

Nonlinear Inflation Dynamics in Menu Cost Economies*

Andres Blanco[†] Corina Boar[‡] Callum Jones[§] Virgiliu Midrigan[¶]

August 2024

Abstract

Existing menu cost models, when parameterized to match the micro-price data, cannot reproduce the extent to which the fraction of price changes increases with inflation. In addition, in the presence of strategic complementarities, they predict implausibly large menu costs and misallocation. We resolve these shortcomings using a multi-product menu cost model that features two key ingredients. First, the products sold by a firm are imperfect substitutes. Second, strategic complementarities are at the firm, not product level. In contrast to existing models, the fraction of price changes increases rapidly with the size of monetary shocks, so our model implies a non-linear Phillips curve.

Keywords: menu costs, inflation, Phillips curve.

*We thank Audrey Azerot and Man Chon Iao for superb research assistance, Qazi Haque, Ernesto Pasten, Francisca Sara-Zaror and Ludwig Straub for valuable discussions, as well as Fernando Alvarez, Francesco Lippi and Mark Gertler for useful feedback. The views expressed are those of the authors and not necessarily those of the Atlanta Fed or the Federal Reserve Board.

[†]Federal Reserve Bank of Atlanta, julioablanco84@gmail.com.

[‡]New York University and NBER, corina.boar@nyu.edu.

[§]Federal Reserve Board, callum.j.jones@frb.gov.

[¶]New York University and NBER, virgiliu.midrigan@nyu.edu.

1 Introduction

A salient feature of the data is that the fraction of price changes increases rapidly in periods of high inflation.¹ Since in sticky price models the fraction of price changes is a key determinant of the slope of the Phillips curve, understanding the causes of high inflation critically depends on these models' ability to reproduce this pattern. In this paper, we study the ability of menu cost models—a class of models in which the fraction of price changes is endogenous²—to reproduce this pattern and the ensuing implications for the shape of the Phillips curve. Importantly, we require these models to be consistent with the distribution of micro-price changes, a key determinant of the real effects of monetary shocks (Midrigan, 2011, Alvarez et al., 2016).

We first show that existing menu cost models with Gaussian idiosyncratic shocks, when calibrated to match the distribution of micro-price changes, cannot reproduce the extent to which the fraction of price changes comoves with inflation. Intuitively, generating the dispersion of price changes in the data requires large idiosyncratic shocks. These idiosyncratic shocks, as opposed to aggregate shocks, drive the bulk of price changes, so the fraction of price changes fluctuates little. An additional shortcoming arises in the presence of microeconomic strategic complementarities which lead to very large menu costs and misallocation from inefficient price dispersion.³

Second, we propose a resolution to these shortcomings by developing a multi-product menu cost model (Midrigan, 2011, Alvarez and Lippi, 2014) with two key ingredients. First, the products sold by a given firm are imperfect substitutes. Second, strategic complementarities are at the firm, rather than at the product level. Both of these assumptions reduce the misallocation from inefficient price dispersion inside the firm, decreasing the relative importance of idiosyncratic shocks in driving repricing decisions and thus allowing the model to reproduce the comovement between inflation and the fraction of price changes in the data. Although it is challenging to directly measure the extent of misallocation inside the firm, the observation that firms change prices infrequently and by large amounts, yet are very responsive to aggregate shocks, suggests that within-firm price dispersion is likely not very costly.

¹See Gagnon (2009), Nakamura et al. (2018), Alvarez et al. (2018) and Karadi and Reiff (2019).

²See, for example, Dotsey et al. (1999), Golosov and Lucas (2007), Gertler and Leahy (2008), Midrigan (2011), Vavra (2013), Alvarez et al. (2016), Alvarez et al. (2022a), Auclert et al. (2022).

³See Dotsey and King (2005) and Klenow and Willis (2016), who document this problem, and Leahy (2011) for a discussion of the importance of strategic complementarities for the slope of the Phillips curve and the real effects of monetary policy.

Third, we use our model to revisit the classic question of how large are the real effects of monetary policy shocks. We show that output responses in our model are very different than those in existing models that we argue are inconsistent with the data. Specifically, in our model output responds non-linearly to shocks of various sizes. The larger the shock is, the stronger the response of the fraction of price changes and therefore the smaller are the real effects. Thus, our model predicts that the Phillips curve is non-linear. In contrast, we show that existing models predict linear output responses, echoing the findings of [Auclert et al. \(2022\)](#).

We start by motivating our analysis using micro-price data that underlies the construction of the Consumer Price Index in the United Kingdom. Since aggregate inflation has not been volatile prior to the recent rise in inflation, we use sectoral data to study the high-frequency comovement between inflation and the fraction of price changes. In line with previous work, we show that the fraction of price changes increases in periods of high inflation. For example, when inflation is close to zero, the fraction of price changes is approximately 10% per month. In contrast, when inflation increases to 10%, the fraction of price changes averages 16%. To assess how important are movements in the fraction of price changes for inflation dynamics, we follow [Klenow and Kryvtsov \(2008\)](#) in decomposing inflation into an intensive margin term that keeps the fraction of price changes constant, and an extensive margin term. We find that the intensive margin term accounts for most of the movements in inflation in periods of low inflation, as in [Klenow and Kryvtsov \(2008\)](#), but for only half of these movements when inflation is relatively high. Thus, the extensive margin of price changes plays an important role at high levels of inflation.

We first confront this evidence using a single-product menu cost model in which firms are subject to Gaussian idiosyncratic and sectoral shocks. To ensure that the model can reproduce the distribution of micro-price changes, we assume that firms face random menu costs of adjusting prices, as in [Dotsey et al. \(1999\)](#), and can occasionally change their price for free, as in [Nakamura and Steinsson \(2010\)](#). A single state variable – the gap between the firm’s price and its flexible-price counterpart, in short *the price gap*, summarizes the history of shocks received by each firm. This price gap determines the hazard that the firm resets its price. In turn, the distribution of price gaps across firms and the adjustment hazard determines the distribution of price changes and the responses of the economy to aggregate shocks.

We calibrate this model to match the fraction and the distribution of price changes

in the data, as well as the volatility of sectoral inflation. A robust prediction of the model is that the fraction of price changes is nearly constant: as sectoral inflation increases from 0% to 10% in our simulations, the monthly fraction of price changes only increases from 11% to 12%, nowhere near as much as in the data. To understand why this is the case, we leverage the result in [Alvarez and Lippi \(2014\)](#) that characterizes how the fraction of price changes responds to an aggregate shock using a continuous time version of the model. This response depends only on the probability of free price changes and the size of the shock relative to the standard deviation of price changes. Since the standard deviation of price changes in the data is high relative to that of sectoral shocks, 19% vs. 1%, the fraction of price changes fluctuates little.

Our findings may appear to contradict a number of existing studies that show that the menu cost model generates substantial fluctuations in the fraction of price changes.⁴ There is, in fact, no contradiction. These papers study a simple menu cost model in the tradition of [Golosov and Lucas \(2007\)](#) which is inconsistent with the higher-order moments of the distribution of price changes, a key determinant of the real effects of monetary policy shocks. This model therefore predicts small real effects of monetary policy shocks, at odds with the VAR evidence ([Christiano et al., 2005](#)).⁵ In contrast, the model we study reproduces the higher-order moments of the distribution of price changes but predicts that the fraction of price changes is nearly constant. We therefore conclude that the single-product menu cost model cannot simultaneously reproduce the distribution of micro-price changes, a key determinant of the real effects of monetary policy, and the extent to which the fraction of price changes comoves with inflation, a key determinant of how non-linear these effects are.

We also characterize the size of the menu costs and losses from misallocation from price dispersion. Our calibration, which features moderate micro strategic complementarities, requires menu costs that represent 8.3% of sales, much larger than the 1% direct estimates in the literature ([Levy et al., 1997](#), [Zbaracki et al., 2004](#)). This calibration also predicts implausibly large losses from misallocation: menu costs reduce aggregate productivity by 20%. We show that even though these shortcomings can be remedied by assuming away strategic complementarities, the model still fails to reproduce the comovement between inflation and the fraction of price changes.

⁴See, for example, [Golosov and Lucas \(2007\)](#), [Gagnon \(2009\)](#), [Nakamura et al. \(2018\)](#), [Alvarez et al. \(2018\)](#) and [Alvarez and Lippi \(2022\)](#).

⁵We illustrate this point in a companion paper, [Blanco et al. \(2024\)](#).

We next develop a model that overcomes these shortcomings. We build on a standard multi-product menu cost model in which each firm sells a continuum of products, each subject to idiosyncratic Gaussian quality shocks, and there are economies of scope in price adjustment in that the firm can change the entire menu of its prices by paying a single menu cost. Two state variables are now necessary to summarize the history of shocks experienced by a firm: the firm's price gap, a weighted average of its product-level price gaps, as well as the duration of the firm's price spells. The latter determines the amount of within-firm misallocation: the older prices are, the larger the misallocation, and thus the larger the losses from leaving prices unchanged. On its own, this model has the same shortcomings as the single-product model.

We therefore introduce two ingredients that allow the model to reproduce the comovement between inflation and the fraction of price changes, while remaining consistent with the distribution of price changes. These ingredients decrease the misallocation from price dispersion within the firm, reducing the importance of idiosyncratic shocks and elevating that of aggregate shocks in repricing decisions. First, individual products sold by a given firm are imperfect substitutes. Our notion of a product is a collection of highly substitutable goods that are subject to correlated shocks. For example, various flavors of tea sold at Starbucks are highly substitutable but also face correlated shocks, so we think of tea at Starbucks as a product that is distinct from pastries, another product sold at Starbucks. Because a firm's products are imperfect substitutes, the losses that the firm faces from its inability to change prices in response to idiosyncratic shocks are small. Second, strategic complementarities, which arise due to decreasing returns to scale, are at the firm, not at the product level. Specifically, there is a firm-specific factor of production that is fixed at the firm level, but perfectly mobile across the products the firm sells, which further reduces the losses from price gap dispersion induced by idiosyncratic shocks inside the firm.

Our multi-product model reproduces the relationship between the fraction of price changes and inflation in the data. As in the data, the extensive margin of price changes accounts for half of the fluctuations in inflation when inflation is high. This is because our model is able to generate large dispersion in price changes with a narrow inaction region, so firms are more responsive to shocks that generate inflation. Moreover, even though our model features moderate strategic complementarities, it requires small menu costs to reproduce the micro-price statistics, in line with the 1% estimates in the data, and implies small losses from misallocation due to inefficient price dispersion.

We further clarify the mechanism of our model by zooming in on a special case in which the elasticity of substitution between the products sold by a firm is zero. This special case is identical to a single-product model provided we adjust the trend money growth rate to ensure that the firm’s price gap drifts at the same rate in both models. Even though these two models have identical implications for the distribution of firm price gaps, decision rules, and aggregate outcomes, the single-product equivalent generates a much smaller dispersion in price changes. Thus, our multi-product model behaves identically to a single-product model with narrower inaction regions, which implies that firm repricing decisions are much more sensitive to aggregate shocks.

We use our model to revisit the real effects of monetary shocks of various sizes. We show that in our model impulse responses are very different than those predicted by existing models. In particular, while the two models respond similarly to small shocks, in our model output responses are non-linear because the fraction of price changes increases rapidly with the size of the shock. In contrast, in existing models the output response scales linearly with the size of the shock, even for monetary shocks as large as 15%, because the fraction of price changes varies little with the size of the shock. In addition, the effects of monetary policy changes are asymmetric in our model: expansionary shocks are more inflationary because asymmetries in the profit function imply that prices are more likely to increase than to fall. We summarize our findings by tracing out the Phillips curve implied by monetary shocks. In our model, in contrast to existing models, the Phillips curve is highly non-linear and even vertical at inflation rates exceeding 5%.

2 Motivating Evidence

This section uses micro price data from the UK to show that the fraction of price changes increases with inflation, and that this is important for inflation fluctuations when inflation is relatively high. Though these facts have been previously documented for other countries, we use the evidence for the UK to quantitatively evaluate the ability of menu cost models to reproduce this pattern. We focus on data for individual sectors rather than the aggregate because inflation is considerably more volatile in individual sectors. Sectoral variation thus allows us to better understand the relationship between inflation and the fraction of price changes but, as we discuss in Section 3, none of our findings hinge on our focus on sectoral variation.

2.1 Data

We use the data that underlie the construction of the Consumer Price Index (CPI) in the UK. The data are collected by the United Kingdom Office for National Statistics (ONS). We use publicly available monthly product-level price quotes from January 1996 to August 2022. Goods and services are classified into 71 classes following the Classification of Individual Consumption by Purpose (COICOP 6). We exclude centrally-collected items, which account for approximately 26% of total consumer expenditures, as well as the energy categories. See Appendix A for details.

In computing price statistics, we use *regular price* series constructed by filtering V-shaped sales that last less than three months.⁶ Kehoe and Midrigan (2015) show that in theory temporary price changes do not contribute much to inflation dynamics. We corroborate their argument empirically by showing that excluding V-shaped sales from the calculation of inflation does not visibly change its time path.

To that end, consider the following decomposition of inflation. Let p_{it} be the price of good i and ω_{it} the weight of that good in the CPI. Aggregate inflation is then

$$\pi_t = \sum_{i \in \mathcal{A}_t} \omega_{it} \log p_{it}/p_{it-1},$$

where $\mathcal{A}_t = \mathcal{R}_t \cup \mathcal{S}_t$ is the set of goods that experience a price change in period t , \mathcal{R}_t is the set of goods that experience a regular price change and \mathcal{S}_t is the set of goods who experience a price change associated with a V-shaped sale. We construct an alternative inflation series based on regular price changes by calculating

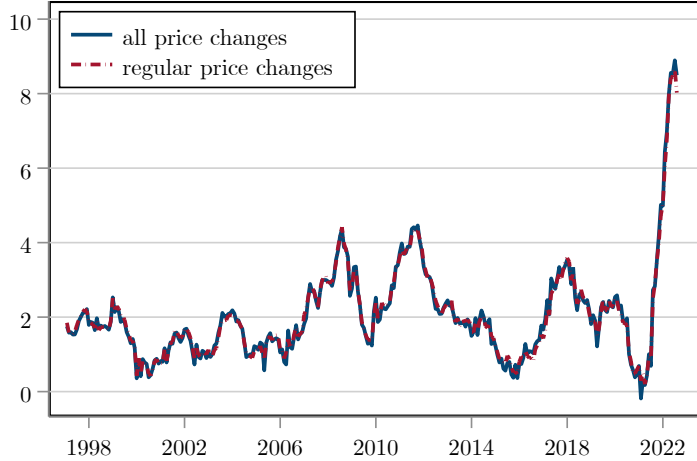
$$\pi_t^R = \sum_{i \in \mathcal{R}_t} \omega_{it} \log p_{it}/p_{it-1},$$

and thus excluding price changes that either initialize or end a V-shaped sale.

Figure 1 compares the inflation series computed using all price changes with that computed using only regular price changes. The figure reports the cumulative inflation in the previous 12 months, that is, the year-to-year percent change in the consumer price index. The two series are nearly indistinguishable, consistent with the theoretical predictions of Kehoe and Midrigan (2015). Motivated by this observation, from now on we focus our analysis on regular price changes only.

⁶We define V-shaped sales as temporary price cuts that return exactly to the original level.

Figure 1: Inflation Calculated Using on All vs. Regular Price Changes



2.2 Inflation and the Fraction of Price Changes

To assess the role of the extensive margin of price changes for inflation dynamics, we follow [Klenow and Kryvtsov \(2008\)](#) in decomposing movements in inflation into an extensive margin component that captures changes in the fraction of price adjustments and an intensive margin that captures movements in the average price change of firms that adjust. Specifically, letting $n_t(s)$ denote the fraction of products in sector s that experience a change in their regular price in period t and $\Delta_t(s)$ denote the average price change conditional on adjustment, we have⁷

$$\pi_t(s) = \Delta_t(s)n_t(s).$$

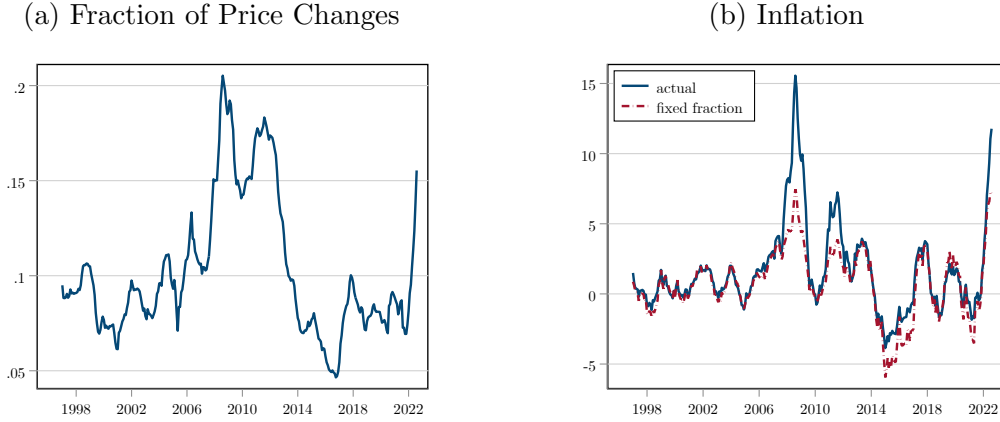
We gauge the role of the extensive margin by constructing a counterfactual inflation series that replaces the observed fraction of price changes $n_t(s)$ with that sector's average fraction of price changes, $\bar{n}(s) = \frac{1}{T} \sum_t n_t(s)$. That is, we calculate

$$\pi_t^c(s) = \Delta_t(s)\bar{n}(s)$$

and compare the dynamics of the actual inflation series $\pi_t(s)$ with the counterfactual inflation $\pi_t^c(s)$ that shuts down movements in the fraction of price changes. Throughout the paper, to mitigate the concern that our results are driven by sampling noise,

⁷All statistics are weighted using item-level consumption expenditure weights.

Figure 2: Inflation Decomposition: Bread and Cereals



Notes: The left panel plots the fraction of price changes $n_t(s)$. The right panel plots $\pi_t(s) = \Delta_t(s)n_t(s)$ and $\pi_t^c(s) = \Delta_t(s)\bar{n}(s)$.

