The Inflation Accelerator*

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

March 2025

Abstract

We develop a tractable sticky price model in which the fraction of price changes evolves endogenously over time and, consistent with the evidence, increases with inflation. Because we assume that firms sell multiple products and choose how many, but not which, prices to adjust in any given period, our model admits exact aggregation and reduces to a one-equation extension of the Calvo model. This additional equation determines the fraction of price changes. The model features a powerful *inflation accelerator* – a feedback loop between inflation and the fraction of price changes – which significantly increases the slope of the Phillips curve during periods of high inflation. Applied to the U.S. time series, our model predicts that the slope of the Phillips curve ranges from 0.02 in the 1990s to 0.12 in the 1970s and 1980s.

Keywords: Phillips curve, inflation, price rigidities.

*We thank Mark Gertler and John Leahy for useful feedback, Hugh Montag and Daniel Villar for sharing the data on the frequency of price changes, and Suk Joon Kim for excellent research assistance. The views expressed are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve Board.

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

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

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

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

1 Introduction

The recent rise in inflation in many economies has spurred considerable interest in further understanding the dynamics of prices. Identifying the causes of high inflation hinges critically on the shape of the Phillips curve. Our goal in this paper is to measure how the slope of the Phillips curve fluctuates in the U.S. macroeconomic time series. Since a key determinant of this slope is the fraction of price changes, it is imperative that we use a model that reproduces the widely documented evidence that the fraction of price changes increases in times of high inflation.[1]

There are two classes of sticky-price models that are typically used. The more popular time-dependent model[2] has the advantage of being tractable, but assumes a time-invariant fraction of price changes, making it less suitable for analyzing episodes in which inflation and the fraction of price changes increase considerably, such as the 1980s. The class of state-dependent (menu cost) models is a natural framework to endogenize the fraction of price changes,[3] but it is considerably more difficult to use for empirical and policy analysis. This difficulty arises because characterizing the general equilibrium dynamics of a model subject to aggregate shocks requires keeping track of the distribution of prices across firms. This challenge is especially pronounced when the fraction of price changes responds to aggregate shocks, which gives rise to important non-linearities.

Our paper proposes a model in which, as in state-dependent models, the fraction of price changes evolves endogenously over time, but which is as tractable as time-dependent models. The main challenge in allowing the fraction of price changes to evolve endogenously over time is that a firm's price adjustment decision depends on how far from the optimum its price is: firms whose prices are further from the optimum have stronger incentives to adjust. Equilibrium outcomes are therefore a function of the entire distribution of price changes, an infinite-dimensional object. We circumvent this challenge by assuming that firms sell a continuum of products and choose how many, but not which, prices to adjust in any period,

[1] See Gagnon (2009), Nakamura et al. (2018), Alvarez et al. (2018), Karadi and Reiff (2019), Montag and Villar (2023) and Blanco et al. (2024) for evidence that the frequency of price changes increases with inflation, as well as Hazell et al. (2022) and Fitzgerald et al. (2024) for evidence on the slope of the Phillips curve using state-level data.

[2] See Calvo (1983) and Taylor (1980).

[3] See, for example, Barro (1972), Sheshinski and Weiss (1977), Dotsey et al. (1999), Golosov and Lucas (2007), Gertler and Leahy (2008), Midrigan (2011), Alvarez and Lippi (2014), Alvarez et al. (2016), Alvarez et al. (2018), Auclert et al. (2022). However, Blanco et al. (2024) show that the canonical menu cost model has difficulties reproducing the extent to which the frequency of price changes comoves with inflation.

subject to an adjustment cost.[4] Because firms cannot choose which prices to adjust, the distribution of prices is no longer necessary to describe adjustment incentives, so the economy admits exact aggregation.[5] We show that our model reduces to a one-equation extension of the Calvo model, with the additional equation pinning down how many prices change in a given period. Our model nests the Calvo model in the limiting case when the adjustment cost goes to infinity. More generally, up to a first-order approximation around a steady state with no inflation, the dynamics of our model are identical to those of the Calvo model.

