

Inflation Distorts Relative Prices: Theory and Evidence*

Klaus Adam (University of Mannheim & CEPR)

Andrey Alexandrov (Princeton University)

Henning Weber (Deutsche Bundesbank)

October 4, 2023

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. In the cross-section of products, the variance of price distortions comoves positively with aggregate inflation over time. In contrast, overall cross-sectional dispersion fails to comove with inflation over time. We show that it is predominantly driven by movements in the dispersion of flexible prices.

JEL Class. No.: E31, E58

*We thank our discussants Hassan Afrouzi, Viacheslav Sheremirov and Alex Wolman, as well as Markus Brunnermeier, Jordi Galí, Loretta Mester, Krzysztof Pytka, Frank Schorfheide, Frank Smets and Carsten Trenkler for useful input and discussions. Klaus Adam gratefully 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 exist no structural empirical evidence validating the notion that inflation has a distorting impact on relative prices.

This gap forms the focal point of our investigation. We first derive a novel theory-consistent empirical approach that allows estimating the marginal effect of inflation on relative price distortions. Subsequently, we apply this methodology to the micro price data underpinning the U.K. Consumer Price Index. We document that inflation is associated - at the level of individual products - with economically significant amounts of price distortions, in line with what sticky price theories predict. Furthermore, in the cross-section of products, price distortions turn out to covary positively with aggregate inflation over time.

Documenting the relationship between inflation and relative price distortions proved challenging and the present paper makes progress on a number of fronts.

First, it is challenging to recover inflation-induced distortions in relative prices from actual price observations. To see why, let p_{jt} denote the relative price *actually* charged for product j in period t , and p_{jt}^* the corresponding *flexible* relative price.² The price gap gap_{jt} due to price setting frictions is then given by

$$\ln p_{jt} = \ln p_{jt}^* + \ln gap_{jt}, \quad (1)$$

and price distortions for product j are conveniently summarized by the variance of the product-specific gaps over time:³

$$dist_j = Var(\ln gap_{jt}).$$

Monetary models postulate that the price distortions ($dist_j$) depend on inflation, but empirically documenting this relationship is challenging: while the actual relative price in equation (1) can be observed, the flexible relative price is unobserved, so that the price gap remains also unob-

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

²The flexible relative price is the price that would be charged for product j in the absence of price setting frictions. It may itself be distorted, e.g., due to market power.

³We discuss the variance of price gaps in the cross-section of products below.

served.⁴ In fact, we formally show in the paper that the flexible relative price and thus the price gap *cannot* be identified from actual relative prices, whenever the flexible price contains a stationary stochastic component.

Given this difficulty, the previous literature does not attempt to identify how price gaps depend on inflation (Wulfsberg (2016) and Nakamura, Steinsson, Sun and Villar (2018)), instead highlights the difficulties associated with empirically recovering price gaps.⁵

An important contribution of the present paper is to show that the *marginal* effect of inflation on price distortions, i.e., how $dist_j$ varies with inflation, can be identified from observed actual prices, even though the *level* of the price gaps (gap_{jt}) are not identified. This is feasible because the variance of the flexible relative price p_{jt}^* in equation (1) is *independent* of inflation, so that the variance of the actual relative price $Var(\ln p_{jt})$ is informative about the variance of price gaps, i.e., price distortions.

We show that time and state-dependent pricing models make identical predictions (up to a second-order approximation) about how the marginal effect of inflation on price distortions can be estimated from observed actual prices: in a first step, one computes residual price variation around the life-cycle trend of a product’s actual relative price time series. In a second step, one relates this residual variation - in the cross-section of products - to a measure of inflation.⁶ This structural approach is valid without restrictions on the behavior of the cross-sectional distribution of flexible prices over time.

A second challenge with identifying the relationship between inflation and relative price distortions is that sticky price theory implies that a marginally higher inflation rate can either *increase* or *decrease* relative price distortions. The direction of the effect depends on whether the observed inflation rate, $\ln \Pi$, lies above or below the optimal inflation rate, $\ln \Pi_j^*$, for product j .⁷ Existing work tends to ignore this issue and often assumes that the optimal inflation rate is zero. Yet, the optimal inflation rate differs from zero for most products and has been found to vary systematically in the cross-section of products (Adam and Weber

⁴In rare cases, additional information about marginal costs and the desired markup 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.

⁵See section IV.A in Nakamura et al. (2018).

⁶More precisely, it needs to be related to a measure of suboptimal inflation, as we explain in the next paragraph.

⁷If the optimal level of inflation lies above (below) the observed level of inflation, then marginally higher rates of inflation decrease (increase) price distortions, according to sticky price models.

(2022), Adam, Gautier, Santoro and Weber (2022)).

To address this issue, the present paper considers price distortions generated by *suboptimal* inflation ($\ln \Pi - \ln \Pi_j^*$), i.e., by the difference between observed and the product-specific *optimal* level of inflation. We thus obtain estimates of the marginal effect of *suboptimal* inflation on relative price distortions.

A third challenge this paper addresses is that it is generally difficult to establish a *causal* relationship between inflation and price distortions by 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, so that movements in inflation are often hard to distinguish from a slow-moving time trend.⁸ As a result, trends in price dispersion over time might either reflect the trend in inflation or others trends which operate concurrently but are unrelated to inflation, e.g., a secular change in the variety of products over time.

Our empirical approach overcomes this identification issue by exploiting cross-sectional variation in the product-specific optimal inflation rate $\ln \Pi_j^*$ during a period in which aggregate inflation Π was relatively stable in the U.K. economy. Variation in the optimal rate Π_j^* in the cross-section of products j is driven by product-specific fundamentals, such as the different rates of productivity progress or a different evolution of monopoly power over time. According to sticky-price theory, such cross-sectional heterogeneity in product-specific fundamentals is unrelated to inflation and thus induces quasi-exogenous variation in the gap between actual and optimal inflation, $(\ln \Pi - \ln \Pi_j^*)$, that we exploit to estimate the *causal* effects of suboptimal inflation on price distortions.

Addressing these three challenges, we show that suboptimal inflation causes relative price distortions in the U.K. economy. Specifically, we find that price distortions at the level of individual products depend on the squared value of *suboptimal* inflation, in line with the theoretical predictions of time or state-dependent pricing models. The squared value of suboptimal inflation has the sign predicted by theory and is statistically significant in 94% of the expenditure categories underlying the U.K. consumer price index. It also has surprisingly high explanatory power across individual products within the typical expenditure category. And in line with the underlying sticky price theories, the distortionary effects of suboptimal inflation are estimated to be stronger in the presence of stronger price rigidities.

Having established that suboptimal inflation gives rise to price dis-

⁸To avoid the possibility that our results are driven by energy price or other special shocks, we exclude the Covid and post-Covid period from our analysis.

tortions over time at the level of individual products, we turn consideration to the dispersion of prices in the cross-section of products. Cross-sectional price dispersion has been the point of departure for much of the earlier literature.

We show that cross-sectional dispersion of actual relative prices, $Var^j(\ln p_{jt})$, is strongly increasing over the sample period (1996 - 2016). Interestingly, cross-sectional price dispersion increases despite the fact that U.K. inflation displays no time trend and only moderate fluctuations.

To explain why overall cross-sectional price dispersion fails to covary with inflation in the data, let $p_{jt}^{*\det}$ denote the deterministic component of a products' flexible relative price and p_{jt}^{*stoch} its stochastic component. We then obtain from equation (1)

$$Var^j(\ln p_{jt}) = Var^j(\ln p_{jt}^{*\det}) + Var^j(\ln p_{jt}^{*stoch} + \ln gap_{jt}). \quad (2)$$

We show that the deterministic component of the flexible price $p_{jt}^{*\det}$ can be identified in the data and that its cross-sectional variance, $Var^j(\ln p_{jt}^{*\det})$, accounts for 99% of observed dispersion of actual prices $Var^j(\ln p_{jt})$. It also explains the overwhelming part of the observed increase in the cross-sectional dispersion of actual prices.

According to sticky price theory, the last term on the right-hand-side of equation (2), which accounts for about 1% of the dispersion of actual prices, should vary with inflation. In fact, we show that inflation-induced movements in the last term reflect *exclusively* inflation-induced movements in cross-sectional price distortion $Var^j(\ln gap_{jt})$.⁹ Specifically, the theory implies that $Var^j(\ln gap_{jt})$ should increase with inflation, provided the optimal inflation rate Π_j^* lies below observed inflation for most products.

This is what we find: $Var^j(\ln gap_{jt})$ comoves positively with inflation over time with a correlation equal to +0.67 that is statistically significant at the 1% level. Cross-sectional price distortions thus increase with inflation over the sample period. And we show that this positive comovement is predominantly driven by products with a low optimal inflation rate Π_j^* , as predicted by the underlying sticky price theories. We also compute an upper and lower bound on the contribution of inflation to cross-sectional price dispersion. Doing so, we find that the peak contribution of inflation to the cross-sectional standard deviation of price gaps, $Std^j(\ln gap_{jt})$, ranges between 3.8% and 5.0% over the sample period.

⁹Note that $Var^j(\ln p_{jt}^{*\det})$ is - by definition - independent of inflation. We show that the covariance $Cov^j(\ln p_{jt}^{*stoch}, \ln gap_{jt})$ also does not depend on inflation, according to sticky price theory.

Our finding that an increase in aggregate inflation leads to an increase in cross-sectional price distortions in the U.K. aligns well with key assumptions made in monetary models and should thus increase confidence in the economic relevance of key policy recommendations derived from these models, e.g., the desirability of targeting low and stable inflation rates. It also aligns well with recent findings in Ascari, Bonmolo and Haque (2022), who show 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 the infers price-induced misallocations arising from product specific mark-ups (Baqae, 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.

Section 2 illustrates the empirical approach developed in this paper using the simplest possible setup. Section 3 introduces the full theory and shows how sticky price models with time or state-dependent pricing frictions imply a regression approach that allows estimating the causal effect of suboptimal inflation on product-level price distortions. Section 4 introduces the U.K. micro price data and section 5 presents our baseline empirical results and a large number of robustness exercises. Section 6 presents a more involved estimation approach that allows to relax some of the identifying assumptions underlying the baseline approach. Section 7 discusses the decomposition of cross-sectional dispersion of actual prices and the comovement of cross-sectional price distortions with inflation over time. A conclusion briefly summarizes.

¹⁰Sara-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 inflation rates.

2 The Approach in a Nutshell

This section illustrates how one can empirically identify from micro price data the marginal contribution of suboptimal inflation on price distortions. The approach differs from the one pursued in Nakamura et al. (2018), who proposed considering the absolute size of price changes as a proxy for relative price distortions. We show that this proxy can sometimes provide misleading signals about the relationship between inflation and relative price distortions.

In our baseline approach, identification is achieved by considering a set of products for which (i) price stickiness and (ii) the shock process driving the idiosyncratic component of the flexible price is homogeneous across products.¹¹ One can then exploit variation in the optimal inflation rate across products to identify the marginal effect of inflation on price distortions. This holds true even if the actual inflation rate is constant over time.

To provide a simple example, suppose that idiosyncratic shocks are simply absent, so that the flexible relative price evolves deterministically, and that prices get adjusted in regular intervals every $N > 1$ periods (Taylor (1979)).¹² Consider product j , which is a physical object or service sold in a specific location.¹³ The flexible optimal relative price $p_{jt}^* = P_{jt}^*/P_t$ of product j is the price the firm would like to charge in the absence of any price setting frictions and evolves deterministically according to

$$\ln p_{jt}^* = \ln p_j^* - t \cdot \ln \Pi_j^*, \quad (3)$$

where p_j^* is a product-specific intercept and Π_j^* a product-specific time trend, capturing differences in marginal costs (or other factors) across products. Finally, suppose 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 the 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 no matter for tracking the desired 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

¹¹These assumptions can be relaxed further, as we show in section 6.

¹²These assumptions are special because they allow identifying the flexible price from micro price data, which fails to be true under the more general assumptions considered later on, but useful for illustrating the approach.

¹³Objects or services that are sold in different locations are treated as different products. The same holds true when an existing product gets substituted by a new product.

only occasionally, so that suboptimal inflation leads to deviations of the relative price from the flexible relative price.

Figure 1 illustrates the situation. 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 $\Pi_1^* < \Pi_2^* < \Pi_3^*$. Assuming that actual inflation Π is equal to Π_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 too low, which means that the relative price falls insufficiently during non-adjustment periods. To compensate for this effect, it becomes optimal to choose a relative price that is lower than the flexible price in adjustment periods, to reduce the gap between the sticky and the flexible relative price over the lifetime of the sticky price. Suboptimally low inflation thus leads to a deviation of the sticky relative price from the flexible relative price. This deviation is even stronger for product $j = 3$, which has a higher optimal inflation rate and - in adjustment periods - a relative price that is even further below the flexible relative price. A larger gap between inflation and the optimal inflation rate 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_j between the sticky relative price around its time trend, i.e., the price distortion for product j , is a function of the square of suboptimal inflation:¹⁴

$$Var(gap_j) = c \cdot (\ln \Pi - \ln \Pi_j^*)^2 \quad (4)$$

where

$$c = \frac{N \cdot (N - 1) \cdot (N + 1)}{12} > 0$$

depends positively on the degree of price stickiness $N > 1$.