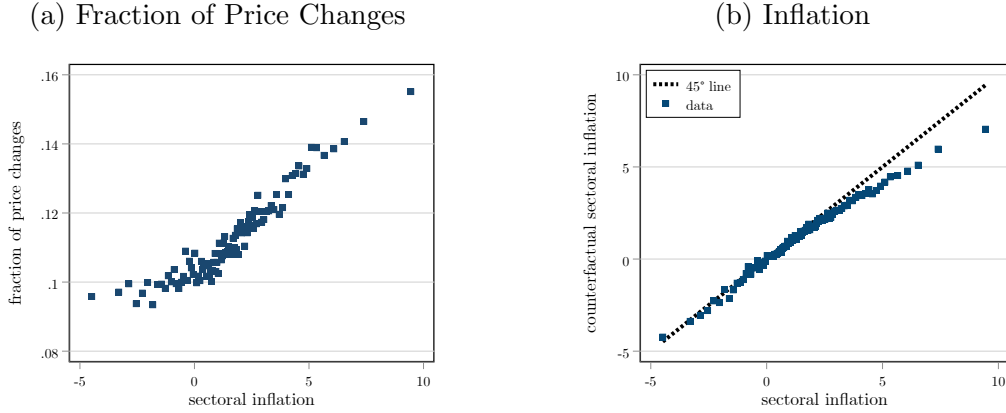
we measure the fraction of price changes as the average monthly fraction of regular price changes in the previous 12 months and inflation as the year-to-year percent change in the sectoral price index.

We first illustrate this decomposition in Figure 2 which shows the fraction of price changes (left panel) and the two inflation series (right panel) for a specific sector – “Bread and Cereals.” The fraction of price changes fluctuates substantially over time, ranging from 5% to 20%. Consequently, the extensive margin of adjustment accounts for a sizable share of movements in inflation in this sector, especially during the 2007-2008 world food crisis, when actual inflation exceeded 15%, while counterfactual inflation only increased to 7%.

Figure 3(a) documents these patterns more systematically by presenting a binned scatterplot of the fraction of price changes in a given sector against sectoral inflation rates pooling data from all sectors and weighting each by its expenditure share. We include sector fixed effects so our results capture high-frequency variation in sectoral inflation rates, not average differences across sectors. The figure shows that the fraction of price changes increases systematically with inflation. For example, as inflation increases from 0 to 10%, the fraction of price changes increases by 6 percentage points, from 10 to 16%. This increase is comparable to that in the U.S. time-series data (Nakamura et al., 2018, Blanco et al., 2024).

To gauge the extent to which these movements in the fraction of price changes

Figure 3: Inflation and the Fraction of Price Changes



Notes: The plot controls for sectoral fixed effects and weights individual sectors by their expenditure share. We measure the fraction of price changes as the 12-month moving average of the fraction of monthly price changes in the preceding year and inflation as the year-to-year percent change in the sectoral price index.

matter for inflation dynamics, Figure 3(b) shows a binned scatterplot of the counterfactual inflation series that keeps the fraction of price changes constant at its historical average against realized inflation. At low levels of inflation the extensive margin accounts for little of the movements in inflation: the counterfactual inflation series increases one-for-one with actual inflation. In contrast, for high inflation, above 4%, ignoring the extensive margin systematically underpredicts inflation.

Table 1 further corroborates these patterns. The slope coefficient in a regression of $\pi_t^c(s)$ on $\pi_t(s)$ is equal to 0.80 for the entire sample, but falls to 0.48 when we restrict attention to periods in which sectoral inflation exceeds its 75th percentile. We therefore conclude that at high rates of inflation, the extensive margin accounts for half of the changes in inflation.

Table 1: Slope of $\pi_t^c(s)$ on $\pi_t(s)$

all observations	0.80
$\pi_t(s) > 75^{th}$ pct.	0.48

Notes: We compute the slope coefficients using an OLS regression that weights observations for each sector using that sector's expenditure weights.

3 Single-Product Menu Cost Model

We first show that existing menu cost models, when calibrated to match the distribution of micro-price changes, cannot reproduce the comovement between inflation and the fraction of price changes in the data. In addition, parameterizations with moderate micro-level strategic complementarities in price setting result in implausibly large menu costs and losses from misallocation due to price dispersion.

Because we study the comovement between sectoral inflation and the fraction of price changes, we consider an economy with a continuum of measure one of ex-ante identical sectors. The output of each sector is used to produce a final good. Each sector consists of a continuum of measure one of firms, each producing a differentiated variety. Firms are subject to idiosyncratic and sectoral shocks. We follow [Golosov and Lucas \(2007\)](#) in assuming that preferences are logarithmic in consumption and linear in hours worked, which allows us to characterize inflation dynamics in each sector in isolation, greatly simplifying computations.

3.1 Consumers

A representative consumer has preferences over consumption and derives disutility from work. The consumer maximizes life-time utility, given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\log c_t - h_t),$$

subject to

$$M_t + \frac{1}{1+i_t} B_t = W_t h_t + D_t + M_{t-1} - P_{t-1} c_{t-1} + B_{t-1} + T_t,$$

where c_t is consumption, h_t is hours worked, P_t is the aggregate nominal price level, M_t is the money supply, B_t is the amount of government bonds, D_t denotes profits and T_t represents government transfers. We assume a constant money growth rate g_m and a cash-in-advance constraint

$$P_t c_t \leq M_t.$$

The optimal labor supply choice implies that $W_t = P_t c_t = M_t$.

3.2 Technology

We next describe the assumptions we make on technology.

3.2.1 Final Goods Producers

Final output is produced by aggregating sectoral output $y_t(s)$ using a technology

$$y_t = \exp \left(\int \log y_t(s) ds \right). \quad (1)$$

The final output is used for consumption only, so the aggregate resource constraint is $c_t = y_t$. The aggregate price index P_t satisfies

$$P_t = \exp \left(\int \log P_t(s) ds \right),$$

where $P_t(s)$ is the price index in sector s . The assumption of a unit elasticity of substitution across sectors implies that sectoral expenditures are proportional to nominal spending, the money supply and nominal wages

$$P_t(s)y_t(s) = P_t y_t = M_t = W_t.$$

3.2.2 Intermediate Goods Producers

Firm f in sector s produces output using a labor-only technology with decreasing returns to scale determined by $\eta \leq 1$

$$y_t(f, s) = e_t(s) z_t(f, s) l_t(f, s)^\eta,$$

where $e_t(s)$ is productivity in sector s , $z_t(f, s)$ is the quality of the product of firm f in that sector and $l_t(f, s)$ the amount of labor the firm uses in production. Sectoral output is produced using a CES aggregator with elasticity of substitution σ

$$y_t(s) = \left(\int \left(\frac{y_t(f, s)}{z_t(f, s)} \right)^{\frac{\sigma-1}{\sigma}} df \right)^{\frac{\sigma}{\sigma-1}}. \quad (2)$$

In addition to affecting the firm's productivity, the quality index $z_t(f, s)$ also affects demand. If prices were flexible, firms would respond to an increase in $z_t(f, s)$

by reducing prices one-for-one, leaving quality-adjusted prices, $z_t(f, s)P_t(f, s)$, and firm revenues unchanged. These quality shocks therefore change firms' desired prices and have the advantage of not requiring keeping track of $z_t(f, s)$ as a state variable.⁸ We assume that the logarithms of $e_t(s)$ and $z_t(f, s)$ follow independent random walk processes with i.i.d. innovations drawn from mean zero normal distributions with standard deviation σ_e and σ_z .⁹

Letting $P_t(f, s)$ denote an individual firm's price, the demand function for the firm's output is given by

$$y_t(f, s) = z_t(f, s) \left(\frac{z_t(f, s) P_t(f, s)}{P_t(s)} \right)^{-\sigma} y_t(s),$$

where the price index in sector s is

$$P_t(s) \equiv \int P_t(f, s) \frac{y_t(f, s)}{y_t(s)} df = \left(\int (z_t(f, s) P_t(f, s))^{1-\sigma} df \right)^{\frac{1}{1-\sigma}}.$$

3.3 Firm Objective

The aggregate implications of menu cost models are shaped by the distribution of price changes (Caballero and Engel, 2007, Midrigan, 2011, Alvarez et al., 2016). We thus assume a flexible menu cost specification that allows the model to reproduce key moments of the distribution of price changes in the data. Specifically, we follow Nakamura and Steinsson (2010) and assume that with probability $1 - \lambda$ firms can change their price for free. With probability λ a price change requires a fixed cost $\xi_t(f, s)$. We assume that the fixed cost is an i.i.d. draw from a uniform distribution $U[0, \bar{\xi}]$ which gives rise to a smooth adjustment hazard (Costain and Nakov, 2011,

⁸An alternative approach, which we discuss in Appendix C, is to assume that $z_t(f, s)$ only affects productivity and scale the menu costs accordingly. An additional advantage of our approach is that the model does not require extremely large idiosyncratic shocks to reproduce the dispersion of price changes. See Klenow and Willis (2016) for an illustration of this problem and Aruoba et al. (2023) for an alternative resolution.

⁹As Karadi and Reiff (2019) show, a model with fat-tailed idiosyncratic shocks can match the distribution of micro-price changes and also predicts a strong response of the fraction of price changes to aggregate shocks. They also show that micro-price data alone cannot identify the shape of the distribution of shocks by presenting several models with fat-tailed shocks that match the same micro moments yet have very different aggregate implications. Because of this identification issue, we follow most of the menu cost literature and assume Gaussian shocks.

Alvarez et al., 2021, Costain et al., 2022).¹⁰

The firm's objective is to maximize the present value of its profits, given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{P_t c_t} \left[(1 + \tau) P_t(f, s) y_t(f, s) - W_t \left(\frac{y_t(f, s)}{e_t(s) z_t(f, s)} \right)^{\frac{1}{\eta}} - \xi_t(f, s) W_t \mathbb{I}_t(f, s) \right],$$

where $\tau = 1/(\sigma - 1)$ is an output subsidy that corrects the markup distortion that would arise even in the absence of menu costs. Letting $\mathbb{I}_t(f, s)$ denote a price adjustment indicator, the last term represents the menu cost of changing prices, denominated in units of labor.

In order to write the firm's problem recursively, we next express its objective as a function of its *price gap*: the ratio of its actual price relative to what the firm would charge under flexible prices. To that end we first define the real marginal cost index in sector s as

$$a_t(s) \equiv \frac{W_t}{P_t(s) y_t(s)} \left(\frac{y_t(s)}{e_t(s)} \right)^{\frac{1}{\eta}},$$

and define the firm's price gap as

$$x_t(f, s) \equiv \bar{a}^\eta \frac{e_t(s) z_t(f, s) P_t(f, s)}{M_t},$$

where \bar{a} is the steady state value of $a_t(s)$. Similarly, we define the sectoral price gap as the CES weighted average of firm price gaps

$$x_t(s) \equiv \left[\int x_t(f, s)^{1-\sigma} df \right]^{\frac{1}{1-\sigma}} = \bar{a}^\eta \frac{e_t(s) P_t(s)}{M_t}.$$

This sectoral price gap is equal to one in steady state, and is inversely related to the sector's real marginal cost

$$x_t(s) = \left(\frac{a_t(s)}{\bar{a}} \right)^{-\eta}. \quad (3)$$

¹⁰In Appendix D we consider a richer specification in which ξ_{it} is drawn from $F(\xi) = \left(\frac{\xi}{\bar{\xi}} \right)^\nu$, a distribution that collapses to the uniform when $\nu = 1$ and is degenerate at $\bar{\xi}$ when $\nu \rightarrow \infty$. Calibrating this richer specification to the distribution of micro-price changes yields a value of ν equal to 1.005, so for parsimony we focus on the uniform distribution. A model with a degenerate distribution of menu costs cannot reproduce the unimodal distribution of micro-price changes in the data, but rather predicts a multi-modal distribution with sharp spikes near the adjustment thresholds, as illustrated in Blanco et al. (2024).

Moreover, under flexible prices $x_t(f, s) = x_t(s) = 1$ and $a_t(s) = \eta$.

With this notation in place, we can write the firm's objective as

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(x_t(s)^{\sigma-1} \left[(1 + \tau) x_t(f, s)^{1-\sigma} - \bar{a} x_t(s)^{\left(\frac{1}{\eta}-1\right)(\sigma-1)} x_t(f, s)^{-\frac{\sigma}{\eta}} \right] - \xi_t(f, s) \mathbb{I}_t(f, s) \right). \quad (4)$$

As in [Burstein and Hellwig \(2008\)](#), decreasing returns to scale introduce a micro-level strategic complementarity in price setting, dampening the response of individual prices to aggregate and sectoral shocks. To see this, notice that the price gap that maximizes the firm's flow profits is

$$x_t(f, s) = \left(\frac{\bar{a}}{\eta} \right)^{\frac{1}{1+\sigma\left(\frac{1}{\eta}-1\right)}} x_t(s)^{\frac{(\sigma-1)\left(\frac{1}{\eta}-1\right)}{1+\sigma\left(\frac{1}{\eta}-1\right)}}.$$

The exponent of $x_t(s)$ determines the strength of strategic complementarities, that is, the extent to which an individual firm's price depends on the price of its competitors. The lower η is or the higher σ , the stronger are strategic complementarities. To see why this is the case, consider an increase in the money supply. If most firms do not adjust their prices, $x_t(s)$ falls. A firm that resets its price recognizes that if it were to increase it, it would experience an output drop, more so the more elastic demand is, that is, the larger is σ . This drop in output would reduce the firm's marginal cost, more so the lower the returns to scale, that is, the lower is η . This decrease in marginal cost dampens the firm's desired price increase.¹¹

Equation (4) shows that the problem of a firm in a given sector only depends on current and future sectoral price gaps $x_t(s)$ and not on other sectoral and aggregate variables. To see how the sectoral price gap $x_t(s)$ is determined in equilibrium, let

$$\hat{x}_t(f, s) \equiv \bar{a}^\eta \frac{e_t(s) z_t(f, s) P_{t-1}(f, s)}{M_t}$$

denote the firm's individual state variable: its last period's price scaled by the money

¹¹An alternative way to introduce micro-level strategic complementarities is through variable markups that depend on firm market shares. As [Alvarez et al. \(2022b\)](#) show, up to a second-order approximation, these approaches are equivalent.

supply and productivity. The firm's state evolves according to

$$\hat{x}_t(f, s) = x_{t-1}(f, s) \frac{e_t(s)}{e_{t-1}(s)} \frac{z_t(f, s)}{z_{t-1}(f, s)} \frac{M_{t-1}}{M_t}. \quad (5)$$

If the firm does not adjust, its price gap is $x_t(f, s) = \hat{x}_t(f, s)$. If the firm adjusts its price, it resets the price gap to $x_t(f, s) = x_t^*(s)$, which is common to all firms that adjust. Letting $v_t^a(s)$ denote the value of adjusting the price and $v_t^n(\hat{x}, s)$ the value of not adjusting, the firm adjusts with probability

$$h_t(\hat{x}, s) = 1 - \lambda + \lambda \min \left\{ \frac{v_t^a(s) - v_t^n(\hat{x}, s)}{\bar{\xi}}, 1 \right\}.$$

Finally, letting $F_t(\hat{x}; s)$ denote the distribution of firms, the sectoral price gap satisfies

$$x_t(s) = \left(\int [h_t(\hat{x}; s) x_t^*(s)^{1-\sigma} + (1 - h_t(\hat{x}; s)) \hat{x}^{1-\sigma}] dF_t(\hat{x}; s) \right)^{\frac{1}{1-\sigma}}.$$

We use the [Krusell and Smith \(1998\)](#) approach to characterize how $x_t(s)$ evolves over time as a function of a single moment of the distribution of $F_t(\hat{x}; s)$. This approach works well in this setting, with an R^2 in the perceived law of motion in excess of 0.999. See [Appendix B](#) for details.

3.4 Losses from Misallocation

Menu costs generate inefficient price dispersion and misallocation. To see this, let

$$l_t(s) = \int l_t(f, s) df$$

denote the amount of labor firms in sector s use in production. We can derive a sectoral production function

$$y_t(s) = e_t(s) \phi_t(s) l_t(s)^\eta,$$

where

$$\phi_t(s) = \left(\int \left(\frac{x_t(f, s)}{x_t(s)} \right)^{-\frac{\sigma}{\eta}} df \right)^{-\eta}$$

captures the losses from misallocation. When prices are flexible $\phi_t(s) = 1$. More generally, dispersion in relative prices reduces $\phi_t(s)$, more so the larger σ/η is. Intuitively, efficiency requires that all firms in a given sector use the same amount of labor, $l_t(f, s) = l_t(s)$. The more elastic demand is, or the stronger the decreasing returns, the larger the dispersion in firm employment implied by a given amount of relative price dispersion, and thus the larger the productivity losses from misallocation.

3.5 Parameterization

Table 2 reports the result of the parameterization. A period is a month and we set the discount factor β to an annualized value of 0.96. We start with an economy with moderate micro-level strategic complementarities by setting σ equal to 6 and the elasticity of labor in the production function η to $2/3$. We show however that even a model without strategic complementarities cannot reproduce the comovement between inflation and the fraction of price changes in the data.

In our baseline calibration, reported in the column labeled “Free price changes,” we choose the money growth rate g_m , the standard deviation of idiosyncratic shocks σ_z , the probability of free price changes $1 - \lambda$, the upper bound of the menu cost distribution $\bar{\xi}$ and the standard deviation of sectoral shocks σ_e to match key moments of the distribution of non-zero price changes in the UK data. Specifically, we target a monthly fraction of price changes of 0.116, a mean price change of 0.018, a standard deviation of price changes of 0.188, a kurtosis of price changes of 3.609, and a volatility of sectoral inflation of 0.029.

We report the data moments in the first column of Panel A in Table 2. We compute these statistics using regular price changes. To mitigate measurement error concerns, we drop the top and bottom 2% of the price change distribution.¹² We calculate the sectoral fraction of price changes in the data as the harmonic weighted average of the fraction of price changes of individual product categories (items) that belong to that sector. To mitigate the concern that dispersion in the size of price changes is driven by ex-ante heterogeneity, we follow Klenow and Kryvtsov (2008) in standardizing the distribution of price changes by the respective item-level mean and variance. See Appendix A for details.

As Panel A shows, the model matches the targeted moments perfectly. As Panel B

¹²See Appendix A for price statistics computed using alternative truncations.