Our key finding is that the slope of the Phillips curve fluctuates considerably in the U.S. time series and increases in times of high inflation due to a feedback loop between inflation and the fraction of price changes. On one hand, an increase in the fraction of price changes increases inflation, more so the higher inflation is to begin with. On the other hand, an increase in inflation increases the firms' incentives to adjust prices, further increasing the fraction of price changes. We refer to this feedback loop as the *inflation accelerator* and show that it is responsible for the bulk of the steepening of the Phillips curve in periods of high inflation. Our model therefore suggests that reducing inflation is much less costly when inflation is high than when it is low.

We study an economy in which multi-product firms sell a continuum of goods and choose what fraction of their prices to change each period, subject to an adjustment cost that is increasing and convex in the number of prices that the firm adjusts. We assume decreasing returns to scale in production, which introduce strategic complementarities in price setting and dampen the slope of the Phillips curve. For clarity, we start by assuming that monetary policy targets nominal spending and that monetary shocks are the only source of aggregate uncertainty. We show, however, in a robustness section, that our results are robust to assuming a conventional Taylor rule and introducing additional sources of aggregate uncertainty.[6]

Relative to the standard Calvo model, endogenizing the fraction of price changes adds a single additional equation that balances the marginal cost of changing prices against the marginal benefit. The marginal benefit increases with inflation, implying that the fraction of price changes increases with inflation. Because endogenizing the fraction of price changes introduces non-linearities in the dynamics of output and inflation, we solve the model using

[4]This assumption is reminiscent of that in Greenwald (2018) who uses a large family construct to endogenize refinancing decisions.

[5]In addition to the menu cost literature, Romer (1990) also endogenizes the frequency of price changes in a Calvo model, but circumvents the curse of dimensionality by assuming that firms choose the frequency of price changes once and for all. In that model the frequency of price changes is constant over time.

[6]In that section, we also show that our findings are robust to eliminating strategic complementarities in price setting and introducing idiosyncratic shocks.

global projection methods, but show that a third-order perturbation provides an accurate approximation, so the model can be solved using readily-available solution techniques.

We start by building intuition for the workings of the model by studying impulse responses to expansionary monetary shocks in environments with low and high trend inflation. We show that the real effects of monetary shocks are considerably smaller in environments with high inflation for two reasons. First, the steady-state fraction of price changes is higher in environments with high inflation. Second, the fraction of price changes increases in response to shocks. Though this increase is relatively small, it has a large impact on the price level because adjusting firms respond to the underlying trend inflation and increase prices by large amounts, an effect reminiscent of Caplin and Spulber (1987).

We build additional intuition for the dynamics of inflation and output by deriving the Phillips curve in our model. The slope of the Phillips curve is equal to the sum of two terms, one identical to the slope in the Calvo model, which increases mechanically with the frequency of price changes, and another which captures the inflation accelerator. Importantly, we show that this second term increases much more rapidly with inflation and thus accounts for the bulk of the increase in the slope of the Phillips curve in high-inflation environments.

We use our model to characterize how the slope of the Phillips curve evolves in the post-war U.S. time series. We do so by first identifying the sequence of monetary shocks that allows the model to reproduce the path of inflation in the data. We show that the model reproduces well the path of the fraction of price changes, notably the sharp rise during the high inflation in the 1970s and 1980s. We then consider log-linear perturbations around the equilibrium point at each date and derive the slope of the Phillips curve. We find that the slope of the Phillips curve varies considerably, ranging from 0.02 in low-inflation periods such as the 1990s to 0.12 in high-inflation periods such as the 1970s and the 1980s.[7] The inflation accelerator accounts for the bulk of this increase: in its absence the higher fraction of price changes in the 1970s and 1980s would only increase the slope of the Phillips curve to 0.04.

That the slope of the Phillips curve varies over time has important implications for the tradeoff between inflation and output stabilization. We gauge how this tradeoff varies over time by calculating a measure of the sacrifice ratio: the fall in output required to achieve a one percentage point reduction in inflation. The sacrifice ratio varies considerably, from 1.4%

[7]That the slope of the Phillips curve is high in times of high inflation and low in times of low inflation is consistent with Morales-Jimenez and Stevens (2024), who use a model with menu costs and information frictions, and Flynn et al. (2024), who use a model with information frictions, to show that output responds less to monetary shocks when inflation is high. It is also consistent with the empirical evidence in Ball and Mazumder (2011) and Hazell et al. (2022).

in the low-inflation period in the 1990s to 0.4% in the high-inflation periods in the 1970s and 1980s. We therefore conclude that our model implies that if inflation is high to begin with, bringing it down requires a smaller drop in output than if inflation is low. Our model thus rationalizes the view that reducing inflation from 10% to 9% is a lot less costly than reducing it from 3% to 2%.