An important insight developed in this paper is the fact that the relationship between suboptimal inflation and price distortion in equation (4) 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 1, and (ii) price distortions, i.e., the gaps between the actual and the flexible price, are identified by the residuals of a regression of actual prices on a time trend, as illustrated in figure 1. Thus micro price data suffices to test whether price distortions vary with suboptimal inflation rates, i.e., whether $c > 0$, as predicted by sticky price theory.

While property (ii) fails to be true when the flexible price also depends on unobserved idiosyncratic shocks, we show in the next section

¹⁴See appendix A for a proof.

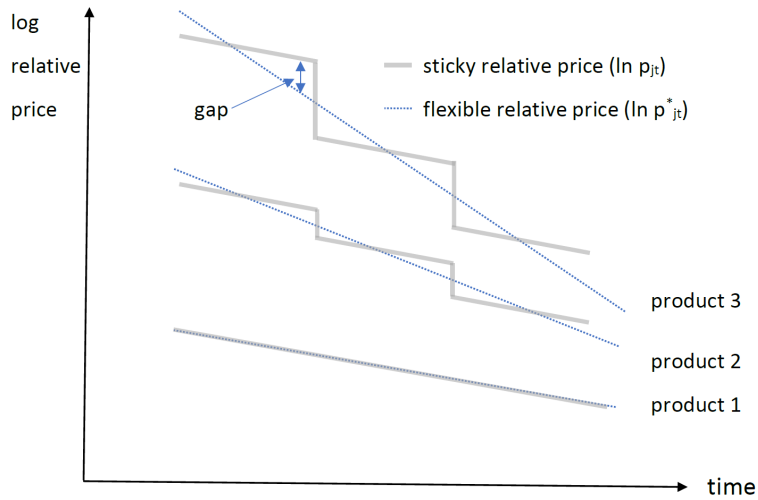


Figure 1: Relative price trends and price gaps

that the presence of such shocks only requires adding a constant to equation (4). This holds true even when considering more plausible pricing setting frictions, such as Calvo or menu-cost frictions.

Interestingly, using the absolute size of price changes as a measure of relative price distortion can lead to misleading conclusions about the relationship between relative price distortions and suboptimal inflation. The absolute size of price changes may respond to inflation in a setting where price distortions fail to do so and it may fail to respond to inflation in a setting where relative price distortions do indeed respond.

To see the first point, consider the example discussed above. The absolute size of log nominal price changes per unit of time is simply a function of price stickiness and suboptimal inflation and equal to $N \cdot |\Pi - \Pi_j^*|$. In the limit where prices become fully flexible ($N \rightarrow 1$), the absolute size of nominal price changes is thus given by $|\Pi - \Pi_j^*|$ 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 absent.¹⁵

This contrasts with the detrended residuals u proposed in figure 1:

¹⁵This 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^* .

in the limit with flexible prices, relative prices follow the dotted lines in the figure, so that the residuals u 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 $N \rightarrow 1$, the coefficient c in equation (4) converges to zero: one arrives at the correct conclusion that suboptimal inflation does *not* lead to relative price distortions.

For the case with sticky prices, the absolute size of price changes may actually fail to respond at all to changes in suboptimal inflation, even in a setting where price distortions do change with suboptimal inflation. Appendix presents an example with sticky prices and idiosyncratic shocks where this is the case and where the detrended residual variance, $Var(u_j)$, again accurately capture how relative price distortions vary with suboptimal inflation.

3 Inflation and Price Distortions: Theory

This section uses sticky price theory to derive a regression equation 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.

We will consider the price setting problem of a firm facing a demand structure that closely matches the implicit demand structure underlying the way how the U.K. Office of National Statistics (ONS) aggregates prices across products. In particular, aggregate consumption C_t is made up of Z different expenditure items (in the language of the ONS), where an expenditure item $z \in \{1, \dots, Z\}$ is a narrow product category, e.g., "Flatscreen TV, 30-inch display" or "CD-player, portable". Expenditure items 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 ,

$$C_{zt} = \left(\int_0^1 C_{jzt}^{\frac{\theta-1}{\theta}} dj \right)^{\frac{\theta}{\theta-1}}, \quad (5)$$

where C_{jzt} denotes the consumed physical units of product j in item z in period t and $\theta > 1$ the elasticity of substitution between products within the item. Aggregate consumption is given by

$$C_t = \prod_{z=1}^Z (C_{zt})^{\psi_z}, \quad (6)$$

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

$$C_{jzt} = \psi_z \left(\frac{P_{jzt}}{P_{zt}} \right)^{-\theta} \left(\frac{P_{zt}}{P_t} \right)^{-1} C_t, \quad (7)$$

where the item price level is defined as $P_{zt} = \left(\int_0^1 P_{jzt}^{1-\theta} dj \right)^{\frac{1}{1-\theta}}$ 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

$$Y_{jzt} = \frac{A_{zt}}{G_{jzt} X_{jzt}} L_{jzt}, \quad (8)$$

where for simplicity L_{jzt} denotes labor input and A_{zt} the level of productivity common to all producers of products in item z at time t .¹⁶ G_{jzt} is a product-specific factor capturing idiosyncratic productivity components that are deterministic from the perspective of the firm, while X_{jzt} is a stochastic idiosyncratic productivity component. In equilibrium, the quantity of products consumed C_{jzt} must be equal to the quantity produced Y_{jzt} .

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.¹⁷

3.1 Time-Dependent Price Setting Frictions

The price setting problem. Consider some product j which is a physical object or service sold in a specific location over time. Otherwise identical objects or services that are sold in different locations are treated as different products in our approach. The same holds true whenever an existing product gets substituted by a new product.

Let $p_{jzt} \equiv P_{jzt}/P_{zt}$ denote the relative price charged for product j , where P_{jzt} denotes the nominal product price and P_{zt} the price index for products in expenditure item z . Similarly, let p_{jzt}^* denote the flexible relative price, i.e., the price the firm would like to charge for product j in period t in the absence of price setting frictions. The flexible price

¹⁶The 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.

¹⁷Deriving analytic results for a unified setup with time *and* state-dependent adjustment frictions (Calvo plus) is analytically difficult, as both price setting frictions require imposing slightly different assumptions for the shock process X_{jzt} .

can differ from the (socially) efficient relative price.¹⁸ Given the demand structure introduced above, appendix D.1 derives the following second-order approximation to the nonlinear optimal price setting problem with Calvo price adjustment frictions:

$$\max_{\ln p_{jzt}} -E_t \sum_{i=0}^{\infty} (\alpha_z \beta)^i (\ln p_{jzt} - i \ln \Pi_z - \ln p_{jzt+i}^*)^2, \quad (9)$$

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 this item. 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 over time by inflation, as long as prices fail to adjust. Deviations of the firm's relative from its flexible optimal 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 (9) 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:

$$\ln p_{jzt}^* = \ln p_{jz}^* - t \cdot \ln \Pi_{jz}^* + \ln x_{jzt}. \quad (10)$$

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 unobserved location-specific effects such as difference 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 variable Π_{jz}^* in equation (10) 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 relative prices 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. It is well-known that the strength of these effects varies across products¹⁹ and

¹⁸This may be due the presence of product-specific monopoly mark-ups. 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.

¹⁹Adam and Weber (2022) document this for the U.K and Adam, Gautier, Santoro and Weber (2022) for France, Germany and Italy.

we will exploit the variation in Π_{jz}^* below to identify the distortionary effects of inflation. We consider a linear time trend in relative prices because the relative price dynamics of newly introduced products are well-approximated by a linear trend.²⁰ Yet, in our empirical analysis we shall also consider nonlinear time trends.

Finally, there is an idiosyncratic stochastic component $\ln x_{jzt}$ in equation (10), which captures idiosyncratic fluctuations induced by changes in productivity or service components at the product level. The absence of a common component in these shocks is justified on the grounds that the left-hand side of equation (10) features the log *relative* price, thus absorbs common components in the nominal price (at the level of a narrowly-defined expenditure category). The stochastic process governing these idiosyncratic components is assumed to be the same for all products within a narrowly-defined expenditure category and satisfies the following restriction:

Assumption 1: Idiosyncratic shocks $\ln x_{jzt}$ are stationary and Markov.

Assumption 1 effectively rules 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 C.²¹ We can thus normalize idiosyncratic shocks so that $E[\ln x_{jzt}] = 0$.

Note that the cross-sectional distribution of flexible prices in expenditure category z is allowed 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 Π_{jz}^* induces changes in the cross-sectional distribution of the flexible relative prices; (ii) as products exit and enter the market, newly entering products may have different product-specific intercepts p_{jz}^* and time trends Π_{jz}^* than exiting products. Since the parameters (p_{jz}^*, Π_{jz}^*) of newly incoming products are drawn from arbitrary time-varying distributions, the setup imposes no restrictions on the evolution of the cross-sectional distribution of flexible relative prices over time.

The optimal reset price. Considering the limit $\beta \rightarrow 1$, the optimal reset price $\ln p_{jzt}^{opt}$ solving problem (9) is given by²²

$$\ln p_{jzt}^{opt} = (\ln p_{jzt}^* - \ln x_{jzt}) + \left(\frac{\alpha_z}{1 - \alpha_z} \right) (\ln \Pi_z - \ln \Pi_{jz}^*) + f(x_{jzt}), \quad (11)$$

²⁰See figure A.XI in the November 2018 working paper version of Argente and Yeh (2022), which depicts the relative price dynamics of newly introduced products using scanner data.

²¹This finding does not depend on assuming Calvo frictions.

²²See appendix D.1 for a derivation.

where

$$f(x_{jzt}) \equiv (1 - \alpha_z) E_t \sum_{i=0}^{\infty} \alpha_z^i \ln x_{jzt+i}. \quad (12)$$

The first term on the r.h.s. of equation (11), $\ln p_{jzt}^* - \ln x_{jzt}$, captures the deterministic component of the flexible price (10). 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 (11) captures effects due to the presence of time-varying idiosyncratic components. Equation (12) shows that it is the expected 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 (11) depends on inflation. If actual 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 actual 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 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 (exceeds) of 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 (11) determines the optimal reset price in periods where prices adjust, the dynamics of the actual relative price for product j in expenditure category z are given by

$$\ln p_{jzt} = \xi_{jzt} (\ln p_{jzt-1} - \ln \Pi_z) + (1 - \xi_{jzt}) \ln p_{jzt}^{opt}, \quad (13)$$

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 (13) 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 (11). 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 effects of suboptimal inflation on price distortions.

The first-stage regression. The first step in estimating the marginal effect of inflation on price distortions consists of running OLS regressions of the form

$$\ln p_{jzt} = \ln a_{jz} - (\ln b_{jz}) \cdot t + u_{jzt}, \quad (14)$$

which regress the relative product price on a product-specific intercept and time trend. To simplify the exposition, we abstract from small sample issues and focus on population regressions.²³ Regression (14) is of interest for two reasons. First, the coefficient estimates deliver²⁴

$$\begin{aligned} \widehat{\ln a_{jz}} &\rightarrow \ln p_{jz}^* \\ \widehat{\ln b_{jz}} &\rightarrow \ln \Pi_{jz}^*, \end{aligned} \quad (15)$$

which shows that the regression allows recovering the deterministic components of the flexible relative price, i.e., the intercept term p_{jz}^* and the product-specific optimal inflation rate Π_{jz}^* . Since the actual relative price follows - in terms of its level and time trend - these deterministic dynamics, the effects of price distortions must be contained in the residuals of regression (14). In fact, these residuals are the second reason why regression (14) is of interest. They are asymptotically given by²⁵

$$u_{jzt} = \xi_{jzt}(u_{jzt-1} - (\ln \Pi_z - \ln \Pi_{jz}^*)) + (1 - \xi_{jzt})(f(x_{jzt}) + \frac{\alpha_z}{1 - \alpha_z}(\ln \Pi_z - \ln \Pi_{jz}^*)) \quad (16)$$

where $\xi_{jzt} = 0$ captures periods in which the price gets adjusted and $\xi_{jzt} = 1$ captures periods without adjustment, and where $f(x_{jzt})$ is defined in equation (12). We next discuss the properties of the regression residuals (16).

The level of price distortions is not identified. Due to price stickiness ($\alpha_z > 0$), the regression residuals u_{jzt} in (16) fail to be very informative about the idiosyncratic shocks, as previously emphasized by Nakamura, Steinsson, Sun and Villar (2018). The underlying intuition is straightforward: in periods where prices do not get adjusted, they reveal no new information about idiosyncratic shocks; and in periods, where prices get adjusted, their adjustment gives considerable weight

²³Small sample effects are discussed in detail in appendix F.

²⁴See appendix D.3 for a formal derivation.

²⁵See appendix D.4 for a derivation.

to expected future values of the idiosyncratic shock, particularly when prices are sticky, see equation (12).

Due to the influence of expected future shock values, the information that becomes available upon a price adjustment, i.e. the term $f(x_{jzt})$ defined in equation (12), fails to identify the underlying process of idiosyncratic shocks $\ln x_{jzt}$. Appendix B proves the following result:

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 processes for $\ln x_{jzt}$ with N states or an infinite number of different Markov processes with $M \geq N$ 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 expectations in equation (12). Since the process for $\ln x_{jzt}$ cannot be identified from observed prices, it is impossible to estimate ‘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 responds to (suboptimal) inflation.

It is worth emphasizing that the result in proposition 1 applies more generally to 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 $\ln y_{jzt}$ plus 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.

One way to deal with the identification problem 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 available.

An alternative approach to handle the identification problem is to impose additional identification 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. 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 hypothesis of a pure random walk in $\ln x_{jzt}$ is strongly rejected in our data, as we show in appendix C.