Table 2: Parameterization of Single-Product Model

A. Moments			
	Data	Free price changes	No free price changes
I. Targeted			
fraction Δp	0.116	0.116	0.116
mean Δp	0.018	0.018	0.018
std. dev. Δp	0.188	0.188	0.188
kurtosis Δp	3.609	3.609	<i>1.998</i>
std dev. $\pi_t(s)$	0.029	0.029	0.029
II. Not targeted			
<i>distribution of Δp</i>			
10 th percentile	0.018	0.021	0.079
25 th percentile	0.045	0.053	0.116
50 th percentile	0.104	0.118	0.165
75 th percentile	0.204	0.214	0.220
90 th percentile	0.334	0.315	0.272
B. Calibrated Parameter Values			
		Free price changes	No free price changes
g_m	mean money growth rate	0.020	0.021
σ_z	s.d. idios. shocks	0.064	0.064
λ	1 - prob. free change	0.911	
$\bar{\xi}$	upper bound menu cost	39.08	3.376
σ_e	s.d. sectoral shocks	0.011	0.010

Note: The money growth rate is annualized and the menu cost is relative to average sales. In the models without free price changes λ is set to 1. In the “No free price changes” calibration we do not target the kurtosis and italicize its implied value. In all calibrations we set $\sigma = 6$, $\eta = 2/3$ and $\beta = 0.96$ (annualized).

shows, the model implies a high dispersion of idiosyncratic shocks relative to sectoral shocks, $\sigma_z = 0.064$ vs. $\sigma_e = 0.011$, a large upper bound on the distribution of menu costs, $\bar{\xi} = 39$ times average monthly sales, and a large probability of free price changes, $1 - \lambda = 0.089$. Free price changes thus account for 77% (0.089/0.116) of all

price changes. As shown in Panel A, the model is also able to replicate the distribution of the size of price changes, which we do not explicitly target. For example, the 10th percentile is 0.018 in the data and 0.021 in the model, whereas the 90th percentile is 0.334 in the data and 0.315 in the model. Thus the model generates nearly as many small and large price changes as in the data. In Appendix F we plot the histogram of price changes predicted by the model against the data.

3.6 Model Implications

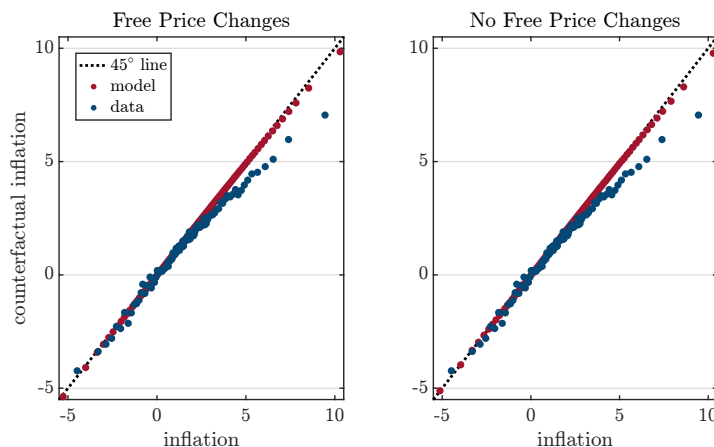
We next evaluate the model’s ability to reproduce the relationship between inflation and the fraction of price changes in the data. We also report the model’s implications for the size of menu costs and losses from misallocation.

Importance of the Extensive Margin. Figure 4 assesses the model’s ability to reproduce the empirical comovement between actual inflation $\pi_t(s)$ and the counterfactual inflation series $\pi_t^c(s)$ that shuts down fluctuations in the fraction of price adjustments. As in the data, we compute inflation as the year-to-year percent change in the sectoral price level, and we compute the fraction of price changes as the 12-month moving average in the preceding year. The left panel of the figure shows that, in contrast to the data, $\pi_t(s)$ and $\pi_t^c(s)$ comove nearly one-for-one in the model. Thus, fluctuations in the fraction of price changes play almost no role in driving inflation dynamics, even when sectoral inflation rates are high. In particular, as inflation increases from 0% to 10% in our simulations, the monthly fraction of price changes only increases by 0.7 percentage points (from 11.3% to 12%), nowhere near the 6 percentage points increase in the data.

Panel A of Table 3 reinforces this point in the second column. The slope coefficients from regressing $\pi_t^c(s)$ on $\pi_t(s)$ are close to one, and fall to only 0.94 when we restrict the sample to periods when sectoral inflation exceeds its 75th percentile. The patterns are at odds with the data, where these coefficients fall below one-half at high rates of inflation.

Size of Menu Costs and Misallocation. We also show that the model requires implausibly large menu costs to reproduce the data and predicts implausibly large losses from misallocation. As the second column of Panel A in Table 3 shows, the average amount of resources spent on adjusting prices in a given period is equal to

Figure 4: Importance of the Extensive Margin



Notes: The figures are based on the calibration reported in Table 2. We compute inflation as the year-to-year percent change in the sectoral price level, and the fraction of price changes as the 12-month moving average in the preceding year.

8.3% of average firm sales, a number much larger than the estimates reported in [Levy et al. \(1997\)](#) and [Zbaracki et al. \(2004\)](#), which are in the neighborhood of 1% of firm revenues. Moreover, the model predicts that aggregate productivity is 19.9% lower than under flexible prices. This number is comparable to the estimates of misallocation reported by [De Loecker et al. \(2020\)](#) and [Baqae and Farhi \(2018\)](#) which encompass *all* distortions that lead to misallocation (taxes, factor adjustment costs, financial frictions, markup variation arising from differences in demand elasticities etc.), as well as differences in production function elasticities across producers.¹³ It is unlikely that menu costs alone account for all observed misallocation in the data.

3.7 Understanding the Results

To understand what drives our findings above, we consider in Appendix E a continuous time version of our model. For simplicity, we assume a fixed, rather than random menu cost, and consider a quadratic approximation to the firm’s objective function and no growth in the money supply, as in [Alvarez and Lippi \(2014\)](#). We use this setting to derive three formulas that describe i) how the fraction of price changes responds to one-time sectoral shocks,¹⁴ as well as the determinants of ii) the

¹³See, for example, [Foster et al. \(2022\)](#) for a discussion.

¹⁴Our technology and preference assumptions imply that the responses of sectoral prices to a sectoral shock are identical to the responses of aggregate prices to an aggregate shock.

Table 3: Model Implications

	Data	$\sigma = 6$		$\sigma = 3$	
		$\eta = 2/3$	$\eta = 1$	$\eta = 2/3$	$\eta = 1$
A. Free Price Changes					
slope of $\pi_t^c(s)$ on $\pi_t(s)$					
all observations	0.80	0.99	0.99	0.99	0.99
$\pi_t(s) > 75^{th}$ pct.	0.48	0.94	0.94	0.96	0.96
menu costs/sales		0.083	0.020	0.023	0.009
losses from misallocation		0.199	0.053	0.071	0.029
B. No Free Price Changes					
slope of $\pi_t^c(s)$ on $\pi_t(s)$					
all observations	0.80	0.98	0.97	0.98	0.98
$\pi_t(s) > 75^{th}$ pct.	0.48	0.94	0.93	0.94	0.93
menu costs/sales		0.065	0.017	0.022	0.009
losses from misallocation		0.089	0.023	0.030	0.012

size of menu costs and iii) misallocation from price dispersion. We then illustrate these theoretical predictions using the quantitative model introduced above.

Importance of the Extensive Margin. As Alvarez and Lippi (2014) show, the response of the fraction of price changes to a one-time shock of size δ is

$$\Delta n = \frac{1-l}{2} \frac{\delta^2}{\mathbb{V}[\Delta p]} + O(\delta^4), \quad (6)$$

where l is fraction of price changes that are free and $\mathbb{V}[\Delta p]$ is the variance of price changes conditional on adjustment.¹⁵ If the size of the sectoral shock is small relative to the variance of price changes, the fraction of price changes responds little. In our calibration, the standard deviation of sectoral shocks is 0.011 and that of price changes is 0.188, so $\delta^2/\mathbb{V}[\Delta p] \approx 0.003$. Therefore, the fraction of price changes responds little. This response is further dampened by the presence of free price changes (in our calibration $l = 0.089/0.116 \approx 3/4$), but since $\delta^2/\mathbb{V}[\Delta p]$ is nearly zero, the fraction of

¹⁵See Appendix E for a derivation.

price changes responds little to aggregate shocks even absent free price changes.

To further illustrate this, we calibrate a version of the model in which we set $\lambda = 1$. As the column labeled “No free price changes” in Table 2 shows, this model can no longer reproduce the kurtosis of price changes and generates a distribution of $|\Delta p|$ that is much less dispersed than in the data. As shown in the right panel of Figure 4 and Panel B of Table 3, this model also predicts that most fluctuations in inflation are driven by the intensive margin of price changes, not changes in the fraction of prices that adjust. Once again, this is because even in this version of the model, sectoral shocks are small relative to the idiosyncratic shocks needed to generate the dispersion in price changes observed in the data.

Relationship to Existing Work. The result above may seem to contradict a number of existing studies that show that the menu cost model generates substantial fluctuations in the fraction of price changes.¹⁶ There is, in fact, no contradiction. These papers study a simple menu cost model in the tradition of Golosov and Lucas (2007) which is inconsistent with the higher-order moments of the distribution of price changes, a key determinant of the real effects of monetary policy shocks. This model therefore predicts small real effects of monetary policy shocks, at odds with the empirical evidence (Christiano et al., 2005). In contrast, the model we study in this section reproduces the higher-order moments of the distribution of price changes but predicts that the fraction of price changes is nearly constant.

An exception is the work of Cavallo et al. (2023) who study a model with a generalized adjustment hazard and can therefore match the distribution of price changes in the data. These authors, however, consider a very large 20% aggregate shock, an order of magnitude larger than the shocks we consider here that are necessary to reproduce the volatility of inflation. In their paper, with a 20% shock, the term $\delta^2/\mathbb{V}[\Delta p]$ in equation (6) is much larger than in our setting. Thus, while it is indeed the case that in menu cost models the fraction of price changes responds significantly to very large shocks, it does not in response to more moderately sized shocks that do trigger movements in the fraction of price changes in the data.

¹⁶See, for example, Golosov and Lucas (2007), Gagnon (2009), Nakamura et al. (2018), Alvarez et al. (2018), Alexandrov (2020) and Alvarez and Lippi (2022) who study the relationship between inflation and the fraction of price changes either in the cross-section, by comparing environments with different levels of steady-state inflation, or in the time-series, as we do, by subjecting the economy to aggregate shocks.

We discuss these points at length using U.S. data in a companion paper, [Blanco et al. \(2024\)](#). There we identify the sequence of aggregate shocks that allow the model to exactly reproduce the path of inflation and show that even though the [Golosov and Lucas \(2007\)](#) model generates fluctuations in the fraction of price changes, it predicts too little dispersion in the size of price changes and therefore little monetary non-neutrality. In contrast, versions of the model that reproduce the dispersion in the size of price changes and generate stronger monetary non-neutrality predict a nearly constant fraction of price changes.

Size of Menu Costs and Misallocation. We also show in [Appendix E](#) that the total expenditures on menu costs in a given period, relative to total revenues,¹⁷ is

$$\mathcal{C} \approx (\sigma - 1) \left[1 + \sigma \left(\frac{1}{\eta} - 1 \right) \right] \frac{\mathbb{V}[\Delta p] \Psi(\mathbb{K}[\Delta p])}{12}, \quad (7)$$

where $\Psi(\cdot)$ is a hump-shaped function that satisfies $\Psi(1) = 1$ and $\Psi(6) = 0$. For our baseline parameterization, $\Psi(3.609) = 1.33$. We also show that the losses from misallocation from price dispersion are

$$\log(\phi) \approx -\frac{\sigma}{2} \left[1 + \sigma \left(\frac{1}{\eta} - 1 \right) \right] \frac{\mathbb{V}[\Delta p] \mathbb{K}[\Delta p]}{6}. \quad (8)$$

Both of these objects depend on the two moments we target in the calibration, the variance and the kurtosis of price changes, as well as the demand elasticity σ and the degree of returns to scale η . As discussed above, our choice of $\sigma = 6$ and $\eta = 2/3$, which generates micro-level strategic complementarities, implies large menu costs and losses from misallocation. In the last three columns of [Table 3](#) we consider alternative values of these two parameters which greatly reduce the calibrated menu costs and losses from misallocation.¹⁸ However, these alternative parameterizations imply weaker or no micro-level strategic complementarities and therefore smaller real effects of monetary shocks, at odds with the evidence ([Christiano et al., 2005](#)).

Importantly, though eliminating micro-level strategic complementarities reduces the size of menu costs and the losses from misallocation, it does not affect the importance of the extensive margin of price changes. This is consistent with equation

¹⁷See [Alvarez et al. \(2016\)](#) for a similar formula which expresses menu costs as a fraction of profits.

¹⁸We re-calibrate the models for each of these alternative parameterizations to match the same targets in [Table 2](#). See [Appendix F](#) for details.

(6) which shows that the response of the fraction of price changes to sectoral shocks does not depend on the demand elasticity σ and the returns to scale η .

To summarize, the single-product menu cost model, when calibrated to match the distribution of micro price changes in the data, cannot reproduce the relationship between inflation and the fraction of price changes. Moreover, parameterizations of the model with moderate micro-level strategic complementarities also imply very large menu costs and losses from misallocation from price dispersion. As we show below, these shortcomings also apply to the canonical multi-product menu cost model with economies of scope in price setting, but can be remedied by introducing features that reduce the degree of misallocation from price gap dispersion inside the firm.

4 Multi-Product Menu Cost Model

We next extend the model to a multi-product setting in which there are economies of scope in the price adjustment technology, as in [Midrigan \(2011\)](#) and [Alvarez and Lippi \(2014\)](#).¹⁹ Each firm sells a unit measure of products. Each product is subject to independent quality shocks. The firm can change the entire menu of prices by paying a random menu cost drawn from a uniform distribution. Since economies of scope in price setting allow us to match the large number of small price changes observed in the data, we no longer need to assume that some price changes are free.

We show that economies of scope, on their own, do not remedy the shortcomings discussed above. We therefore introduce two ingredients in the multi-product model, both of which reduce the misallocation from price dispersion within the firm and narrow the inaction region, thus allowing the model to reproduce the relationship between inflation and the fraction of price changes we document in the data. Additionally, the model implies small menu costs and losses from misallocation even in the presence of strategic complementarities.

The first ingredient is a nested CES aggregator in which the elasticity of substitution between the products sold by a given firm is lower than that across firms. Our notion of a product is a collection of highly substitutable goods, subject to correlated shocks. For example, we think of tea sold by Starbucks as representing a product because different flavors or sizes of tea are close substitutes that experience correlated

¹⁹See [Bhattarai and Schoenle \(2014\)](#) and [Bonomo et al. \(2022\)](#) for evidence on multi-product pricing.

shocks. In contrast, different pastries sold by Starbucks, while highly substitutable among themselves, are much less substitutable with tea. The second ingredient is that decreasing returns to scale arise due to a specific factor of production that is fixed at the firm level, but is perfectly mobile across the products a firm sells. This implies that there are decreasing returns and therefore strategic complementarities at the firm level, but the losses from misallocation within the firm are lower than in a model where the decreasing returns to scale are also at the product level.

Since the multi-product model shares many elements with the single-product model above, we only discuss the new ingredients that we introduce here.

4.1 Technology

The technology for producing the final good is the same as in equation (1) and that for producing sectoral output is

$$y_t(s) = \left(\int y_t(f, s)^{\frac{\sigma-1}{\sigma}} df \right)^{\frac{\sigma}{\sigma-1}}.$$

A firm produces a unit mass of products that are aggregated into a firm-level composite using

$$y_t(f, s) = \left(\int \left(\frac{y_{it}(f, s)}{z_{it}(f, s)} \right)^{\frac{\gamma-1}{\gamma}} di \right)^{\frac{\gamma}{\gamma-1}},$$

where γ is the elasticity of substitution between different products and $z_{it}(f, s)$ is the quality of product i , which follows a random walk process

$$\log z_{it+1}(f, s) = \log z_{it}(f, s) + \sigma_z \varepsilon_{it+1}^z(f, s),$$

where σ_z is the volatility of innovations and $\varepsilon_{it+1}^z(f, s)$ is an i.i.d. draw from a standard normal distribution.²⁰ The demand for an individual product is given by

$$y_{it}(f, s) = z_{it}(f, s) \left(\frac{z_{it}(f, s) P_{it}(f, s)}{P_t(f, s)} \right)^{-\gamma} y_t(f, s),$$

²⁰In an earlier draft, we considered an alternative specification with firm-level, in addition to product level quality shocks. While that economy also reproduces the relationship between inflation and the fraction of price changes, it also features stronger selection effects and thus less non-neutrality for small aggregate shocks. Since we would like to argue that the non-linearity in our model is not driven by selection effects, here we substantially reduce these by eliminating firm level shocks.

where the composite price of the bundle of products of firm f is

$$P_t(f, s) \equiv \int P_{it}(f, s) \frac{y_{it}(f, s)}{y_t(f, s)} di = \left(\int (z_{it}(f, s) P_{it}(f, s))^{1-\gamma} di \right)^{\frac{1}{1-\gamma}}.$$

Individual products are produced with a technology that uses labor and an input, say managerial, that is in fixed supply at the firm level but perfectly mobile across individual products. Specifically, letting $m_{it}(f, s)$ denote the amount of the fixed input used for product i , the production function is

$$y_{it}(f, s) = e_t(s) z_{it}(f, s) m_{it}(f, s)^{1-\eta} l_{it}(f, s)^\eta. \quad (9)$$

We normalize the supply of the fixed factor to 1, so the choice of $m_{it}(f, s)$ satisfies

$$\int m_{it}(f, s) di = 1.$$

This technology exhibits constant returns to scale at the individual product level but decreasing returns at the firm level. Assuming instead that the fixed factor is immobile across products, so that $m_{it}(f, s) = 1$, the technology in equation (9) also features decreasing returns to scale at the product level. As we show below, under this alternative assumption, the losses from misallocation within the firm are larger.

4.2 Firm Objective

We next discuss the problem of the firm. Its life-time value is

$$V_0(f, s) = \mathbb{E}_0 \sum_{t=0}^{\infty} \frac{\beta^t}{P_t c_t} \left[(1 + \tau) \int P_{it}(f, s) y_{it}(f, s) di - W_t l_t(f, s) - \xi_t(f, s) W_t \mathbb{I}_t(f, s) \right],$$

where $\mathbb{I}_t(f, s)$ is an indicator for whether the firm changes its menu of prices. Letting

$$x_{it}(f, s) = \bar{a}^\eta \frac{e_t(s) z_{it}(f, s) P_{it}(f, s)}{M_t}$$

denote the price gap of product i , the firm's price gap is

$$x_t(f, s) = \left(\int x_{it}(f, s)^{1-\gamma} di \right)^{\frac{1}{1-\gamma}} = \bar{a}^\eta \frac{e_t(s) P_t(f, s)}{M_t}.$$

Letting $l_t(f, s) = \int l_{it}(f, s) di$ denote the total amount of labor a firm uses in production, we can derive a firm-level production function

$$y_t(f, s) = e_t(s) \phi_t(f, s) l_t(f, s)^\eta,$$

where $\phi_t(f, s)$ is firm productivity which can fall below one because of losses from misallocation from price gap dispersion inside the firm. Notice that this firm-level production function features decreasing returns to scale, which generates strategic complementarities, as in the single-product menu cost model.