The rest of the paper proceeds as follows. Section 2 presents the model. Section 3 describes the parameterization. Section 4 analyzes the steady state of the model. Section 5 applies the framework to the time-series U.S. data. Section 6 discusses several robustness exercises. Section 7 concludes.

2 Model

We study an economy in which firms adjust prices infrequently. In contrast to the standard New Keynesian model, we allow the frequency of price changes to fluctuate endogenously over time by assuming that multi-product firms choose what fraction of their prices to adjust in any given period. We circumvent the need to keep track of the distribution of prices by assuming that firms choose how many, but not which, prices to change. Owing to this assumption, our model reduces to a one-equation extension of the standard Calvo model, with the additional equation describing how the fraction of price changes, and therefore the slope of the Phillips curve, is pinned down each period.

For clarity, we start by assuming that monetary policy targets nominal spending, which evolves over time according to a random walk process. Shocks to the growth rate of nominal spending are the only source of aggregate fluctuations. We then show in the robustness section below that assuming instead that monetary policy follows a Taylor rule does not change our key findings.

2.1 Consumers

A representative consumer has preferences over consumption c_t and hours worked h_t and maximizes life-time utility

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

subject to the budget constraint

$$P_t c_t + \frac{1}{1+i_t} B_{t+1} = W_t h_t + D_t + B_t,$$

where P_t is the nominal price level, B_t are holdings of a risk-free bond which pays nominal interest i_t, D_t are the dividends from the firms the representative consumer owns, and W_t is the nominal wage rate.

2.2 Monetary Policy

We assume that monetary policy targets nominal spending, $M_t \equiv P_t c_t$, which follows a random walk with drift

$$\log \frac{M_{t+1}}{M_t} = \mu_{t+1} = \mu + \varepsilon_{t+1},$$

where μ is the average growth rate of nominal spending and ε_{t+1} are Gaussian innovations with standard deviation σ. As Afrouzi and Yang (2021) point out, this specification of the monetary policy rule is equivalent to an interest rate rule in which the central bank assigns the same weight to inflation and output growth.

2.3 Technology

There is a continuum of intermediate goods firms indexed by i. Each firm produces a continuum of products k with technology

$$y_{ikt} = (l_{ikt})^{\eta},$$

where y_{ikt} is the output of product k produced by firm i, l_{ikt} is the labor used in production and $\eta \leq 1$ is the span-of-control parameter which, as in Burstein and Hellwig (2008), introduces a micro-level strategic complementarity in price setting.

A perfectly competitive final goods sector aggregates the intermediate goods y_{ikt} into a composite final good using a CES aggregator

$$y_t = \left(\int_0^1 \int_0^1 (y_{ikt})^{\frac{\theta-1}{\theta}} \, \mathrm{d}k \mathrm{d}i \right)^{\frac{\theta}{\theta-1}},$$

where θ is the elasticity of substitution, which we assume to be the same both across products and across firms. This implies that the demand for an individual product is

$$y_{ikt} = \left(\frac{P_{ikt}}{P_t} \right)^{-\theta} y_t, \tag{1}$$

where P_{ikt} is the price of an individual product and

$$P_t = \left(\int_0^1 \int_0^1 (P_{ikt})^{1-\theta} \, \mathrm{d}k \mathrm{d}i \right)^{\frac{1}{1-\theta}}$$

is the aggregate price index.

2.4 Problem of Intermediate Goods Producers

We next describe the profit maximization problem of intermediate goods producers.

Period Profits. The nominal profits of firm i from producing product k are

$$P_{ikt}y_{ikt} - \tau W_t l_{ikt},$$

where $\tau = 1 - 1/\theta$ is a subsidy that removes the markup distortion that would arise even in the absence of price rigidities. Using the demand function (1), we can express real profits as

$$\left(\frac{P_{ikt}}{P_t}\right)^{1-\theta} y_t - \tau \frac{W_t}{P_t}\left(\frac{P_{ikt}}{P_t}\right)^{-\frac{\theta}{\eta}} y_t^{\frac{1}{\eta}}. \tag{2}$$