We now show that it is simply not necessary to identify the *level* of price distortions to estimate the marginal effects of suboptimal inflation on price distortions. We discuss this point in the next subsection.

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 (11) highlights that any non-zero gap $\ln \Pi_z - \ln \Pi_{jz}^*$ generates front-loading of prices upon price adjustment times, 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 increasing the variance of the regression u_{jzt} in the first-stage regression (16).

Therefore, the variance of first-stage residuals satisfies the following relationship:²⁶

Proposition 2 *The variance of the first-stage residual in equation (14) is given by*

$$\text{Var}(u_{jzt}) = v_z + c_z \cdot (\ln \Pi_z - \ln \Pi_{jz}^*)^2, \quad (17)$$

where the intercept

$$v_z \equiv \text{Var} \left((1 - \alpha_z) E_t \sum_{i=0}^{\infty} \alpha_z^i \ln x_{jzt+i} \right) \quad (18)$$

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

$$c_z \equiv \frac{\alpha_z}{(1 - \alpha_z)^2}. \quad (19)$$

The intercept term v_z in equation (17) contains both efficient price components, e.g., the presence of idiosyncratic fundamental shocks, and price distortions that arise due to price stickiness, see equation (18). In particular, price stickiness causes the loading on the current idiosyncratic shocks to be too low relative to the flexible price case. Without additional information, it is impossible to further decompose to what extent v_z reflects efficient or inefficient forces, which is precisely the feature

²⁶See appendix D.4 for a derivation.

preventing identification of the level of price distortions from observed actual prices. The second term on the r.h.s. of equation (17) captures the effects of suboptimal inflation on price distortions. According to the theory, the coefficient c_z is an increasing function of the degree of the Calvo price stickiness parameter α_z .

Equation (17) is a second-stage regression equation and a key equation we shall exploit in the present paper. It uses the residual variance from the first-stage equation (14) 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 (15). Equation (17) 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 section 6).

Appendix F describes in detail the two stage estimation approach that allows estimating the coefficient c_z . It shows that the second-stage estimate for c_z is biased towards zero, due to the presence of first-stage estimation error. The second-stage estimate of c_z thus provide 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 *not* creating price distortions ($H_0 : c_z = 0$), this works against our main finding.

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

3.2 State-Dependent Price Setting Frictions

We now present a model with state-dependent pricing. To be able to get closed-form solutions, we consider a continuous-time setup and a slightly more restrictive process for the idiosyncratic shocks. Within this setup, we derive continuous-time analogue to proposition 2. The firm's objective (9) becomes:

$$\max_{\{\tau_{jzi}, \Delta \ln p_{jzi}\}_{i=1}^{\infty}} -E \left[\int_t^{\infty} e^{-\rho(s-t)} (\ln p_{jzt+s} - \ln p_{jzt+s}^*)^2 ds + \kappa_z \sum_{i=1}^{\infty} e^{-\rho(\tau_{jzi}-t)} \right] \quad (20)$$

The parameter $\rho > 0$ is the discount rate, τ_{jzi} 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 $\tau_{jzi} + s$ is given by $\ln p_{jz\tau_{jzi} + s} - s \ln \Pi_z$ between adjustment periods, reflecting relative price

erosion due to inflation.

The flexible relative price $\ln p_{jzt}^*$ follows a continuous-time analogue of (10) with an additional restriction on the idiosyncratic process $\ln x_{jzt}$, namely that it assumes values from a finite grid $\{\ln x_{1z}, \dots, \ln x_{Nz}\}$ and switches from grid point i to grid point j with Poisson intensity λ_{izj}^X .²⁷

Appendix E shows that under $\rho \rightarrow 0$ and for sufficiently small adjustment cost κ_z ,²⁸ the OLS regression (14) recovers the exact same coefficients as in the time-dependent model. Furthermore, the variance of residuals depends on product-specific suboptimal inflation:

$$\text{Var}(u_{jzt}) = \text{Var}(\ln x_z) + c_z^{MC} \cdot (\ln \Pi_z - \ln \Pi_{jz}^*)^2 + O((\ln \Pi_z / \Pi_{jz}^*)^4), \quad (21)$$

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_z^{MC} is now a function of the shock process parameters $\lambda_{iz}^X = \sum_{j \neq i} \lambda_{izj}^X$:

$$c_z^{MC} \equiv E \left[\frac{1}{(\lambda_{iz}^X)^2} \right].$$

If λ_{iz}^X is constant across states, then

$$c_z^{MC} = \frac{1}{\Lambda_z^2} \quad (22)$$

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 c_z^{MC} differs slightly from the one in the discrete time setup with Calvo friction, see equation (19), for which $\Lambda_z = 1 - \alpha_z$. This is so because multiple price adjustments can happen per unit of time under continuous time modeling. Notice also that 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 (21). This is the reason why equation (21) 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 objective), see equation (17).

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

²⁷The 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.

²⁸Note that we *do not* consider a limiting case $\kappa_z \rightarrow 0$, instead our result holds for all $\kappa_z \leq \bar{\kappa}$ for some $\bar{\kappa} > 0$.

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.²⁹ This is essential for our identification approach, which relies on cross-product variation in relative price trends.³⁰

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, etc.). Given the sample selection described further below, we consider 1071 different expenditure items and 15.5 million price observations over the sample period.

A product within an item is a sequence of price observations for a particular object or service sold in a particular store. Otherwise identical objects or services that are 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.

We then estimate the first-stage equation (14) for every product in the sample and estimate the second-stage equation (17) at the level of the expenditure item $z = 1, \dots, 1071$, considering all products j belonging to the item, i.e., we estimate

$$\widehat{Var}(u_{jzt}) = v_z + c_z \cdot (\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*)^2 + \varepsilon_{jz} \quad (23)$$

where $\widehat{Var}(u_{jzt})$ is the variance of first-stage residuals of product j in item z and $\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*$ the corresponding first-stage estimate of the gap between the item-level inflation rate and product-specific optimal infla-

²⁹See Adam and Weber (2022) who estimate the optimal aggregate inflation rate for the U.K. from relative price trends.

³⁰For these reasons, micro price data is more attractive for our analysis than supermarket scanner data, which covers fewer product categories and also contains many categories, e.g., food, for which relative price trends tend to be less strongly pronounced.

Total number of price quotes used	15.4 million			
	mean	median	min	max
Number of products per item	726	560	50	3,201
Number of price quotes per item	14,415	10,763	407	73,313

Table 1: Basic product and price statistics

tion.³¹ Estimation of equation (23) delivers 1071 estimates c_z , one for each expenditure item. We focus in our analysis on the item-level rather than on the aggregate level because doing so increases the chances that our two key identifying assumptions (identical degrees of price rigidity & identical stochastic processes driving idiosyncratic shocks) are satisfied. These assumptions will be relaxed in section 6, where we consider a more demanding estimation approach.

The data methodology follows the one used in Adam and Weber (2022), who provide further details. Starting from the raw micro price data, we delete products with duplicate price observations in a given month³² and also delete all price observations flagged by ONS 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 only include 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. This leads to a selection of 1093 expenditure items from the 1233 contained in the raw data. Furthermore, we only consider products with a minimum length of six price observations after eliminating sales prices from the sample.³³ We account for outliers by eliminating the 5% of products with the highest values for $\widehat{Var}(u_{jzt})$ and for $(\ln \widehat{\Pi}_z / \Pi_{jz}^*)^2$ within each expenditure item.

³¹See appendix F for details of the estimation approach, including arguments showing why two-stage estimation approach only biases the coefficient c_z towards zero, i.e., against finding a role for suboptimal inflation on inefficient price dispersion.

³²Duplicate price quotes can arise because the U.K. Office of National Statistics (ONS) does not disclose all available locational information underlying the data, so that in rare cases we cannot uniquely identify the product price.

³³We identify sales prices using the sales flag recorded by the ONS price collectors.

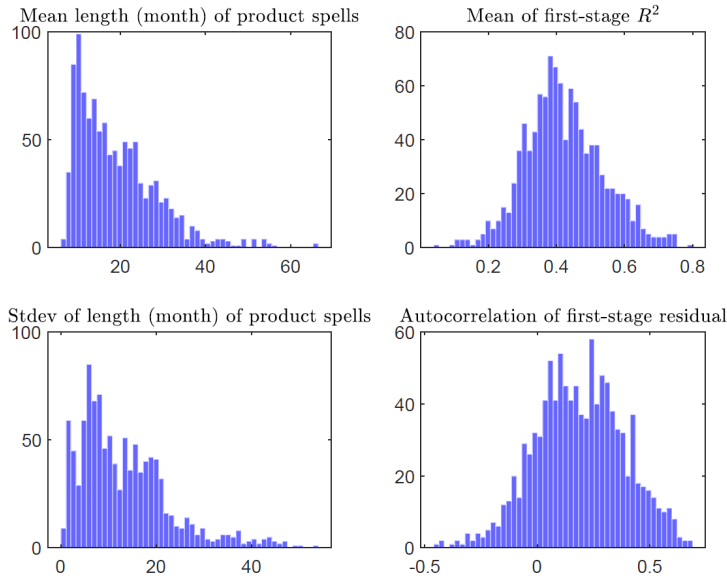


Figure 2: Descriptive statistics: first-stage regression

We then consider all expenditure items containing at least 50 products.³⁴ This leads us to the 1071 expenditure items that we use in our empirical analysis.³⁵ Table 1 reports basic statistics on the number of products and price observations per item. Given our approach, the average number of price observations per product is equal to 20 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 2 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 and 30 months, which is long enough to estimate an intercept and slope parameter in our first-stage.

³⁴We 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.

³⁵Not all these items are present throughout the sample period, as expenditure items get added and removed. For the average year, we have 503 expenditure items.

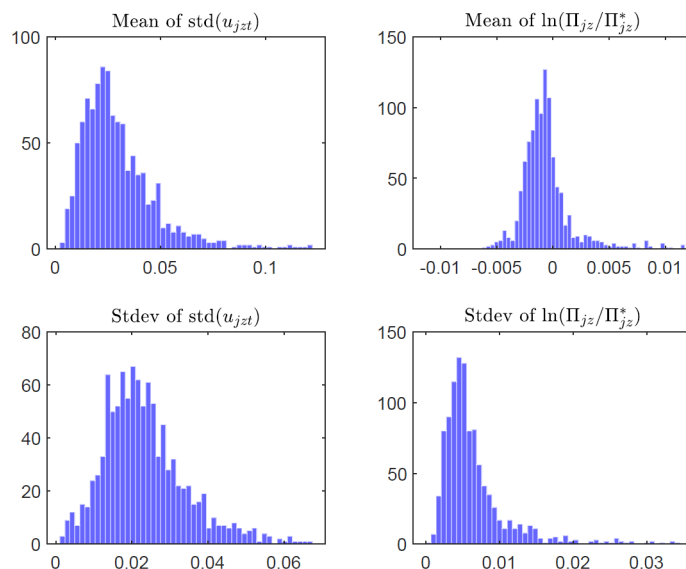


Figure 3: Descriptive statistics: second-stage regression

The bottom left panel in figure 2 highlights that there is a considerable amount of variation in the length of product lives within each item. We exploit this feature below to present estimates that are based on products whose price can be observed for at least 12 or 25 months (instead of 6 in our baseline) .

The top right panel in figure 2 reports the distribution of the mean R^2 values of the first-stage regression (61) across items. For most items, the intercept and time trend tend to capture on average between 30% and 50% 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 2 depicts the distribution of the mean autocorrelation of these residuals. The autocorrelation is significantly below one, showing that the assumption of a random walk is implausible given our data.³⁶

The top left panel of figure 3 reports the mean standard deviation of the regression residual across items.³⁷ For most items, the average standard deviation ranges between 2% and 4%. The standard deviation of the standard deviation of residuals is shown in the bottom left panel

³⁶See appendix C for formal tests of the random walk hypothesis, which are based on price observations from price adjustment periods.

³⁷We report moments of the non-squared variables entering the second-stage regression to increase readability of the figures.




Figure 4: Baseline results from estimating equation (23)

of figure 3. 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 3 depicts the distribution of item-level means of the suboptimal inflation rate.³⁸ For the vast majority of items, the average suboptimal inflation rate lies between $\pm 0.5\%$ per month. The lower right panel in figure 3 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 (23), which captures how suboptimal inflation distorts relative prices. To increase chances that our key identifying assumptions (identical degrees of price rigidity and identical stochastic processes for idiosyncratic disturbances) are satisfied, the estimation is carried out separately for each of the 1071 U.K. expenditure items in our sample. Section 6 presents an alternative estimation approach that allows relaxing these assumptions.

5.1 Baseline Results

Figure 4 presents our baseline estimation outcome. The left panel depicts the distributions of the estimated coefficients c_z , obtained from estimating equation (23) for each expenditure category z . We find that 97% of the estimated coefficients are positive, in line with the predictions of sticky price theories. The right panel in figure 4 depicts the distribution of t -statistics: 94% of the estimates have a t -statistic larger than two, 82% a t -statistic larger than five, and only 1% of the coefficients have a t -statistic below minus two. Figure 4 thus provides overwhelming support for the notion that suboptimal inflation gives rise to price distortions at the product level. Row 1 in table 2 reports further details of the regression outcome.

Interestingly, the median adjusted R^2 value of the second-stage regression (23) is 15%. 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.³⁹ 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 (23).