In our model with a mobile specific input, firm productivity is

$$\phi_t(f, s) = \left(\int \left(\frac{x_{it}(f, s)}{x_t(f, s)} \right)^{-\gamma} di \right)^{-1}.$$

In contrast, if the specific input is fixed at the product level, firm productivity is

$$\phi_t(f, s) = \left(\int \left(\frac{x_{it}(f, s)}{x_t(f, s)} \right)^{-\frac{\gamma}{\eta}} di \right)^{-\eta},$$

and is lower than in our model for a given dispersion in price gaps.

With this notation, the firm's objective can be expressed in terms of the firm-level price gap and the losses from misallocation as

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(x_t(s)^{\sigma-1} \left[(1 + \tau) x_t(f, s)^{1-\sigma} - \bar{a} x_t(s)^{\left(\frac{1}{\eta}-1\right)(\sigma-1)} x_t(f, s)^{-\frac{\sigma}{\eta}} \phi_t(f, s)^{-\frac{1}{\eta}} \right] - \xi_t(f, s) \mathbb{I}_t(f, s) \right).$$

Thus the firm's flow profits depend on the amount of misallocation from price gap dispersion inside the firm. The lower this misallocation is, the smaller are the profit losses from not adjusting prices, so the lower the incentives to adjust in response to idiosyncratic shocks.

Since the firm's objective only depends on its price gap $x_t(f, s)$, the productivity $\phi_t(f, s)$ and the sectoral price gap $x_t(s)$, we can write the firm's problem recursively by summarizing the distribution of the firm's price gaps with two idiosyncratic state variables. To derive these, consider a firm that does not adjust the prices $P_{it-1}(f, s)$

it inherits from the previous period. Its composite price index is then equal to

$$P_t(f, s) = \left(\int (z_{it}(f, s) P_{it-1}(f, s))^{1-\gamma} di \right)^{\frac{1}{1-\gamma}}.$$

Since $z_{it}(f, s)$ follows a geometric random walk with independent innovations, this composite price index evolves according to

$$P_t(f, s) = \exp \left((1 - \gamma) \frac{\sigma_z^2}{2} \right) P_{t-1}(f, s).$$

If $\gamma > 1$, the composite price drifts down over time at a rate that increases with the volatility of idiosyncratic shocks. Intuitively, the composite price index is a quantity-weighted average of individual product prices, so even though individual prices are constant, consumers reallocate demand towards products with cheaper quality-adjusted prices.

The first state variable we keep track of is then

$$\hat{x}_t(f, s) = \bar{a}^\eta \frac{e_t(s) \exp \left((1 - \gamma) \frac{\sigma_z^2}{2} \right) P_{t-1}(f, s)}{M_t},$$

the firm's price gap in the absence of price changes. If the firm resets its prices, the gap is equal to $x_t(f, s) = x_t^*(s)$, the optimal reset price gap, otherwise it is $x_t(f, s) = \hat{x}_t(f, s)$. This state variable evolves over time according to

$$\hat{x}_t(f, s) = \exp \left((1 - \gamma) \frac{\sigma_z^2}{2} \right) x_{t-1}(f, s) \frac{e_t(s)}{e_{t-1}(s)} \frac{M_{t-1}}{M_t}.$$

The second state variable we keep track of is the duration of a firm's price spell, as this determines the losses from misallocation within the firm. To see why, notice that when the firm resets its prices, it sets $x_{it}(f, s) = x_t^*(s)$ and eliminates the misallocation inside the firm, so $\phi_t(f, s) = 1$. Over time, the losses from misallocation increase because the distribution of price gaps becomes more dispersed. In our model with a mobile specific factor, the productivity of a firm whose prices are d periods old is

$$\phi_t(f, s) = \exp \left(-d\gamma \frac{\sigma_z^2}{2} \right). \tag{10}$$

For a given duration, the losses from misallocation are increasing in the elasticity of substitution γ and the volatility of idiosyncratic shocks. Assuming instead that the specific factor is immobile across products implies

$$\phi_t(f, s) = \exp\left(-d\gamma\frac{\sigma_z^2}{2}\left(1 + \gamma\left(\frac{1}{\eta} - 1\right)\right)\right), \quad (11)$$

so the losses from misallocation are larger than in our baseline model whenever $\eta < 1$. Thus, by assuming that γ is relatively low and the strategic complementarities are at the firm rather than the product level, we reduce the losses from misallocation caused by idiosyncratic productivity shocks and therefore reduce their importance relative to aggregate shocks in determining price adjustment decisions.

As in the single-product model, we use the [Krusell and Smith \(1998\)](#) approach to characterize how the sectoral price gap $\hat{x}_t(s)$ evolves over time in response to sectoral shocks. We find that the method works well even in the multi-product setting, with a R^2 in the perceived law of motion for the sectoral price gap in excess of 0.997.

Lastly, we note that the distribution of price changes for a firm with a price gap $\hat{x}_t(f, s)$ that last changed its price d periods ago and adjusts in period t is

$$\log \frac{P_{it}^*(f, s)}{P_{it-d}(f, s)} \sim N\left(\log \frac{x_t^*(s)}{\hat{x}_t(f, s)} + d(1 - \gamma)\frac{\sigma_z^2}{2}, d\sigma_z^2\right).$$

The older the firm's prices are, the more dispersed its price gaps and therefore the more dispersed its price changes. In turn, the distribution of overall price changes is equal to a mixture of the normal distributions above and is therefore fat-tailed as long as there is randomness in the menu costs, which generates dispersion in the duration of price spells conditional on adjustment.

4.3 Parameterization

Table 4 shows the parameterization of the two variants of the multi-product model. In both of these we set $\sigma = 6$ and $\eta = 2/3$, as in the single-product model, so they also feature moderate strategic complementarities. In our baseline economy, which we refer to as *our model*, we assume that the specific factor is mobile across products and set $\gamma = 1$. As we show below, this value of γ allows the model to reproduce

that menu costs are approximately 1% of firm revenues.²¹ In the *standard* multi-product economy, we assume that the specific factor is immobile across products and set $\gamma = \sigma = 6$.

Our calibration strategy is similar to that in the single-product model, except that we no longer explicitly target the kurtosis of price changes since we have one fewer parameter. Panel A of Table 4 shows that both multi-product models are able to perfectly match the targeted moments, namely the fraction of price changes, the mean and standard deviation of price changes, and the volatility of sectoral inflation. The models also reproduce well the untargeted statistics: the kurtosis of price changes and the distribution of the size of price changes. Panel B reports the calibrated parameter values. With the exception of the upper bound of the menu cost distribution, which is much smaller in our model, both models imply similar parameter values.

4.4 Model Implications

We next discuss the ability of the multi-product models to reproduce the relationship between inflation and the fraction of price changes, as well as their implications for the size of menu costs and the losses from misallocation.

Importance of the Extensive Margin. Our model predicts that the fraction of price changes increases much more in periods of high inflation. For example, when inflation increases from 0 to 10%, our model predicts that the fraction of price changes increases by 5 percentage points, from 10.5 to 15.5%, close to the 6 percentage points increase in the data. In contrast, the standard multi-product model predicts a much smaller increase of 1.2 percentage points.

Figure 5 plots the relationship between sectoral inflation $\pi_t(s)$ and the counterfactual sectoral inflation $\pi_t^c(s)$ that eliminates fluctuations in the fraction of price changes. Our model reproduces the non-linear relationship between the two, whereas the standard multi-product model predicts that the extensive margin of price changes plays almost no role in driving inflation dynamics, even at high rates of inflation.

Table 5 further corroborates this point. In our model, the slope coefficient from regressing $\pi_t(s)$ on $\pi_t^c(s)$ is 0.83, close to the 0.80 in the data, and falls to 0.57 when sectoral inflation is above its 75th percentile, close to its 0.48 empirical counterpart. In

²¹In the robustness section we report results for alternative values of γ .

Table 4: Parameterization of Multi-Product Model

A. Moments			
	Data	Our model	Standard
I. Targeted			
fraction Δp	0.116	0.116	0.116
mean Δp	0.018	0.018	0.018
std. dev. Δp	0.188	0.188	0.188
std dev. $\pi_t(s)$	0.029	0.029	0.029
II. Not targeted			
kurtosis Δp	3.609	3.873	3.612
<i>distribution of Δp</i>			
10 th percentile	0.018	0.020	0.022
25 th percentile	0.045	0.053	0.055
50 th percentile	0.104	0.115	0.118
75 th percentile	0.204	0.206	0.209
90 th percentile	0.334	0.311	0.312

B. Calibrated Parameter Values

		Our model	Standard
g_m	mean money growth rate	0.023	0.022
σ_z	s.d. idios. shocks	0.063	0.063
$\bar{\xi}$	upper bound menu cost	0.960	26.56
σ_e	s.d. sectoral shocks	0.011	0.010

Note: The money growth rate is annualized and the menu cost is relative to average sales.

contrast, in the standard model, regressing $\pi_t(s)$ on $\pi_t^c(s)$ yields a slope coefficient of 0.97. This slope coefficient only falls to 0.89 when inflation exceeds its 75th percentile.

Size of Menu Costs and Misallocation. The last two rows of Table 5 show that our model requires small menu costs, in line with the 1% estimates from the data, to reproduce the micro-price statistics. In contrast, the standard model requires menu costs that are 30 times larger. Additionally, our model implies that the productivity losses from misallocation from menu costs are 1.3%, whereas the standard model

Figure 5: Importance of the Extensive Margin

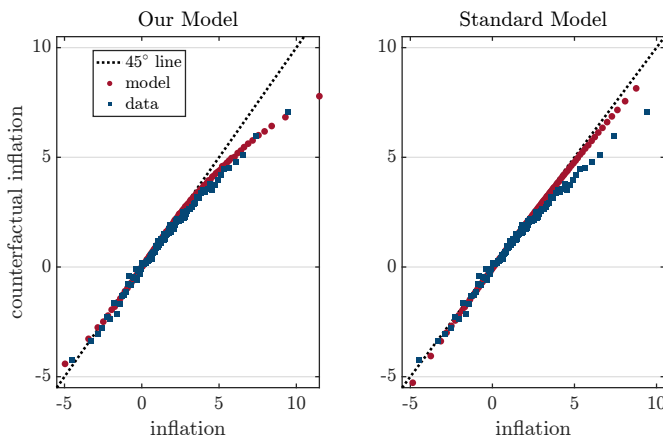


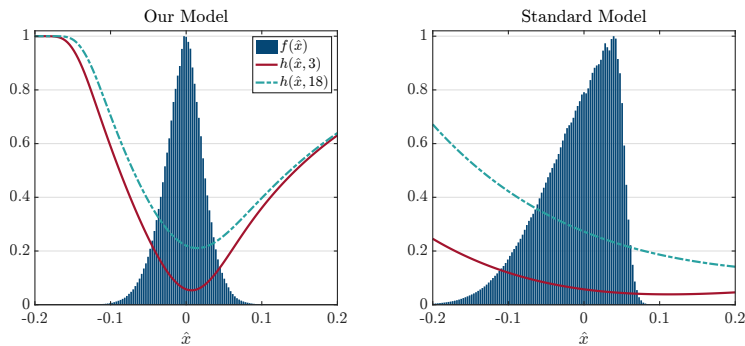
Table 5: Multi-product Model Implications

	Data	Our model	Standard
slope of $\pi_t^c(s)$ on $\pi_t(s)$			
all observations	0.80	0.83	0.97
$\pi_t(s) > 75^{th}$ pct.	0.48	0.57	0.89
menu costs/sales		0.010	0.297
losses from misallocation		0.013	0.367

implies much larger losses, equal to 36.7%. That menu costs and misallocation are smaller in our model follows from the fact that the same volatility of quality shocks σ_z translates into much lower within-firm losses from misallocation, as shown in equations (10) and (11).

As we showed in the single-product model, one could reduce the size of menu costs and losses from misallocation by simply eliminating strategic complementarities in price setting and reducing the demand elasticity faced by individual firms. Doing so would render the aggregate price index much more flexible, reducing the real effects of monetary shocks, at odds with the empirical evidence. By assuming that the elasticities of substitution between and across firms are different and that complementarities are across but not within firms, our model breaks the tight link between micro-level complementarities and the size of menu costs and misallocation.

Figure 6: Adjustment Hazards in Multi-Product Model



4.5 Understanding the Results

To understand why the fraction of price changes fluctuates much more in our model relative to the standard model, Figure 6 plots the distribution of firm price gaps $\hat{x}_t(f, s)$ across firms and sectors, as well as the adjustment hazard for firms that last adjusted prices 3 and 18 months ago. We note that the adjustment hazards are much flatter in the standard model compared to our model. This reflects the much larger menu costs required to match the micro-price statistics in the standard model. The flatter the adjustment hazard, the smaller is the fraction of firms that adjusts in response to a shock that shifts the distribution of price gaps $\hat{x}_t(f, s)$, and therefore the lower the response of the fraction of price changes. We also note that the adjustment hazard increases with the duration of prices because the resulting increase in misallocation inside the firm given by equations (10) and (11) makes it more likely for firms with older prices to adjust.

To sharpen the intuition behind our findings, recall that in the continuous-time version of the single-product model without free price changes the response of the fraction of price changes to a one-time sectoral shock of size δ is

$$\Delta n = \frac{1}{2} \frac{\delta^2}{\mathbb{V}[\Delta p]} + O(\delta^4). \quad (12)$$

This result relies on the one-to-one relationship between the variance of price changes $\mathbb{V}[\Delta p]$ and the width of the inaction region, S , namely $\mathbb{V}[\Delta p] = S^2$. Thus, the extent to which the fraction of price changes increases after a shock of size δ depends on how large the shock is *relative* to the width of the inaction region.

We argue next that in the multi-product model there is no longer a tight relation-

ship between the variance of price changes and the width of the inaction region: as shown in Figure 6, the two multi-product models have different adjustment hazards despite matching the same variance of price changes. To that end, consider a version of our multi-product model in which the elasticity of substitution between a firm's products is equal to $\gamma = 0$, so there are no losses from misallocation inside the firm and $\phi_t(f, s) = 1$. The firm's objective then becomes

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(x_t(s)^{\sigma-1} \left[(1 + \tau) x_t(f, s)^{1-\sigma} - \bar{a} x_t(s)^{\left(\frac{1}{\eta}-1\right)(\sigma-1)} x_t(f, s)^{-\frac{\sigma}{\eta}} \right] - \xi_t(f, s) \mathbb{I}_t(f, s) \right), \quad (13)$$

where, absent a price change, the firm price gap evolves according to

$$x_t(f, s) = \exp\left(\frac{\sigma_z^2}{2}\right) x_{t-1}(f, s) \frac{e_t(s)}{e_{t-1}(s)} \frac{M_{t-1}}{M_t}. \quad (14)$$

Comparing equations (13) and (14) to their counterparts (4) and (5) from the single-product model, reveals that the multi-product model with $\gamma = 0$ is equivalent to a single-product model in which there are no idiosyncratic shocks and in which the money growth rate satisfies

$$g_m^{\text{single}} = g_m^{\text{multi}} - \frac{\sigma_z^2}{2}.$$

This equivalent single-product model has identical value functions, adjustment thresholds and optimal reset prices as its multi-product counterpart, and therefore has the same aggregate implications. However, the two models have very different implications for the distribution of price changes.

We illustrate this point by calibrating the multi-product model with $\gamma = 0$ using the same strategy as above. The second column of Table 6 shows that this model reproduces the targeted moments perfectly. The last column of the table reports the moments implied by evaluating the equivalent single-product model at the calibrated parameters from our multi-product model with $\gamma = 0$. The equivalent single-product model implies a standard deviation of price changes of only 1.9%, much smaller than the 18.8% in the multi-product analog and in the data.

Since the two models have identical aggregate implications, they predict the same response of the fraction of price changes to a sectoral shock of size δ . We can therefore leverage equation (12) to express the response of the fraction of price changes in our multi-product model as a function of the variance of price changes

in the equivalent single-product model. For example, for a sectoral shock of 0.01, $\delta^2/\mathbb{V}[\Delta p] = 0.01^2/0.019^2 \approx 0.25$, much larger than in the single-product model we considered in Section 3. Intuitively, our multi-product model behaves identically to a single-product model with narrower inaction regions, which implies that firm repricing decisions are much more sensitive to aggregate shocks.

Table 6: Equivalent Single-product Model

	Data	Our model $\gamma = 0$	Equivalent single product
I. Targeted			
frequency Δp	0.116	0.116	0.116
mean Δp	0.018	0.018	0.000
std. dev. Δp	0.188	0.188	0.019
std dev. $\pi_t(s)$	0.029	0.029	0.029
II. Not targeted			
kurtosis Δp	3.609	4.480	2.211
<i>distribution of Δp</i>			
10 th percentile	0.018	0.019	0.007
25 th percentile	0.045	0.050	0.010
50 th percentile	0.104	0.109	0.016
75 th percentile	0.204	0.199	0.022
90 th percentile	0.334	0.309	0.030

Note: The moments reported in the column “Equivalent single product” are computed by evaluating the equivalent single product model at the calibrated parameters from our model with $\gamma = 0$.

