

Nonlinear Inflation Dynamics in Menu Cost Economies*

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

December 2022

Abstract

We show that canonical menu cost models, when parameterized to match the distribution of price changes, suffer three important shortcomings: they require implausibly large menu costs, they predict a large amount of misallocation, and they cannot reproduce the comovement between the frequency of price changes and inflation in the data. These shortcomings are amplified in the presence of microeconomic strategic complementarities. We resolve them by extending the standard multi-product menu cost model along two dimensions. First, we assume that strategic complementarities are at the firm, not the product level. Second, we assume that the products sold by a firm are imperfect substitutes. In contrast to standard models, the frequency of price changes increases rapidly with the size of monetary shocks, so our model implies nonlinear output responses. Even for small shocks, our model predicts stronger selection effects and therefore more flexible price responses and smaller real effects.

Keywords: menu costs, inflation, Phillips curve.

*We thank Audrey Azerot and Man Chon Iao for superb research assistance. The views expressed are those of the authors and not necessarily those of the Federal Reserve Board or the Federal Reserve System.

[†]University of Michigan, jablanca@umich.edu.

[‡]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

Macroeconomists often invoke menu costs as an important source of price rigidities. In menu cost economies, firms are more likely to adjust prices in response to large shocks, so Phillips curves are potentially non-linear. However, most work using menu cost models studies the responses to small aggregate shocks, often times using linear methods.¹ The recent rise in inflation in many modern economies suggests a need to understand how menu cost economies respond to large shocks. Understanding the causes of high inflation hinges critically on whether Phillips curves are as flat when inflation is high as they are when inflation is low.

In this paper, we study the importance of non-linearities in menu cost economies that reproduce the distribution of micro price changes. As recent work by [Alvarez et al. \(2016\)](#) showed, the latter is a critical determinant of the real effects of small monetary shocks in a large class of menu cost economies. We show that standard menu cost models with Gaussian idiosyncratic productivity shocks, both single- and multi-product, have three important shortcomings when calibrated to match this distribution. First, they require implausibly large menu costs and imply considerable profit losses from frictions to price adjustment. Second, they predict a large amount of misallocation from dispersion in prices. Third, they cannot reproduce the strong empirical comovement between inflation and the frequency of price changes. These shortcomings are amplified in the presence of microeconomic strategic complementarities that make a firm’s optimal reset price depend on the price of its competitors ([Klenow and Willis, 2016](#)).

We propose a resolution to these shortcomings by extending the standard multi-product menu cost model ([Alvarez and Lippi, 2014](#)) along two dimensions. First, we assume that strategic complementarities are at the firm, not the product level. Second, we assume that the products sold by a given firm are imperfect substitutes. Both of these assumptions reduce the amount of misallocation from inefficient price dispersion inside the firm and thus require lower menu costs to reproduce the frequency of price changes. These two extensions go a long way towards remedying the three shortcomings of canonical menu cost models.

We use our model to revisit the classic question of how large are the real effects of monetary policy. We find that the output responses to monetary shocks in our model are very different than those in the standard model. First, even in response to a small monetary shock, our

¹See, for example, [Dotsey et al. \(1999\)](#), [Goloso and Lucas \(2007\)](#), [Midrigan \(2011\)](#), [Vavra \(2013\)](#), [Alvarez et al. \(2016\)](#), [Alvarez et al. \(2022\)](#), [Auclert et al. \(2022\)](#). An exception is the work of [Karadi and Reiff \(2019\)](#) who study the response to large shocks in an economy similar to [Midrigan \(2011\)](#).

model predicts a much larger degree of price flexibility owing to a stronger selection effect (Golosov and Lucas, 2007). Second, output responds non-linearly to shocks of various sizes. The larger the shock is, the stronger the response of the frequency of price changes and therefore the smaller the real effects. Thus, our model predicts that the Phillips curve is highly non-linear.

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 disaggregated sectoral data to study the high-frequency comovement between sectoral inflation and the sectoral frequency of price changes. In line with previous work, we show that the frequency of price changes increases in periods of high inflation.² For example, when sectoral inflation is close to zero, the frequency of price changes is approximately 10% per month. In contrast, when inflation increases to 5%, the frequency of price changes averages 14%. We follow Klenow and Kryvtsov (2008) in decomposing movements in inflation into an intensive margin term that keeps the frequency of price changes constant, and an extensive margin term. As in Klenow and Kryvtsov (2008), the intensive margin term accounts for most of the movements in inflation in periods of low inflation, but for only half of these movements when inflation is relatively high. Thus, the frequency of price changes plays an important role at high levels of inflation.

We begin our analysis using a standard single-product menu cost model in which firms are subject to Gaussian idiosyncratic, sectoral and aggregate shocks. Firms face random menu costs of adjusting prices, as in Dotsey et al. (1999), and have occasional opportunities to change their price for free, as in Nakamura and Steinsson (2010). As in existing work, 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 idiosyncratic 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 frequency and distribution of price changes in the data. The calibrated model displays three shortcomings. First, it requires very large menu costs: the average amount of resources spent on changing prices in any given period is equal to 8.8% of GDP, much larger than the 1% direct estimate in the literature (Levy et al., 1997, Zbaracki et al., 2004). Since most price changes in this economy are free, the average menu

²This pattern has been documented before by Gagnon (2009), Nakamura et al. (2018), Alvarez et al. (2018) and Karadi and Reiff (2019) using aggregate data from episodes of high inflation in other countries.

cost does not fully reflect the losses firms face from frictions to price adjustment. We find that these frictions reduce the value of the firm by 44%, a sizable amount and at odds with the view that menu costs generate small losses to individual firms (Mankiw, 1985). Second, the model implies substantial misallocation: aggregate productivity is 21.6% lower than under flexible prices. Though this number is comparable to the estimates of misallocation reported by De Loecker et al. (2020) and Baqaee and Farhi (2018), these estimates encompass losses from all distortions, not just those due to menu costs. Third, the model predicts that the frequency of price changes is nearly constant, even in times of high inflation, at odds with the patterns we document in the data.

We next consider a multi-product setting in which each firm sells a continuum of products, each subject to idiosyncratic and firm-specific Gaussian quality shocks. There are economies of scope in price adjustment in that the firm can change the entire menu of its prices by paying a single fixed cost, assumed constant for all firms and time periods. We show that 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 the product-level price gaps of the firm, 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.

We show that economies of scope, on their own, do not remedy the three shortcomings enumerated above. We therefore extend the model along two dimensions, both of which reduce the misallocation from price dispersion within the firm. Our first assumption is that 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. Second, we assume that 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 product specific shocks are small. We show that the re-calibrated model goes a long way towards remedying the shortcomings of the standard menu cost models. For example, our model predicts that menu costs represent 2.4% of average firm revenue, much closer to the empirical estimates,

and that the losses from misallocation are only 2%. Importantly, our model predicts that the frequency of price changes increases much more in times of high inflation, half as much as it does in the data.

We use our model to revisit the real effects of monetary shocks by studying impulse response functions to one-time, unanticipated and permanent shocks of various sizes. We show that in our model impulse responses are very different than those predicted by standard models. First, output responds non-linearly to shocks because the frequency of price changes increases rapidly with the size of the shock. In contrast, in the standard model the output response scales nearly linearly with the size of the shock, even for monetary shocks as large as 15%, because the frequency of price changes is nearly constant. Second, even for small shocks, for which the frequency of price changes responds little, the real effects in our model are smaller than in the standard model. This is due to a much stronger selection effect.

We provide intuition for these two results by zooming in on a special case of our multi-product model, one in which the elasticity of substitution between the products sold by a given firm is equal to zero. In this case the duration of price spells is no longer a state variable because there is no misallocation inside the firm. We show that this version of the multi-product model is identical to a single-product model provided we adjust the trend growth rate of the money supply 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 and decision rules, they have different implications for the distribution of price change. The single-product economy generates a bi-modal distribution with most mass near the (s, S) thresholds. In contrast, the multi-product model matches the distribution of price changes in the data well due to the large product-specific shocks. Despite the different distributions of price changes, the two models respond identically to aggregate shocks. Thus, the multi-product model without misallocation inside the firm inherits the properties of the single-product model which, as [Golosov and Lucas \(2007\)](#) pointed out, features strong selection effects. Moreover, since our model requires less volatile firm-level shocks to reproduce the dispersion of price changes, the menu cost is lower, implying narrower (s, S) bands, and therefore more non-linear output responses.

We summarize our findings by tracing out the Phillips curve implied by monetary shocks. In our model, in contrast to the standard model, the Phillips curve is highly non-linear and becomes nearly vertical at inflation rates exceeding 10%.

2 Motivating Evidence

This section uses sectoral micro price data from the UK to corroborate that the frequency of price changes systematically increases with inflation, and that increases in the frequency of price changes account for a significant share of movements in inflation when inflation is relatively high. Though these facts have been documented in previous work using data for other countries,³ we use the numbers for the UK to quantitatively evaluate the ability of several variants of the menu cost model to reproduce this pattern. In contrast to existing work, we focus on data for individual sectors rather than the aggregate. As we show below, inflation is considerably more volatile in individual sectors compared to the aggregate, so sectoral variation allows us to better understand the relationship between inflation and frequency of price changes in periods of high inflation.

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 6-digit Classification of Individual Consumption by Purpose (COICOP 6). Each item in a given class is constructed with product-level price quotes by either sampling individual outlets or by collecting prices centrally (for example, university tuition fees). We exclude centrally-collected items, which account for approximately 26% of total consumer expenditure.

In computing price statistics, we use *regular price* series constructed by filtering V-shaped sales that last less than three months.⁴ Kehoe and Midrigan (2015) show that in theory temporary price changes do not contribute much to inflation dynamics. We next 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

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

³See, for example, Gagnon (2009) using data for Mexico, Nakamura et al. (2018) using data for the US, Alvarez et al. (2018) using data for Argentina, Karadi and Reiff (2019) using data for Hungary.

⁴We 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



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 denote 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.

2.2 Inflation and Frequency of Adjustment

We follow Klenow and Kryvtsov (2008) in decomposing movements in inflation into an extensive margin component that captures changes in the frequency of price adjustments and an intensive margin which captures movements in the average price change of firms that adjust.

Specifically, letting $f_t(s)$ 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 have⁵

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

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

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

and compare the dynamics of the actual inflation series $\pi_t(s)$ with the counterfactual $\pi_t^c(s)$ that shuts down movements in the frequency of price changes.

Figure 2 shows the fraction of price changes (left panel) and the two inflation series (right panel) for a specific COICOP 6 sector – “Bread and Cereals.” As earlier, our measure of the frequency of price changes is the average monthly fraction of (regular) price changes in the previous 12 months. Our measure of inflation is the year-to-year percent change in the sectoral price index.