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 level⁴⁰ increases the standard deviation of first-stage residuals by 3.5 percentage points.⁴¹

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 7 will assess the quantitative importance of relative price distortions using directly the (unbiased) first-stage residuals.

Sticky price theories suggest that the coefficient c_z is determined by the adjustment rate for prices, see equations (19) and (22). 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 (19) to solve for the regression-implied share of non-adjusting

³⁸See appendix H for information on the cross-sectional distribution of the product-specific optimal inflation rate Π_{jz}^* .

³⁹Recall that item-specific constants do not contribute to the R^2 values of the second-stage regressions (23).

⁴⁰The 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 3.

⁴¹The predicted increase in the variance is $0.12\% = 12 \cdot (0.01)^2$ and the reported 3.5% number is the square root of 0.12%.




Figure 5: Observed versus regression-implied price stickiness (α_z)

products α_z , we obtain:⁴²

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

Figure 5 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 strongly positive correlation equal to +0.5. This shows that price distortions 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 5, while theory predicts the two measures to align along this line. A downward bias in the regression-implied value for α_z can arise from downward bias in our estimated coefficients c_z , which emerges due to first-stage estimation error, see appendix F.

Overall, our baseline results provide strong support for the notion that suboptimal inflation distorts relative prices. The next section explores the robustness of this finding.

⁴²We perform the inversion only for items with strictly positive estimated c_z , which is true for 97% 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.

Row	Specification	Share of positive point estimates ($c_z > 0$)	Share of estimates c_z with t -stat < -5 < -2 $> +2$ $> +5$	Median c_z	Mean c_z	Median adj. R^2	No. of price obs. (millions)
1	Baseline	97%	0% 1% 94% 82%	8.81	11.85	15%	15.4
2	Alternative 2nd Stage	96%	0% 1% 90% 74%	8.32	11.64	17%	15.4
3	Add linear term to r.h.s.	53%	1% 14% 26% 8%	0.27	-3.16	0%	15.4
4	Only squared inflation on r.h.s. Sign-dependent coefficients: $c^+, \ln \Pi_z / \Pi_{jz}^* > 0$ $c^-, \ln \Pi_z / \Pi_{jz}^* < 0$	97% 98%	0% 1% 89% 68% 0% 1% 93% 73%	8.11 11.41	10.4 16.11	16% 16%	14.7 14.7
	Alternative Selection						
5	Using only products with ≥ 12 price observations	98%	0% 1% 92% 75%	12.72	15.99	15%	12.6
6	≥ 24 price observations	95%	0% 1% 85% 62%	23.42	27.97	15%	8.7
7	≥ 2 nominal price changes	91%	0% 3% 71% 30%	6.6	8.74	5%	8.4
8	≥ 4 nominal price changes	85%	0% 4% 53% 14%	6.77	10.81	3%	5
9	Include sales prices	98%	0% 1% 93% 75%	8.76	13.06	12%	16.8
	Alternative 1st Stage						
10	Nonlinear 1st stage	97%	0% 1% 93% 78%	6.41	9.23	12%	15.4
11	Trend stability test	97%	0% 0% 90% 67%	7.1	11.09	13%	6.3
12	- with p -value $\geq 10\%$ - with p -value $\geq 20\%$	97%	0% 1% 88% 61%	7.27	11.37	13%	4.9
	Within Product Regression						
13	All products	99%	0% 1% 97% 90%	3.15	3.32	21%	15.4
14	Only products with ≥ 1 price changes <i>per</i> half life	97%	0% 1% 86% 49%	3.38	4.19	10%	6.7
15	≥ 2 price changes <i>per</i> half life	91%	1% 2% 70% 20%	4.28	5.49	5%	3.7

Table 2: Regression outcomes: baseline estimation and robustness exercises




Figure 6: Robustness to adding a linear term (equation (24), estimates for c_z are reported in table 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 2, which also reports the baseline outcome.

Adding Linear Terms. Sticky price theories predict that only the squared deviation of inflation from its optimal level should explain the variance of first-stage regression residuals. In particular, the linear gap between actual and optimal inflation should have a zero coefficient when added to the right-hand side of equation (23). One can test this over-identifying restriction and run regressions of the form

$$\widehat{Var}(u_{jzt}) = v_z + v_z^l \cdot \ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^* + c_z \cdot (\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*)^2 + \varepsilon_{jzt}, \quad (24)$$

to 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 6 reports the distribution of the estimated v_z^l (left panel) and the associated distribution of t -statistics (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 2 shows that estimates for c_z are hardly affected by the presence of a linear term, again in line with what sticky price theory suggests.

Inflation versus Suboptimal Inflation as Regressor. It turns out to be important for our empirical results that the right-hand side of

equation (23) features the squared value of *suboptimal* inflation rather than simply the squared value of inflation. To illustrate this point, let $\widehat{\ln \Pi_z}(j)$ denotes the average item-level inflation rate prevailing *over the lifetime of product j* and consider the following alternative formulation of the second-stage regression

$$\widehat{Var}(u_{jzt}) = v_z + c_z \cdot (\widehat{\ln \Pi_z}(j))^2 + \varepsilon_{jz}, \quad (25)$$

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

Row 3 in table 2 shows that outcomes differ radically from the baseline: (i) about half of the point estimates for c_z are then positive with the other half being negative; (ii) 60% of the coefficients are statistically insignificant, and (iii) the R^2 value of the regression drops almost to zero.

This shows that one would wrongly conclude that inflation is *not* associated with price distortions, if one assumes product-specific optimal inflation to be equal to zero, as suggested by textbook sticky price models. This finding also highlights that the baseline findings are predominantly due to differences in the *product-specific optimal inflation rate* Π_{jz}^* in the cross-section of products j within an item z .

Positive versus Negative Deviations from Optimal Inflation.

We now explore whether the direction of the deviation from optimal inflation makes a difference for observed 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 price have to *increase* when $\ln \Pi_z / \Pi_{jz}^* > 0$. If price rigidities depend on the direction of the price adjustment, then positive versus negative deviations from the optimal inflation rate generate price distortions of different strength. We can test whether this is the case by estimating the baseline equation (23) using coefficients that depend on the sign of the deviation:

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

where $I_{\{x\}}$ is an indicator function that is equal to 1 if x is true and zero otherwise.

Row 4 in table 2 shows that one obtains similar outcomes in terms of the share of positive point estimates and for the statistical significance of the estimates, independently of the sign of the deviation. Yet, the magnitude of the two coefficients differs notably: the mean and median estimate for c^+ is significantly smaller than the corresponding numbers for c^- . Price distortions are thus smaller when firms have to increase prices to counteract the effects of suboptimal inflation. This suggests that downward rigidity of prices is somewhat stronger than upward rigidity.

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 or perhaps even to spurious results. We address these concerns by selecting products for which estimation errors are likely going to be 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 2. While results barely change in terms of the share of positive coefficients c_z and their statistical significance, the magnitudes 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 .

In a second approach, we use the number of nominal price changes a selection criterion for including products in the regressions. The idea behind this approach is that we would like to exclude products with only few price changes, so as to avoid 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 price changes and a regression using only products with 4 or more price changes. Rows 7 and 8 in table 2 show that one obtains again a very large number of positive point estimates and high statistical significance levels, albeit less strongly than in the baseline. Also, the R^2 value of regression falls to about one half or one-third of the baseline level and the magnitudes of coefficient estimates decline relative to the baseline. Despite this, support for the notion that suboptimal inflation distorts relative prices remains very 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 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 (14). Since the presence of nonlinear time trends may be source of concern, we recompute the first-stage residuals allowing also a quadratic time trend and then use the resulting residuals in our second-stage regression (23). Row 10 in table 2 reports

the regression outcomes. We obtain again 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%.⁴³ The estimation outcomes are reported in rows 11 and 12 of table 2 and hardly change compared to the baseline.

6 Exploiting Within-Product Variation

This section explores an alternative estimation strategy that allows to significantly relax key identifying assumptions underpinning the baseline estimation approach. The strategy consists of exploiting *within-product* variation and can deal with settings in which idiosyncratic shock processes and price rigidities *both differ* across products within the *same* expenditure item. It thus addresses key concerns one might have with the baseline approach but comes at the cost of increased second-stage attenuation bias.

The key idea underlying the alternative approach consists of splitting the sample life of products into two equally long subsamples and to exploit variation in the inflation rate over time 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 and $\ln \Pi_{jz2} - \ln \Pi_{jz}^*$ the suboptimal rate in the second half.⁴⁴

Consider first the case with Calvo frictions: equation (17) then applies separately in the first and in the second half life of each product.⁴⁵ This allows taking - for each product j in item z - the time differences across the product half lives, which delivers

$$\begin{aligned} & \text{Var}_1(u_{jzt}) - \text{Var}_2(u_{jzt}) \\ & = c_z((\ln \Pi_{jz1} - \ln \Pi_{jz}^*)^2 - (\ln \Pi_{jz2} - \ln \Pi_{jz}^*)^2), \end{aligned} \quad (26)$$

where $\text{Var}_1(u_{jzt})$ and $\text{Var}_2(u_{jzt})$ denote the residual variances in the first and second half of the product lifetime, respectively.⁴⁶ The key

⁴³As 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. This is not a problem here, as it only increases the strictness of selecting products featuring for a constant trend.

⁴⁴The suboptimal rates can be estimated using equation (62) in appendix F separately for the first and second half of product lifetime.

⁴⁵This assumes that the change in inflation are not anticipated, which is the case whenever *changes* in inflation are unpredictable.

⁴⁶These are estimated using the same regression as in the baseline approach.

feature of equation (26) is that it eliminates the constant present in the baseline regression specification (17). This allows testing whether the coefficient $c_z = \alpha_z / (1 - \alpha_z)^2$ in equation (26) is positive *without* requiring that idiosyncratic shock processes are identical across products. Moreover, when 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 (26) will recover the average price distortion across products, i.e.,

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

provided the product-specific coefficients $\alpha_{jz} / (1 - \alpha_{jz})^2$ display conditional mean independence from the regressor in (26).⁴⁷ 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.

The within-product estimation approach generalizes the baseline approach, but the second-stage (26) will likely feature stronger right-hand side measurement error: one now has to estimate (in the first stage) how suboptimal inflation changes over the product life, rather than just the level of suboptimal inflation. One can thus expect increased second-stage attenuation bias in the estimated coefficient c_z .

Next, consider the case with menu cost frictions, for which similar arguments apply. Taking differences across the first and second half of product life using equation (21), one obtains (up to a second-order approximation) again equation (26), but with the regression coefficient now given by $c_z = E[1 / (\lambda_{iz}^X)^2]$, where λ_{iz}^X is the switching intensity in the i -th idiosyncratic state of the idiosyncratic shock process. The regression coefficient is now independent of the menu-cost parameter κ , so that the estimation approach (26) remains valid in a menu-cost setting in the presence of product-specific menu-costs.⁴⁸ If the expected switching intensities $E[1 / (\lambda_{iz}^X)^2(j)]$ also differ across products j within the *same* item, but display conditional-mean independence from the regressor in equation (26), then OLS estimation of equation (26) again

⁴⁷See appendix G for a proof and further details.

⁴⁸Heterogeneity 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.

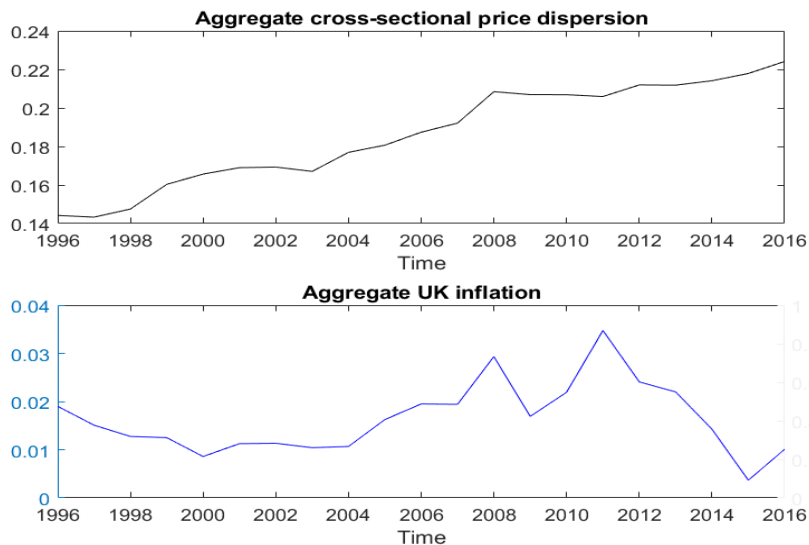


Figure 7: Aggregate cross-sectional price dispersion and inflation

recovers the average distortion coefficient

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

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 (26) is positive.

Row 13 in table 2 reports the outcomes from estimating equation (26). It shows that 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 than in the baseline. 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 2 repeat the within-estimation approach using only products that have at least 1 or at least 2 nominal price changes *per half life*. One then still obtains very strong support for the notion that suboptimal inflation is associated with relative price distortions.

7 Understanding Cross-Sectional Price Dispersion

The analysis focused thus far on price distortions over time at the level of individual products. This section shifts focus and considers the cross-sectional dispersion of prices at a given point in time and its comovement with inflation over time.

The top panel in figure 7 depicts an aggregate measure of cross-sectional price dispersion. The aggregate measure is constructed by computing first the cross-sectional variance of relative prices, $Var^j(\ln p_{jzt})$, for each item z and in each year. Item-level variances are then aggregated across items using expenditure weights. The figure shows that the cross-sectional dispersion of prices has increased by more than 50% over the sample period. At the same time, aggregate inflation, depicted in the bottom panel of figure 7, displays no clear time trend.