4.6 The Role of γ

We next explore how the conclusions of our model are shaped by the within firm elasticity of substitution γ . To that end, we recalibrate two versions of our model in which we set $\gamma = 0$ and $\gamma = 3$, respectively. We report the results of the calibration in Appendix F and summarize the main predictions here. Table 7 shows that when $\gamma = 0$ the model predicts too large of a role for the extensive margin of price changes in driving inflation fluctuations relative to the data: when sectoral inflation exceeds its 75th percentile, 90% of the movements in inflation are due to variation in the

Table 7: Alternative Values of γ

	Data	$\gamma = 0$	$\gamma = 1$	$\gamma = 3$
slope of $\pi_t^c(s)$ on $\pi_t(s)$				
all observations	0.80	0.60	0.83	0.91
$\pi_t(s) > 75^{th}$ pct.	0.48	0.10	0.57	0.71
menu costs/sales		0.001	0.010	0.034
losses from misallocation		0.001	0.013	0.043

fraction of price changes relative to the 52% in the data. This version of the model also predicts very small menu costs, 0.1% of total revenues, and insignificant losses from misallocation. When $\gamma = 3$ the extensive margin fluctuates too little. For example, when sectoral inflation exceeds its 75th percentile, 29% of the movements in inflation are due to the extensive margin. This version of the model requires too large menu costs, 3.4% of revenues, and predicts losses from misallocation that are three times larger than in our baseline with $\gamma = 1$. We therefore conclude that allowing for a relatively low value of γ and therefore small losses from misallocation inside the firm is critical for the model to reproduce the role of the extensive margin of price changes for inflation dynamics, as well as reconcile the evidence that menu costs are relatively small in an economy that features moderate strategic complementarities.²²

5 Real Effects of Monetary Shocks

We next use the model developed above to revisit the classic question in the menu cost literature: how large are the real effects of monetary policy? That is, how much does output respond to monetary shocks of various sizes? We show that output responses in our model are very different than those in the existing models that we showed are inconsistent with the data. Specifically, in our model output responds non-linearly to shocks of various sizes. The larger the shock is, the stronger the response of the

²²Alvarez and Lippi (2014) also consider a nested CES specification with different elasticities. They conclude that when the number of products sold by a firm goes to infinity and there are no common shocks, allowing for different elasticities does not affect the dynamics of the economy. In contrast, in our setting with common shocks (i.e. shocks to sectoral productivity), allowing for different elasticities significantly changes the dynamics of the economy in response to large, but not to small shocks.

fraction of price changes and therefore the smaller the real effects. Thus, our model predicts a non-linear Phillips curve.

5.1 Impulse Response to Monetary Shocks

We start by reporting impulse responses to a one-time, unanticipated and permanent increase of 1% and 10% in the money supply M_t , starting from the ergodic distribution of price gaps induced by sectoral and idiosyncratic productivity shocks. Figure 7 plots the impulse response of aggregate output y_t in the top row and that of the fraction of price changes in the bottom row. We compare the impulse responses in our model and the standard single-product model with $\sigma = 6$, $\eta = 2/3$ and free price changes.²³ To ease comparison, we rescale the y-axis in the output impulse responses by the size of the shock. We make two points. First, for a small shock of 1% the real effects of monetary policy are nearly the same in the two models, echoing the findings of Alvarez and Lippi (2014) who show that these effects are pinned down by the kurtosis of the distribution of price changes. Second, while in the single-product model the impulse response scales linearly with the shock, consistent with the findings of Auclert et al. (2022), in our model a large shock of 10% implies a disproportionately smaller output response. Not only is the output response on impact smaller in our model (6.3% vs. 8.9%), but it also shorter lived.

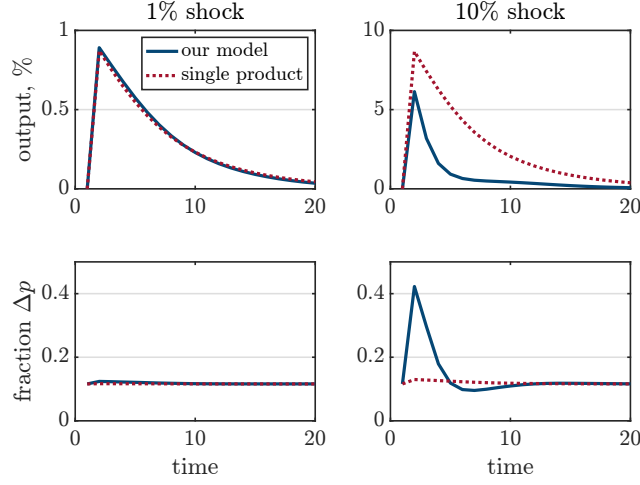
To see why this is the case, the bottom row of Figure 7 shows that, in contrast to the single-product model, in which the fraction of price changes responds very little to both small and large shocks, in our model it increases considerably after a large money shock. For example, though the fraction of price changes responds little to a money shock of 1%, it jumps on impact to 40% after a 10% money shock.

In Table 8 we zoom in on the impact response of inflation to a money shock Δm . We calculate the pass-through of the shock to inflation $\Delta\pi/\Delta m$ and decompose it into three channels. Our decomposition, in the spirit of Caballero and Engel (2007) and Costain and Nakov (2011), starts from the observation that, up to a first-order approximation, inflation in the absence of the shock is equal to

$$\pi = \int \omega h(\omega) df(\omega),$$

²³In Appendix F we show that the responses in the single-product model are similar to those in the standard multi-product model.

Figure 7: Response of Output and Fraction of Price Changes to Money Shock



where ω is the desired price change, $h(\omega)$ is the adjustment hazard and $f(\omega)$ is the ergodic distribution of desired price changes across firms and sectors. The money shock increases all firms' desired price changes to $\omega + \alpha$, where

$$\alpha = \tilde{x}^* - x^* + \Delta m,$$

and where \tilde{x}^* is the average across sectors of the log reset price in the first period after the money shock and x^* is the average across sectors of the log reset price in the absence of the shock. The money shock changes the inflation rate to

$$\tilde{\pi} = \int (\omega + \alpha) \tilde{h}(\omega) df(\omega),$$

where $\tilde{h}(\omega)$ is the new adjustment hazard as a function of ω , the desired price change absent the money shock. The change in inflation $\Delta\pi \equiv \tilde{\pi} - \pi$ can then be decomposed into the following three terms

$$\Delta\pi = \underbrace{\alpha \int h(\omega) df(\omega)}_{\text{Calvo}} + \underbrace{\alpha \int (\tilde{h}(\omega) - h(\omega)) df(\omega)}_{\text{frequency}} + \underbrace{\int \omega (\tilde{h}(\omega) - h(\omega)) df(\omega)}_{\text{selection}}.$$

The first term, which we refer to as the Calvo term, captures the price increase that the shock generates if the frequency of price changes were to remain constant at its steady state level $\int h(\omega) df(\omega)$. The second term, which we refer to as the frequency

term, captures the price increase resulting from the increase in the frequency of price changes from its steady state level to $\int \tilde{h}(\omega) df(\omega)$. The final term is the Golosov and Lucas (2007) selection effect that captures the change in mix of firms that adjust prices. We note that this is purely an accounting decomposition, as all of these effects are interdependent. For example, a stronger selection effect leads to more price flexibility and thus a smaller reduction in the optimal reset price \tilde{x}^* and therefore a larger Calvo effect.

We make two observations based on the decomposition results in Table 8. First, for small shocks the models have a similar pass-through to inflation, mostly driven by the Calvo effect. Moreover, the relative importance of the Calvo, frequency and selection effects is similar across the two models: in both changes in the frequency play a negligible role and selection accounts for a quarter of the pass-through. Second, in our model the pass-through increases rapidly with the size of the shock: from 0.114 for a 1% shock to 0.441 for a 10% shock vs. from 0.128 to 0.145 in the single-product model. This increase is primarily accounted for by the increase in the frequency of price changes, while the strength of the Calvo and selection effects remains comparable to that in the single-product model.

Table 8: Inflation Pass-through to Monetary Shock on Impact

	Single-product		Our model	
	1%	10%	1%	10%
total pass-through	0.128	0.145	0.114	0.441
<i>Calvo</i>	0.092	0.096	0.079	0.105
<i>frequency</i>	0.001	0.012	0.006	0.286
<i>selection</i>	0.035	0.037	0.030	0.050

5.2 Non-Linear Phillips Curve

We next investigate how non-linear are the real effects of changes in monetary policy for a wider range of shocks. Specifically, we consider money shocks that range from -15% to 15% and report the impact response of the frequency of price changes and output, as well as the cumulative response of output.

Figure 8: Impact Response of the Fraction of Price Changes

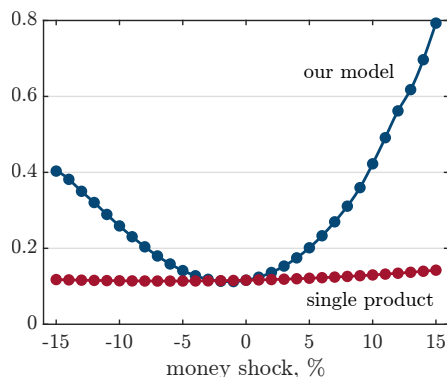


Figure 8 shows that though the fraction of price changes is relatively insensitive to the size of the money shock in the single product model, it responds much more in our model. Specifically, in the neighborhood of zero the frequency does not respond to money shocks, but increases fast away from zero. Moreover, the frequency response is asymmetric: 42% of firms change prices after a 10% increase in the money supply, and 26% do so after a 10% fall in the money supply. This asymmetry is driven by the asymmetry in the profit function: sub-optimally low prices are much more costly than sub-optimally high prices.

Figure 9(a) displays the impact response of output to monetary shocks of various sizes. In the single product model, the response is nearly linear, reflecting the relatively constant fraction of price changes. In our model, the output response is highly non-linear. Though the output response is similar to that in the single product model for shock sizes between -7% and 5%, our model predicts lower real effects for larger shocks. The output response on impact is also asymmetric in our model, reflecting the asymmetry in the response of the fraction of price changes: a 15% money shock leads to only a 2.8% increase in output, whereas a -15% leads to a 8.9% fall in output.

The cumulative impulse responses of output, depicted in Figure 9(b), exhibit a similar pattern, but the non-linearity of the response in our model is even more pronounced, a consequence of the lower persistence of the output response in our model. As can be inferred from Figure 7, the response of output in our model has a lower half-life. The asymmetry is also present in the cumulative output response: while a 15% increase in money supply shock has almost no cumulative real effects, a 15% fall in money supply leads to a cumulative fall in output of 36%.

Figure 10 summarizes this discussion by depicting the Phillips curve implied by

Figure 9: Output Response to Money Shocks

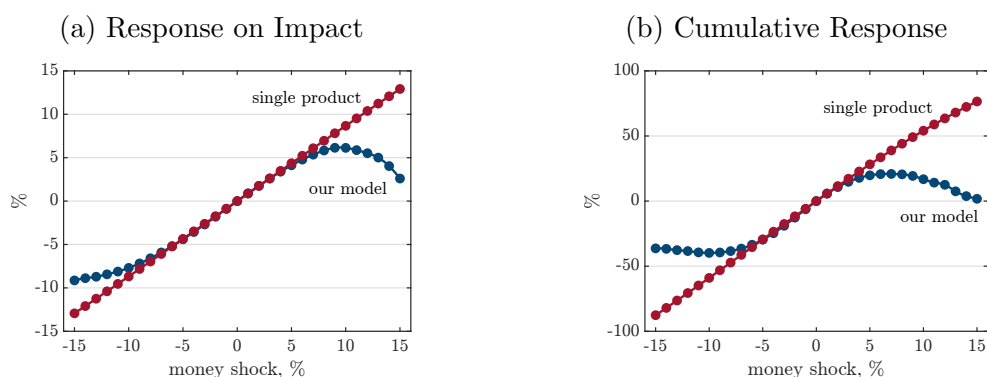
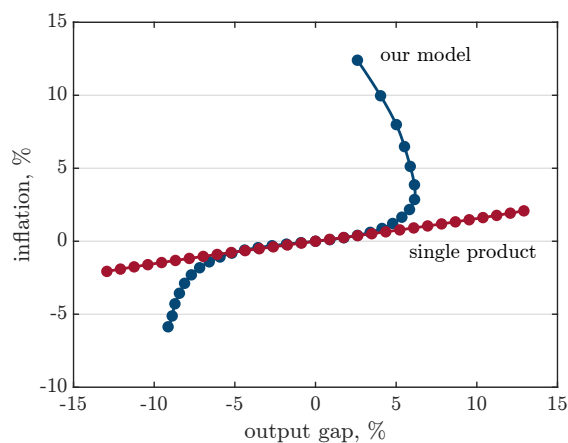


Figure 10: Phillips Curve, Impact Responses



the impact responses of output and inflation to the money shocks considered above. While the Phillips curve is approximately linear in the single product model, it is highly non-linear in our model. In particular, at low levels of inflation the Phillips curve in our model has a similar slope as that in the single model, but it becomes much steeper and eventually vertical at high rates of inflation.

6 Conclusions

We show that canonical menu cost models, when calibrated to match the distribution of micro price changes, cannot reproduce the strong comovement between the fraction of price changes and inflation observed in the data. They therefore predict linear inflation dynamics even in response to large shocks, as in time-dependent pricing

models with a constant frequency of price changes. Moreover, these menu cost models require large menu costs and predict large losses from misallocation in the presence of microeconomic strategic complementarities in price setting.

We propose a resolution to these shortcomings by extending a multi-product menu cost model along two dimensions. First, we assume that individual products sold by a given firm are imperfect substitutes. Second, we assume that strategic complementarities are at the firm rather than the product level. Both these assumptions limit the losses from misallocation from price dispersion within the firm and allow the model to reproduce the importance of the extensive margin of price changes at high levels of inflation. The model also implies small menu costs and losses from misallocation, even in the presence of strategic complementarities.

The key force that drives our results is that within-firm price gap dispersion is not too costly to the firm. The firm is therefore willing to tolerate large idiosyncratic shocks, reducing their importance relative to aggregate shocks in determining repricing decisions. We conjecture that any mechanism that reduces within-firm misallocation will allow the model to simultaneously reproduce the large dispersion in price changes observed in the data and the sensitivity of repricing decisions to the level of inflation. Although it is challenging to directly measure the extent of within-firm misallocation, the observation that firms change prices infrequently and by large amounts, yet are very responsive to aggregate or sectoral shocks, suggests that within-firm price dispersion is likely not very costly.

We use the model to study the real effects of monetary policy. In contrast to existing models, our model predicts a highly non-linear and asymmetric output response to shocks, owing to a more responsive frequency of price changes. The model implies that the Phillips curve is nearly vertical when inflation is sufficiently high.

References

- Alexandrov, Andrey**, “The Effects of Trend Inflation on Aggregate Dynamics and Monetary Stabilization,” 2020.
- Alvarez, Fernando and Francesco Lippi**, “Price Setting With Menu Cost for Multiproduct Firms,” *Econometrica*, 2014, *82* (1), 89–135.
- **and** –, “The Analytic Theory of a Monetary Shock,” *Econometrica*, 2022, *90* (4), 1655–1680.
- Alvarez, Fernando E, Francesco Lippi, and Takis Souganidis**, “Price Setting with Strategic Complementarities as a Mean Field Game,” Working Paper 30193, National Bureau of Economic Research July 2022.
- , – , **and** –, “Price setting with strategic complementarities as a mean field game,” Technical Report, National Bureau of Economic Research 2022.
- Alvarez, Fernando, Francesco Lippi, and Aleksei Oskolkov**, “The Macroeconomics of Sticky Prices with Generalized Hazard Functions,” *The Quarterly Journal of Economics*, 11 2021, *137* (2), 989–1038.
- , **Hervé Le Bihan, and Francesco Lippi**, “The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach,” *American Economic Review*, October 2016, *106* (10), 2817–51.
- , **Martin Beraja, Martín Gonzalez-Rozada, and Pablo Andres Neumeyer**, “From Hyperinflation to Stable Prices: Argentina’s Evidence on Menu Cost Models,” *The Quarterly Journal of Economics*, 09 2018, *134* (1), 451–505.
- Aruoba, S. Boragan, Eugene Oue, Felipe Saffie, and Jonathan L. Willis**, “Real Rigidities, Firm Dynamics and Monetary Nonneutrality: The Role of Demand Shocks,” Technical Report 2023.
- Auclert, Adrien, Rodolfo D Rigato, Matthew Rognlie, and Ludwig Straub**, “New Pricing Models, Same Old Phillips Curves?,” Working Paper 30264, National Bureau of Economic Research July 2022.
- Baley, Isaac and Andrés Blanco**, “Aggregate dynamics in lumpy economies,” *Econometrica*, 2021, *89* (3), 1235–1264.

- Baqae, David Rezza and Emmanuel Farhi**, “Productivity and Misallocation in General Equilibrium,” 2018. LSE working paper.
- Bhattarai, Saroj and Raphael Schoenle**, “Multiproduct Firms and Price-Setting: Theory and Evidence from U.S. Producer Prices,” *Journal of Monetary Economics*, 2014, *66*, 178–192.
- Blanco, Andres**, “Optimal Inflation Target in an Economy with Menu Cost and Zero Lower Bound,” *American Economic Journal: Macroeconomics*, 2021, *13* (3), 108–141.
- Blanco, Andrés, Corina Boar, Callum J Jones, and Virgiliu Midrigan**, “Nonlinear Dynamics in Menu Cost Economies? Evidence from U.S. Data,” 2024. NBER Working Paper 32748.
- Bonomo, Marco, Carlos Carvalho, Oleksiy Kryvtsov, Sigal Ribon, and Rodolfo Rigato**, “Multi-Product Pricing: Theory and Evidence From Large Retailers,” 2022.
- Burstein, Ariel and Christian Hellwig**, “Welfare Costs of Inflation in a Menu Cost Model,” *American Economic Review*, May 2008, *98* (2), 438–43.
- Caballero, Ricardo J. and Eduardo M.R.A. Engel**, “Price Stickiness in Ss Models: New Interpretations of Old Results,” *Journal of Monetary Economics*, 2007, *54*, 100–121.
- Cavallo, Alberto, Francesco Lippi, and Ken Miyahara**, “Large Shocks Travel Fast,” Working Paper 31659, National Bureau of Economic Research September 2023.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans**, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 2005, *113* (1), 1–45.
- Costain, James and Anton Nakov**, “Distributional Dynamics Under Smoothly State-Dependent Pricing,” *Journal of Monetary Economics*, 2011, *58* (6), 646–665.
- , – , and **Borja Petit**, “Flattening of the Phillips Curve with State-Dependent Prices and Wages,” *The Economic Journal*, 2022, *132* (642), 546–581.