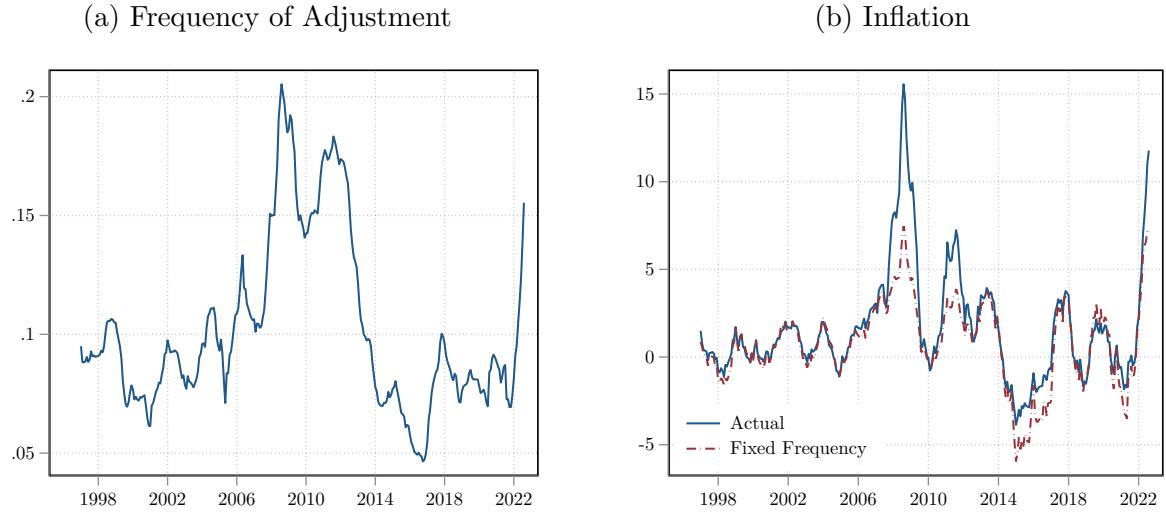
As the figure shows, the frequency of price changes in this sector fluctuates substantially over time, ranging from 5% to 20%. Consequently, the extensive margin of adjustment accounts for a sizable fraction of movements in inflation, especially during the 2007-2008 world food crisis, when actual inflation exceeded 15%, while counterfactual inflation only increased to 7%.

Figure 3(a) documents these patterns more systematically by presenting a binned scatterplot of the sectoral frequency of price changes against sectoral inflation rates pooling data from all sectors and weighing each by its expenditure share. We include sectoral fixed effects so our results capture high-frequency variation in sectoral inflation rates, not trend differences across sectors. The figure shows that the frequency of price changes systematically increases with inflation. For example, when inflation is in the neighborhood of zero, the frequency of price changes is approximately equal to 10% per month and when inflation is 5%, the frequency of price changes averages 14%.

To assess the importance of movements in the frequency of price changes for the dynamics of inflation, Figure 3(b) shows a binned scatterplot of the counterfactual inflation series that keeps the frequency of price changes constant at its historical average against realized

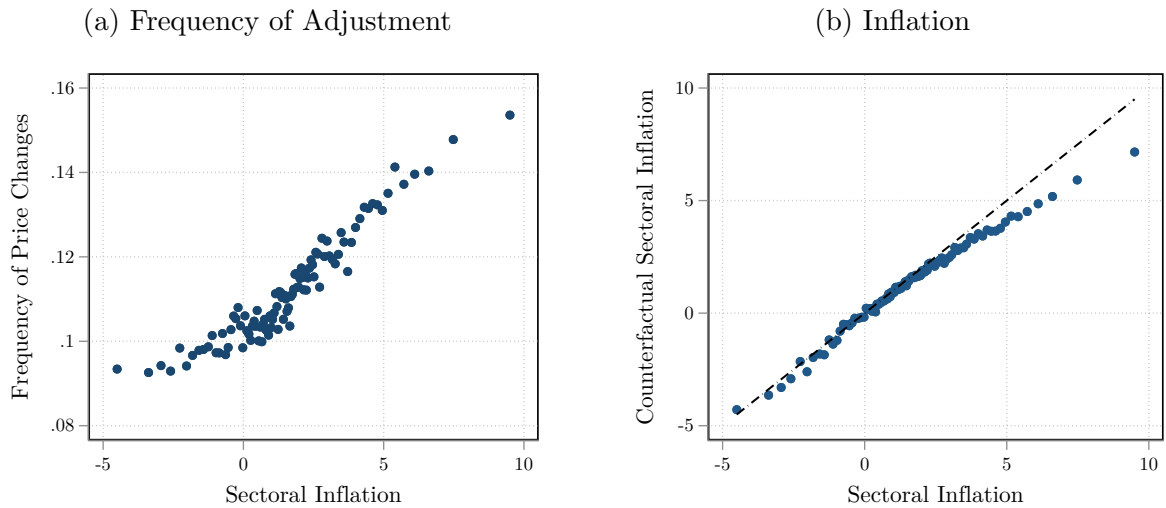
⁵All statistics are weighted using item-level consumption expenditure weights.

Figure 2: Inflation Decomposition: Bread and Cereals



Notes: The left panel plots the frequency of adjustment $f_t(s)$. The right panel plots $\pi_t(s) = \Delta_t(s)f_t(s)$ and $\pi_t^c(s) = \Delta_t(s)\bar{f}(s)$.

Figure 3: Inflation and the Frequency of Price Changes



Notes: The binned scatter controls for sectoral fixed effects and weighs individual sectors by the corresponding expenditure weights.

Table 1: Role of Extensive Margin

A. Inflation Volatility	
s.d. $\pi_t(s)$	2.87
s.d. $\pi_t^c(s)$	2.51
ratio	0.87

B. Slope of $\pi_t^c(s)$ on $\pi_t(s)$	
all observations	0.80
$\pi_t(s) > 75^{th}$ pct.	0.48
$\pi_t(s) > 90^{th}$ pct.	0.39

Notes: We compute the statistics in Panel A by first calculating the standard deviation of the two inflation series for each sector and then calculating the expenditure-weighted average of the sector-level standard deviations. We compute the slope coefficients in Panel B using an OLS regression that weighs observations for each sector using that sector’s expenditure weights.

inflation. Note that for 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 rates, above 4%, ignoring the extensive margin systematically underpredicts inflation.

Table 1 further corroborates these patterns. Panel A shows that the standard deviation of the counterfactual inflation series is equal to 2.51%, representing 87% of the standard deviation of actual inflation. This is consistent with [Klenow and Kryvtsov \(2008\)](#), who document that most movements in inflation are due to the intensive margin. However, Panel B shows that the extensive margin becomes much more important at higher rates of inflation. 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 and 0.39 when sectoral inflation exceeds its 75th and 90th percentile, respectively. We therefore conclude that at higher rates of inflation, the extensive margin accounts for more than one half of the variability of inflation.

3 Single-Product Menu Cost Model

We next show that the standard single-product menu cost, calibrated to match the distribution and frequency of price changes in the data, has three shortcomings. First, the model requires implausibly large menu costs to reproduce the frequency of price changes in the data. Second, it implies implausibly large losses from misallocation due to price dispersion. Third, it predicts that the frequency of price changes comoves little with inflation. As we show below, these shortcomings also apply to the standard multi-product model with economies of scope in price adjustment.

Because we are interested in studying the comovement between sectoral inflation and frequency of price changes, we assume an economy that consists of a continuum of ex-ante identical sectors. The output of each sector is used to produce a final consumption good using a Cobb-Douglas technology. Each sector consists of a continuum of firms, each producing a differentiated variety. Following [Midrigan \(2011\)](#) and much of the subsequent menu cost literature, we assume that idiosyncratic firm-level shocks are shocks to *quality*: they change both the costs of producing the good, as well as the consumers' demand for it, but absent menu costs leave the firm's profits unchanged. In addition to idiosyncratic shocks, we allow for shocks to sectoral productivity, as well as shocks to monetary policy. We follow [Golosov and Lucas \(2007\)](#) in assuming that preferences are logarithmic in consumption and linear in hours worked. This assumption is widely used in menu cost models and implies that the dynamic problem of firms in a given sector depends solely on that sector's inflation and marginal cost. This 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

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

where c_t is consumption, h_t is hours worked, P_t is the aggregate nominal price level, M_{t+1} is the money supply, B_{t+1} is the amount of government bonds, D_t denotes profits and T_t represents government transfers.

For simplicity, we assume that monetary policy is a money growth rule and that nominal spending is subject to a cash-in-advance constraint,

$$P_t c_t \leq M_t.$$

The Euler equation for bond holdings is then

$$\frac{1}{c_t} = \beta \mathbb{E}_t (1 + i_t) \frac{P_t}{P_{t+1}} \frac{1}{c_{t+1}},$$

and the cash-in-advance constraint binds if $i_t > 0$, as in our numerical experiments. Since the optimal labor supply choice satisfies

$$\frac{W_t}{P_t} = c_t,$$

the nominal wage is equal to the money supply

$$W_t = P_t c_t = M_t.$$

Note that the timing assumptions we make here are as in [Rotemberg \(1987\)](#) and imply that the cash-in-advance constraint does not distort the labor supply choice because labor income in period t can be used for consumption immediately.

3.2 Technology

We next describe the assumptions we make on technology.

3.2.1 Final Goods Producers

Final output is produced using a Cobb-Douglas aggregator across sectoral output $y_t(s)$

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

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 and the demand for a given sector's output is

$$y_t(s) = \left(\frac{P_t(s)}{P_t} \right)^{-1} y_t.$$

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 $\eta \leq 1$

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

where $e_t(s)$ is a productivity component common to all firms in sector s , $u_t(f, s)$ represents the quality of an individual firm f in that sector and $l_t(f, s)$ the amount of production labor the firm hires. As in [Burstein and Hellwig \(2008\)](#), decreasing returns to scale introduce a form of strategic complementarity across price setters, dampening the response of individual prices to aggregate and sectoral shocks.

Sectoral output is obtained by aggregating firm output using a Dixit-Stiglitz aggregator with elasticity of substitution σ

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

Notice that in addition to shifting the firm's productivity, the quality index $u_t(f, s)$ also acts as a demand shifter. If prices were flexible, firms would respond to an increase in $u_t(f, s)$ by reducing prices one-for-one, leaving quality adjusted prices, namely $u_t(f, s)P_t(f, s)$, and firm revenues unchanged. These quality shocks therefore provide a simple mechanism that changes firm's desired prices without requiring us to keep track of $u_t(f, s)$ as a state variable.⁶ For tractability, we assume that $e_t(s)$ and $u_t(f, s)$ follow random walk processes

$$\log e_{t+1}(s) = \log e_t(s) + \sigma_e \varepsilon_{t+1}^e(s) \quad (3)$$

and

$$\log u_{t+1}(f, s) = \log u_t(f, s) + \sigma_u \varepsilon_{t+1}^u(f, s),$$

where $\varepsilon_{t+1}^e(s)$ and $\varepsilon_{t+1}^u(f, s)$ are i.i.d innovations drawn from a standard normal distribution.

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

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

where

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

is the price index in sector s .

⁶An alternative approach, which we discuss in the Appendix, would be to assume that $u_t(f, s)$ only affects productivity and scale the menu costs accordingly.