At first glance, this suggests that inflation is not associated with price dispersion in the cross-section of products. Yet, this conclusion turns out to be wrong: we show that the increase in cross-sectional price dispersion over time, depicted in the top panel of figure 7, is due to an increase in the dispersion of *flexible* prices over time, which masks an underlying positive relationship between inflation and cross-sectional price distortions. Higher inflation is thus associated in the data with higher cross-sectional price *distortions*, even though it displays no relationship with cross-sectional price *dispersion* over time.

Section 7.1 shows how one can decompose cross-sectional price dispersion into a component capturing identifiable components of the flexible price distribution and a remainder, whose variation over time identifies variation in price distortions over time. Section 7.2 uses this result to analyze (i) the relationship between cross-sectional price distortions and inflation at the item level and (ii) the relationship between aggregate inflation and an expenditure-weighted average of item-level cross-sectional price distortions. Section 7.3 documents that the increase in cross-sectional price dispersion is almost fully accounted by an increase in the dispersion of flexible prices. It also derives quantitative bounds on the amount of cross-sectional price distortion generated by inflation over the sample period.

7.1 Decomposing Cross-Sectional Price Dispersion

From the sticky price theories analyzed in section 3, it follows that the price of product j in expenditure category z evolves *over time* according

to⁴⁹

$$\ln p_{jzt} = \ln p_{jz}^* - \ln \Pi_{jz}^* \cdot t + u_{jzt}, \quad (27)$$

where the residuals u_{jzt} have mean zero, are independent across j and z , and have variance *over time* equal to

$$\text{Var}(u_{jzt}) = v_z + c_z \cdot (\ln \Pi_z - \ln \Pi_{jz}^*)^2. \quad (28)$$

We are now interested in decomposing the *cross-sectional* variance of prices, denoted by $\text{Var}^j(\ln p_{jzt})$, of the products j present at some time t in some item z . In particular, we would like to evaluate how this measure of *cross-sectional* price dispersion depends on the item-level inflation rate Π_z .

Suppose there is a unit mass of products j in item z and that each month a share of products randomly exits the sample and gets replaced by newly sampled products. In general, 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.⁵⁰

Appendix H shows that time variation in the cross-sectional distribution of optimal inflation rates $\{\Pi_{jz}^*\}$ is minor. This allows restricting consideration to a setting with a time-invariant cross-sectional distribution of optimal inflation rates. Specifically, we assume that upon product entry, the optimal inflation rate Π_{jz}^* is an i.i.d. draw from $\{\Pi_z^{*1}, \Pi_z^{*2}, \dots, \Pi_z^{*I}\}$, where Π_z^{*i} is chosen with probability $m_z^i \geq 0$ for $i = 1, \dots, I$ and $\sum_i m_z^i = 1$.

In contrast, the distribution of estimated intercepts $\{p_{jz}^*\}$ strongly moves with time in the data. To account for this, we allow for arbitrary time variation in the cross-sectional distribution of intercepts for newly incoming products.⁵¹ Given this setup, we derive the following decomposition result.⁵²

Proposition 3 *Let $\text{Var}^j(\cdot)$ denote the variance in the cross-section of products j . The cross-sectional variance of relative prices in expenditure*

⁴⁹See equation (53) in appendix D.3 for the case with Calvo frictions and equation (54) in appendix E for the case with menu costs.

⁵⁰We assume that upon the time of entry, the 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, instead products that are newly sampled by the Office of National Statistics.

⁵¹The covariance between the distribution of intercepts $\{p_{jz}^*\}$ and optimal inflation rates $\{\Pi_{jz}^*\}$ is also left unrestricted.

⁵²See appendix I for the proof.

category z at time t is then given by

$$\text{Var}^j(\ln p_{jzt}) = \text{Var}^j(\ln p_{jz}^* - \ln \Pi_{jz}^* \cdot t) + \text{Var}^j(u_{jzt}), \quad (29)$$

where

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

Equation (29) decomposes the cross-sectional price dispersion into two components. The first component, $\text{Var}^j(\ln p_{jz}^* - \ln \Pi_{jz}^* \cdot t)$, captures identifiable elements of the flexible price distribution that are deterministic from the viewpoint of an individual product j . We can identify this component using our first stage estimates of (a_{jz}, b_{jz}) from equation (14). The second component in equation (29), $\text{Var}^j(u_{jzt})$, is given by (30) and depends on the constant v_z , which captures variation induced by stochastic components of the flexible price, and $c_z \cdot E^j[(\ln \Pi_z - \ln \Pi_{jz}^*)^2]$, which captures price distortions induced by inflation.⁵³

The decomposition in proposition 3 holds at each point in time in a setting where steady-state inflation Π_z is constant. Yet, the decomposition also applies in a setting where inflation is slowly changing over time. For instance, suppose that inflation changes from year to year according to a random walk, with inflation being equal to Π_{zt} in year t . Price setters then expect future inflation to be equal to the current inflation rate Π_{zt} , so that our steady-state pricing results continue to apply. And since the vast majority of prices have adjusted over the course of a year, the cross-sectional dispersion of prices at the end of each year will depend largely only on the inflation rate Π_{zt} prevailing during year t . Equation (30) thus provides a theory-implied relationship linking (yearly) inflation rates Π_{zt} to the cross-sectional distribution of first-stage residuals $\text{Var}^j(u_{jzt})$ at the end of the year.

Depending on the distribution of optimal inflation rates, an increase in Π_{zt} can lead to either an increase or a decrease in price distortions: if average optimal inflation rate ($E^j[\Pi_{jz}^*]$) lies below actual inflation, then price distortions are predicted to increase with inflation. The opposite is true if the average optimal inflation rate lies above actual inflation. Using our first-stage residuals, we can estimate $\text{Var}^j(u_{jzt})$ as $\text{Var}^j(\hat{u}_{jzt})$ and test whether the predicted relationship with inflation is actually present in the data.⁵⁴ We investigate this issue in the next section.

⁵³In the absence of price stickiness, we have $c_z = 0$ so that price dispersion does not depend on inflation.

⁵⁴The estimated first-stage residuals \hat{u}_{jzt} are unbiased but contaminated with measurement error. To the extent that measurement error does not vary over time, $\text{Var}^j(\hat{u}_{jzt})$ will correctly capture the time variation of $\text{Var}^j(u_{jzt})$.




Figure 8: Correlation between inflation and price distortions at the item level

7.2 Inflation and Cross-Sectional Price Distortions over Time

The present section investigates the comovement between inflation and cross-sectional price distortions over time. It considers first comovement at the level of expenditure items and then comovement at the aggregate level, i.e., for the expenditure-weighted average across items.

7.2.1 Item Level Results

This section test whether the correlation over time between the item-level inflation rate Π_{zt} and the cross-sectional price distortion, as measured by $Var^j(\hat{u}_{jzt})$, behaves in line with the predictions of proposition 3

To test this prediction, we compute the correlation between Π_{zt} and $Var^j(\hat{u}_{jzt})$ over time, using all items z for which we have at least three years of data.⁵⁵ The top panel in figure 8 depicts the resulting distribution of correlations across items, using all correlations with a p -value less or equal to 10%. The bottom panel depicts the distribution of p -values.⁵⁶ Figure 8 shows that there are significantly positive and significantly negative correlations, but more positive than negative ones.⁵⁷

Proposition 3 implies that positive (negative) correlations emerge whenever optimal average inflation ($E^j[\Pi_{jz}^*]$) lies above the average ac-

⁵⁵This is the case for 696 expenditure items.

⁵⁶206 of the 696 correlations have p -values smaller than 10%.

⁵⁷This result is robust to choosing tighter p -values, e.g., a value of 5%, or to considering all correlations, independently of their p -value.

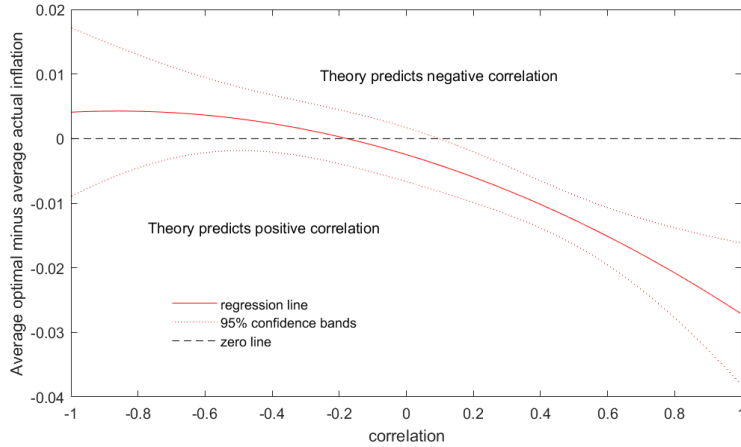


Figure 9: The relationship between average optimal minus actual inflation (y-axis) and the correlation between inflation and cross-sectional price distortions (x-axis)

tual inflation rate in the item. Figure 9 shows that this is indeed the case: 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 in line with the predictions of proposition 3. This is particularly true for the statistically significant parts of the regression line.⁵⁸

7.2.2 The Expenditure-Weighted Average Item

We now consider an economy-wide measure of cross-sectional price distortions and its comovement with aggregate inflation. In particular, we aggregate the cross-sectional item-level variances $Var^j(\hat{u}_{jzt})$, considered in the previous section, across items using item-level expenditure weights. According to equation (30), time-variation in this measure reflects time variation in cross-sectional price distortions.

Figure 10 depicts the resulting aggregate cross-sectional distortion measure together with the aggregate inflation rate.⁵⁹ Unlike aggregate price dispersion, which is trending upward, see figure 7, aggregate price distortions do not show much of a time trend. In addition, aggregate

⁵⁸This 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 significant.

⁵⁹Note that aggregate inflation is also an expenditure-weighted average of item-level inflation rates. Figure 10 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.

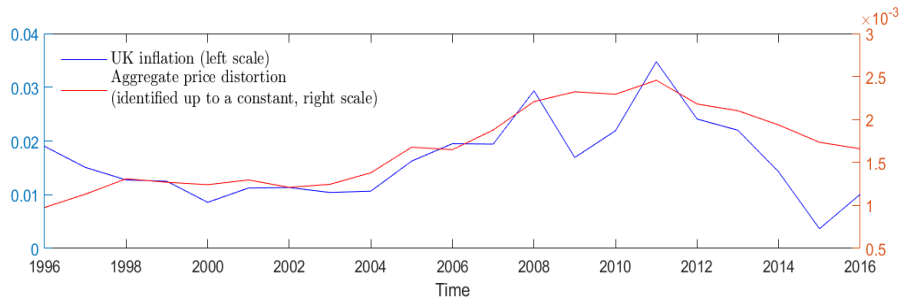


Figure 10: Aggregate inflation and aggregate cross-sectional price distortions

price distortions covary positively with aggregate inflation: the correlation between both measures is equal to $+0.67$ and is significant at the 1% level. This shows that higher aggregate inflation rates are associated with larger amounts of cross-sectional price distortions in the data.

Importantly, this result is not driven by outliers in the distribution of first-stage residuals. For instance, results are similar when removing the 2.5% highest and 2.5% lowest residuals in each item before computing the variance of first-stage residuals. Likewise, computing instead a robust dispersion measure leads to very similar outcomes.⁶⁰

Proposition 3 predicts that the positive correlation between aggregate inflation and cross-sectional price distortions is driven by products for which the optimal inflation rate Π_{jz}^* lies below the actual inflation rate. This theoretical prediction can again be tested. To do so, we group individual products according to their optimal inflation rate. Specifically, we consider the 1/3 of products with the highest and the 1/3 of products with the lowest optimal inflation rate in each expenditure item and then recompute price distortions for these two sub-groups.⁶¹

The top group of products has an (unweighted) average optimal inflation rate that varies between $+2.75\%$ and $+5.0\%$ over time, which roughly covers the range in which actual inflation moves. The bottom group, however, has a deeply negative optimal inflation rate that ranges

⁶⁰Following Nakamura et al. (2018), we computed the interquartile range (IQR) of first-stage residuals at the level of each expenditure item and then use the expenditure-weighted median to aggregate across items. This leads to very similar conclusions.

⁶¹We split products within each expenditure item, rather splitting products across all items combined, to avoid that results are driven by compositional effects. As is well-known, the average optimal inflation rates varies systematically in the cross section of items.

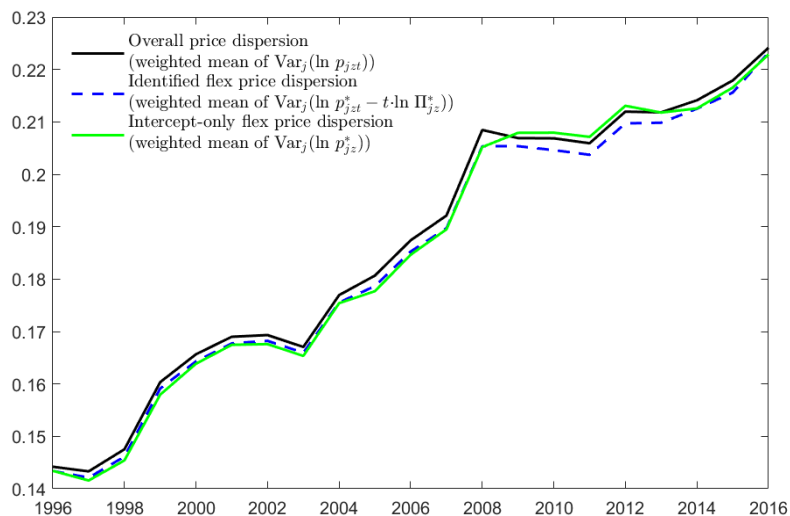


Figure 11: Overall versus flexible cross-sectional price dispersion (various identified components)

between -6% and -9% over time. According to the theory, this group should display a strong positive correlation between price distortions and inflation over time. In contrast, the top group should display no or only a weak correlation with inflation.