- De Loecker, Jan, Jan Eeckhout, and Gabriel Unger**, “The Rise of Market Power and the Macroeconomic Implications,” *The Quarterly Journal of Economics*, 01 2020, *135* (2), 561–644.
- Dotsey, Michael and Robert G. King**, “Implications of State-Dependent Pricing for Dynamic Macroeconomic Models,” *Journal of Monetary Economics*, 2005, *52* (1), 213–242.
- , – , and **Alexander L. Wolman**, “State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output,” *The Quarterly Journal of Economics*, 1999, *114* (2), 655–690.
- Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo, and Josephine Smith**, “How Frequent Are Small Price Changes?,” *American Economic Journal: Macroeconomics*, April 2014, *6* (2), 137–55.
- Foster, Lucia S, John C Haltiwanger, and Cody Tuttle**, “Rising Markups or Changing Technology?,” Technical Report 30491 September 2022.
- Gagnon, Etienne**, “Price Setting during Low and High Inflation: Evidence from Mexico,” *The Quarterly Journal of Economics*, 08 2009, *124* (3), 1221–1263.
- Gali, Jordi**, *Monetary Policy, Inflation, and the Business Cycle*, Princeton University Press, 2008.
- Gertler, Mark and John Leahy**, “A Phillips Curve with an Ss Foundation,” *Journal of Political Economy*, 2008, *116* (3), 533–572.
- Golosov, Mikhail and Robert Lucas**, “Menu Costs and Phillips Curves,” *Journal of Political Economy*, 2007, *115* (2), 171–199.
- Karadi, Peter and Adam Reiff**, “Menu Costs, Aggregate Fluctuations, and Large Shocks,” *American Economic Journal: Macroeconomics*, July 2019, *11* (3), 111–46.
- Kehoe, Patrick and Virgiliu Midrigan**, “Prices Are Sticky After All,” *Journal of Monetary Economics*, 2015, *75*, 35–53.
- Klenow, Peter J. and Jonathan L. Willis**, “Real Rigidities and Nominal Price Changes,” *Economica*, July 2016, *83*, 443–472.

- **and Oleksiy Kryvtsov**, “State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?,” *The Quarterly Journal of Economics*, 08 2008, *123* (3), 863–904.
- Krusell, Per and Anthony Smith Jr.**, “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 1998, *106* (5), 867–896.
- Leahy, John**, “A Survey of New Keynesian Theories of Aggregate Supply and Their Relation to Industrial Organization,” *Journal of Money, Credit and Banking*, 2011, *43*, 87–110.
- Levy, Daniel, Mark Bergen, Shantanu Dutta, and Robert Venable**, “The Magnitude of Menu Costs: Direct Evidence From Large U. S. Supermarket Chains,” *The Quarterly Journal of Economics*, 1997, *112* (3), 791–825.
- Midrigan, Virgiliu**, “Menu Costs, Multi-product Firms, and Aggregate Fluctuations,” *Econometrica*, 2011, *79* (4), 1139–1180.
- Nakamura, Emi and Jon Steinsson**, “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.
- **and Jón Steinsson**, “Monetary Non-neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, 08 2010, *125* (3), 961–1013.
- , – , **Patrick Sun, and Daniel Villar**, “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation,” *The Quarterly Journal of Economics*, 08 2018, *133* (4), 1933–1980.
- Vavra, Joseph**, “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *The Quarterly Journal of Economics*, 09 2013, *129* (1), 215–258.
- Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen**, “Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets,” *The Review of Economics and Statistics*, 2004, *86* (2), 514–533.

Appendix

For Online Publication

A Data

A.1 Overview

We use the data that underlie the construction of the Consumer Price Index (CPI) in the UK. The data are collected by the United Kingdom’s Office for National Statistics (ONS).¹ We use public monthly product-level price quotes and item-level price indexes from January 1996 to August 2022.

Goods and services are classified following the 6-digit Classification of Individual Consumption by Purpose (COICOP 6).² The CPI is produced in stages, with indexes derived at each stage weighted together to give higher level indexes. A sample of prices of items which are representative of UK consumer expenditure are collected in line with the COICOP classification system. There are currently around 650 representative items in the UK CPI price basket of goods. The items usually have fairly broad specifications (such as a roll of wallpaper or women’s jeans). Price collectors choose a selection of products which conform to that item specification. Product-level price quotes are collected by either sampling individual outlets or are collected centrally (for example, university tuition fees).

A.2 Weights

Class-level The COICOP class-level weights are largely calculated from household final consumption expenditure data which covers the relevant population and range of goods and services and are classified by COICOP. This is supplemented by other data sources, including the Living Costs and Food Survey (LCF) data, International Passenger Survey data, and data from Public Sector Branch. The weights used in compiling the measures of consumer price inflation are updated annually following

¹The descriptions in this section are taken from the Consumer Price Indices Technical Manual published by the ONS and available [here](#).

²COICOP is a hierarchical classification system comprising: Divisions e.g. 01 Food & non-alcoholic beverages, Groups e.g. 01.1 Food, and Classes (the lowest published level) e.g. 01.1.1 Bread and cereals. See [here](#) for a description of the COICOP classification.

ONS reviews of the representative items in the basket, so that the weights reflect the introduction of new items and the deletion of others. In addition, using up-to-date expenditure data ensures that the indexes remain representative of current expenditure patterns over time.

Item-level Some items within a class represent themselves while others represent a subclass of expenditure within a section. However, other items represent price changes for a set of items, which are not priced, so for these the weight reflects total expenditure on all items in the set. The expenditure figures for all items in a section are expressed as a percentage of the section weight. Each percentage is rounded to the nearest unit, except where percentages are less than 0.5 which are rounded up to 1. Manual adjustments are then made by the ONS to constrain the sum of each section’s item weights to 100.

The item weights are updated twice each year—with the January index when the new COICOP weights are introduced, and in February when the representative items that make up the basket of goods and services are updated. When the basket of goods and services is updated in February, item weights are updated by drawing on data from a variety of sources. These include detailed National Accounts expenditure data, LCF data, market research data and other sources including administrative data. For each COICOP class, the sum of the new item weights introduced in February is constrained to be equal to the updated class weight introduced in the previous month.

A.3 Sources

We use several datasets published by the ONS to construct our master panel dataset.

1. Price quotes. The price quote data is sourced from the ONS website, which contains the [latest data](#) and the [historical data](#).
2. Item identifier, COICOP classification, and COICOP weights. The [item index data](#), the [classification](#) of items into COICOP classifications and the COICOP weights are provided by the ONS.
3. Aggregated price indexes. We also use the [price indexes](#) published by the ONS at COICOP-6 and above levels of disaggregation.

A.4 Compiling the Dataset

To compile the dataset, we use the following steps.

1. Import data. In this step we generate a dataset of unprocessed price quotes and a dataset of item-COICOP classifications and CPI weights.
2. Process item-level data and price quotes. In this step we correct for recording errors and drop price quotes that are invalidated by the ONS. We also use the algorithm in [Blanco \(2021\)](#) to recover unique price trajectories for price quotes with the same product-outlet identifier.
3. Merge price quotes data with item identifiers and weights.

Our final master panel dataset is comprised of around 38 million unique price quote observations from 1996m1 to 2022m8. All statistics and analyses are produced with this dataset.

A.5 Data Checks

To confirm that the disaggregated price indexes generate the published aggregate CPI index using the sector-level weights in our dataset we construct

$$\pi_t = \sum_s w_t(s) \pi_t(s) \tag{15}$$

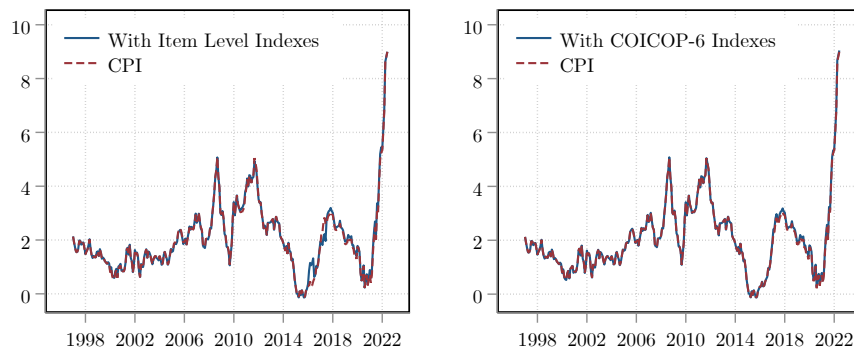
where π_t is inflation, $\pi_t(s)$ is inflation at the s -sector level of disaggregation, and $w_t(s)$ is the corresponding weight in our dataset. [Figure A.1](#) plots the constructed inflation rate against the published inflation series.

A.6 Constructing Micro-Price Statistics

We use our master panel dataset to construct the micro-price statistics that we use to calibrate the model. We apply the following steps in sequence to the dataset:

1. **Filters:** We drop price changes with a sale flag and noncomparable product substitution flag. We also drop prices that are centrally collected by the ONS. We next remove quotes for products that are not observed in the dataset for at least 6 months. We next drop prices that are not rounded to the nearest cent,

Figure A.1: Inflation Rates



and which could indicate recording errors (see [Eichenbaum et al., 2014](#)). We next drop observations if the number of observations for the item-category is less than 20. Out of our initial number 37,708,793 unique price quotes, these filtering steps eliminates 3,639,521 observations.

2. **Removing energy products:** We drop observations that are classified as “energy” at the COICOP-6 level, following the ONS classification.³ This removes 381,134 observations.
3. **Product-level weights:** The weights in our dataset are observed at the item-date level. We construct product-date level weights by dividing the item-date weight by the number of quotes observed for that item-date.
4. **Regular prices:** Since the sales flag is unlikely to cover all sales observed, we next construct regular prices by applying an algorithm that filters out V-shaped price series that last less than three months. As we plot in the text, the aggregate inflation series computed from regular prices is close to the series computed using posted prices.
5. **Standardization:** We follow [Klenow and Kryvtsov \(2008\)](#) and calculate for each price change $\Delta p_{it}(j)$ of quote i that belongs to product category j the

³There are five COICOP-6 classifications that are grouped as “energy”: Electricity (04.5.1), Gas (04.5.2), Liquid fuels (04.5.3), Solid fuels (04.5.4), Fuels and lubricants (07.2.2).

standardized price change

$$\hat{\Delta}p_{it}(j) = \frac{\Delta p_{it}(j) - \mu_{\Delta}(j)}{\sigma_{\Delta}(j)}\sigma_{\Delta} + \mu_{\Delta},$$

where $\mu_{\Delta}(j)$ and $\sigma_{\Delta}(j)$ are the item-level mean and standard deviation of non-zero price changes and μ_{Δ} and σ_{Δ} are the overall ones.

6. **Remove outliers:** Our final step is to remove the top 2% and bottom 2% of observations based on the normalized price changes.

Table A.1 shows price statistics for different sets of prices from our dataset.⁴ Removing energy prices drops the frequency of adjustment from 0.18 to 0.17, and using regular prices drops it further to 0.12. The remaining price statistics are broadly similar.

We next show in Table A.2 a subset of commonly reported statistics under different filtering assumptions to understand how the procedures we use to process the UK data affects the price statistics. We report the mean absolute value of price changes, the standard deviation of price changes, the kurtosis of price changes, and the frequency of price changes. In all cases, we remove energy prices, as categorized by the ONS at the COICOP-6 level. The first row shows the price statistics for all prices. The remaining rows report the price statistics for regular prices computed under our algorithm that filters out V-shaped price series that last less than three months (the baseline filter), and alternatively computed using the regular price algorithm of Kehoe and Midrigan (2015) (the KM filter). When considering regular prices, we present statistics for different treatments of observations at the top and bottom of the price change distribution. In the set of statistics that we target in the baseline estimation, we remove the top and bottom 2% of observations ordered by the normalized price changes. In Table A.2 we present results. The bottom rows of Table A.2 show these statistics, where reported, in other papers (that use different datasets). The results indicate that the key statistics do not depend on the filter used to produce regular prices.

⁴In all cases, we standardize prices at the item level and remove the top 2% and bottom 2% of outliers.

Table A.1: Micro Price Statistics

	All Prices	No Energy	Regular Prices No Energy
frequency Δp	0.180	0.168	0.116
<i>distribution of Δp</i>			
mean	0.014	0.014	0.018
std. dev.	0.188	0.190	0.188
kurtosis	3.121	3.145	3.609
<i>distribution of Δp</i>			
10 th percentile	0.019	0.019	0.018
25 th percentile	0.050	0.050	0.045
50 th percentile	0.115	0.115	0.104
75 th percentile	0.216	0.217	0.204
90 th percentile	0.327	0.330	0.334

Notes: In all cases, we remove the top and bottom 2% outliers.

Table A.2: Price Statistics Across Filters

	Mean $ \Delta p $	Std Dev Δp	Kurtosis Δp	Freq Δp
All prices	0.177	0.255	12.270	0.164
Regular prices				
- Baseline filter	0.168	0.250	13.987	0.116
- Baseline filter, 2% outliers	0.142	0.188	3.609	0.116
- KM filter	0.167	0.248	13.889	0.093
- KM filter, 2% outliers	0.141	0.186	3.566	0.093
Literature				
- Klenow and Kryvtsov (2008)	0.113	–	–	–
- Nakamura and Steinsson (2008)	0.085	–	–	0.087–0.111
- Midrigan (2011)	0.11	–	4.02	0.116
- Karadi and Reiff (2019)	0.099	–	3.98	0.126
- Blanco (2021)	0.153	0.205	3.809	0.126

Notes: For regular prices, the ‘baseline filter’ removes V-shaped sales that last less than 3 months, and the ‘KM filter’ follows the algorithm in [Kehoe and Midrigan \(2015\)](#). Removing $x\%$ outliers drops the top and bottom x -th percentiles of the data ordered by normalized price changes. [Nakamura and Steinsson \(2008\)](#) report the median $|\Delta p|$. [Blanco \(2021\)](#) reports statistics for the UK. [Karadi and Reiff \(2019\)](#) reports statistics for Hungary.

B Solution Method

In the presence of strategic complementarities in price setting the firms' optimal pricing decisions depend on the price of their competitors. Recall that the assumptions we make on preferences imply that the problem of a firm in a given sector depends only on the current and future sectoral price gaps $x_t(s)$. We follow the approach of [Krusell and Smith \(1998\)](#) and postulate that $x_t(s)$ is a function of a single moment of the distribution of firm state variables. Specifically, the moment we use is

$$\hat{x}_t(s) \equiv \bar{a}^\eta \frac{e_t(s)}{M_t} P_{t-1}(s),$$

the previous period's price level, scaled by the money supply and sectoral productivity.⁵ We postulate that the sectoral price gap depends on this state variable according to

$$x_t(s) = \mathcal{X}(\hat{x}_t(s)).$$

Given the function $\mathcal{X}(\cdot)$, $\hat{x}_t(s)$ evolves according to

$$\hat{x}_{t+1}(s) = \frac{e_{t+1}(s)}{e_t(s)} \frac{M_t}{M_{t+1}} \mathcal{X}(\hat{x}_t(s)).$$

We use the [Krusell and Smith \(1998\)](#) approach to pin down $\mathcal{X}(\cdot)$. In particular, we parameterize $\mathcal{X}(\cdot)$ using Chebyshev polynomials in order to capture potential non-linearities. For any given guess of $\mathcal{X}(\cdot)$, we solve the firm's decision rules, simulate histories of sectoral productivity shocks, and find the sectoral price gap $x_t(s)$ that is consistent with the firms' decision rules. We then use projection methods to update our guess of $\mathcal{X}(\cdot)$ using simulated data on $x_t(s)$ and $\hat{x}_t(s)$, and iterate until convergence. We find that the [Krusell and Smith \(1998\)](#) approach works well in both the single- and multi-product models, with an R^2 in the perceived law of motion in excess of 0.999.

⁵Equivalently, $\hat{x}_t(s) = \left(\int \left(\frac{z_{t-1}(f,s)}{z_t(f,s)} \right)^{1-\sigma} \hat{x}_t(f,s)^{1-\sigma} df \right)^{\frac{1}{1-\sigma}}$.

C An Economy With Idiosyncratic Productivity Shocks

We describe an economy in which idiosyncratic shocks are shocks to productivity, as opposed to quality. We show that the firm's problem is nearly isomorphic to the problem of a firm in our baseline model provided one rescales the menu cost appropriately.

We now suppose that the technology for aggregating individual products into a final sector good is

$$y_t(f, s) = \left(\int y_{it}(f, s)^{\frac{\gamma-1}{\gamma}} di \right)^{\frac{\gamma}{\gamma-1}}$$

and

$$y_t(s) = \left(\int y_t(f, s)^{\frac{\sigma-1}{\sigma}} df \right)^{\frac{\sigma}{\sigma-1}}.$$

Notice that we no longer have taste shifters in these aggregators. The demand functions are therefore

$$y_{it}(f, s) = \left(\frac{P_{it}(f, s)}{P_t(f, s)} \right)^{-\gamma} y_t(f, s)$$

$$y_t(f, s) = \left(\frac{P_t(f, s)}{P_t(s)} \right)^{-\sigma} y_t(s).$$

As earlier, the production function is

$$y_{it}(f, s) = e_t(s) z_{it}(f, s) m_{it}(f, s) l_{it}(f, s)^\eta,$$

and the optimal choice of the specific factor m_{it} implies that the total amount of labor the firm needs to produce the bundle $y_{it}(f, s)$ is

$$l_t(f, s) = \left(\int \frac{y_{it}(f, s)}{e_t(s) z_{it}(f, s)} di \right)^{\frac{1}{\eta}}.$$

Notice that now $z_{it}(f, s)$ represents a product-specific productivity shock. Firm profits are

$$\sum_{t=0}^{\infty} \frac{\beta^t}{P_t c_t} \left[(1 + \tau) \int P_{it}(f, s) y_{it}(f, s) di - W_t \left(\int \frac{y_{it}(f, s)}{e_t(s) z_{it}(f, s)} di \right)^{\frac{1}{\eta}} - W_t \xi_t(f, s) \mathbb{I}_t(f, s) \right],$$

where we now assume that the fixed cost of changing prices $\xi_t(f, s)$ depends on the firm's productivity, as we discuss below. Absent such rescaling, firms whose productivity grows over time would face smaller menu costs relative to their profits and no longer be subject to pricing frictions.

To show that in this environment the problem of the firm is similar to that in our baseline model, let us define the several objects. First, the first-best level of a firm's productivity is

$$z_t(f, s) = \left(\int z_{it}(f, s)^{\gamma-1} di \right)^{\frac{1}{\gamma-1}}.$$

This evolves over time according to

$$z_t(f, s) = z_{t-1}(f, s) \exp\left((\gamma - 1) \frac{\sigma_z^2}{2}\right),$$

given our assumption that individual productivity evolves according to a geometric random walk process with Gaussian innovations. We can write the firm's production function as

$$y_t(f, s) = e_t(s) z_t(f, s) \phi_t(f, s) l_t(f, s)^\eta,$$

where $\phi_t(f, s)$ represents the losses from misallocation inside the firm, given by

$$\phi_t(f, s) = \left(\int \frac{z_t(f, s)}{z_{it}(f, s)} \left(\frac{P_{it}(f, s)}{P_t(f, s)} \right)^{-\gamma} di \right)^{-1}.$$

Also let

$$z_t(s) = \left(\int z_t(f, s)^{\frac{(\sigma-1)\frac{1}{\eta}}{1+\sigma(\frac{1}{\eta}-1)}} df \right)^{\frac{1+\sigma(\frac{1}{\eta}-1)}{(\sigma-1)\frac{1}{\eta}}}$$

denote the sectoral weighted average of individual firm's composite productivities. This term also evolves over time according to a deterministic trend.