3.3 Menu Costs and Firm Objective

As recent work emphasized, the aggregate implications of menu cost models are shaped by the distribution of individual firms' desired price changes as well as the shape of the adjustment hazards, which in turn determine the distribution of firm price changes (Caballero and Engel, 2007, Midrigan, 2011, Alvarez et al., 2016). We therefore 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 and with probability λ a price change requires a fixed cost $\xi_t(f, s)$. We follow Khan and Thomas (2008) in assuming that the fixed cost is an i.i.d. draw from a uniform distribution $U[0, \bar{\xi}]$ which gives rise to a smoothly increasing adjustment hazard, as in Costain and Nakov (2011) and Alvarez et al. (2021).

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

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

where $\tau = \sigma / (\sigma - 1)$ 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.

It is convenient to rewrite the firm's 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 do so, 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

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

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

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

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

$$x_t(s) = \left(\frac{a_t(s)}{\bar{a}} \right)^{-\eta}. \quad (4)$$

We note that 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

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[(1 + \tau) \left(\frac{x_t(f, s)}{x_t(s)} \right)^{1-\sigma} - a_t(s) \left(\frac{x_t(f, s)}{x_t(s)} \right)^{-\frac{\sigma}{\eta}} - \xi_t(f, s) \mathbb{I}_t(f, s) \right], \quad (5)$$

or using equation (4),

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(x_t(s)^{\sigma-1} \left[(1 + \tau) x_t(f, s)^{1-\sigma} - \bar{a} x_t(s)^{\left(\frac{1}{\eta}-1\right)(\sigma-1)} x_t(f, s)^{-\frac{\sigma}{\eta}} \right] - \xi_t(f, s) \mathbb{I}_t(f, s) \right). \quad (6)$$

Given a sectoral price gap $x_t(s)$, the firm level price gap that maximizes the firm's flow profits is equal to

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

so the exponent

$$\theta \equiv \frac{(\sigma-1)\left(\frac{1}{\eta}-1\right)}{1+\sigma\left(\frac{1}{\eta}-1\right)}$$

determines the strength of strategic complementarities, that is, the extent to which an individual firm's price depends on the price of its competitors.⁷ In particular, the lower η is or the higher σ is, the larger the value of θ , that is, the stronger strategic complementarities. Intuitively, if after an increase in the money supply most firms do not adjust their prices and therefore $x_t(s)$ is low, a firm that resets its price recognizes that if it were to increase its price it would experience an output drop. This output drop is increasing in the demand elasticity σ . Since marginal cost is upward sloping, more so the lower η is, the output drop would lead to a decline in marginal cost. The firm would therefore not respond to the money shock fully and instead keep its price gap close to $x_t(s)$. This complementarity thus amplifies the output responses to monetary shocks as it increases the effective degree of price stickiness in the economy.

Equation (6) 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 any other sectoral and aggregate variables. To see how the sectoral price gap $x_t(s)$ is determined in equilibrium, let

$$\hat{x}_t(f, s) = x_{t-1}(f, s) \frac{e_t(s)}{e_{t-1}(s)} \frac{u_t(f, s)}{u_{t-1}(f, s)} \frac{M_{t-1}}{M_t}$$

⁷Notice that we define θ to be the negative of the corresponding parameter in Alvarez et al. (2022).

denote the firm's individual state variable. If a firm with this state variable does not adjust, its price gap is $x_t(f, s) = \hat{x}_t(f, s)$ and therefore falls if the money supply expands and increases if sectoral or firm productivity increase. If, in contrast, 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 for a firm with state \hat{x} , the probability that the firm adjusts is given by

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

Given the optimal reset price $x_t^*(s)$, the adjustment hazard $h_t(\hat{x}; s)$ and the firm distribution $F_t(\hat{x}; s)$, the sectoral price gap satisfies

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

3.4 Aggregation

The frictions to price adjustment give rise to a non-degenerate distribution of price gaps $x_t(f, s)$ and therefore dispersion in markups across firms and time. This, in turn, leads to productivity losses from misallocation and inefficient fluctuations in sectoral and aggregate employment. To characterize these, note first that we can recover the firm's markup

$$\mu_t(f, s) = \eta(1 + \tau) \frac{P_t(f, s) y_t(f, s)}{W_t l_t(f, s)},$$

which can be rewritten as a function of its own price gap and the sectoral price gap,

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

Following [Edmond et al. \(2018\)](#), we can then calculate the sectoral markup as the sales-weighted harmonic average of individual firm markups,

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

We can also derive a sectoral production function that determines how much output $y_t(s)$ a given sector produces using a total amount of labor

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

Specifically, the sectoral production function is

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

where $\phi_t(s)$ summarizes the losses from misallocation due to dispersion in relative price gaps

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

When prices are flexible, $x_t(f, s) = x_t(s) = 1$ and $\phi_t(s) = 1$. More generally, dispersion in relative prices reduces $\phi_t(s)$ below 1, 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 losses from misallocation.

In addition to reducing sectoral productivity, menu costs also generate inefficient fluctuations in sectoral employment. In particular, the amount of labor used by a sector is

$$l_t(s) = \frac{(1 + \tau) \eta}{\mu_t(s)} \tag{7}$$

and is decreasing in the sectoral markup. In contrast, if prices were flexible, sectoral employment would be $l_t(s) = \eta$.

Lastly, we can derive the aggregate implications of frictions to price adjustment in a similar fashion by noting that the aggregate markup is the harmonic average of sectoral markups

$$\mu_t = \left(\int \mu_t(s)^{-1} ds \right)^{-1},$$

and that aggregate productivity is equal to

$$\bar{e}_t \equiv \frac{y_t}{l_t^\eta} = \left(\int \left(\frac{x_t}{x_t(s) \phi_t(s)} \right)^{\frac{1}{\eta}} ds \right)^{-\eta},$$

where

$$x_t = \exp \left(\int \log \frac{x_t(s)}{e_t(s)} ds \right).$$

3.5 Parameterization

Table 2 reports the result of the parameterization. A period is a month. We set the discount factor β to an annualized value of 0.96, the demand elasticity σ to 6, implying a flexible price

markup of 20%, and the elasticity of labor in the production function η to 2/3. These are values commonly used in the literature. We calibrate the remaining parameters to match moments on the frequency and size of price changes in the UK micro data reported in Panel A. We compute the counterparts of these statistics in the steady-state of the model, with no sectoral or aggregate shocks.

We calculate the sectoral frequency of price changes in the data as the harmonic weighted average of the frequency of price changes of individual product categories (items) that belong to that sector. We follow [Klenow and Kryvtsov \(2008\)](#) in standardizing the distribution of price changes by the respective item-level mean and variance. That is, we calculate for each price change $\Delta p_{it}(j)$ of quote i that belongs to product category j the 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.

As Panel A shows, the model is able to match the targeted moments well. The frequency of price changes is equal to 11.6% per month. The model reproduces well the kurtosis of non-zero price changes in the data (3.65 vs. 3.61). Despite its parsimony, the model matches well the prevalence of both large and small price changes, as well as the distribution of the size (absolute value) of price changes. For example, 10% of price changes are less than 1.8% in absolute value in the data and 2.0% in the model. As Panel B shows, the model requires very dispersed productivity shocks ($\sigma^u = 0.067$) and a large probability of free price changes ($1 - \lambda = 0.09$). Free price changes thus account for 78% (0.09/0.116) of all price changes.

We next highlight the first shortcoming of the menu cost model: it requires implausibly large menu costs to reproduce the data. The upper bound of the distribution of menu costs, $\bar{\xi}$, is 43 times average monthly firm sales. Since firms only pay the menu cost when they draw a sufficiently small one, the average amount of resources spent on costs of adjusting prices in a given period is equal to 11.6% of the wage bill or 8.8% 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. Because of the prevalence of free price changes, the magnitude of menu costs alone is not informative about the costs of price adjustment. We thus find it useful to compare the life-time value of a firm in our menu cost economy to that in a counterfactual setting in which the firm faces no frictions to price adjustment. The value of the firm in the menu cost economy is 44% lower than in the absence of menu costs.

Thus, menu costs are akin to a 44% tax on firm profits, a sizable amount, and at odds with the view that firms suffer small losses from barriers to changing prices (Mankiw, 1985).

Table 2: Parameterization of Single-Product Model

A. Moments		
	Data	Model
frequency Δp	0.116	0.116
<i>distribution of Δp</i>		
mean	0.018	0.016
std. dev.	0.188	0.195
kurtosis	3.609	3.649
5 th percentile	-0.327	-0.316
10 th percentile	-0.226	-0.227
25 th percentile	-0.081	-0.099
50 th percentile	0.026	0.017
75 th percentile	0.119	0.143
90 th percentile	0.247	0.266
95 th percentile	0.340	0.333
<i>distribution of Δp</i>		
mean	0.142	0.152
std. dev.	0.125	0.124
5 th percentile	0.009	0.011
10 th percentile	0.018	0.020
25 th percentile	0.045	0.055
50 th percentile	0.104	0.120
75 th percentile	0.204	0.220
90 th percentile	0.334	0.326
95 th percentile	0.413	0.392


B. Parameter Values

<i>Assigned</i>		<i>Calibrated</i>		
β	discount factor	0.96	g_m mean money growth rate	0.021
σ	demand elasticity	6	σ^u s.d. idios. shocks	0.067
η	labor elasticity	2/3	λ 1 - prob. free change	0.910
			$\bar{\xi}$ upper bound menu cost	43.23

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

Figure 4 provides some intuition for why the model requires such large menu costs and such a large probability of free price changes to match the micro data. The left panel plots the distribution of desired price changes, $f(\Delta p)$, where $\Delta p = \log x^*/\hat{x}$ is the firm's

Figure 4: Distribution of Desired Price Changes and Adjustment Hazard



desired price change, and superimposes the probability of adjustment, $h(\Delta p)$, as a function of the desired price change. The right panel plots the distribution of actual price changes conditional on adjustment, $g(\Delta p) \sim h(\Delta p)f(\Delta p)$. Since many price changes in the data are in the neighborhood of zero, the adjustment hazard at zero, $1 - \lambda$, must be quite high. In turn, reproducing the tails of the price change distribution requires a relatively flat hazard, that is, large menu costs. This allows the model to reproduce the fraction of producers whose prices are far away from the optimum and thus desire a large price change.

We next highlight the second shortcoming of the menu cost model: it generates implausibly large losses from misallocation. Table 3 summarizes the model's implications for the distribution of markups across firms within a sector. The average cost-weighted markup, the object that determines the sector's labor share, is equal to 1.195, in the middle of the range of existing estimates. The average sales-weighted markup is equal to 1.592. As pointed out by Edmond et al. (2018), the sales-weighted average markup is equal to the cost-weighted markup plus a term that captures the dispersion of markups across firms. Thus, the model features a great deal of markup dispersion. Indeed, the table shows that markups range from a 10th percentile of 0.50 to a median of 1.04 and a 90th percentile of 1.94, so a sizable fraction of producers sell at a price substantially below marginal cost. Overall, the model predicts a great deal of misallocation: aggregate productivity is 21.63% lower than under flexible prices. This number is comparable to the estimates of misallocation reported by De Loecker et al.

Table 3: Distribution of Markups

cost-weighted average markup	1.195
sales-weighted average markup	1.592
<i>cost-weighted distribution of markups</i>	
10 th percentile	0.496
25 th percentile	0.691
50 th percentile	1.043
75 th percentile	1.585
90 th percentile	1.940
misallocation losses, %	21.63

(2020) and Baqaee and Farhi (2018). However, the latter 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.⁸ It is implausible that menu costs alone account for all observed misallocation in the data.