Losses from Misallocation. Differences in the price of products sold by a given firm generate losses from misallocation, reducing firm productivity. To see this, let

$$y_{it} = \left(\int (y_{ikt})^{\frac{\theta-1}{\theta}} \, \mathrm{d}k\right)^{\frac{\theta}{\theta-1}}$$

denote the composite output produced firm i and let

$$l_{it} = \int l_{ikt} \mathrm{d}k$$

denote the total amount of labor the firm uses. We can then derive a firm-level production function

$$y_{it} = \left(\frac{X_{it}}{P_{it}}\right)^{\theta} l_{it}^{\eta},$$

where

$$P_{it} = \left(\int (P_{ikt})^{1-\theta} \, \mathrm{d}k\right)^{\frac{1}{1-\theta}} \tag{3}$$

denotes the price index of firm i and

$$X_{it} = \left(\int (P_{ikt})^{-\frac{\theta}{\eta}} \, \mathrm{d}k\right)^{-\frac{\eta}{\theta}} \tag{4}$$

determines the extent of misallocation. Absent dispersion in prices, $X_{it}/P_{it} = 1$ and productivity is maximized. With price dispersion, $X_{it}/P_{it} < 1$ and productivity is reduced.

Price Adjustment Cost. We assume that the firm has a convex cost of changing prices denominated in units of labor. This cost is increasing in the number of prices n_{it} the firm resets and is equal to

$$\frac{\xi}{2}\left(n_{it}-\bar{n}\right)^2, \quad \text{if } n_{it}>\bar{n}$$

and zero otherwise. Here, ξ determines the size of the adjustment cost and $\bar{n}$ is the fraction of free price changes. The key assumption we make is that although the firm can choose how many prices to change in a given period, it cannot choose which prices to change. By endogenizing the frequency of price changes, the model can capture the evidence that firms are more likely to adjust prices in times of high inflation, as in menu cost models, but in a much more tractable way. When $\xi \to \infty$, the model collapses to the Calvo model with a constant frequency $\bar{n}$.

Our model shares similarities with that in Romer (1990) which endogenizes the frequency of prices changes in the Calvo model.[8] In Romer (1990) firms choose a once-and-for-all price adjustment probability, balancing the gains from more frequent adjustment against the costs of repricing. Extending that model to allow for a time-varying adjustment probability would require keeping track of the distribution of prices because the gains from adjusting would be higher for prices further away from the optimum, just like in menu cost models.[9] In contrast, our assumption that firms sell a continuum of products and choose how many, but not which, prices to change, implies that firms are ex-post identical and that a small number of state variables are sufficient to characterize a firm's incentives to adjust prices. This feature allows exact aggregation and renders our model very tractable.

Price Setting. We next describe the firms' problem in detail. The value of the firm is the present discounted sum of its flow profits (2). The log-linear specification of preferences implies that $c_t = \frac{W_t}{P_t} = y_t$ and allows to write the value of the firm as

$$\mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \left(\int \left[\left(\frac{P_{ikt+s}}{P_{t+s}} \right)^{1-\theta} - \tau \left(\frac{P_{ikt+s}}{P_{t+s}} \right)^{-\frac{\theta}{\eta}} y_{t+s}^{\frac{1}{\eta}} \right] \mathrm{d}k - \frac{\xi}{2} \left(n_{it+s} - \bar{n} \right)^2 \right),$$

or, using the definitions of P_{it} and X_{it} in equations (3) and (4),

$$\mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \left[\left(\frac{P_{it+s}}{P_{t+s}} \right)^{1-\theta} - \tau \left(\frac{X_{it+s}}{P_{t+s}} \right)^{-\frac{\theta}{\eta}} y_{t+s}^{\frac{1}{\eta}} - \frac{\xi}{2} \left(n_{it+s} - \bar{n} \right)^2 \right]. \tag{5}$$

[8] See also Kiley (2000), Devereux and Yetman (2002) and Bakhshi et al. (2007).

[9] See also Alvarez et al. (2021) and Cavallo et al. (2024) for variants of the menu cost model in which firms choose the price adjustment probability subject to a convex adjustment cost.

The firm chooses what fraction of prices n_{it} to reset every period and the reset price P_{it}^*. Because all products are identical, $P_{ikt} = P_{it}^*$ for all products whose price is reset.