We define the price gaps as follows. The sectoral price gap is given by

$$x_t(s) = \bar{a}^\eta \frac{e_t(s) z_t(s) P_t(s)}{M_t(s)}.$$

The firm-level price gap is given by

$$x_t(f, s) = \bar{a}^\eta \frac{e_t(s) u_t(s) \left(\frac{\tilde{u}_t(f, s)}{u_t(s)} \right)^{\frac{1}{1+\sigma(\frac{1}{\eta}-1)}} P_t(f, s)}{M_t(s)}.$$

The product-level price gap is given by

$$x_{it}(f, s) = \bar{a}^\eta \frac{e_t(s) z_t(s) \left(\frac{z_t(f, s)}{z_t(s)} \right)^{\frac{1}{1+\sigma(\frac{1}{\eta}-1)}} \frac{z_{it}(f, s)}{z_t(f, s)} P_{it}(f, s)}{M_t(s)}.$$

We assume that the menu cost scales with the firm's productivity

$$\bar{\xi}_t(f, s) = \left(\frac{z_t(f, s)}{z_t(s)} \right)^{\frac{(\sigma-1)\frac{1}{\eta}}{1+\sigma(\frac{1}{\eta}-1)}}.$$

This assumption ensures that the menu cost is equal to a constant fraction of the firm's (flexible price) profits, so they do not vanish for firms that grow increasingly large. We can then rewrite the firm's objective as

$$\sum_{t=0}^{\infty} \beta^t \left(\frac{z_t(f, s)}{z_t(s)} \right)^{\frac{(\sigma-1)\frac{1}{\eta}}{1+\sigma(\frac{1}{\eta}-1)}} \left[(1 + \tau) \left(\frac{x_t(f, s)}{x_t(s)} \right)^{1-\sigma} - a_t(s) \phi_t(f, s)^{-\frac{1}{\eta}} \left(\frac{x_t(f, s)}{x_t(s)} \right)^{-\frac{\sigma}{\eta}} - \bar{\xi}_t(f, s) \mathbb{I}_t(f, s) \right].$$

This objective is nearly identical to that in the baseline model with quality shocks, except that we have an additional term due to firm productivity growth affecting the discount factor. In addition, since we scale prices by different terms involving productivity, the law of motion for price gaps changes accordingly.

We finally note that if the firm does not adjust prices, misallocation inside the firm is equal to

$$\phi_t(f, s) = \left(\int \frac{z_t(f, s)}{z_{it}(f, s)} \left(\frac{P_{it}(f, s)}{P_t(f, s)} \right)^{-\gamma} di \right)^{-1} = \frac{\left(\int z_{it}(f, s)^{-1} di \right)^{-1}}{z_t(f, s)},$$

and evolves over time according to the same law of motion as in our baseline model with quality shocks

$$\phi_t(f, s) = \phi_{t-1}(f, s) \exp \left(-\gamma \frac{\sigma_z^2}{2} \right).$$

D A Flexible Distribution of Menu Costs

In our baseline analysis we assumed that menu costs are drawn from a uniform distribution. Here, we consider a more flexible distribution

$$F(\xi) = \left(\frac{\xi}{\bar{\xi}} \right)^\nu,$$

which collapses to the uniform distribution when $\nu = 1$ and is degenerate at $\bar{\xi}$ when $\nu \rightarrow \infty$. We calibrate the parameters of the model to match the same targets as in the baseline and also target the percentiles of the distribution of the size of price changes that were untargeted in the baseline.

Table D.3 reports the results of this calibration. For comparison, we also calibrate a model with a uniform menu cost distribution to match the same targets. As the table shows, the value of ν that best reproduces the data is very close to one, so the uniform distribution fits the data just as well as the flexible one. For this reason, we focus on a uniform menu cost distribution throughout the paper.

E Single-Product Model in Continuous Time

This section analyzes the model presented in Section 3 in continuous time. We assume $g_m = 0$, so the money supply is constant.

E.1 Environment

Household and Final Goods Producers. The problem of households and final goods producers is as in Section 3.

Intermediate Goods Producers. The intermediate goods producer faces the same production technology as in Section 3. The technology for changing prices is as follows: in a period of length dt , the firm faces a non-random menu cost $\bar{\xi}$ denominated in units of labor with probability $1 - \varphi dt$ and a zero menu cost with probability φdt . The firm chooses a sequence of price adjustment dates $\{T_h\}_{h=1}^\infty$ and log-price adjustments $\{\Delta p_h\}_{h=1}^\infty$. For a given firm-level log price gap \hat{x} and a sectoral

Table D.3: Parameterization with A Flexible Menu Cost Distribution

A. Targeted Moments

	Data	Flexible	Uniform
fraction Δp	0.116	0.115	0.116
mean Δp	0.018	0.018	0.019
std. dev. Δp	0.188	0.192	0.192
kurtosis Δp	3.609	3.613	3.613
std dev. $\pi_t(s)$	0.029	0.029	0.029
<i>distribution of Δp</i>			
10 th percentile	0.018	0.021	0.021
25 th percentile	0.045	0.054	0.054
50 th percentile	0.104	0.120	0.120
75 th percentile	0.204	0.219	0.218
90 th percentile	0.334	0.321	0.321

B. Calibrated Parameter Values

	Flexible	Uniform
g_m mean money growth rate	0.021	0.022
σ_z s.d. idios. shocks	0.066	0.066
ξ upper bound menu cost	40.41	40.64
σ_e s.d. sectoral shocks	0.011	0.011
ν menu cost curvature	1.005	

Note: The money growth rate is annualized and the menu cost is relative to average sales.

log price gap \hat{X} , the sequential formulation of the firms' problem is

$$V^{\hat{X}}(\hat{x}) = \max_{\{T_h, \Delta p_h\}_{h=1}^{\infty}} \mathbb{E}_0 \left[\int_0^{\infty} e^{-\rho t} \Pi_t(s, f) dt - \sum_{i=1}^{\infty} \xi_{T_h}(f, s) \right], \quad (16)$$

where the flow profit function is given by

$$\Pi_t(s, f) = e^{\hat{X}_t(s)(\sigma-1)} \left((1 + \tau) e^{(\hat{x}_t(s, f) - \hat{X}_t(s))(1-\sigma)} - \bar{a} e^{\hat{X}_t(s)(\frac{1}{\eta}-1)(\sigma-1)} e^{-\frac{\sigma}{\eta} \hat{x}_t(f, s)} \right),$$

subject to the law of motion of firm price gap

$$\hat{x}_t(f, s) = \text{constant} + \hat{x} + \log(e_t(s)) + \log(z_t(f, s)) + \sum_{h: T_h \leq t} \Delta p_h$$

and the law of motion of the sectoral price gap $\hat{X}_t(s)$ with initial condition \hat{X} .

E.2 Cost of Price Rigidity

We compute the cost of price rigidity under the assumption that $\sigma_e = 0$, i.e., there are no sectoral shocks. Moreover, since sectors are identical, without loss of generality we omit the sectoral index. Let $V^{\hat{X}}(\hat{x})$ be the optimal firm value defined in equation (16). The flow profit function $\Pi^{\hat{X}}(\hat{x})$ in a sector with sectoral (log) price gap \hat{X} and firm-level (log) price gap \hat{x} is given by

$$\Pi^{\hat{X}}(\hat{x}) := e^{\hat{X}(\sigma-1)} \left((1 + \tau) e^{(\hat{x} - \hat{X})(1-\sigma)} - \bar{a} e^{\hat{X}(\frac{1}{\eta}-1)(\sigma-1)} e^{-\frac{\sigma}{\eta}\hat{x}} \right). \quad (17)$$

Given the definition of the flow profits in equation (17), we can write the firm's value as

$$V^{\hat{X}}(\hat{x}) = \max_{\tau} \mathbb{E} \left[\int_0^{\tau} e^{-\rho t} \Pi^{\hat{X}}(\hat{x}_t) dt + e^{-\rho\tau} \left[\xi_{\tau} + \max_{\hat{x}^*} V^{\hat{X}}(\hat{x}^*) \right] \mid \hat{x}_0 = \hat{x} \right]. \quad (18)$$

The next propositions characterize the costs of price rigidity. First, we characterize the losses from misallocation due to price dispersion. Second, we characterize the size of the menu cost.

Let $\mathbb{E}[\hat{x}^m]$ denote the m-th moment of the log price gap distribution and $\mathbb{V}[\hat{x}]$ denote its variance. Similarly, $\mathbb{E}[\Delta p^m]$ denotes the m-th moment of the log price change distribution and $\mathbb{V}[\Delta p] = \mathbb{E}[\Delta p^2] - \mathbb{E}[\Delta p]^2$ and $\mathbb{K}[\Delta p] = \frac{\mathbb{E}[\Delta p^4]}{\mathbb{E}[\Delta p^2]^2}$ denotes the variance and kurtosis of the price change distribution, respectively. Finally, n denotes the fraction of price changes and $\mathbb{E}[\tau]$ the average duration between price changes.

From now on, we consider a quadratic approximation of firms' flow profits around the optimal static log price gap. This approximation implies a symmetric value function around the optimal price gap and therefore a symmetric policy function. Let $(\hat{x}^-, \hat{x}^*, \hat{x}^+)$ be the optimal policy, characterized by the lower and upper adjustment thresholds \hat{x}^- and \hat{x}^+ , and the reset price \hat{x}^* . We normalize the units in which we express the price gap to ensure that the optimal reset price is zero. Moreover, the

symmetry of the value function implies that $-\hat{x}^- = \hat{x}^+ = \bar{x}$. The following proposition characterizes the quadratic approximation of the profit function.

Proposition 1. *Define*

$$\hat{x}^*(\hat{X}) = \arg \max_{\hat{x}} \Pi^{\hat{X}}(\hat{x}). \quad (19)$$

Then, up to a second-order approximation,

$$\Pi^{\hat{X}}(\hat{x}) = \Pi^{\hat{X}}(\hat{x}^*(\hat{X})) + \frac{1}{2} \frac{\partial^2 \Pi^{\hat{X}}(\hat{x})}{\partial \hat{x}^2} \Bigg|_{\hat{x}=\hat{x}^*(\hat{X})} (\hat{x} - \hat{x}^*(\hat{X}))^2 + O\left((\hat{x} - \hat{x}^*(\hat{X}))^3\right), \quad (20)$$

where the optimal reset price is given by

$$\hat{x}^*(\hat{X}) = \frac{\eta}{\eta + \sigma(1 - \eta)} \log \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right) + \frac{(\sigma - 1)(1 - \eta)}{\eta + \sigma(1 - \eta)} \hat{X}. \quad (21)$$

and the level and curvature of the profit function are given by

$$\Pi^{\hat{X}}(\hat{x}^*(\hat{X})) = e^{\left(\sigma - 1 - \frac{(\sigma - 1)^2(1 - \eta)}{\eta + \sigma(1 - \eta)}\right) \hat{X}} \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right)^{\frac{\eta(1 - \sigma)}{\eta + \sigma(1 - \eta)}} (1 + \tau) \frac{\eta + \sigma(1 - \eta)}{\sigma} \quad (22)$$

$$\frac{\partial^2 \Pi^{\hat{X}}(\hat{x})}{\partial \hat{x}^2} \Bigg|_{\hat{x}=\hat{x}^*(\hat{X})} = -e^{\left(\sigma - 1 - \frac{(\sigma - 1)^2(1 - \eta)}{\eta + \sigma(1 - \eta)}\right) \hat{X}} \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right)^{\frac{\eta(1 - \sigma)}{\eta + \sigma(1 - \eta)}} (1 + \tau)(\sigma - 1) \frac{\eta + \sigma(1 - \eta)}{\eta} \quad (23)$$

and the revenue is equal to

$$e^{\hat{X}(\sigma - 1)} (1 + \tau) e^{(1 - \sigma)\hat{x}^*(\hat{X})} = e^{\left(\sigma - 1 - \frac{(\sigma - 1)^2(1 - \eta)}{\eta + \sigma(1 - \eta)}\right) \hat{X}} \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right)^{\frac{\eta(1 - \sigma)}{\eta + \sigma(1 - \eta)}} (1 + \tau) \quad (24)$$

Proof. The first-order condition for the profit function

$$\Pi^{\hat{X}}(\hat{x}) = e^{\hat{X}(\sigma - 1)} \left[(1 + \tau) e^{(1 - \sigma)\hat{x}} - e^{(\sigma - 1)\frac{1 - \eta}{\eta} \hat{X}} e^{-\frac{\sigma}{\eta} \hat{x}} \right] \quad (25)$$

is given by

$$0 = (1 + \tau)(1 - \sigma) e^{(1 - \sigma)\hat{x}} + \frac{\sigma}{\eta} e^{(\sigma - 1)\frac{1 - \eta}{\eta} \hat{X}} e^{-\frac{\sigma}{\eta} \hat{x}}, \quad (26)$$

which implies an optimal markup

$$\hat{x}^*(\hat{X}) = \frac{\eta}{\eta + \sigma(1 - \eta)} \log \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right) + \frac{(\sigma - 1)(1 - \eta)}{\eta + \sigma(1 - \eta)} \hat{X}. \quad (27)$$

Therefore,

$$\begin{aligned} \Pi^{\hat{X}}(x^*(\hat{X})) &= e^{\hat{X}(\sigma-1)} \left[(1 + \tau) e^{(1-\sigma)\hat{x}^*(\hat{X})} - e^{(\sigma-1)\frac{1-\eta}{\eta}\hat{X}} e^{-\frac{\sigma}{\eta}\hat{x}^*(\hat{X})} \right] \\ &= e^{\left((\sigma-1) - \frac{(\sigma-1)^2(1-\eta)}{\eta + \sigma(1-\eta)} \right) \hat{X}} \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right)^{\frac{\eta(1-\sigma)}{\eta + \sigma(1-\eta)}} \frac{\eta + \sigma(1 - \eta)}{\sigma} (1 + \tau) \end{aligned} \quad (28)$$

and

$$\frac{\partial^2 \Pi^{\hat{X}}(\hat{x}^*(\hat{X}))}{\partial \hat{x}^2} = -e^{\left((\sigma-1) - \frac{(\sigma-1)^2(1-\eta)}{\eta + \sigma(1-\eta)} \right) \hat{X}} \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right)^{\frac{\eta(1-\sigma)}{\eta + \sigma(1-\eta)}} (1 + \tau)(\sigma - 1) \frac{\eta + \sigma(1 - \eta)}{\eta}. \quad (29)$$

Finally, the revenue is given by

$$\begin{aligned} e^{\hat{X}(\sigma-1)} (1 + \tau) e^{(1-\sigma)\hat{x}^*(\hat{X})} &= e^{\hat{X}(\sigma-1)} (1 + \tau) \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right)^{\frac{\eta(1-\sigma)}{\eta + \sigma(1-\eta)}} e^{\left(-\frac{(\sigma-1)^2(1-\eta)}{\eta + \sigma(1-\eta)} \right) \hat{X}} \\ &= e^{\left((\sigma-1) - \frac{(\sigma-1)^2(1-\eta)}{\eta + \sigma(1-\eta)} \right) \hat{X}} \left(\frac{\sigma}{(1 + \tau)(\sigma - 1)\eta} \right)^{\frac{\eta(1-\sigma)}{\eta + \sigma(1-\eta)}} (1 + \tau) \end{aligned} \quad (30)$$

□

Proposition 2. *Aggregate productivity is approximately equal to*

$$\log(\phi) \approx -\frac{\sigma(\eta + (1 - \eta)\sigma)}{\eta} \mathbb{V}[\hat{x}] = -\frac{\sigma}{2} \left(1 + \sigma \left(\frac{1}{\eta} - 1 \right) \right) \frac{\mathbb{V}[\Delta p] \mathbb{K}[\Delta p]}{6} \quad (31)$$

Proof. Recall that we can write aggregate productivity as

$$\phi = \left(\int_f e^{-\frac{\sigma}{\eta}(\hat{x}_f - \hat{X})} df \right)^{1/\eta}. \quad (32)$$

Using the same approximation as in [Gali \(2008\)](#), we have that

$$\log(\phi) \approx -\frac{\sigma(\eta + (1 - \eta)\sigma)}{2\eta} \mathbb{V}[\hat{x}]. \quad (33)$$

Because the profit function is symmetric, we have that $\hat{x}^*(\hat{X}) = \mathbb{E}[\hat{x}]$. With this result, we can use Corollary 3 of [Baley and Blanco \(2021\)](#), and the definition of kurtosis, we have that

$$\mathbb{V}[\hat{x}] = \frac{1}{6} \frac{\mathbb{E}[\Delta p^4]}{\mathbb{E}[\Delta p^2]} = \frac{1}{6} \mathbb{E}[\Delta p^2] \mathbb{K}[\Delta p]. \quad (34)$$

Putting these two results together, we have

$$\log(\phi) \approx -\frac{\sigma}{2} \left(1 + \sigma \left(\frac{1}{\eta} - 1 \right) \right) \frac{\mathbb{E}[\Delta p^2] \mathbb{K}[\Delta p]}{6} = -\frac{\sigma}{2} \left(1 + \sigma \left(\frac{1}{\eta} - 1 \right) \right) \frac{\mathbb{V}[\Delta p^2] \mathbb{K}[\Delta p]}{6}, \quad (35)$$

where the last equality uses that when the drift is zero, $\mathbb{E}[\Delta p] = 0$ and $\mathbb{E}[\Delta p^2] = \mathbb{V}[\Delta p^2]$ ([Baley and Blanco, 2021](#)).

□

Proposition 3. *The total amount paid on menu costs, expressed as a fraction of revenue, is approximately equal to*

$$\frac{n\mathbb{E}[\xi_\tau]}{e^{\hat{X}(\sigma-1)}(1+\tau)e^{(1-\sigma)\hat{x}^*(\hat{X})}} = (\sigma - 1) \left[1 + \sigma \left(\frac{1}{\eta} - 1 \right) \right] \frac{\mathbb{V}[\Delta p] \Psi(\mathbb{K}[\Delta p])}{12} \quad (36)$$

where $\Psi(1) = 1$, $\Psi(6) = 0$, and increasing-decreasing in its argument.