3.6 Role of Extensive Margin of Adjustment

We next illustrate the third shortcoming of the menu cost model: in contrast to the data, the frequency of price changes comoves little with inflation, rendering the extensive margin of price adjustments unimportant for inflation dynamics.

To that end, we subject firms to sectoral productivity shocks $e_t(s)$ that evolve according to equation (3), and we choose the standard deviation σ_e to match the average standard deviation of sectoral inflation.⁹ We use the Krusell and Smith (1998) approach to solve the firm’s problem in the presence of sectoral shocks. Recall that the firm’s problem depends only on the current and future sectoral price gaps $x_t(s)$. We postulate that the firm’s perceived law of motion for the sectoral price gap is a function of the current gap and the sectoral productivity shock

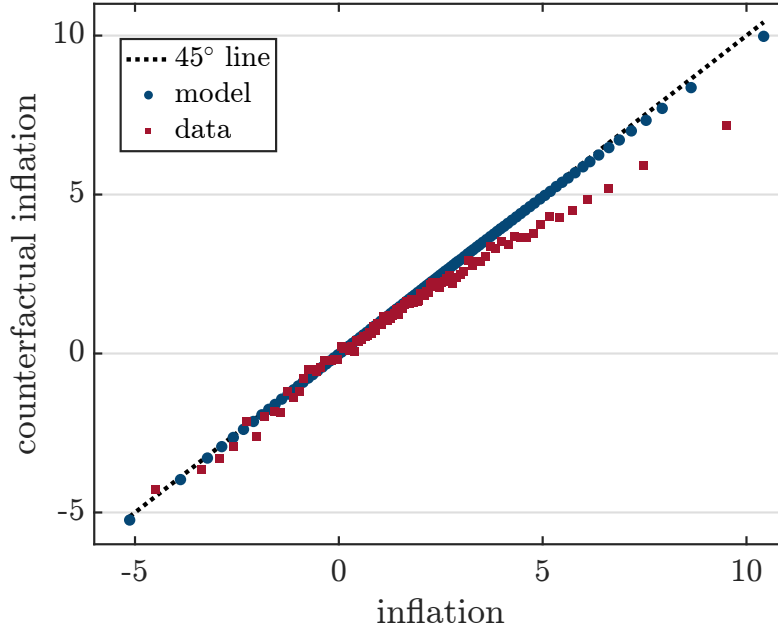
$$x_{t+1}(s) = \mathcal{X}(x_t(s), \varepsilon_{t+1}(s)).$$

Since we are interested in characterizing potentially nonlinear responses to shocks, we parameterize $\mathcal{X}(\cdot)$ using Chebyshev polynomials. For any given guess of $\mathcal{X}(\cdot)$, we solve the firm’s

⁸See, for example, Foster et al. (2022) for a discussion.

⁹When we add sectoral shocks, we reduce the standard deviation of idiosyncratic shocks σ_u to ensure that the model reproduces the moments in Table 2. For now, we assume there are no aggregate monetary shocks.

Figure 5: Importance of the Extensive Margin



decision rules, simulate histories of sectoral productivity shocks, and find the sectoral price gap $x_t(s)$ that is consistent with the firm’s decision rules. We then use projection methods to update our guess of $\mathcal{X}(\cdot)$ using simulated data on $x_t(s)$ and $\varepsilon_t(s)$, and iterate until convergence. We find that the [Krusell and Smith \(1998\)](#) approach works well in this setting, in that the R^2 in the perceived law of motion exceeds 0.9999.

Figure 5 assesses the model’s ability to reproduce the empirical comovement between the counterfactual inflation series that shuts down fluctuations in the frequency of price adjustments $\pi_t^c(s)$ and actual inflation $\pi_t(s)$. In contrast to the data, in the model $\pi_t^c(s)$ and $\pi_t(s)$ comove one-for-one, suggesting that the extensive margin of price adjustment plays no role in inflation fluctuations, even at high rates of inflation.

Table 4 corroborates this point. Panel A shows that the standard deviation of counterfactual inflation is nearly as large as that of actual inflation, 2.83 vs. 2.87. Panel B shows that the slope coefficients from regressing $\pi_t^c(s)$ on $\pi_t(s)$ are close to one, even when we restrict the sample to periods when sectoral inflation exceeds its 75th and 90th percentile. The patterns are at odds with the data.

To summarize, the single-product menu cost model, when calibrated to match the distribution of micro price changes in the data, suffers three important shortcomings. First, the model requires implausibly large menu costs and implies considerable firm losses from frictions to price adjustment. Second, the model predicts a large amount of misallocation

Table 4: Importance of the Extensive Margin

Inflation Volatility		
	Data	Model
s.d. $\pi_t(s)$	2.87	2.87
s.d. $\pi_t^c(s)$	2.51	2.83
ratio	0.87	0.99

Slope of $\pi_t^c(s)$ on $\pi_t(s)$		
	Data	Model
all observations	0.80	0.99
$\pi_t(s) > 75^{th}$ pct.	0.48	0.94
$\pi_t(s) > 90^{th}$ pct.	0.39	0.92

from dispersion in marginal revenue products. Third, the model cannot reproduce the strong comovement between the frequency of price changes and inflation observed in the data. As the robustness section shows, even when $\eta = 1$, so there are no strategic complementarities, the single-product model continues to predict that the frequency of price changes is approximately constant, even in times of high inflation, and that menu costs and misallocation are large, albeit less so than with strategic complementarities. Moreover, as we show below, these shortcomings also apply to the canonical multi-product menu cost model with economies of scope in price setting.

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 \(2014\)](#).¹⁰ Specifically, we assume that each firm sells a unit measure of products, each subject to idiosyncratic and firm specific quality shocks. Moreover, we assume that the firm can change the entire menu of prices by paying a single fixed cost. Intuitively, this model can match the large number of both small and large price changes because the price adjustment decision is determined by the distribution of price gaps across all products, not individual price gaps.

We show that economies of scope, on their own, do not remedy the three shortcomings of

¹⁰See [Bhattarai and Schoenle \(2014\)](#) and [Bonomo et al. \(2022\)](#) for evidence on multi-product pricing.

the menu cost model enumerated above. We therefore extend the model along two dimensions, both of which reduce the misallocation from price dispersion within the firm and allow us to substantially remedy the shortcomings of the standard menu cost model. Intuitively, since in a multi-product economy firms adjust when misallocation within the firm is large, reducing it lowers the menu costs that the model requires to reproduce the frequency of price changes and thus the amount of misallocation. Moreover, since menu costs are small, the extensive margin of price adjustment becomes more important, as in the data.

Our first assumption 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. As we show below, this implies that though there are decreasing returns to scale at the firm level and therefore strategic complementarities across firms, the losses from misallocation within the firm are lower relative to a multi-product model where the decreasing returns to scale are at the product level. Second, we assume that individual products sold by a given firm are imperfect substitutes, an assumption that further reduces the losses from misallocation within a firm. Our notion of a product is a collection of highly substitutable goods that are 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 shocks. In contrast, different pastries sold by Starbucks, while highly substitutable among themselves, are much less substitutable with tea.

To conserve space, we present the general version of our model only and discuss in passing various special cases. Since the model shares many elements with the single-product model above, we only discuss the new ingredients of the model we introduce here.

4.1 Technology

The output $y_t(f, s)$ of individual firms is aggregated into a final good using the same aggregators as in (2) and (1). A given firm produces a continuum 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 $z_{it}(f, s)$, the quality of product i , 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. The demand for an individual product is given by

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

where

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

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

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

$$y_{it}(f, s) = e_t(s) u_t(f, s) z_{it}(f, s) m_{it}(f, s)^{1-\eta} l_{it}(f, s)^\eta. \quad (8)$$

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

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

For a given amount of labor used in the production of product i , the optimal choice of the fixed factor is

$$m_{it}(f, s) = \frac{l_{it}(f, s)}{l_t(s)}, \quad (9)$$

where

$$l_t(f, s) = \int l_{it}(f, s) di = \left(\int \frac{y_{it}(f, s)}{e_t(s) u_t(f, s) z_{it}(f, s)} di \right)^{\frac{1}{\eta}} \quad (10)$$

is the total amount of labor used by the firm. Substituting equation (9) into equation (8) reveals that at the product level the technology is linear in the amount of labor used by product i .

We make two remarks. First, eliminating dispersion in $z_{it}(f, s)$ across products allows us to nest the single-product menu cost model as a special case. Second, assuming instead that the fixed factor is immobile across products, so that $m_{it}(f, s) = 1$ and equation (8) features decreasing returns to scale at the product level, implies that the amount of labor the firm uses is given by

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

As we show below, under this alternative assumption, the losses from misallocation within the firm are larger compared to our baseline.

4.2 Menu Costs and Firm Objective

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 menu costs are random nor that some price changes are free. We let $\bar{\xi}$ denote the deterministic cost a firm must incur to change the entire menu of prices. The firm's life-time value is then given by

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

Let

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

denote the price gap of product i , which now also scales the price by the product's idiosyncratic quality $z_{it}(f, s)$. The firm's price gap is simply

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

Aggregating the labor used by different products allows us to derive a firm-level production function

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

where

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

captures the losses from misallocation arising from price gap dispersion inside the firm. Notice that this firm-level production function features decreasing returns to scale, which imply that the firm's optimal price gap $x_t(f, s)$ depends on the price gap of its competitors $x_t(s)$, as in the single-product menu cost model.

We note that if the specific input is instead fixed at the product level, the firm-level production function is unchanged and the expression for misallocation is instead

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

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

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

where $\mathbb{I}_t(f, s)$ is an indicator for whether the firm changes its menu of prices. Thus, the only difference between this objective and that of a single-product firm is that in a single-product firm there is no misallocation inside the firm.

4.3 Recursive Formulation

We build on the insights of [Alvarez and Lippi \(2014\)](#) and summarize the distribution of price gaps inside the firm $x_{it}(f, s)$ with sufficient statistics. In contrast to their work, which considers a quadratic approximation to the firm's objective, here we solve the model non-linearly and have to keep track of two 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, we have that the firm's composite price index evolves over time 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 goods with cheaper quality-adjusted prices and the firm's composite price falls.

The first state variable we keep track of is

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

the price gap the firm would have in the absence of price changes. If the firm resets its prices, its price 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

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

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 therefore there is no misallocation inside the firm, $\phi_t(f, s) = 1$. Over time, the losses from misallocation increase because the distribution of price gaps becomes more dispersed

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

Thus, the losses from misallocation for a firm whose prices are d periods old are

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

For a given duration, these losses are increasing in the elasticity of substitution γ and the volatility of idiosyncratic shocks.

We note that assuming instead that the specific factor is immobile across products implies

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

so the losses from misallocation are larger than in our baseline model whenever $\eta < 1$.