This is indeed what we find: for the top group, the correlation between inflation and price distortions is weaker ($+0.38$) and only marginally significant (p -value of 0.09); for the bottom group, the correlation is strongly positive ($+0.69$) and highly significant (p -value of 0.001). In line with sticky price theory, the positive correlation between inflation and inefficient price dispersion at the aggregate level is thus driven by products with optimal inflation rates that lie below actual inflation.

7.3 Bounds on Price Distortions and Changes in the Dispersion of Flexible Prices

This section derives upper and lower bounds on the amount of relative price distortions due to inflation over the sample period and discusses the drivers of the upward trend in aggregate price dispersion in the top panel of figure 7.

Proposition 3 implies that the (aggregated) variance of first stage residuals represents an upper bound on the amount of price distortions

that is due to inflation.⁶² The upper bound of the variance reached in figure 10 is approximately $2.5 \cdot 10^{-3}$. Therefore, absent any flexible price dispersion, price distortions gives rise to a standard deviation of prices of at most $\sqrt{2.5 \cdot 10^{-3}} = 5\%$. While this is quantitatively large, price distortions account only for about 1% of aggregate 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 the theory - solely due to inflation. This range is approximately equal to $1.5 \cdot 10^{-3}$ in figure 10 and implies (in the absence of flexible price dispersion) that inflation would induce variation in the standard deviation of prices of up to $\sqrt{1.5 \cdot 10^{-3}} = 3.87\%$ over time. Again, this appears sizable in absolute terms.

Aggregate price dispersion, however, is overwhelmingly driven by price dispersion that is also present under flexible prices. Figure 11 depicts the aggregate price dispersion, previously shown in the top panel of figure 7), together with the dispersion of the identifiable components of the flexible price dispersion (the expenditure-weighted item level variances $Var^j(\ln p_{jz}^* - \ln \Pi_{jz}^* \cdot t)$ from proposition 3).

Figure 11 shows that the identifiable component of flexible price dispersion accounts for the vast majority of aggregate price dispersion and also closely tracks it over time. Since time variation in the distribution of optimal inflation rates ($\ln \Pi_{jz}^*$) is very limited, virtually all time-series variation is due to time-series variation in the cross-sectional dispersion of the intercepts ($\ln p_{jz}^*$).⁶³

This shows that time series variation in aggregate price dispersion is to a large extent driven by time series variation in flexible price dispersion, which strongly rises over time. The increase in flexible price dispersion may reflect a number of economic forces, such as a widening cross-sectional distribution of mark-ups, productivities, or unobserved product qualities/variety. The large increase in flexible price dispersion over time is also the predominant reason why aggregate inflation fails to covary with aggregate price dispersion over time, see figure 7.

8 Conclusions

In summary, our research derives three key insights:

1. We establish a robust link between deviations of inflation from its

⁶²This is so because the intercept $v_z \geq 0$ in equation (30) is not due to inflation.

⁶³To make comparisons meaningful over time, figure 11 reports the dispersion coming from intercepts using the normalized intercepts $\ln p_{jz}^* - \Pi_{jz}^* \cdot t_{jz}^0$, where t_{jz}^0 is the time period in which the product first enters the sample.

product-specific optimal level and an increase in price distortions at the product level.

2. At the aggregate level, we find a positive association between variations in aggregate inflation and cross-sectional price distortion, which is mainly driven by products with optimal rates of inflation below actual inflation.
3. The dynamics of aggregate cross-sectional price dispersion over time are largely driven by identifiable components of the flexible price dispersion. This suggests that factors beyond inflation are the main driver of the dynamics in aggregate price dispersion.

Collectively, these findings offer substantial empirical support for the theoretical foundations of sticky price models and the monetary policy implications they engender.

In future research, we intend to explore the relationship between price distortions and demand misallocation. This requires the observation of product quantities alongside prices, which is feasible only for a narrower range of products. Yet, it holds significant implications for understanding the real implications of suboptimal inflation.

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.
- (2022): “Estimating the Optimal Inflation Target from Trends in Relative Prices,” forthcoming in: *American Economic Journal: Macroeconomics*.
- ALVAREZ, F., M. BERAJA, M. GONZALEZ-ROZADA, AND P. A. NEUMEYER (2019): “From Hyperinflation to Stable Prices: Argentina’s Evidence on Menu Cost Models,” *The Quarterly Journal of Economics*, 134(1), 451–505.
- ALVAREZ, F., F. LIPPI, AND A. OSKOLOV (2022): “The Macroeconomics of Sticky Prices and Generalized Hazard Functions,” *Quarterly Journal of Economics*, 137, 989–1038.
- ARGENTE, D., AND C. YEH (2022): “Product Life Cycle, Learning, and Nominal Shocks,” *The Review of Economic Studies*, 89(6), 2992–3054.

- 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 Keynesian Framework 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), 108–113.
- 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.

A 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:

$$\max_{\ln p_{jt}} - \sum_{i=0}^{N-1} (\ln p_{jt+i} - \ln p_{jt+i}^*)^2 = \max_{\ln p_{jt}} - \sum_{i=0}^{N-1} (\ln p_{jt} - \ln p_{jt}^* - i \ln(\Pi/\Pi_j^*))^2$$

The first order condition yields:

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

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

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

Since the flex 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 $0 \leq i < N$:

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

Note that adjustment size is given by:

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

A.1 Absolute Price Changes May Miss Price Distortions

Suppose now that 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:

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

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:

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

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

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

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

B 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 (12):

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

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

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

where λ_{ij}^f is the probability of observing f_j in the subsequent period, conditional on observing f_i in the previous period.⁶⁴ 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 $\ln \mathbf{x} \equiv [\ln x_1, \dots, \ln x_N]$ 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 $\ln \tilde{\mathbf{x}} \equiv [\ln \tilde{x}_1, \dots, \ln \tilde{x}_N, \ln \tilde{x}_{N+1}]$ 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 $\varepsilon > 0$.⁶⁵ We now construct the transition matrix $\tilde{\Lambda}^x$ in the following way:

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

All elements in black are borrowed directly from the Λ^x matrix, whereas elements in red are to be solved for.⁶⁶ The first $(N-1)$ rows of $\tilde{\Lambda}^x$ ensure

⁶⁴This can be achieved by conditioning on price spells of length one.

⁶⁵One 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.

⁶⁶We order states such that $\lambda_{N1}^x > 0$ and $\lambda_{NN}^x > 0$. 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 , $\lambda_{ij_1}^x > 0$ and $\lambda_{ij_2}^x > 0$.

that for all $i < N$:

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

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

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

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

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

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

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

and the remaining elements $\tilde{\lambda}_N^x$ and $\tilde{\lambda}_{N+1}^x$ 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 $\tilde{\lambda}_{N1}^x$, $\tilde{\lambda}_N^x$, $\tilde{\lambda}_{(N+1)1}^x$ and $\tilde{\lambda}_{N+1}^x$ are all $\in [0, 1]$. Such ε always exists since we have ordered the states to ensure $\lambda_{N1}^x > 0$ and $\lambda_{NN}^x > 0$ 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 $\frac{\lambda_{iN}^x}{2} + \frac{\lambda_{iN}^x}{2} = \lambda_{iN}^x$. Finally, note that states $\ln x_N$ and $\ln x_{N+1}$ have the same unconditional probability,⁶⁷ and therefore the probability of moving from f_N to f_i is equal to $\frac{1}{2} \left(\tilde{\lambda}_{Ni}^x + \tilde{\lambda}_{(N+1)i}^x \right) = \lambda_{Ni}^x$ for all

⁶⁷The 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} .

$i < N$. This implies that the probability of staying in f_N is also the same as in the original process (λ_{NN}^x).

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.

C 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, whenever $\ln x_{jzt}$ is a random walk. This holds true with Calvo frictions, see equation (11), but also for the case with menu cost frictions.

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

$$\ln p_{jzt_{n+1}}^{opt} - \ln p_{jzt_n}^{opt} = -\ln \Pi_{jz}^* \cdot (t_{n+1} - t_n) + \ln e_{jzn+1} \quad (31)$$

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 and have adjustment-time-dependent variance $(t_{n+1} - t_n)\sigma_z^2$, where σ_z^2 denotes the innovation variance in the random walk in expenditure item z . These two properties can be tested.

To test for the adjustment-time-dependent variance, we use all observations $(t_{n+1} - t_n, e_{jzn+1})$ within some item z to run the regression

$$(\ln e_{jzn+1})^2 = a_z + b_z(t_{n+1} - t_n) \quad (32)$$

and check whether $b_z = \sigma_z^2 > 0$ as predicted by the random walk. Figure 12 reports the distribution of the estimated b_z and the associated t -statistics using all products with $N_{jz} > 3$. It shows that the random walk hypothesis $b_z > 0$ is strongly rejected by the data.

Second, we can also test if the residuals $\ln e$ in (31) are uncorrelated over time. 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




Figure 12: Random walk test, equation (32)

estimates for all products with $N_{jz} > 3$:

$$\widehat{Var}_z = \sum_j \left(\frac{N_{jz} - 2}{\sum_k (N_{kz} - 2)} \sum_{n=2}^{N_{jz}} \frac{\left(\frac{\ln e_{jzn}}{\sqrt{t_n - t_{n-1}}} \right)^2}{N_{jz} - 2} \right)$$

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

The left panel in figure 13 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 $\ln x_{jzt}$. While the estimated coefficients b_z in (32) are then symmetrically centered around zero (but still not predominantly positive), the evidence on the auto-correlation of the residuals remains almost identical to the one shown in figure 13 for the case with an estimated time trend Π_{jz}^* in the first-stage regression.

Based on these findings, we conclude that unobserved shocks in our data do *not* follow a pure random walk.




Figure 13: Autocorrelation of residuals and bootstrapped p-values (random walk implies autocorrelation of zero)

D Details of the Calvo Model

D.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,

$$\max_{P_{jzt}} E_t \sum_{i=0}^{\infty} \alpha_z^i \frac{\Omega_{t,t+i}}{P_{t+i}} \left[(1 + \tau) P_{jzt} - \frac{W_{t+i}}{A_{zt+i}} G_{jzt+i} X_{jzt+i} \right] Y_{jzt+i} \quad (33)$$

$$s.t. \quad Y_{jzt+i} = \psi_z \left(\frac{P_{jzt}}{P_{zt+i}} \right)^{-\theta} \left(\frac{P_{zt+i}}{P_{t+i}} \right)^{-1} Y_{t+i}, \quad (34)$$

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 (8), and the nominal wage given by W_{t+i} . The parameter τ is a sales subsidy (tax if negative). Maximization is subject to equation (34), which is derived from the cost-minimizing household demand function (7) using market clearing conditions.

D.1.1 Balanced growth path

We approximate the profit maximization problem (33) 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_{zt}^e allocated to production in item z is also constant over time. All idiosyncratic shocks continue 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 (5) subject to the production function (8) and the feasibility constraint that $L_{zt}^e = \int L_{jzt}^e dj$. This implies that the efficient level of output in item z is

$$Y_{zt}^e = \frac{A_{zt}}{\Delta_{zt}^e} L_{zt}^e, \quad (35)$$

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

$$1/\Delta_{zt}^e \equiv \left(\int_0^1 (1/(G_{jzt} X_{jzt}))^{\theta-1} dj \right)^{\frac{1}{\theta-1}}. \quad (36)$$

We consider a balanced growth path in which $1/\Delta_{zt}^e = 1/\Delta_z^e$, so that equation (35) implies that item-level productivity is given by⁶⁸

$$\Gamma_{zt}^e \equiv A_{zt}/\Delta_z^e. \quad (37)$$

Using equation (6), aggregate productivity Γ_t^e of the economy is given by⁶⁹

$$\Gamma_t^e \equiv \prod_{z=1}^Z (\Gamma_{zt}^e)^{\psi_z}. \quad (38)$$

Equation (37) and the previous equations imply that the steady-state growth rate of aggregate output and consumption along the balanced growth path, $\gamma^e \equiv \Gamma_t^e/\Gamma_{t-1}^e$, is given by

$$\gamma^e = \prod_{z=1}^Z a_z^{\psi_z} \quad (39)$$

where a_z denotes the steady-state growth rate of item-level productivity A_{zt} . From equation (37), we also obtain that the steady-state growth rate of item-level output and consumption, $\gamma_z^e \equiv \Gamma_{zt}^e/\Gamma_{zt-1}^e$, is given by

$$\gamma_z^e = a_z.$$

⁶⁸It is straightforward to accommodate also a trend in $1/\Delta_{zt}^e$ in the balanced growth path, but this does not yield any additional insights.

⁶⁹To see why, substitute equilibrium output for equilibrium consumption in equation (6) and detrend all output variables in the resulting equation by their growth trends. This yields

$$\frac{Y_t^e}{\Gamma_t^e} = \left[\frac{\prod_{z=1}^Z (\Gamma_{zt}^e)^{\psi_z}}{\Gamma_t^e} \right] \prod_{z=1}^Z \left(\frac{Y_{zt}^e}{\Gamma_{zt}^e} \right)^{\psi_z},$$

so that the aggregate growth trend is given by equation (38).