Proof. Under the assumption of the CalvoPlus model, we can write $n\mathbb{E}[\xi_\tau]$ as

$$n\mathbb{E}[\xi_\tau] = \frac{n - \varphi}{n} (n\bar{\xi}). \quad (37)$$

The normalized upper adjustment threshold \bar{x} is

$$\bar{x} = \left(\frac{12\sigma_u^2 \bar{\xi}}{\Pi_{\hat{x}^2}(\hat{x}^*(\hat{X}))} \right)^{1/4}. \quad (38)$$

Using that $\sigma_z^2 = n\mathbb{E}[\Delta p^2]$ if the drift is equal to zero (see Proposition 1 in [Alvarez et](#)

al., 2016) we have

$$n\bar{\xi} = \mathbb{E}[\Delta p^2] \frac{\Pi_{\hat{x}^2}^{\hat{X}}(\hat{x}^*(\hat{X}))}{12} \left(\frac{\bar{x}^2}{\mathbb{E}[\Delta p^2]} \right)^2, \quad (39)$$

we have that

$$n\mathbb{E}[\xi_\tau] = \frac{n-\varphi}{n} (n\bar{\xi}) = \mathbb{E}[\Delta p^2] \frac{\Pi_{\hat{x}^2}^{\hat{X}}(\hat{x}^*(\hat{X}))}{12} \left[\frac{n-\varphi}{n} \left(\frac{\bar{x}^2}{\mathbb{E}[\Delta p^2]} \right)^2 \right]. \quad (40)$$

We next show that the last term is only a function of $\frac{\varphi\bar{x}^2}{\sigma_z^2}$. Using the result that $\frac{\varphi\bar{x}^2}{\sigma_z^2}$ is a function of kurtosis (see Proposition 6 in Alvarez et al., 2016), implies that the last term is only a function of the kurtosis, which will complete the proof of the Proposition.

Let $\mathcal{P}(\hat{x})$ and $\mathcal{D}(\hat{x})$ be the solutions of the following differential equations

$$\varphi\mathcal{P}(\hat{x}) = \frac{\sigma_z^2}{2}\mathcal{P}''(\hat{x}), \quad \mathcal{P}(\bar{x}) = \mathcal{P}(-\bar{x}) = 1, \quad (41)$$

$$\varphi\mathcal{D}(\hat{x}) = \varphi\left(\frac{-\hat{x}}{\bar{x}}\right)^2 + \frac{\sigma_z^2}{2}\mathcal{D}''(\hat{x}), \quad \mathcal{D}(\bar{x}) = 1, \mathcal{D}(-\bar{x}) = 1. \quad (42)$$

It is easy to check that $\mathcal{P}(0) = \frac{n-\varphi}{n}$ and $\mathcal{D}(0) = \frac{\mathbb{E}[\Delta p^2]}{\bar{x}^2}$. Using these definitions and letting $\hat{z} = \hat{x}/\bar{x}$, we can normalize $\mathcal{P}(\cdot)$ and $\mathcal{D}(\cdot)$ as $\tilde{\mathcal{P}}(\hat{z}) := \mathcal{P}(\hat{z}\bar{x})$ and $\tilde{\mathcal{D}}(\hat{z}) := \mathcal{D}(\hat{z}\bar{x})$. Doing a change of variable, we have that

$$\varphi\tilde{\mathcal{P}}(\hat{z}) = \frac{\sigma_z^2}{2\bar{x}^2}\tilde{\mathcal{P}}''(\hat{z}), \quad \tilde{\mathcal{P}}(1) = \tilde{\mathcal{P}}(-1) = 1, \quad (43)$$

$$\varphi\tilde{\mathcal{D}}(\hat{z}) = \varphi\hat{z}^2 + \frac{\sigma_z^2}{2\bar{x}^2}\tilde{\mathcal{D}}''(\hat{z}), \quad \tilde{\mathcal{D}}(1) = \tilde{\mathcal{D}}(-1) = 1. \quad (44)$$

Letting $\Phi = \frac{\varphi 2\bar{x}^2}{\sigma_z^2}$

$$\Phi\tilde{\mathcal{P}}(\hat{z}) = \tilde{\mathcal{P}}''(\hat{z}), \quad \tilde{\mathcal{P}}(1) = \tilde{\mathcal{P}}(-1) = 1, \quad (45)$$

$$\Phi\tilde{\mathcal{D}}(\hat{z}) = \Phi\hat{z}^2 + \tilde{\mathcal{D}}''(\hat{z}), \quad \tilde{\mathcal{D}}(1) = \tilde{\mathcal{D}}(-1) = 1, \quad (46)$$

with $\tilde{\mathcal{P}}(0) = \frac{n-\lambda}{n}$ and $\tilde{\mathcal{D}}(0) = \frac{\mathbb{E}[\Delta p^2]}{\bar{x}^2}$. Since Φ is a strictly decreasing function of the kurtosis of price changes, we have that $\frac{\tilde{\mathcal{P}}(0)}{\tilde{\mathcal{D}}(0)^2} = \Psi(\mathbb{K}[\Delta p])$. It is easy to check that $\Psi(1) = 1$ and $\Psi(6) = 0$ and that the function is increasing-decreasing in its argument.

Thus, we have that

$$n\mathbb{E}[\xi_\tau] = \mathbb{E}[\Delta p^2] \frac{\Pi_{\hat{x}^2}^{\hat{X}}(\hat{x}^*(\hat{X}))}{12} \Psi(\mathbb{K}[\Delta p]). \quad (47)$$

Dividing by the revenue, we have that

$$\begin{aligned} \frac{n\mathbb{E}[\xi_\tau]}{e^{\hat{X}(\sigma-1)}(1+\tau)e^{(1-\sigma)\hat{x}^*(\hat{X})}} &= \mathbb{E}[\Delta p^2] \frac{\Pi_{\hat{x}^2}^{\hat{X}}(\hat{x}^*(\hat{X}))}{12e^{\hat{X}(\sigma-1)}(1+\tau)e^{(1-\sigma)\hat{x}^*(\hat{X})}} \Psi(\mathbb{K}[\Delta p]) \\ &= \mathbb{E}[\Delta p^2](\sigma-1) \frac{\eta + \sigma(1-\eta)}{12\eta} \Psi(\mathbb{K}[\Delta p]) \\ &= \mathbb{E}[\Delta p^2] \frac{\sigma-1}{12} \left[1 + \sigma \left(\frac{1}{\eta} - 1 \right) \right] \Psi(\mathbb{K}[\Delta p]). \end{aligned} \quad (48)$$

Finally, since $\mathbb{E}[\Delta p] = 0$, we have that $\mathbb{E}[\Delta p^2] = \mathbb{V}[\Delta p]$ and

$$\frac{n\mathbb{E}[\xi_\tau]}{e^{\hat{X}(\sigma-1)}(1+\tau)e^{(1-\sigma)\hat{x}^*(\hat{X})}} = \mathbb{V}[\Delta p] \frac{\sigma-1}{12} \left[1 + \sigma \left(\frac{1}{\eta} - 1 \right) \right] \Psi(\mathbb{K}[\Delta p]). \quad (49)$$

□

E.3 The Effect of a Sectoral Shock on the Fraction of Price Changes

We next characterize the impact effect of a sectoral productivity shock of size δ on the fraction of price changes n . Let $\Delta n(\delta)$ be the change in the fraction of price changes upon such a shock. The next proposition characterizes this object under the assumption that equilibrium policies are unchanged.

Proposition 4. *Let l be the fraction of free price changes in the steady state, i.e., $l = \frac{\varphi}{n}$. Then*

$$\Delta n(\delta) = \frac{1-l}{2} \frac{\delta^2}{\mathbb{V}[\Delta p]} + O(\delta^4) \quad (50)$$

Proof. Let $S_t(\delta)$ be the mass of firms without a price change following a sectoral shock of size $-\delta$ which decreases the price gap by δ . We refer to $S_t(\delta)$ as the survival function t periods following the shock. Then, on impact,

$$\Delta n(\delta) = 1 - S_0(\delta), \quad (51)$$

since there is a positive mass of firms changing their price on impact. The survival function at time 0 is given by

$$S_0(\delta) = \int_{-\bar{x}}^{\bar{x}} f(\hat{x} + \delta) d\hat{x}, \quad (52)$$

where $f(\hat{x})$ satisfies the Kolmogorov forward equation with boundary conditions

$$\varphi f(\hat{x}) = \frac{\sigma_z^2}{2} \frac{d^2 f(\hat{x})}{d\hat{x}^2} \quad \forall \hat{x} \in (-\bar{x}, \bar{x}) / \{0\}, \quad (53)$$

$$f(\pm\bar{x}) = 0 \quad \forall \hat{x} \notin [-\bar{x}, \bar{x}], \quad (54)$$

$$\int_{\mathbb{R}} f(\hat{x}) = 1, \quad (55)$$

$$f(\hat{x}) \in \mathbb{C}(\mathbb{R}) \cap \mathbb{C}^2((-\bar{x}, \bar{x})). \quad (56)$$

Notice that above we wrote the price gap as the difference from its mean. Using a fourth-order Taylor approximation we have

$$S_0(\delta) = S_0(0) + \frac{dS_0(0)}{d\delta} \delta + \frac{1}{2!} \frac{d^2 S_0(0)}{d\delta^2} \delta^2 + \frac{1}{3!} \frac{d^3 S_0(0)}{d\delta^3} \delta^3 + O(\delta^4) \quad (57)$$

We next characterize each term in this approximation. From now on, without loss of generality, we assume that $\delta > 0$ and evaluate these terms by taking the limit as $\delta \downarrow 0$.

Order-zero: Using the boundary condition (55)

$$S_0(0) = \int_{-\bar{x}}^{\bar{x}} f(\hat{x}) d\hat{x} = 1. \quad (58)$$

First-order: Using a change of variables and imposing that $f(\hat{x}) = 0$ for all $\hat{x} < -\bar{x}$, we have that

$$S_0(\delta) = \int_{-\bar{x}}^{\bar{x}} f(\hat{x} + \delta) d\hat{x} = \int_{-\bar{x}-\delta}^{\bar{x}-\delta} f(\hat{x}) d\hat{x} = \int_{-\bar{x}}^{\bar{x}-\delta} f(\hat{x}) d\hat{x}. \quad (59)$$

Applying the Leibniz rule

$$S'_0(\delta) = -f(\bar{x} - \delta), \quad (60)$$

$$S'_0(0) = -\lim_{\hat{x} \uparrow \bar{x}} f(\hat{x}) = 0, \quad (61)$$

where the last equation uses the boundary conditions (54) and (56).

Second-order: Observe that

$$S''_0(\delta) = f'(\bar{x} - \delta) \quad (62)$$

and taking the limit, $S''_0(0) = \lim_{\hat{x} \uparrow \bar{x}} f'(\hat{x})$. We now characterize $\lim_{\hat{x} \uparrow \bar{x}} f'(\hat{x})$. Since the fraction of price change satisfies

$$n = \varphi + \frac{\sigma_z^2}{2} \left(\lim_{\hat{x} \downarrow \bar{x}} f'(\bar{x}) - \lim_{\hat{x} \uparrow \bar{x}} f'(\bar{x}) \right), \quad (63)$$

using symmetry $f(\hat{x}) = f(-\hat{x})$ if and only if $f'(\hat{x}) = -f'(-\hat{x})$, we have that

$$-\frac{(n - \varphi)}{\sigma_z^2} = S''_0(0). \quad (64)$$

Using the fact that $\sigma_z^2 = n\mathbb{E}[\Delta p^2]$

$$S''_0(0) = -\frac{(n - \varphi)}{n\mathbb{E}[\Delta p^2]}. \quad (65)$$

Third-order: Using the boundary conditions (53) and (56)

$$S'''_0(\delta) = -f''(\bar{x} - \delta) = -\frac{2\varphi}{\sigma_z^2} f(\bar{x} - \delta) \quad (66)$$

with $S'''_0(0) = 0$.

Taking all the results together,

$$S_0(\delta) = 1 - \frac{n - \varphi}{2n} \frac{\delta^2}{\mathbb{E}[\Delta p^2]} + O(\delta^4). \quad (67)$$

Thus,

$$\Delta n(\delta) = \frac{1 - l}{2} \frac{\delta^2}{\mathbb{V}[\Delta p]} + O(\delta^4). \quad (68)$$

□

F Additional Figures and Tables

Figure F.2 shows the distribution of price changes implied by the baseline single-product model and compares it to the data. Tables F.4 and F.5 report the parameterization of the single-product models with and without free price changes, for different values of σ and η . Table F.6 reports the calibration results when we consider alternative values of γ in our multi-product model. Figure F.3 plots the fraction of price changes in our model and the standard multi-product model for money shocks ranging from -15% to 15%. Figures F.4(a) and F.4(b) display the impact and cumulative output response to these shocks.

Figure F.2: Distribution of Price Changes

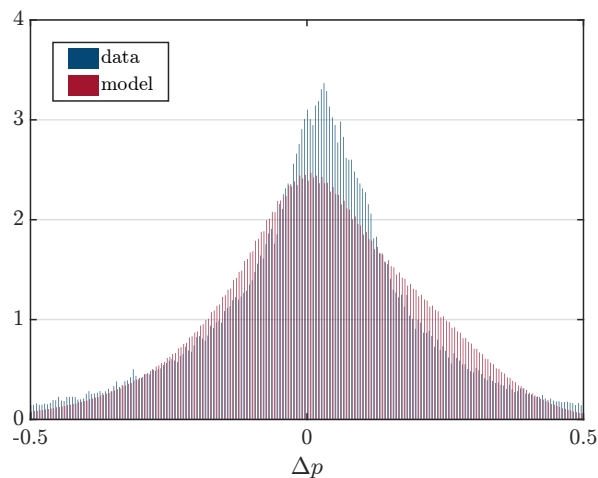


Table F.4: Parameterization of Single-Product Model with Free Price Changes

A. Moments				
	Data	$\sigma = 6$ $\eta = 1$	$\sigma = 3$ $\eta = 2/3$	$\eta = 1$
I. Targeted				
frequency Δp	0.116	0.116	0.116	0.116
mean Δp	0.018	0.018	0.018	0.018
std. dev. Δp	0.188	0.188	0.188	0.188
kurtosis Δp	3.609	3.609	3.609	3.609
std dev. $\pi_t(s)$, %	2.870	2.870	2.870	2.870
II. Not targeted				
<i>distribution of Δp</i>				
10 th percentile	0.018	0.020	0.020	0.019
25 th percentile	0.045	0.052	0.051	0.051
50 th percentile	0.104	0.116	0.114	0.114
75 th percentile	0.204	0.213	0.211	0.210
90 th percentile	0.334	0.316	0.317	0.318
B. Calibrated Parameter Values				
		$\sigma = 6$ $\eta = 1$	$\sigma = 3$ $\eta = 2/3$	$\eta = 1$
g_m	mean money growth rate	0.021	0.020	0.021
σ_z	s.d. idios. shocks	0.064	0.064	0.064
λ	1 - prob. free change	0.909	0.907	0.907
$\bar{\xi}$	upper bound menu cost	12.11	18.91	7.924
σ_e	s.d. sectoral shocks	0.010	0.011	0.010

Note: The money growth rate is annualized. In all calibrations, we set $\beta = 0.96$ (annualized).

Table F.5: Parameterization of Single-Product Model with no Free Price Changes

A. Moments				
	Data	$\sigma = 6$ $\eta = 1$	$\sigma = 3$ $\eta = 2/3$	$\eta = 1$
I. Targeted				
frequency Δp	0.116	0.116	0.116	0.116
mean Δp	0.018	0.018	0.018	0.018
std. dev. Δp	0.188	0.188	0.188	0.188
kurtosis Δp	3.609	<i>1.931</i>	<i>1.839</i>	<i>1.819</i>
std dev. $\pi_t(s)$, %	2.870	2.870	2.870	2.870
II. Not targeted				
<i>distribution of Δp</i>				
10 th percentile	0.018	0.079	0.079	0.079
25 th percentile	0.045	0.116	0.116	0.115
50 th percentile	0.104	0.165	0.165	0.165
75 th percentile	0.204	0.221	0.221	0.221
90 th percentile	0.334	0.274	0.277	0.277
B. Calibrated Parameter Values				
		$\sigma = 6$ $\eta = 1$	$\sigma = 3$ $\eta = 2/3$	$\eta = 1$
g_m	mean money growth rate	0.022	0.021	0.021
σ_z	s.d. idios. shocks	0.064	0.064	0.064
$\bar{\xi}$	upper bound menu cost	0.908	1.228	0.501
σ_e	s.d. sectoral shocks	0.009	0.010	0.009

Note: The money growth rate is annualized. We do not target the kurtosis and italicize its implied value. In all calibrations, we set $\beta = 0.96$ (annualized).

Table F.6: Alternative Parameterizations of Our Multi-Product Model

A. Moments			
	Data	$\gamma = 0$	$\gamma = 3$
I. Targeted			
fraction Δp	0.116	0.116	0.116
mean Δp	0.018	0.018	0.018
std. dev. Δp	0.188	0.188	0.188
std dev. $\pi_t(s)$	0.029	0.029	0.029
II. Not targeted			
kurtosis Δp	3.609	4.480	3.722
<i>distribution of Δp</i>			
10 th percentile	0.018	0.019	0.021
25 th percentile	0.045	0.050	0.054
50 th percentile	0.104	0.109	0.116
75 th percentile	0.204	0.199	0.208
90 th percentile	0.334	0.309	0.311
B. Calibrated Parameter Values			
		Our model	Standard
g_m	mean money growth rate	0.022	0.023
σ_z	s.d. idios. shocks	0.063	0.063
$\bar{\xi}$	upper bound menu cost	0.057	3.178
σ_e	s.d. sectoral shocks	0.010	0.011

Note: The money growth rate is annualized and the menu cost is relative to average sales.

Figure F.3: Impact Response of the Fraction of Price Changes

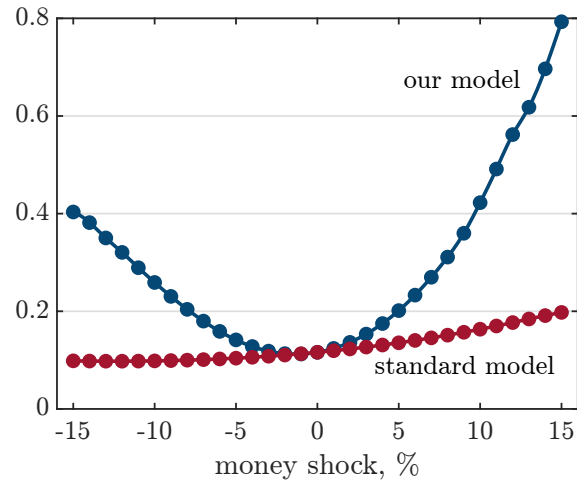


Figure F.4: Output Response to Money Shocks