Dropping the dependence on s and letting x_t denote the sectoral price gap, the value of not adjusting is given by

$$v_t^n(\hat{x}, d) = x_t^{\sigma-1} \left[(1 + \tau) \hat{x}^{1-\sigma} - \bar{a} x_t^{(\frac{1}{\eta}-1)(\sigma-1)} \hat{x}^{-\frac{\sigma}{\eta}} \exp\left(\frac{\gamma d \sigma_z^2}{\eta \frac{\sigma_z^2}{2}}\right) \right] + \beta \mathbb{E}_t \max(v_{t+1}^n(\hat{x}', d+1), v_{t+1}^a - \bar{\xi}),$$

where the law of motion for the firm's price gap is given by

$$\hat{x}' = \exp\left(\left(1 - \gamma\right) \frac{\sigma_z^2}{2} + \sigma_e \varepsilon_{t+1}^e + \sigma_u \varepsilon_{t+1}^u(f) - g_m\right) \hat{x}_t. \quad (11)$$

The value of adjusting is

$$v_t^a = \max_{x_t^*} x_t^{\sigma-1} \left[(1 + \tau) (x_t^*)^{1-\sigma} - \bar{a} x_t^{(\frac{1}{\eta}-1)(\sigma-1)} (x_t^*)^{-\frac{\sigma}{\eta}} \right] + \beta \mathbb{E}_t \max(v_{t+1}^n(\hat{x}', d+1), v_{t+1}^a - \bar{\xi}).$$

The law of motion for the price gap \hat{x}' is similar to equation (11), with x_t^* replacing \hat{x}_t .

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, with weights given by the distribution over state variables conditional on adjustment $g_t(\hat{x}, d; s) \sim h_t(\hat{x}, d; s) f_t(\hat{x}, d; s)$.

4.4 Parameterization

Table 5 shows the parameterization of the model. In our baseline economy, which we refer to as *our model*, we set $\gamma = 1$, $\sigma = 6$ and assume that the specific factor is mobile across products.¹¹ To illustrate the role of the two departures from the standard multi-product model, we also show results from a *standard* multi-product economy in which $\gamma = \sigma = 6$ and the specific factor is immobile across products.

Panel A shows that our model reproduces the distribution of price changes well. Panel B shows that the volatility of firm-level shocks σ_u is lower than that of product-level shocks, σ_z . Intuitively, the relative volatility of these two shocks is identified by the distribution of price changes. Specifically, as σ_z/σ_u goes to zero, the model becomes the single-product model which cannot reproduce the large number of small price changes given the assumption of a constant fixed cost. In contrast, as σ_z/σ_u goes to infinity, the law of motion for \hat{x}_t is deterministic, so all firms change prices after the same number of periods, as in the Taylor model. In this case, the distribution of prices changes conditional on adjustment is normal, so the model cannot match the excess kurtosis in the data. In this sense, the relative volatility of the two shocks in our model plays the same role as the number of products does in Alvarez and Lippi (2014). Notice that the standard multi-product model also reproduces the distribution of price changes well.

We next discuss the ability of our model to remedy the shortcomings of the single-product menu cost model. Consider first the implied menu costs. As Panel B of Table 5 shows, our model requires much smaller costs of changing prices: $\bar{\xi} = 0.21$ of average firm sales. Since firms adjust infrequently, the average amount of resources spent on costs of price adjustment in a given period is equal to 2.4% of average firm sales. Though this is still larger than the 1% estimate in the literature, it is substantially smaller compared to what the single-product model predicts. As before, we find it useful to compare the life-time value of a firm in the economy with menu costs with that of a firm that faces no frictions to price adjustment. In our model menu costs reduce firm value by as much as a 7.9% tax on firm profits would. This is much smaller than the implicit tax on profits in the single-product model. Notice that the standard multi-product model also implies implausibly large menu costs: $\bar{\xi} = 2.21$ of average sales, implying that the total resource cost of changing prices is 25.8% of average sales.

Consider next our model's implications for dispersion in firm markups and the losses from misallocation. Table 6 shows that markups range from a 10th percentile of 1.02 to a

¹¹In the robustness section we report results for alternative values of γ .

Table 5: Parameterization of Multi-Product Model

A. Moments			
	Data	Our model	Standard
frequency Δp	0.116	0.116	0.117
<i>distribution of Δp</i>			
mean	0.018	0.017	0.025
std. dev.	0.188	0.196	0.196
kurtosis	3.609	3.566	3.508
5 th percentile	-0.327	-0.312	-0.325
10 th percentile	-0.226	-0.232	-0.233
25 th percentile	-0.081	-0.107	-0.093
50 th percentile	0.026	0.023	0.042
75 th percentile	0.119	0.119	0.160
90 th percentile	0.247	0.256	0.260
95 th percentile	0.340	0.330	0.317
<i>distribution of Δp</i>			
mean	0.142	0.154	0.157
std. dev.	0.125	0.123	0.120
5 th percentile	0.009	0.012	0.013
10 th percentile	0.018	0.023	0.025
25 th percentile	0.045	0.060	0.064
50 th percentile	0.104	0.127	0.133
75 th percentile	0.204	0.218	0.224
90 th percentile	0.334	0.322	0.320
95 th percentile	0.413	0.392	0.385

B. Calibrated Parameter Values

		Our model	Standard
g_m	mean money growth rate	0.023	0.035
σ^u	s.d. firm shocks	0.025	0.037
σ^z	s.d. product shocks	0.062	0.058
$\bar{\xi}$	menu cost	0.207	2.207

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


90th percentile of 1.38 and imply that aggregate productivity is only 1.97% lower than under flexible prices. These losses are much smaller than in the standard models, both single- and multi-product, where productivity is 21.63% and 21.24% lower, respectively.

Figure 6 plots the distribution of firm price gaps $\hat{x}_t(f, s)$ in a given sector, as well as the adjustment hazard for firms that last adjusted 6 and 12 months ago. We make three

Table 6: Distribution of Markups


cost-weighted average markup	1.194	1.191
sales-weighted average markup	1.210	1.285
<i>cost-weighted distribution of markups</i>		
10 th percentile	1.019	0.827
25 th percentile	1.088	0.918
50 th percentile	1.195	1.103
75 th percentile	1.266	1.405
90 th percentile	1.382	1.793
misallocation losses, %	1.97	21.24

Figure 6: Distribution of Price Gaps



observations. First, multi-product economies feature less dispersion in the price gaps than the single-product economy. This is because firm level shocks are less dispersed, as now a lot of the dispersion in price changes is due to product specific shocks. Second, because menu costs are smaller in our model, the (s, S) bands are narrower than in the standard model. As we show below, the width of the (s, S) bands has important implications for how the economy responds to large aggregate shocks. The narrower the bands are, the larger the fraction of firms that end up outside the bands and therefore adjust prices after a large shock. Third, the (s, S) bands narrow with the duration of prices because the increase in misallocation inside the firm reduces the value of inaction relative to that of adjustment.

Figure 7: Importance of the Extensive Margin



4.5 Role of Extensive Margin of Adjustment

We next evaluate the ability of our model to reproduce the comovement between the frequency of price changes and inflation. As before, we subject firms to sectoral productivity shocks and use the [Krusell and Smith \(1998\)](#) approach to solve the firm’s problem. 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 0.9998.

Figure 7 illustrates our model’s prediction regarding the relationship between inflation $\pi_t(s)$ and the counterfactual inflation series $\pi_t^c(s)$ that keeps the frequency of price changes constant. At low levels of inflation $\pi_t(s)$ and $\pi_t^c(s)$ comove one-for-one, as in the data. At higher inflation, keeping the frequency of price changes constant underpredicts inflation, although less so in the model compared to the data.

Table 7 further corroborates this point. In our model, counterfactual inflation accounts for 89% of the volatility of actual inflation, close to the 87% in the data. Moreover, as in the data, the extensive margin becomes particularly important at high levels of inflation: the slope coefficient from regressing $\pi_t^c(s)$ on $\pi_t(s)$ is 0.89 for the overall sample, but falls to 0.72 and 0.64 when we restrict the sample to periods in which inflation exceeds its 75th and 90th percentile. Thus, our model accounts for half of the importance of the extensive margin for inflation dynamics. In contrast, the standard multi-product menu cost model predicts a much smaller role for the extensive margin of price adjustment.

Table 7: Importance of the Extensive Margin

Inflation Volatility			
	Data	Our model	Standard
s.d. $\pi_t(s)$	2.87	2.87	2.86
s.d. $\pi_t^c(s)$	2.51	2.55	2.70
ratio	0.87	0.89	0.94


Slope of $\pi_t^c(s)$ on $\pi_t(s)$			
	Data	Our Model	Standard
all observations	0.80	0.89	0.94
$\pi_t(s) > 75^{th}$ pct.	0.48	0.72	0.82
$\pi_t(s) > 90^{th}$ pct.	0.39	0.64	0.78

To summarize, we show that extending the multi-product menu cost model to reduce the amount of misallocation within the firm goes a long way towards remedying the three shortcomings of the standard menu cost models we highlighted.

5 Real Effects of Monetary Shocks

We next 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 a monetary shock? We show that output responses in our model are very different than those in the standard model that we argued is inconsistent with the data. We show two results. First, in our model output responds non-linearly to shocks of various sizes. The larger the shock is, the stronger the response of the frequency of price changes and therefore the smaller the real effects. Thus, our model predicts that the slope of the Phillips curve is non-linear. Second, even for small shocks the real effects are smaller in our model compared to the standard model, even though both match the distribution of price changes equally well. Thus, the distribution of price changes no longer pins down the output responses, in contrast to standard models (Alvarez et al., 2016).

Figure 8: Output Response to Money Shock



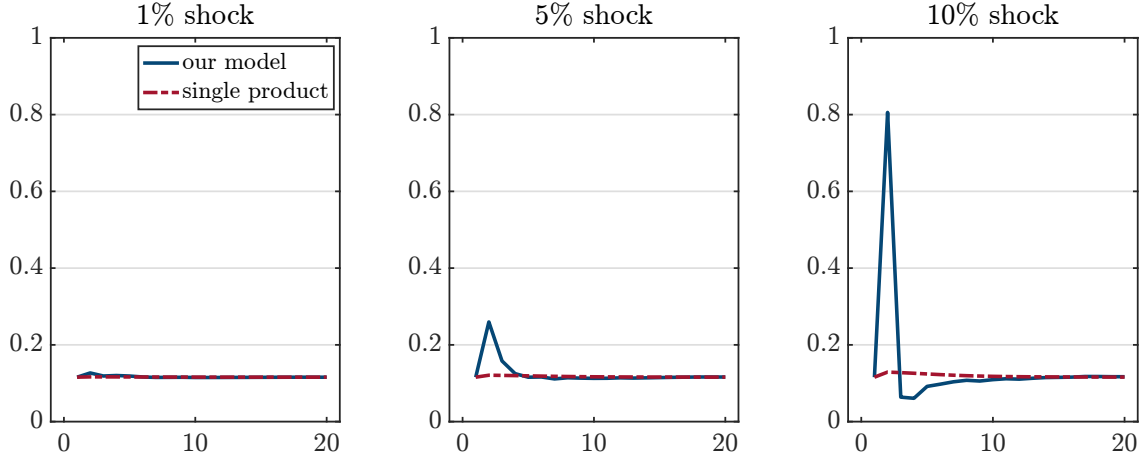
5.1 Impulse Response to Monetary Shocks

We start by reporting the impulse response of aggregate output y_t to a one-time, unanticipated and permanent increase of 1%, 5% and 10% in the money supply M_t , starting from the steady-state of the model without aggregate or sectoral uncertainty. Figure 8 plots these responses in our model and the standard single-product model. To ease comparison, we rescale the y-axis by the size of the shock. We make two points. First, even for a small shock of 1% the real effects of monetary policy are smaller in our model. Since $P_t y_t = M_t$, this simply reflects that the aggregate price level is more flexible in our model. Second, while in the single-product model the impulse response scales linearly with the shock, in our model larger shocks imply disproportionately smaller output responses. Indeed, when the shock is large (10%) the output response on impact is smaller than when the shock is 5%. Thus, the larger the shock is, the larger the discrepancy between the real effects of the shock in our model and the standard model.

