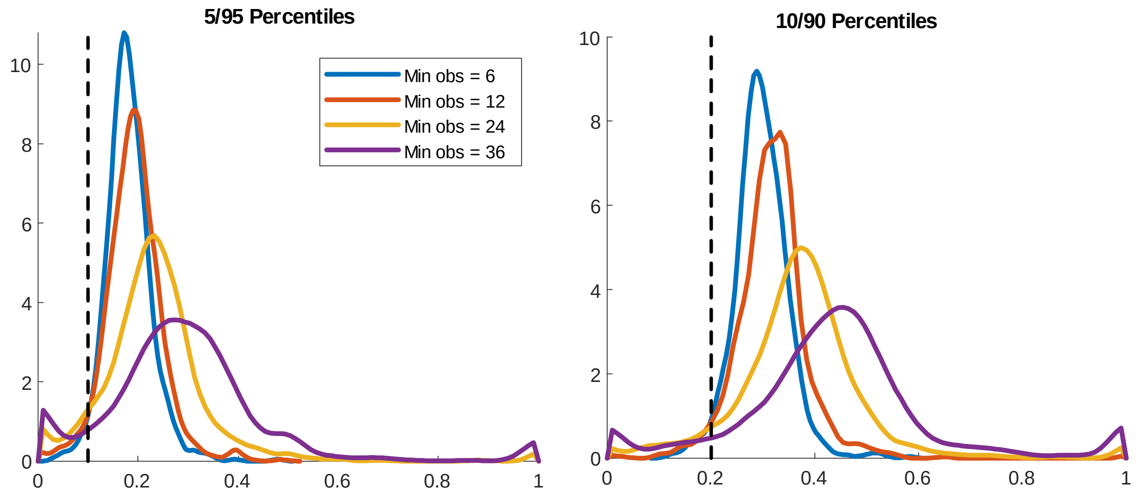

Figure 19: Share of t-statistics outside bootstrapped confidence intervals

thus strengthens the evidence for trend heterogeneity. Naturally, it is easier to detect trend heterogeneity among products with longer life spans, since their price paths are driven to a larger extend by trends and to a smaller extent by idiosyncratic shocks.

J Details of the Within-Product Regression Approach

The within product regression (25) takes the form

$$Y = c_z \cdot X \tag{68}$$

where Y is a $N \times 1$ vector of consisting of $Var_1(u_{jz}) - Var_2(u_{jz})$ for $j = 1, ..., N$, X a vector consisting of $\left(\ln \Pi_{jz1} - \ln \Pi^*_{jz}\right)^2 - \left(\ln \Pi_{jz2} - \ln \Pi^*_{jz}\right)^2$ for $j = 1, .., N$ and c_z is a scalar. The true relationship between Y and X is given by

$$Y = CX + \varepsilon,$$

where Y and X are random variables and

$$C = \begin{pmatrix} c_{1z} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & c_{Nz} \end{pmatrix}$$

is a diagonal coefficient matrix of random coefficients satisfying the conditional mean independence assumption $E[C|X] = E[C] = c \cdot I_{N\times N}$, with the scalar c denoting the expected value of the true coefficient. The residual vector ε a $N \times 1$ vector of (higher-order approximation) residuals satisfying $E[\varepsilon|X] = 0$. The OLS estimate of c_z in equation (68) is given by

$$\widehat{c}_z = (X'X)^{-1} X'Y$$

and under the stated assumptions its expectation satisfies

$$\begin{aligned}
E[\widehat{c}_z] &= E[(X'X)^{-1} X'Y] \\
&= E[E[(X'X)^{-1} X' (CX + \varepsilon) |X]] \\
&= E[(X'X)^{-1} X' \underbrace{E[C|X]}_{=cI} X] + (X'X)^{-1} X' \underbrace{E[\varepsilon|X]}_{=0}] \\
&= E[(X'X)^{-1} X'X]c \\
&= c,
\end{aligned}$$

as claimed in the main text.

J.1 Within-Product Approach with Menu Cost Frictions

We consider here the case with menu cost frictions, for which similar arguments apply as with Calvo frictions. Taking differences across the first and second half of product life using equation (19), one obtains (up to a second-order approximation) again equation (25), but with the regression coefficient now given by $c_z = E[1/\left(\lambda_{iz}^X\right)^2]$, where λ_{iz}^X is the switching intensity in the i-th state of the idiosyncratic shock process. The regression coefficient is now independent of the menu-cost parameter κ, so that the estimation approach (25) remains valid in a menu-cost setting in the presence of product-specific menu-costs.[68] If the expected switching intensities $E[1/\left(\lambda_{iz}^X\right)^2(j)]$ also differ across products j within the *same* item, but display conditional-mean independence from the regressor in equation (25), then OLS estimation of equation (25) again recovers the average coefficient

$$E\left[\widehat{c}_z\right] = E\left[\frac{1}{\left(\lambda_{iz}^X(j)\right)^2}\right].$$

As with Calvo frictions, one can thus test whether suboptimal inflation distorts relative prices *without* having to assume that products have identical menu costs and identical processes governing idiosyncratic shocks. And as with Calvo frictions, the test requires checking whether c_z in equation (25) is positive.