To characterize these optimal choices, we first describe how the choice of P_{it}^* and n_{it} affect firm profits in future periods. Consider first the term $(P_{it+s})^{1-\theta}$ in equation (5). Using the definition of the firm's price index in equation (3) and the assumption that the firm chooses how many, but not which, prices to change, allows us to write this term as a function of the history of previously chosen reset prices and repricing probabilities as follows

$$\begin{aligned}(P_{it+s})^{1-\theta} = {} & n_{it+s}\left(P_{it+s}^*\right)^{1-\theta} + (1-n_{it+s})\, n_{it+s-1}\left(P_{it+s-1}^*\right)^{1-\theta} + \\ & (1-n_{it+s})\,(1-n_{it+s-1})\, n_{it+s-2}\left(P_{it+s-2}^*\right)^{1-\theta} + \cdots + \\ & \prod_{j=1}^{s}(1-n_{it+j}) n_{it}\left(P_{it}^*\right)^{1-\theta} + \prod_{j=1}^{s}(1-n_{it+j})\,(1-n_{it})\,(P_{it-1})^{1-\theta}.\end{aligned} \tag{6}$$

The first term on the right hand side represents the contribution of the n_{it+s} newly reset prices in period $t+s$. The second term represents the contribution of the $(1-n_{it+s})\, n_{it+s-1}$ prices that were reset in period $t+s-1$ and were not reset in period $t+s$. This pattern continues with each subsequent term accounting for the contribution of prices reset in each period leading up to $t+s$, including those reset in period t, captured by the first term in the last line of the expression, as well as those reset prior to period t, captured by the last term of the expression. In writing this last term we used the definition of the price index in equation (3) to express the history of all reset prices prior to period t using a single state variable, P_{it-1}. A similar argument allows us to rewrite the term $(X_{it+s})^{-\frac{\theta}{\eta}}$ as

$$\begin{aligned}(X_{it+s})^{-\frac{\theta}{\eta}} = {} & n_{it+s}\left(P_{it+s}^*\right)^{-\frac{\theta}{\eta}} + (1-n_{it+s})\, n_{it+s-1}\left(P_{it+s-1}^*\right)^{-\frac{\theta}{\eta}} + \\ & (1-n_{it+s})\,(1-n_{it+s-1})\, n_{it+s-2}\left(P_{it+s-2}^*\right)^{-\frac{\theta}{\eta}} + \cdots + \\ & \prod_{j=1}^{s}(1-n_{it+j}) n_{it}\left(P_{it}^*\right)^{-\frac{\theta}{\eta}} + \prod_{j=1}^{s}(1-n_{it+j})\,(1-n_{it})\,(X_{it-1})^{-\frac{\theta}{\eta}}.\end{aligned} \tag{7}$$

We can now characterize the optimal choice of P_{it}^* and n_{it}. To derive the optimality condition with respect to P_{it}^* we note that equations (6) and (7) imply that

$$\frac{\partial\left(P_{it+s}\right)^{1-\theta}}{\partial P_{it}^*} = (1-\theta)\left(P_{it}^*\right)^{-\theta}\prod_{j=1}^{s}(1-n_{it+j})n_{it}$$

and

$$\frac{\partial\left(X_{it+s}\right)^{-\frac{\theta}{\eta}}}{\partial P_{it}^*} = -\frac{\theta}{\eta}\left(P_{it}^*\right)^{-\frac{\theta}{\eta}-1}\prod_{j=1}^{s}(1-n_{it+j})n_{it}.$$

Therefore, the reset price P_{it}^* that maximizes the value of the firm satisfies the first order condition

$$\left(\frac{P_{it}^*}{P_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} = \frac{1}{\eta}\frac{b_{2it}}{b_{1it}},$$

where

$$b_{1it} = \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \prod_{j=1}^{s} (1 - n_{it+j}) \left(\frac{P_{t+s}}{P_t}\right)^{\theta-1}$$

and

$$b_{2it} = \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \prod_{j=1}^{s} (1 - n_{it+j}) \left(\frac{P_{t+s}}{P_t}\right)^{\frac{\theta}{\eta}} (y_{t+s})^{\frac{1}{\eta}}.$$

The terms b_{1it} and b_{2it} capture the present value of revenue and marginal costs in future periods, weighted by the probability that a price reset today in still in effect in that future period. An increase in future inflation reduces the real value of the reset price and increases revenue and marginal costs and therefore the weight that the firm places on that period. Strategic complementarities, captured by the term $1 + \theta\left(\frac{1}{\eta} - 1\right)$, dampen the extent to which the reset price responds to aggregate shocks, generating additional price stickiness. These expressions that determine the optimal reset price are analogous to those obtained in a Calvo model, with the only difference being that in the Calvo model n_t is constant.