We next explain why this is the case. First, Figure 9 plots the corresponding impulse responses of the frequency of price adjustment. In contrast to the standard model, in which the frequency responds very little to shocks of all sizes, in our model the frequency increases after a large money shock. For example, though the frequency responds little to a money shock of 1%, it nearly doubles to 25% for a 5% shock and jumps to 80% after a 10% shock.

Second, 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), starts from the observation that, up to a first-order approximation, inflation in the absence of the shock is

Figure 9: Frequency Response to Money Shock



equal to

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

where ω is the desired price change, $h(\omega)$ is the adjustment hazard and $f(\omega)$ is the steady-state distribution of desired price changes. 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 log reset price in the first period after the money shock and x^* is 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 [Goloso 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.

Table 8 reports the results of this decomposition. We make two observations. First, in our model money shocks have a larger pass-through to inflation than in the single-product model. This is true even for a small shock of 1%, for which the pass-through is 0.32 in our model and only 0.13 in the single-product model. This difference is almost entirely accounted for by a stronger selection effect, which is responsible for two thirds of the overall response in our model. Second, in our model the pass-through increases rapidly with the size of the shock: from 0.32 for a 1% shock to 0.86 for a 10% shock vs. from 0.13 to 0.15 in the single-product model. This increase is primarily accounted for by the increase in the frequency of price changes.

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

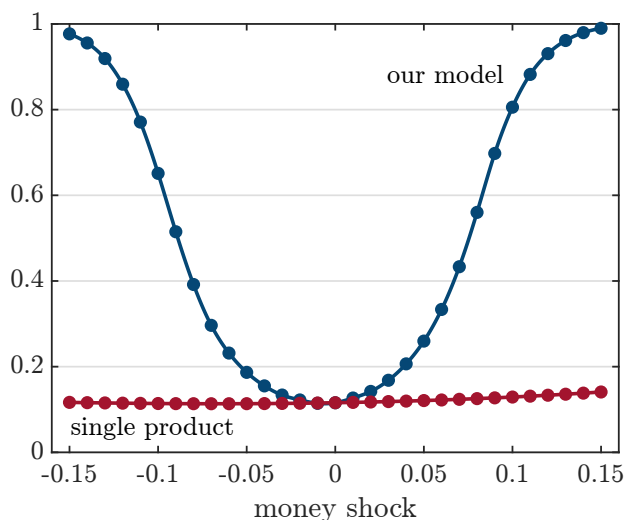
	Single-product			Our model		
	1%	5%	10%	1%	5%	10%
total pass-through	0.129	0.135	0.146	0.323	0.421	0.861
<i>Calvo</i>	0.094	0.096	0.098	0.095	0.099	0.111
<i>frequency</i>	0.001	0.004	0.011	0.009	0.123	0.660
<i>selection</i>	0.035	0.036	0.037	0.219	0.198	0.090

5.2 Non-Linear Phillips Curve

We next investigate how non-linear are the real effects of money shocks 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 10 shows that though the frequency of price changes is relatively insensitive to the size of the money shock in the standard model, it responds very non-linearly in our model. Specifically, in the neighborhood of zero the frequency does not respond to money shocks, but increases fast away from zero and is nearly 100% for shocks exceeding 10% in absolute value. Recall that our model only partially reproduces the importance of the extensive

Figure 10: Frequency Response on Impact



margin of price adjustment in the sectoral data, so the results in Figure 10 likely understate the non-linear response of the frequency.

Figure 11(a) displays the impact response of output to monetary shocks of various sizes. In the single product model, the response is nearly linear, with a slope coefficient of 0.87, reflecting the relatively constant frequency of price changes. In our model, the slope is highly non-linear. For small shocks the output response is comparable to that in the standard model. As the shock gets larger, the output response is smaller than in the standard model and peaks for shocks of approximately 7% in absolute value. Further increases in the size of the shock reduce the output response, owing to the steep rise in the frequency of price changes. Consequently, if the shock is large enough and all firms adjust prices, there are no real effects of monetary policy.

The cumulative impulse responses of output, depicted in Figure 11(b), exhibit a similar pattern, but the discrepancy between our model and the standard model is even larger, owing to the lower persistence of output in our model. As can be inferred from Figure 8, the response of output in our model has a lower half-life.

Figure 12 summarizes this discussion by depicting the Phillips curve implied by the impact responses of output and inflation to the money shocks considered above. While the Phillips curve is approximately linear in the standard model, it is highly non-linear in our model. In particular, at low levels of inflation the Phillips curve has a slope only slightly larger than in the standard model, and becomes vertical at high rates of inflation.

Figure 11: Output Response to Money Shocks

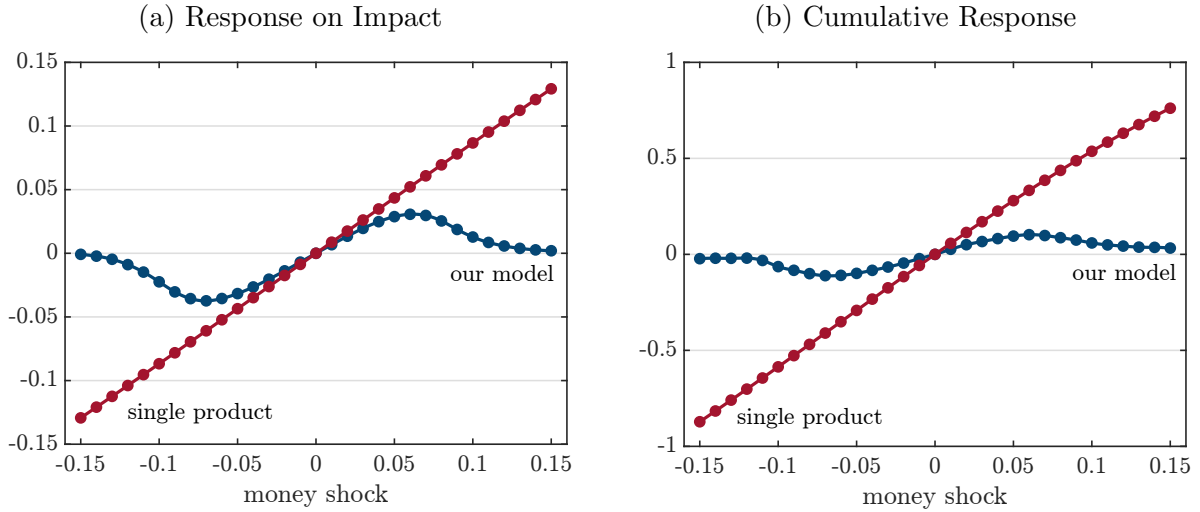
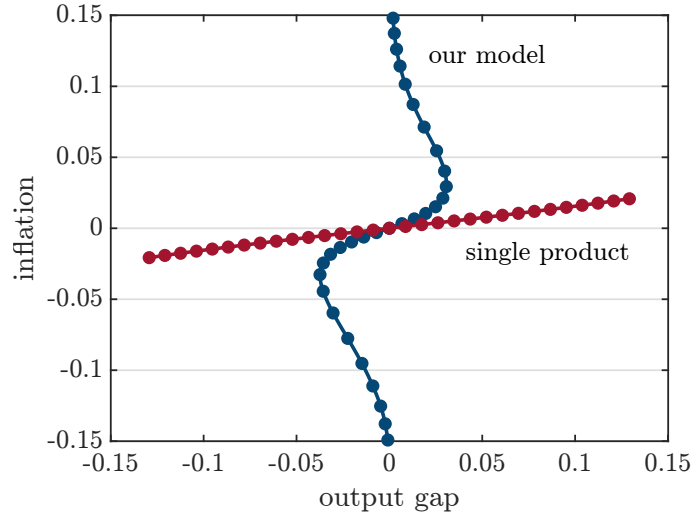


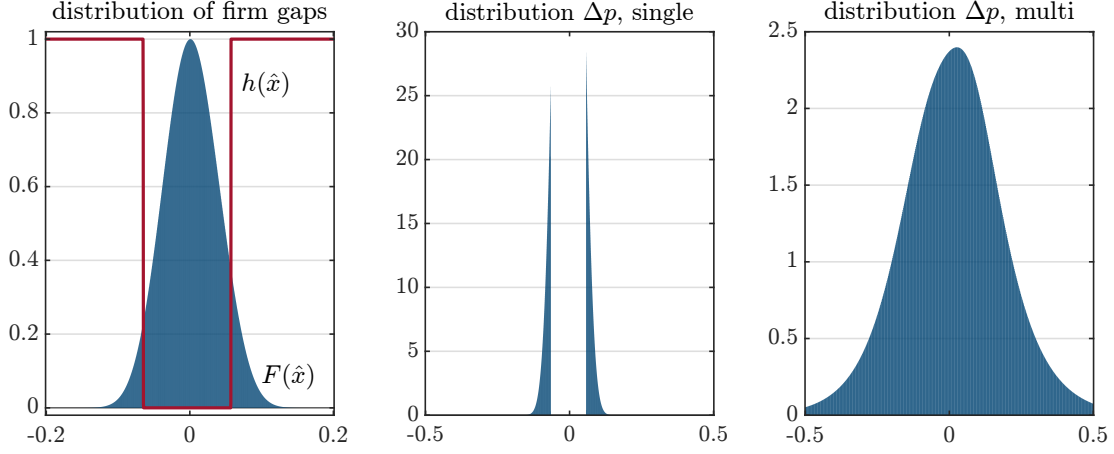
Figure 12: Phillips Curve, Impact Responses



5.3 Intuition

We next provide intuition for why reducing the amount of misallocation inside the firm increases the strength of the selection effect in response to small shocks and the response of the frequency of price changes to large shocks. To this end, we 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. Since the duration of price spells only affects firms' profits through their effect on misallocation, $\phi_t(f, s) = \exp\left(-d\gamma\frac{\sigma_z^2}{2}\right)$, setting $\gamma = 0$ implies that the firm's decisions are no longer a function of the duration d .