D.1.2 Detrended firm problem

With growth-consistent preferences that exhibit constant relative risk aversion, the one-period household discount factor is given by $\Omega = \omega(\gamma^e)^{-\sigma} < 1$, where σ denotes relative risk aversion and ω is the rate of time preference. Using this expression, the firm problem (33)-(34) along the balance growth path can be written as

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

where $y = Y_{t+i}/\Gamma_{t+i}^e$ denotes detrended output. Furthermore, using equation (37) 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

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

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

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

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

$$\frac{y_z}{y} = \psi_z \left(\frac{P_{zt} \Gamma_{zt}^e}{P_t \Gamma_t^e} \right)^{-1}. \quad (40)$$

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

$$\Gamma_t^e E_t \sum_{i=0}^{\infty} (\alpha_z \omega (\gamma^e)^{1-\sigma})^i \left[\psi_z (1 + \tau) \frac{P_{jzt}}{P_{zt}} \Pi_z^{-i} - \frac{w}{\Delta_z^e} \frac{y_z}{y} G_{jzt+i} X_{jzt+i} \right] \left(\frac{P_{jzt}}{P_{zt}} \Pi_z^{-i} \right)^{-\theta} y, \quad (41)$$

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

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

To rewrite the firm objective (41) 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 $\tilde{p}_{jzt+i} \equiv p_{jzt} \Pi_z^{-i}$, 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 (37) 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 (40) to substitute for the last term on the right hand side in the previous equation shows that marginal costs can be expressed as

$$\psi_z mc_{jzt} = \frac{w}{\Delta_z^e} \frac{y_z}{y} G_{jzt} X_{jzt}. \quad (42)$$

Substituting the previous equation and the definition of the relative price \tilde{p}_{jzt+i} into the firm objective in equation (41) yields, after dropping the pre-multiplying constant $\psi_z \Gamma_t^e$:

$$E_t \sum_{i=0}^{\infty} (\alpha_z \omega (\gamma^e)^{1-\sigma})^i [(1+\tau) \tilde{p}_{jzt+i} - mc_{jzt+i}] (\tilde{p}_{jzt+i})^{-\theta} y. \quad (43)$$

D.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 (43), that we seek to quadratically approximate, can then be written as

$$E_t \sum_{i=0}^{\infty} (\alpha \omega (\gamma^e)^{1-\sigma})^i [(1+\tau) \tilde{p}_{jt+i} - mc_{jt+i}] (\tilde{p}_{jt+i})^{-\theta} y \quad (44)$$

where it is understood that α, \tilde{p}_{jt+i} and mc_{jt+i} are item specific objects. From equation (42) follows that

$$\ln mc_{jt} = \ln mc_j - (\ln \Pi_j^*) \cdot t + \ln x_{jt}. \quad (45)$$

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 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 (44), the objective for period $t + i$ is given by

$$D_{jt+i} = [(1 + \tau)e^{\ln \tilde{p}_{jt+i}} - e^{\ln mc_{jt+i}}] (e^{\ln \tilde{p}_{jt+i}})^{-\theta} y. \quad (46)$$

We approximate this objective to second order in the variables $\ln \tilde{p}_{jt+i}$ 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 (46) yields

$$\begin{aligned} D_{jt+i} &= (y\vartheta^{-\theta}) e^{(1-\theta)\ln mc_{jt+i}^{\det}} [-\theta(\ln \tilde{p}_{jt+i} - \ln(\vartheta mc_{jt+i}^{\det}))^2 \\ &\quad + 2\theta(\ln \tilde{p}_{jt+i} - \ln(\vartheta mc_{jt+i}^{\det}))(\ln mc_{jt+i} - \ln mc_{jt+i}^{\det})] + O(3) \\ &= (-\theta y\vartheta^{-\theta}) (mc_{jt+i}^{\det})^{1-\theta} [\ln \tilde{p}_{jt+i} - \ln(\vartheta mc_{jt+i}^{\det}) - (\ln mc_{jt+i} - \ln mc_{jt+i}^{\det})]^2 + \text{t.i.p.} + O(3) \\ &= (-\theta y\vartheta^{-\theta}) (mc_{jt+i}^{\det})^{1-\theta} [\ln \tilde{p}_{jt+i} - \ln(\vartheta mc_{jt+i}^{\det})]^2 + \text{t.i.p.} + O(3), \end{aligned} \quad (47)$$

where t.i.p. collects terms independent of policy and it follows from equation (45) that $mc_{jt+i}^{\det} = mc_j e^{-(\ln \Pi_j^*)(t+i)}$. Thus, we rewrite the Taylor approximation coefficient in the previous equation according to

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

We can now express the expected discounted sum of period profits in equation (44) accurate to second order according to

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

which is proportional to

$$-E_t \sum_{i=0}^{\infty} (\alpha\beta_j)^i [\ln p_{jt} - i \ln \Pi - \ln(p_{jt+i}^*)]^2 + \text{t.i.p.} + O(3) \quad (48)$$

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

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

which implies using equation (45)

$$p_{jt}^* = p_j^* e^{-(\ln \Pi_j^*)t} x_{jt},$$

which is equal to (10) 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 (48), which delivers

$$\max_{\ln p_{jt}} -E_t \sum_{i=0}^{\infty} (\alpha \beta_j)^i (\ln p_{jt} - i \ln(\Pi/\Pi_j^*) - \ln p_{jt}^* - \ln x_{jt+i} + \ln x_{jt})^2. \quad (49)$$

The first-order condition is given by

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

which implies that the optimal price is given by

$$\ln p_{jt}^{opt} = \ln p_{jt}^* - \ln x_{jt} + \left(\frac{\alpha \beta_j}{1 - \alpha \beta_j} \right) \ln(\Pi/\Pi_j^*) + E_t (1 - \alpha \beta_j) \sum_{i=0}^{\infty} (\alpha \beta_j)^i \ln x_{jt+i} \quad (50)$$

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 $\beta_j \rightarrow 1$, this reduces to equation (11).

D.3 Asymptotics of the first-stage regression

To simplify notation, we drop the item-level subscript z in the remainder of this appendix. Starting with equation (13), we substitute $\ln p_{jt}^{opt}$ using equation (11) and also use (10) to obtain

$$\ln p_{jt} = \xi_{jt} (\ln p_{jt-1} - \ln \Pi) + (1 - \xi_{jt}) \left(\ln p_j^* - t \ln \Pi_j^* + \frac{\alpha}{1 - \alpha} \ln(\Pi/\Pi_j^*) + f(x_{jt}) \right), \quad (51)$$

where $f(x_{jt})$ is defined in equation (12).

To derive the OLS estimates of the parameters in equation (14), we rearrange equation (51) to

$$\begin{aligned} \ln p_{jt} + t \ln \Pi_j^* &= \xi_{jt} (\ln p_{jt-1} + (t - 1) \ln \Pi_j^* - \ln(\Pi/\Pi_j^*)) \\ &+ (1 - \xi_{jt}) \left(\ln p_j^* + \frac{\alpha}{1 - \alpha} \ln(\Pi/\Pi_j^*) + f(x_{jt}) \right). \end{aligned} \quad (52)$$

Computing the unconditional expectation yields

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

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

$$\ln p_{jt} = \ln p_j^* - t \ln \Pi_j^* + u_{jt}, \quad (53)$$

where u_{jt} denotes an expectation error with zero mean. This shows that for regression (14) we get

$$\widehat{\ln a_j} \rightarrow \ln p_j^* \\ \widehat{\ln b_j} \rightarrow \ln \Pi_j^*,$$

as the number of price observations becomes large.

D.4 Proof of proposition 2

This appendix derives equations (16) and (17) in the main text. To simplify notation, we drop the item-level subscript z in the remainder of the appendix. We substitute equation (53) into equation (52), which yields directly yields equation (16). Squaring equation (16), taking unconditional expectations, and using independence of ξ_{jt} yields

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

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

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

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 (17).

E Details of the State-Dependent Model

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

E.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:

$$\begin{aligned} dz_{jt} &= d \ln p_{jt} - d \ln p_{jt}^* = - \underbrace{(\ln \Pi - \ln \Pi_j^*)}_{\mu_j} dt - d \ln x_{jt} \\ d \ln x_{jt} &= \sum_{i=1}^N (\ln x_i - \ln x_{jt}) dJ_t^i(\ln x_{jt}) \end{aligned}$$

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:

$$\begin{aligned} \ln p_{jt} &= \ln p_j^* + \ln x_{jt} - t \ln \Pi_j^* + z_{jt} \\ E[\ln p_{jt} + t \ln \Pi_j^*] &= \ln p_j^* + \underbrace{E[\ln x_{jt}]}_{=0} + E[z_{jt}] \end{aligned} \quad (54)$$

And thus the estimates of OLS regression (14) converge to

$$\begin{aligned} \widehat{\ln a_j} &\rightarrow \ln p_j^* \\ \widehat{\ln b_j} &\rightarrow \ln \Pi_j^*, \end{aligned}$$

if $E[z_{jt}] = 0$, which is true in the limiting case as $\rho \rightarrow 0$, as shown below.⁷⁰ Furthermore, residuals and their variance can be written as:

$$\begin{aligned} u_{jt} &= \ln p_{jt} - \ln p_j^* + t \ln \Pi_j^* = z_{jt} + \ln x_{jt} \\ Var(u_{jt}) &= E[z_{jt}^2] + 2E[z_{jt} \ln x_{jt}] + Var(\ln x_{jt}) \end{aligned} \quad (55)$$

⁷⁰While this result is shown formally under the assumption of sufficiently small κ , it holds more generally. As $\rho \rightarrow 0$, 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.

E.2 Solution

The firm's objective is to maximize its value from equation (20), given by:

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

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

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

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

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

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

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

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

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

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:

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

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

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:

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

with

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

and value function satisfies:

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

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:*

$$\underline{z}_i^2 - \hat{z}_i^2 = \lambda_i^X \kappa \tag{56}$$

$$e^{\alpha_i \hat{z}_i} (1 - \alpha_i \hat{z}_i) = e^{\alpha_i \underline{z}_i} (1 - \alpha_i \underline{z}_i) \tag{57}$$

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

Conditional on state x_i , the price gap distribution satisfies:

$$\begin{aligned}\lambda_i^X f_i(z) &= \mu \partial_z f_i(z) \\ \int_{\underline{z}_i}^{\hat{z}_i} f_i(z) dz &= 1\end{aligned}$$

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:

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

$$E[z^2|x_i] = \int_{\underline{z}_i}^{\hat{z}_i} z^2 f_i(z) dz = \frac{\hat{z}_i + \underline{z}_i}{\alpha_i} - \hat{z}_i \underline{z}_i \quad (58)$$

$$E[z^2] = E_x \left[\frac{\hat{z}_i + \underline{z}_i}{\alpha_i} - \hat{z}_i \underline{z}_i \right] \quad (59)$$

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 $E[zx] = E[x_i E[z|x_i]] = 0$ and the main object of interest – the variance of residuals from the OLS regression (14) – is given by:

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

E.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:

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

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

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

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

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

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 $\rho \rightarrow 0$. As we take the limit with $N \rightarrow \infty$, 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 $E[\tau_N] \rightarrow \infty$. Thus $\lim_{\rho \rightarrow 0} \rho V(z, x, \rho) = -E[z^2] - \kappa \lambda_a \equiv A$ 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:

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

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 $\bar{v} = \max_i \tilde{V}(\hat{z}, x_i)$ and $\underline{v} = \min_i \tilde{V}(\hat{z}, x_i)$. Then:

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

Meaning $\underline{v} = \bar{v}$. 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 func-

tion satisfies:

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

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 $\{z(x_i), \hat{z}(x_i), \bar{z}(x_i)\}_{i=1}^N$ 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 , $z - (\ln x_j - \ln x_i) \notin [\underline{z}(x_j), \bar{z}(x_j)]$, 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, $\bar{\kappa}$ is such that $\min_{ij} |\ln x_i - \ln x_j| = \max_i \bar{z}(x_i) - \min_i \underline{z}(x_i)$ in model M . Such $\bar{\kappa} > 0$ always exists since for all i $\lim_{\kappa \rightarrow 0} \bar{z}(x_i) = \lim_{\kappa \rightarrow 0} \underline{z}(x_i) = 0$. ■

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:

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

Similarly:

$$\begin{aligned} -\alpha_i C_i^v e^{-\alpha_i \hat{z}_i} - \frac{2\hat{z}_i}{\lambda_i^X} + \frac{2}{\alpha_i \lambda_i^X} &= -\alpha_i C_i^v e^{-\alpha_i \underline{z}_i} - \frac{2\underline{z}_i}{\lambda_i^X} + \frac{2}{\alpha_i \lambda_i^X} \\ C_i^v e^{-\alpha_i \hat{z}_i} + \frac{2\hat{z}_i}{\alpha_i \lambda_i^X} &= C_i^v e^{-\alpha_i \underline{z}_i} + \frac{2\underline{z}_i}{\alpha_i \lambda_i^X} \\ C_i^v e^{\alpha_i (\underline{z}_i - \hat{z}_i)} &= C_i^v + e^{\alpha_i \underline{z}_i} \frac{2(\underline{z}_i - \hat{z}_i)}{\alpha_i \lambda_i^X} \end{aligned} \quad (60)$$

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 (60). ■

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 (56) with respect to μ . For the second equality,

differentiate equilibrium condition (57) 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 (57) 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 $\mu \rightarrow 0$, $\hat{z}_i \rightarrow 0$, $\underline{z}_i \rightarrow -\sqrt{\lambda_i^X \kappa}$ and $E[z^2] \rightarrow 0$.