J.2 Details of the Across-Item Estimation Approach

Estimation of equation (26) in the main text is based on the following result, which allows for item-specific price stickiness (or menu costs) and item-specific idiosyncratic shock processes:

[68] Heterogeneity in adjustment costs has only fourth order effects on the variance of first-stage residuals. This is also true in the baseline approach with menu cost frictions.

Proposition 12 *Suppose all products j within an expenditure item z have the same optimal inflation rate $\Pi_{jz}^* = \Pi_z^*$. Let Π_z denote the actual inflation rate in item z and $Var(u)$ the variance of the first-stage residuals in item z (obtained under the assumption of a common optimal inflation rate). Consider the second-stage regression equation*

$$Var(u_z) = v_0 + c_0(\ln \Pi_z - \ln \Pi_z^*)^2 \tag{69}$$

and let $\widehat{c}_0$ denote the OLS estimate of c_0 and suppose there is no measurement error in $(\ln \Pi_z - \ln \Pi_z^*)^2$. For the case with Calvo frictions, the OLS estimate recovers the average second stage coefficient, i.e.,

$$E[\widehat{c}_0] = E\left[\frac{\alpha_z}{(1-\alpha_z)^2}\right],$$

whenever the $\frac{\alpha_z}{(1-\alpha_z)^2}$ and the $v_z = Var\left((1-\alpha_z)E_t \sum_{i=0}^{\infty} \alpha_z^i \ln x_{jzt+i}\right)$ are random variables with identical means for all z, and with conditional means that do not depend on $((\ln \Pi_1 - \ln \Pi_1^*)^2, ..., (\ln \Pi_Z - \ln \Pi_Z^*)^2)$. Similarly, with menu cost frictions, we have

$$E[\widehat{c}_0] = E\left[1/\left(\lambda_{iz}^X\right)^2\right],$$

whenever $E\left[1/\left(\lambda_{iz}^X\right)^2\right]$ and $Var(\ln x_z)$ are random variables with identical means for all z and conditional means that do not depend on $((\ln \Pi_1 - \ln \Pi_1^*)^2, ..., (\ln \Pi_Z - \ln \Pi_Z^*)^2)$.

Proof: Equation (69) is of the form

$$Y = X \cdot \begin{pmatrix} v_0 \\ c_0 \end{pmatrix}, \tag{70}$$

where Y is a $Z \times 1$ vector consisting of the variance of first-stage residuals $Var(u_z)$ for all items $z = 1, ..., Z$ and X a $Z \times 2$ vector of regressors containing the intercept and the second-stage regressor

$$X = \left(1_{Z\times 1}, \begin{pmatrix} (\ln \Pi_1 - \ln \Pi_1^*)^2 \\ \vdots \\ (\ln \Pi_Z - \ln \Pi_Z^*)^2 \end{pmatrix}\right).$$

The intercept v_0 and the second-stage coefficient of interest c_0 are scalars.

The true relationship between Y and X, however, is given by

$$Y = v + C \begin{pmatrix} (\ln \Pi_1 - \ln \Pi_1^*)^2 \\ \vdots \\ (\ln \Pi_Z - \ln \Pi_Z^*)^2 \end{pmatrix}, \tag{71}$$

where v is $Z \times 1$ vector containing the item-specific intercepts

$$v = \begin{pmatrix} v_1 \\ \vdots \\ v_Z \end{pmatrix}$$

and

$$C = \begin{pmatrix} c_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & c_Z \end{pmatrix}$$

is a diagonal coefficient matrix containing the item-specific second-stage coefficients for the second column of X. (The precise expression for these coefficients depends on the considered price setting friction.) Given our assumptions about conditional mean independence, we have

$$E[v|X] = E[v] = \overline{v} \cdot 1_{Z\times 1}$$
$$E[C|X] = E[C] = \overline{c} \cdot I_{Z\times Z},$$

where the scalars $\overline{v}$ and $\overline{c}$ denote the true expected value of the intercept and the second-stage coefficients, respectively. (The true expectations of these coefficients depend also on the considered price setting friction.) The residual vector ε is a $Z \times 1$ vector of (higher-order approximation) residuals satisfying $E[\varepsilon|X] = 0$.