Figure 13: Distribution of Price Changes in Single- and Multi-Product Model



Dropping the dependence on s , the Bellman equation that describes the value of inaction is only a function of the firm's overall price gap \hat{x}

$$v_t^n(\hat{x}) = x_t^{\sigma-1} \left[(1 + \tau) \hat{x}^{1-\sigma} - \bar{a} x_t^{(\frac{1}{\eta}-1)(\sigma-1)} \hat{x}^{-\frac{\sigma}{\eta}} \right] + \beta \mathbb{E}_t \max(v_{t+1}^n(\hat{x}'), v_{t+1}^a - \bar{\xi}),$$

where the price gaps evolves according to

$$\hat{x}' = \exp \left(\left((1 - \gamma) \frac{\sigma_z^2}{2} + \sigma_u \varepsilon_{t+1}^u(f, s) - g_m \right) \hat{x} \right).$$

The value of adjustment is

$$v_t^a = \max_{x_t^*} x_t^{\sigma-1} \left[(1 + \tau) (x_t^*)^{1-\sigma} - \bar{a} x_t^{(\frac{1}{\eta}-1)(\sigma-1)} (x_t^*)^{-\frac{\sigma}{\eta}} \right] + \beta \mathbb{E}_t \max(v_{t+1}^n(\hat{x}'), v_{t+1}^a - \bar{\xi}),$$

where the price gap evolves according to the same law of motion, with x_t^* replacing \hat{x} .

The single-product model is a special case of a multi-product model in which there are no product-level shocks, so $\sigma_z = 0$. The equations above then imply that when $\gamma = 0$ the multi-product model with $\sigma_z > 0$ is equivalent to the single-product model with $\sigma_z = 0$, provided we adjust the trend growth rate of money supply to ensure that the price gap drifts at the same rate in the two models. Thus, a single-product model in which

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

has identical value functions, adjustment thresholds and optimal reset prices as its multi-product counterpart. We can therefore obtain intuition for how the multi-product economy works by studying an equivalent single-product model.

The left panel of Figure 13 shows the steady state distribution of price gaps \hat{x} and the adjustment hazard in the multi-product economy with $\gamma = 0$ that we calibrate to match the distribution of micro-price changes in the data (see the robustness section below for details about the calibration). This is also the distribution of price gaps in the single-product economy in which we adjust the trend money growth rate appropriately. Even though these two models have identical implications for the distribution of firm price gaps and decision rules, they have different implications for the distribution of price change.

The middle panel of the figure shows the distribution of price changes in the single-product economy. This distribution is reminiscent of that in Golosov and Lucas (2007): it features neither small nor very large price changes and has a large mass near the (s, S) thresholds. In contrast, in the multi-product version of the model, the bimodal distribution of price gaps conditional on adjustment does not translate into a bimodal distribution of price changes, since individual products experience price changes largely due to product-specific shocks. As the right panel of Figure 13 shows, the latter reproduces the evidence on the distribution of price changes well.

Consider next the impulse response of output to a monetary shock. The shock reduces all firms' price gaps by the same amount. The aggregate price gap, in both the single- and multi-product economy, is given by

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

and therefore responds identically in the two economies. Thus, even though the single- and multi-product economy imply different distributions of price changes, they predict identical responses to monetary shocks.

We conclude that selection effects are stronger when we reduce misallocation inside the firm because the multi-product model inherits the properties of the single-product model. Since the selection effect is strong in the single-product model, as pointed out by Golosov and Lucas (2007), it is also strong in the multi-product model. Moreover, since the model requires less volatile firm-level shocks to reproduce the dispersion of price changes, the menu cost required to reproduce the frequency of price changes is lower, implying narrower (s, S) bands, and therefore more non-linear output responses.¹²

¹²An alternative approach to narrowing the (s, S) bands is to assume fat-tailed productivity shocks, as in Midrigan (2011). Karadi and Reiff (2019) show that model reproduces well the evidence from Hungary that the frequency of price changes responds strongly to large changes in the value added tax.

6 Robustness

We next evaluate the robustness of our findings to alternative parameterizations of the menu cost models we considered above. Specifically, we show that the three shortcomings we highlighted in the single-product model are not driven by our assumption that there are decreasing returns to scale and therefore strategic complementarities. Moreover, we show that our results in the multi-product model are robust to alternative values of γ , the elasticity of substitution between products sold by a given firm.

6.1 Single-Product Model

We investigate the extent to which the shortcomings we highlighted in the single-product model are driven by our assumption that there are decreasing returns to scale. We do so by setting $\eta = 1$ and recalibrating the remaining parameters to match the same set of moments as before. Table 13 in the Appendix reports the calibrated parameters and the targeted moments and shows that this model can also match the distribution of price changes in the data.

The first column of Table 9 shows that the model with constant returns requires smaller menu costs, 2.1% of average sales and predicts smaller losses from misallocation, 5.71%, compared to the single-product model with decreasing returns. Because in the single-product model a large fraction of price changes are free, a more meaningful statistic is the overall losses the firm incurs due to menu costs. We find these losses are equivalent to a 32.4% tax on the firm's profits, a sizable amount. Even though the model without decreasing returns to scale implies smaller menu costs and losses from misallocation, it fails to reproduce the comovement between inflation and the frequency of price changes. In particular, the slope coefficient from regressing counterfactual inflation π_t^c on actual inflation π_t is 0.99 and falls to only 0.93 when we restrict the simulated sample to periods in which inflation is above its 90th percentile, suggesting a much weaker role for the extensive margin of price adjustment than in the data.

6.2 Multi-Product Model

We also investigate the extent to which our results are sensitive to γ , the elasticity of substitution between products within the firm. Recall that a relatively low elasticity of substitution is required to reduce the degree of misallocation within the firm. In our baseline, we set

Table 9: The Three Shortcomings in Alternative Parameterizations

	Single-product	Multi-product	
	$\eta = 1$	$\gamma = 0$	$\gamma = 3$
menu costs/sales	0.021	0.014	0.047
misallocation, %	5.71	0.92	3.97
slope of π_t^c on π_t			
all observations	0.99	0.88	0.89
$\pi_t > 90^{th}$ pct.	0.93	0.61	0.63

$\gamma = 1$. Here, we consider two alternative values: $\gamma = 0$, corresponding to the case in which goods within the firm are perfect complements, and $\gamma = 3$. In both cases, we recalibrate the remaining parameters to match the same moments as before. We report the results of the parameterization in Table 14 in the Appendix.

Table 9 shows that the less substitutable goods within the firm are, the smaller the menu cost and lower the amount of misallocation. In particular, when $\gamma = 0$, menu costs amount to 1.4% of total sales, close to the empirical estimate, and the losses from misallocation are only 0.92%. When $\gamma = 3$, the menu costs are a larger fraction of average sales (4.7%) and the losses from misallocation are larger (3.97%). Both these parameterizations imply a stronger comovement between inflation and the frequency of price changes than the standard model. The slope coefficients in regressions of π_t^c on π_t are similar to those in our baseline.

7 Conclusions

Canonical menu cost models, of both the single-product and multi-product variety, suffer three important shortcomings when calibrated to match the distribution of micro price changes in the data. First, they require implausibly large menu costs and imply considerable profit losses from frictions to price adjustment. Second, they predict a large amount of misallocation from price dispersion. Third, they cannot reproduce the strong comovement between the frequency of price changes and inflation observed in the data. These shortcomings are exacerbated in the presence of microeconomic strategic complementarities.

We propose a resolution to these shortcomings by extending the multi-product menu cost model along two dimensions. First, we assume that strategic complementarities are at

the firm rather than the product level. Second, we assume that individual products sold by a given firm are imperfect substitutes. Both these assumptions limit the losses from misallocation from price dispersion within the firm and go a long way, but not fully, towards remedying the three shortcomings of the canonical menu cost models.

We use the model to study the real effects of monetary policy. We find that, in contrast to standard models, our model predicts smaller and highly non-linear output responses to shocks, owing to a stronger selection effect and a more responsive frequency of price changes. The model implies that the Phillips curve is nearly vertical when inflation exceeds 10%.

References

- Alvarez, Fernando and Francesco Lippi**, “Price Setting With Menu Cost for Multiproduct Firms,” *Econometrica*, 2014, *82* (1), 89–135.
- 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.
- 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, Martn 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.
- 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.
- 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.
- 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.
- Costain, James and Anton Nakov**, “Distributional Dynamics Under Smoothly State-Dependent Pricing,” *Journal of Monetary Economics*, 2011, *58* (6), 646–665.
- Davies, Richard**, “Prices and Inflation in the UK - A New Dataset,” Technical Report 55, Centre for Economic Performance February 2021.

- 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, Robert G. King, 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.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu**, “How Costly Are Markups?,” Working Paper 24800, National Bureau of Economic Research July 2018.
- 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.
- 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.
- Khan, Aubhik and Julia K. Thomas**, “Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics,” *Econometrica*, 2008, *76* (2), 395–436.
- 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.
- 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.
- Mankiw, N. Gregory**, “Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly,” *The Quarterly Journal of Economics*, 05 1985, *100* (2), 529–538.

- Midrigan, Virgiliu**, “Menu Costs, Multi-product Firms, and Aggregate Fluctuations,” *Econometrica*, 2011, 79 (4), 1139–1180.
- Nakamura, Emi 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.
- Petrella, Ivan, Emiliano Santoro, and Lasse de la Porte Simonsen**, “Time-varying Price Flexibility and Inflation Dynamics,” 2018.
- Rotemberg, Julio**, “The New Keynesian Microfoundations,” in “NBER Macroeconomics Annual 1987, Volume 2” 1987, pp. 69–116.
- 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

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 Office for National Statistics (ONS). We use public monthly product-level price quotes and item-level price indexes from January 1996 to August 2022.¹³

Goods and services are classified following the 6-digit Classification of Individual Consumption by Purpose (COICOP 6).¹⁴ The CPI is produced in stages, with indexes derived at each stage weighted together to give higher level indexes.¹⁵ A sample of prices of items which are representative of UK consumer expenditure are collected in line with the COICOP classification system.¹⁶ There are currently around 650 representative items in the CPI price basket of goods. The items usually have fairly broad specifications (such as a roll of wallpaper or womens 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.¹⁷ This is supplemented by other data sources, including the Living Costs and Food Survey (LCF) data, International Passenger Survey data (IPS) and data from Public Sector Branch. The weights used in compiling the measures of consumer price inflation are updated annually following ONS reviews of the representative

¹³See also [Davies \(2021\)](#) and [Petrella et al. \(2018\)](#) who use these data.

¹⁴COICOP 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.

¹⁵This description is taken from the Consumer Price Indices Technical Manual published by the ONS and available [here](#).

¹⁶For example, for the item home-killed lamb, prices are collected for ‘loin chops with bone’ and ‘shoulder with bone’. Other joints, and loin chops and shoulders without bones, are not priced; it is assumed that their price movements are close to those of the joints of lamb that are priced.

¹⁷The descriptions in this subsection are taken from the Consumer Price Indices Technical Manual published by the ONS and available [here](#).

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.¹⁸ 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.¹⁹ 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 sections 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.²⁰
2. Item identifier, COICOP classification, and COICOP weights. The item index data is sourced from the ONS website.²¹ The classification of items into COICOP classifications are also provided by the ONS.²² COICOP weights are provided for each item.

¹⁸For example, within electrical appliances, the electric cooker item represents only itself and not any other kinds of electrical appliances.