Proof. Combine equilibrium conditions (56) and (57) to obtain:

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

Since the LHS is always positive, and so is the exponent on the RHS, $\lim_{\mu \rightarrow 0} \hat{z}_i = 0$. It then follows from (56) that $\lim_{\mu \rightarrow 0} \underline{z}_i = -\sqrt{\lambda_i^X \kappa}$ and from (59) that $\lim_{\mu \rightarrow 0} E[z^2] = 0$. ■

Proof of Proposition 7. From Lemmas 9 and 10, and equation (58) it follows that:

$$\begin{aligned}
\underline{z}'_i &\equiv \frac{\partial \underline{z}_i}{\partial \mu} = \frac{1}{\lambda_i^X} + \frac{\mu \hat{z}_i - \lambda_i^X \hat{z}_i \underline{z}_i}{\mu \lambda_i^X \underline{z}_i} \\
\lim_{\mu \rightarrow 0} \underline{z}'_i &= \frac{1}{\lambda_i^X} - \lim_{\mu \rightarrow 0} \frac{\hat{z}_i}{\mu} = \frac{1}{\lambda_i^X} - \lim_{\mu \rightarrow 0} \hat{z}'_i
\end{aligned}$$

At the same time, by Lemma 9: $\hat{z}_i = \frac{\underline{z}_i \underline{z}'_i}{\hat{z}'_i}$, and by Lemma 10: $\lim_{\mu \rightarrow 0} \frac{\underline{z}'_i}{\hat{z}'_i} = 0$.

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 $\lim_{\mu \rightarrow 0} \hat{z}'_i = \frac{1}{\lambda_i^X}$. From (56) it follows that $\lim_{\mu \rightarrow 0} \underline{z}'_i = 0$ and from (59) that $\lim_{\mu \rightarrow 0} \frac{\partial E[z^2]}{\partial \mu} = 0$. 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 $\mu \rightarrow 0$, 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 $\tau(z) = \frac{1}{\Lambda} (1 - e^{\alpha(\underline{z}-z)})$. It follows from Lemma 6 and equation (59) that:

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

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

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

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

F Details of the Regression Approach

This section discusses econometric details associated with estimating our key equation (17), which relates price distortions to suboptimal inflation at the product level. In our baseline empirical approach, we estimate equation (17) 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 obtain a large number of estimates of the coefficient of interest c in equation (17).

We use a two-step estimation approach, because neither the left-hand side variable nor the right hand-side variables in equation (17) 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 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 (17) can be estimated using the residuals of relative-price regressions of the form

$$\ln p_{jzt} = \ln a_{jz} - (\ln b_{jz}) \cdot t + u_{jzt} \quad (61)$$

where j denotes the product and $z \in \{1, \dots, Z\}$ the expenditure item under consideration, with Z being the total number of expenditure items in our sample.⁷¹

Estimation of the right-hand side variables in equation (17) would require estimating the average inflation rate, $\ln \Pi_z$, and the product specific optimal inflation rate, $\ln \Pi_{jz}^*$. Having two first-stage estimates on the right-hand side of equation (17) is, however, unattractive on econometric grounds.⁷² A more parsimonious way to proceed is to estimate instead directly the gap between the item-level and product-specific optimal inflation rate ($\ln \Pi_z / \Pi_{jz}^*$) in the first stage. This can be achieved by adding the price level equation

$$\ln P_t = \ln P_0 + \ln \Pi \cdot t$$

to equation (10). Adding the item-level subindex z , we obtain for every product *another* first-stage regression of the form

$$\ln P_{jzt} = \ln \tilde{a}_{jz} + (\ln \Pi_z / \Pi_{jz}^*) \cdot t + \tilde{u}_{jzt} \quad (62)$$

where P_{jzt} denotes the *nominal* product price. Equation (62) shows that the time trend in the nominal price of the product directly identifies the gap between item-level inflation and the product-specific optimal inflation rate. Equations (61) and (62) jointly make up our first-stage seemingly unrelated regression (SUR) system.

Since the SUR system (61)-(62) does not feature exclusion restrictions, OLS estimation is identical to GLS estimation, despite the presence of correlated residuals. OLS estimation delivers an unbiased estimate of the gap $\ln \Pi_z / \Pi_{jz}^*$ and an unbiased estimate of the residual variance of interest,

$$\widehat{Var}(u_{jzt}) = \frac{1}{T_{jz} - 2} \sum_t (\hat{u}_{jzt})^2,$$

where T_{jz} denotes the number of price observations for product j in item z .

⁷¹To simplify notation, the previous sections have suppressed the item index z .

⁷²It requires discussing, amongst other things, the covariance in the estimation errors of these two right-hand side variables.

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

$$\widehat{Var}(u_{jzt}) = v_z + c_z \cdot (\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*)^2 + \varepsilon_{jz} \quad (63)$$

using OLS estimation. This delivers an estimate of c_z for each expenditure item $z = 1, \dots, Z$. The error term ε_{jz} in equation (63) 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 (21).

While the first-stage estimates $\widehat{Var}(u_{jzt})$ and $\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*$ 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 now show next.

To illustrate this point, we assume that the first-stage residuals are normally distributed. (The more general case with non-normal errors is discussed in appendix F.1 below.) When estimating the SUR system (61)-(62), the estimation error in $\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*$ 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 (63) is independent of the estimation error in $(\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*)^2$ 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 (63) 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 $(\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*)^2$ biases the second-stage estimate of c_z towards zero. This is so because measurement error in $(\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*)^2$ generates a classic attenuation effect. In addition, estimation error in $\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*$ raises the expected value of $(\ln \widehat{\Pi}_z / \widehat{\Pi}_{jz}^*)^2$, 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 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., ones associated with errors in price collection, we eliminate within each

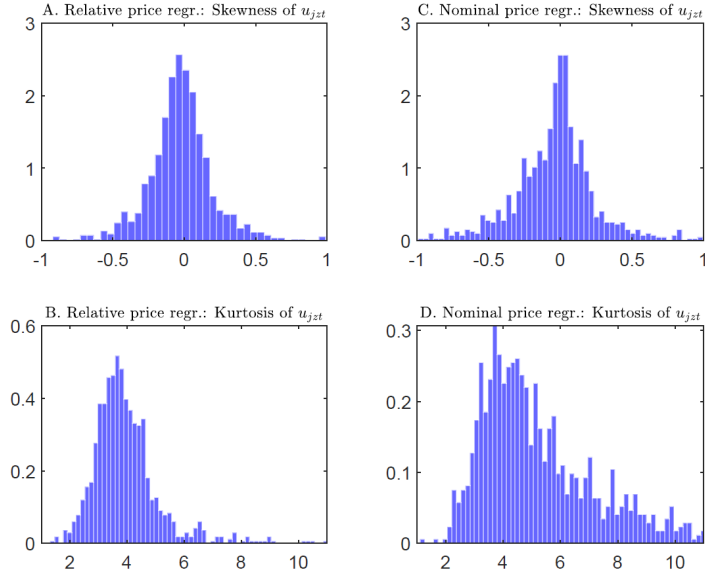


Figure 14: Skewness and kurtosis of the first-stage regression residuals

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.

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

Figure 14 reports the skewness and kurtosis of the first-stage regression residuals of equation (61) (left-hand side panels) and equation (62) (right-hand side panels) across the considered expenditure items.⁷³ 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 14, the situation looks different. Kurtosis values often lie above the value of 3 implied by a normal distribution.

We now show that quite similar arguments as for the normal case 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 to insure that the estimation error in the l.h.s. variable $\widehat{Var}(u_{jzt})$ in equation (63) is orthogonal to

⁷³The measures use outlier trimmed residuals by considering the 2.5%-97.5% quantile of the residual distribution.

(rather than independent of) the estimation error in the r.h.s. regressor $(\ln \widehat{\Pi_z / \Pi_{jz}^*})^2$.

Recall that the errors in $(\ln \widehat{\Pi_z / \Pi_{jz}^*})$ and $\{\widehat{u}_{jzt}\}$ are orthogonal by construction. A violation of orthogonality between $(\ln \widehat{\Pi_z / \Pi_{jz}^*})^2$ and $\widehat{Var}(u_{jzt})$ can thus only arise because these variables are nonlinear rather than linear functions of $\ln \widehat{\Pi_z / \Pi_{jz}^*}$ 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 $(\ln \widehat{\Pi_z / \Pi_{jz}^*})^2$ and $\widehat{Var}(u_{jzt})$ holds whenever the residuals satisfy

$$Cov\left[\left((0, 1) (X'X)^{-1} X'u(0, 1)'\right)^2, (1, 0)'u'Mu(1, 0)|X\right] = 0, \quad (64)$$

where

$$X' \equiv \begin{pmatrix} 1 & 1 & 1 & \dots \\ 0 & 1 & 2 & \dots \end{pmatrix} \quad (65)$$

is the matrix of first-stage regressors and M the matrix defined in (66) below. Condition (64) is a condition on the true residuals u , which is satisfied in the special case with normal errors. Condition (64) holds by construction when replacing the true residuals u by the estimated OLS or GLS residuals \widehat{u} , thus cannot be tested empirically using the regression residuals.⁷⁴

To understand why condition (64) 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:

$$\underbrace{Y}_{Tx2} = \underbrace{X}_{Tx2} \underbrace{\beta}_{2x2} + \underbrace{u}_{Tx2},$$

where X denotes the (deterministic) regressors defined in (65) and Y the stacked vector of the left-hand side variables (p_{jzt}, P_{jzt}) in equations

⁷⁴Using the notation introduced below, this follows from the fact that

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

(61) and (62). Letting u_t denote the residuals at date t and u the stacked residual vector, we have $E[u_t] = 0$ and

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

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

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

and the regression residuals by

$$\underbrace{\widehat{u}}_{Tx2} = MY = Mu \text{ where } M \equiv (I - X(X'X)^{-1}X') \quad (66)$$

We have

$$\begin{aligned} E[\widehat{u}'\widehat{u}|X] &= E[\underbrace{u'}_{2xT} \underbrace{M'M}_{TxT} \underbrace{u}_{Tx2} |X] \\ &= E[\underbrace{u'}_{2xT} \underbrace{M}_{TxT} \underbrace{u}_{Tx2} |X] \\ &= \text{tr}(M)E[u'u|X] \\ &= \frac{1}{T-2} \begin{pmatrix} v_{11}^2 & v_{12} \\ v_{12} & v_{22}^2 \end{pmatrix}, \end{aligned}$$

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

$$\widehat{v}_{11}^2 \equiv \frac{(1,0)'\widehat{u}'\widehat{u}(1,0)}{T-2}. \quad (67)$$

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

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

The last equality holds if and only if

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

which is the case if and only if condition 64 holds, as $E\left[\frac{(1,0)'uM'Mu(1,0)}{\text{tr}(M'M)}\right] = v_{11}^2$.

G Details of the Within Product Regression Approach

The within product regression (26) takes the form

$$Y = c_z \cdot X \quad (68)$$

where Y is a $N \times 1$ vector of consisting of $Var(u_{1jz}) - Var(u_{2jz})$ for $j = 1, \dots, N$, X a vector consisting of $(\ln \Pi_{jz1} - \ln \Pi_{jz}^*)^2 - (\ln \Pi_{jz2} - \ln \Pi_{jz}^*)^2$ for $j = 1, \dots, N$ 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

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

and its expectations satisfies under the stated assumptions

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

as claimed in the main text.

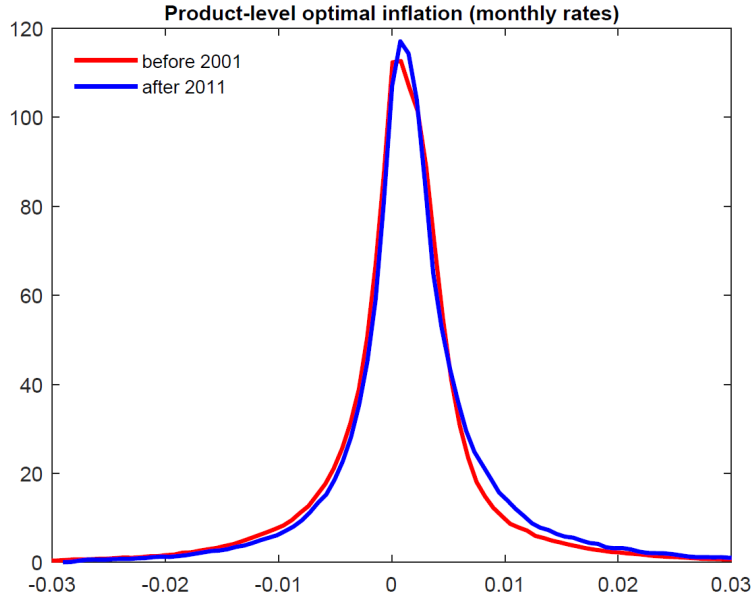


Figure 15: Distribution of product-specific optimal inflation rates Π_{jz}^* in 1996-2000 versus 2012-2016 (monthly rates, unweighted)

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

Figure 15 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.

I Proof of Proposition 3

From equation (27) we get

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

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).