We now show that under the stated assumptions, OLS estimates of $(v_0, c_0)'$ in equation (70) recover the expected value $(\overline{v}, \overline{c})$ of the coefficients:

$$
\begin{aligned}
E\left[\begin{pmatrix} \widehat{v}_0 \\ \widehat{c}_0 \end{pmatrix}\right] &= E[(X'X)^{-1} X'Y] \\
&= E[E[(X'X)^{-1} X'Y|X]] \\
&= E[E[(X'X)^{-1} X' \left(v + C \begin{pmatrix} (\ln \Pi_1 - \ln \Pi_1^*)^2 \\ \vdots \\ (\ln \Pi_Z - \ln \Pi_Z^*)^2 \end{pmatrix} \right) |X]] \\
&= E[(X'X)^{-1} X' \left(\underbrace{E[v|X]}_{=\overline{v} \cdot 1_{Z\times 1}} + \underbrace{E[C|X]}_{=\overline{c} I_{Z\times Z}} \begin{pmatrix} (\ln \Pi_1 - \ln \Pi_1^*)^2 \\ \vdots \\ (\ln \Pi_Z - \ln \Pi_Z^*)^2 \end{pmatrix} \right)] \\
&= E[(X'X)^{-1} X' \left(\overline{v} \cdot 1_{Z\times 1} + \overline{c} \begin{pmatrix} (\ln \Pi_1 - \ln \Pi_1^*)^2 \\ \vdots \\ (\ln \Pi_Z - \ln \Pi_Z^*)^2 \end{pmatrix} \right)] \\
&= E[(X'X)^{-1} X'X] \begin{pmatrix} \overline{v} \\ \overline{c} \end{pmatrix} \\
&= \begin{pmatrix} \overline{v} \\ \overline{c} \end{pmatrix}
\end{aligned}
$$

as claimed.

K Cross Sectional Distribution of Product-Specific Optimal Inflation Rates over Time

Figure 20 depicts the cross-sectional distribution of product-specific optimal inflation rates Π_{jz}^* across all products and all items in the first and last five years in of our sample (1996-2000 and 2012- 2016). It shows that this distribution is remarkably stable over time.

L Proof of Proposition 3

From equation (27) we get

$$
\begin{aligned}
Var^j\left(\ln p_{jzt}\right) = &\, Var^j\left(\ln p_{jz}^* - \ln \Pi_{jz}^* \cdot t\right) + Var^j\left(u_{jzt}\right) \\
&+ Cov^j(\ln p_{jz}^*, u_{jzt}) \\
&- t \cdot Cov^j(\ln \Pi_{jz}^*, u_{jzt}).
\end{aligned}
$$

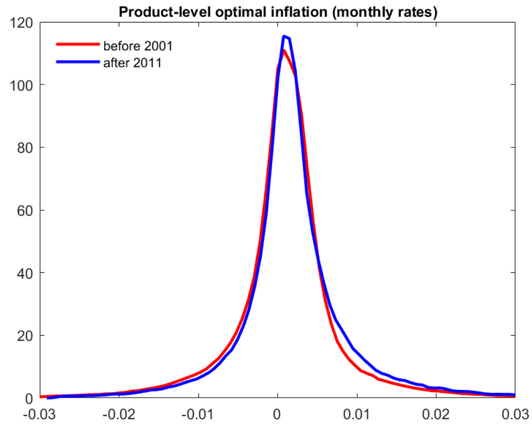

Figure 20: Distribution of product-specific optimal inflation rates Π^*_{jz} in 1996-2000 versus 2012-2016 (monthly rates, unweighted)

We next show that $Cov^j(\ln p^*_{jz}, u_{jzt}) = Cov^j(\ln \Pi^*_{jz}, u_{jzt}) = 0$:

$$\begin{aligned} Cov^j(\ln p^*_{jz}, u_{jzt}) &= E^j[\ln p^*_{jz} u_{jzt}] - E^j[\ln p^*_{jz}]\underbrace{E^j[u_{jzt}]}_{=0} \\ &= E^j[E^j[\ln p^*_{jz} u_{jzt} | p^*_{jz}]] \\ &= E^j[\ln p^*_{jz} \underbrace{E^j[u_{jzt} | p^*_{jz}]}_{=0}] \\ &= 0. \end{aligned}$$

Similarly:

$$\begin{aligned} Cov^j(\ln \Pi^*_{jz}, u_{jzt}) &= E^j[\ln \Pi^*_{jz} u_{jzt}] - E^j[\ln \Pi^*_{jz}]\underbrace{E^j[u_{jzt}]}_{=0} \\ &= E^j[E^j[\ln \Pi^*_{jz} u_{jzt} | \Pi^*_{jz}]] \\ &= E^j[\ln \Pi^*_{jz} \underbrace{E^j[u_{jzt} | \Pi^*_{jz}]}_{=0}] \\ &= 0. \end{aligned}$$

It thus only remains to compute the cross-sectional variance of residuals, $Var^j(u_{jzt})$. These residuals are described by a mixture distribution in which one first draws the relative price trend $\Pi_z^{*(i)}$ with probability m_{zi}. Subsequently, we draw corresponding residuals u_{jzt}. Since the residuals are independent across j, the cross-variance of residuals for any given $\Pi_z^{*(i)}$ is equal to their variance over time, as given in equation (28). Therefore, the variance of the mixture distribution is given by equation (30).