¹⁹For example, a screwdriver is one of several items representing all spending on small tools within DIY materials, and there are other items within the section representing all spending on paint, timber, fittings and so on.

²⁰The link to the latest data is [here](#) and to historical data is [here](#).

²¹The link is [here](#).

²²The link is [here](#).

3. Aggregated price indexes. We also use the price indexes published by the ONS at COICOP-6 and above levels of disaggregation.²³

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, 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.²⁴
3. Merge price quotes data with item identifiers and weights.

Our final master panel dataset is comprised of the variables in Table 10. We have around 38 million unique price quote observations from 1996m1 to 2022m8. All statistics and analyses are conducted with this dataset.

A.5 Data Checks

We do two checks on our panel dataset.

1. First, we confirm that the diagggregated price indexes generate the published aggregate CPI index using the sector-level weights in our dataset. We construct:

$$\pi_t = \sum_s w_t(s) \pi_t(s) \tag{12}$$


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 14 plots the constructed aggregate index against the published aggregate index; Figure 15 plots the corresponding inflation rates.

2. The second check we do is to compare inflation constructed using the final micro price dataset to the published aggregate inflation series. Figure 16 plots inflation rates calculated from the micro data using regular prices, and shows that they have a close correspondence to the published aggregate inflation rate.

²³The link is [here](#).

²⁴These observations can arise because of confidentiality reasons.

Figure 14: Price Indexes



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, 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.²⁵ 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.

²⁵There 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).

Table 10: Variables in Dataset

Variable	Description
date	Date of price quote observation
quote_id	Identifier for the price quote
weight_id	Identifier for the item
weight	Weight
price	Price
coicop_6	COICOP-6 classification
CPIagg	Aggregate CPI index

Figure 15: Inflation Rates

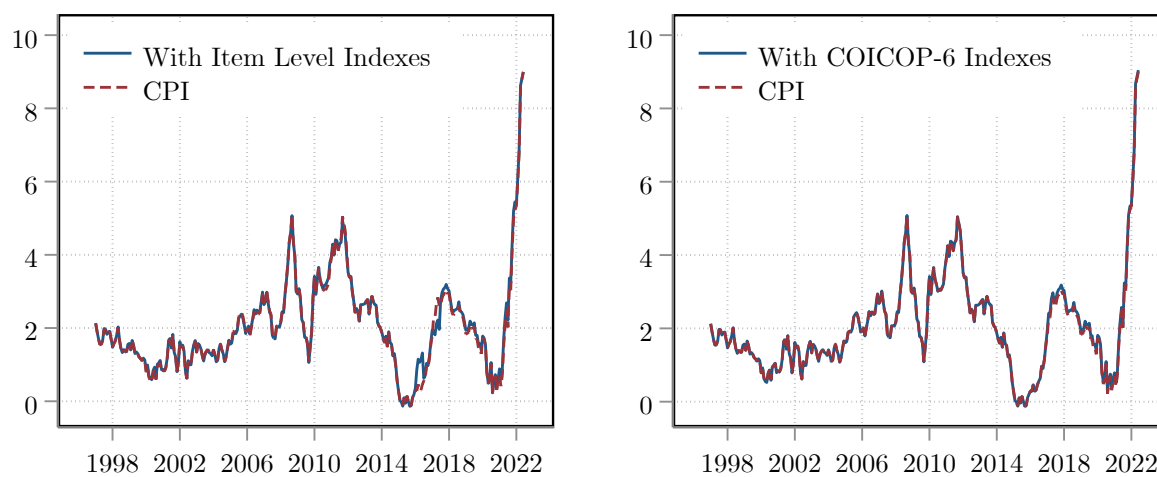
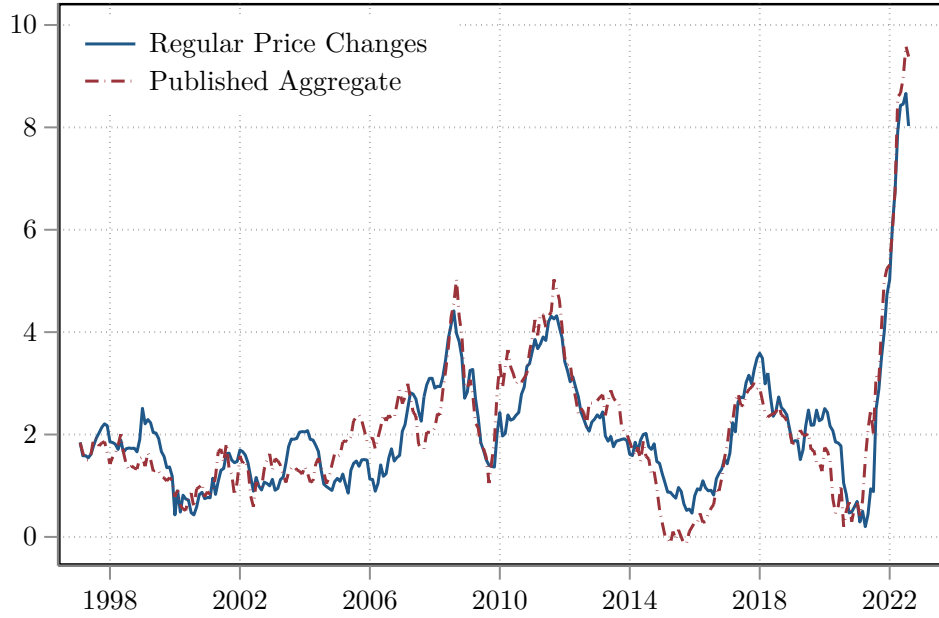


Figure 16: Inflation Calculated Using Regular Price Changes



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. As an example of our algorithm that constructs regular prices, Figure 17 plots the posted price and regular price for an item in the COICOP-6 Bread and Cereals category (01.1.1).

5. **Standardization:** We follow Klenow and Kryvtsov (2008) and compute the standardized price change:

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


where $\mu_{\Delta}(s)$ and μ_{Δ} are the mean of the non-zero log price changes in sector s and in the aggregate, respectively, and $\sigma_{\Delta}(s)$ and σ_{Δ} construct the item-level mean and variance of price changes.

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

Table 11 shows price statistics for different sets of prices from our dataset.²⁶ Removing

²⁶In all cases, we standardize prices at the item level and remove the top 2% and bottom 2% of outliers.

Figure 17: Posted and Regular Price Example: “Bread and Cereals” Product



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.

Table 12 shows the weights at the COICOP-2 level in our sample used to compute micro price statistics and compared to the weights that are published by the ONS in the CPI.

Table 11: 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
5 th percentile	-0.325	-0.328	-0.327
10 th percentile	-0.238	-0.240	-0.226
25 th percentile	-0.100	-0.099	-0.081
50 th percentile	0.024	0.023	0.026
75 th percentile	0.126	0.126	0.119
90 th percentile	0.250	0.252	0.247
95 th percentile	0.329	0.332	0.340
<i>distribution of Δp</i>			
mean	0.146	0.147	0.142
std. dev.	0.119	0.121	0.125
5 th percentile	0.009	0.009	0.009
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
95 th percentile	0.394	0.398	0.413

Table 12: COICOP-2 Level Weights in Sample Versus Published Weights

	Weight, Sample	Weight, CPI
1 Food and Non-Alcoholic Beverages	0.19	0.12
2 Alcoholic Beverages and Tobacco	0.09	0.05
3 Clothing and Footwear	0.08	0.06
4 Housing, Utilities, and Other Fuels	0.03	0.14
5 Furniture, Household Eq./Maintenance	0.12	0.08
6 Health	0.03	0.02
7 Transport	0.08	0.14
8 Communication	0.00	0.03
9 Recreation and Culture	0.14	0.13
10 Education	0.00	0.03
11 Restaurants	0.13	0.11
12 Miscellaneous Goods and Services	0.11	0.09
Aggregate	1.00	1.00

B 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

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

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

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

Notice that now $z_{it}(f, s)$ represents a product-specific productivity shock and $u_t(f, s)$ represents a firm-specific productivity shock. The firm's profits are

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

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 then write the firm's production function as

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

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

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

We also let the composite firm productivity be

$$\tilde{u}_t(f, s) = u_t(f, s) z_t(f, s),$$

which evolves according to a geometric random walk process with Gaussian innovations and a drift equal to $(\gamma - 1) \frac{\sigma_z^2}{2}$. Also let

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

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

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

The firm-level price gap is given by

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

The product-level price gap is given by

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

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

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

This assumption ensures that the menu cost is equal to 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

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

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

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

C Additional Figures and Tables

This section reports the parameter values and targeted moments in the robustness section discussed in the paper.

Table 13: Alternative Parameterization: Single-Product Model

A. Moments		
	Data	$\eta = 1$
frequency Δp	0.116	0.116
<i>distribution of Δp</i>		
mean	0.018	0.013
std. dev.	0.188	0.196
kurtosis	3.609	3.624
5 th percentile	-0.327	-0.318
10 th percentile	-0.226	-0.230
25 th percentile	-0.081	-0.102
50 th percentile	0.026	0.014
75 th percentile	0.119	0.139
90 th percentile	0.247	0.264
95 th percentile	0.340	0.334
<i>distribution of Δp</i>		
mean	0.142	0.152
std. dev.	0.125	0.125
5 th percentile	0.009	0.011
10 th percentile	0.018	0.020
25 th percentile	0.045	0.055
50 th percentile	0.104	0.120
75 th percentile	0.204	0.220
90 th percentile	0.334	0.328
95 th percentile	0.413	0.395
B. Calibrated Parameter Values		
		$\eta = 1$
g_m	mean money growth rate	0.019
σ^u	s.d. firm shocks	0.067
λ	1-prob. free change	0.909
$\bar{\xi}$	upper bound menu cost	12.49

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

Table 14: Alternative Parameterizations: Multi-Product Model

A. Moments			
	Data	$\gamma = 0$	$\gamma = 3$
frequency Δp	0.116	0.115	0.115
<i>distribution of Δp</i>			
mean	0.018	0.012	0.025
std. dev.	0.188	0.178	0.190
kurtosis	3.609	3.889	3.405
5 th percentile	-0.327	-0.275	-0.306
10 th percentile	-0.226	-0.205	-0.224
25 th percentile	-0.081	-0.100	-0.094
50 th percentile	0.026	0.013	0.037
75 th percentile	0.119	0.122	0.152
90 th percentile	0.247	0.229	0.254
95 th percentile	0.340	0.303	0.317
<i>distribution of Δp</i>			
mean	0.142	0.138	0.152
std. dev.	0.125	0.113	0.117
5 th percentile	0.009	0.010	0.012
10 th percentile	0.018	0.021	0.024
25 th percentile	0.045	0.053	0.061
50 th percentile	0.104	0.112	0.128
75 th percentile	0.204	0.193	0.217
90 th percentile	0.334	0.290	0.313
95 th percentile	0.413	0.359	0.377

B. Calibrated Parameter Values

		$\gamma = 0$	$\gamma = 3$
g_m	mean money growth rate	0.019	0.034
σ^u	s.d. firm shocks	0.026	0.028
σ^z	s.d. product shocks	0.057	0.059
$\bar{\xi}$	menu cost	0.122	0.410

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