To build additional intuition, we can alternatively express the firm's optimal reset price as a function of the expected present value of future marginal costs,

$$P_{it}^* = \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s \omega_{it+s} MC_{it+s},$$

where

$$MC_{it+s} = \frac{1}{\eta} W_{t+s} (y_{it+s})^{\frac{1}{\eta}-1} = \frac{1}{\eta} W_{t+s} \left(\frac{P_{it}^*}{P_{t+s}}\right)^{-\theta\left(\frac{1}{\eta}-1\right)} y_{t+s}^{\frac{1}{\eta}-1}$$

is the marginal cost in period $t+s$ of producing a good whose price was last reset in period t and

$$\omega_{it+s} = \frac{\beta^s (P_{t+s})^{\theta-1} \prod_{j=1}^{s} (1 - n_{it+j})}{\mathbb{E}_t \sum_{s=0}^{\infty} \beta^s (P_{t+s})^{\theta-1} \prod_{j=1}^{s} (1 - n_{it+j})}$$

denotes the relative weight of period $t+s$ in determining the firm's optimal price, with the weight reflecting the probability that the price is still in effect at that point, as well as the amount the firm expects to sell given the aggregate price index in that period.

To derive the optimality condition with respect to n_{it}, we first note that equations (6) and (7) imply that

$$\frac{\partial (P_{it+s})^{1-\theta}}{\partial n_{it}} = \prod_{j=1}^{s} (1 - n_{it+j}) \left((P_{it}^*)^{1-\theta} - (P_{it-1})^{1-\theta}\right)$$

and

$$\frac{\partial \left(X_{it+s}\right)^{-\frac{\theta}{\eta}}}{\partial n_{it}} = \prod_{j=1}^{s} \left(1 - n_{it+j}\right) \left(\left(P_{it}^{*}\right)^{-\frac{\theta}{\eta}} - \left(X_{it-1}\right)^{-\frac{\theta}{\eta}}\right).$$

The fraction of price changes n_{it} that maximizes the value of the firm therefore satisfies the first order condition

$$\xi \left(n_{it} - \bar{n}\right) = b_{1it} \left(\left(\frac{P_{it}^{*}}{P_t}\right)^{1-\theta} - \left(\frac{P_{it-1}}{P_t}\right)^{1-\theta}\right) - \tau b_{2it} \left(\left(\frac{P_{it}^{*}}{P_t}\right)^{-\frac{\theta}{\eta}} - \left(\frac{X_{it-1}}{P_t}\right)^{-\frac{\theta}{\eta}}\right).$$

In choosing what fraction of prices to adjust, the firm balances the price adjustment costs against the benefits resulting from changing its price index, captured by the first term on the right hand side, and reducing misallocation inside the firm, captured by the second term on the right hand side. Notice that the terms b_{1it} and b_{2it} that determine the optimal reset price also determine the firm's incentive to adjust prices. For example, the higher is output y_{t+s} in future periods, the larger is b_{2it} and therefore the stronger the incentive to reduce misallocation and thus economize on labor costs. Similarly, the more likely the firm is to adjust its prices in the future, the lower are b_{1it} and b_{2it} and therefore the benefits from adjusting prices today. Note that the firms' incentives to adjust are shaped by two state variables, the firm's price index, P_{it-1}, and the amount of misallocation inside the firm, X_{it-1}, as in the multi-product menu cost model of Blanco et al. (2024). Nevertheless, because firms are ex-post identical, we do not need to keep track of the joint distribution of these two state variables.

2.5 Equilibrium

Since all firms are identical, $n_{it} = n_t$ and $P_{it}^{*} = P_t^{*}$. Consequently all firms have the same price indices and losses from misallocation. Let $p_t^{*} = P_t^{*}/M_t$, $p_t = P_t/M_t$, $x_t = X_t/P_t$ and $\pi_t = P_t/P_{t-1}$ and recall that

$$w_t = c_t = y_t = \frac{1}{p_t},$$

where $w_t = \frac{W_t}{P_t}$ is the real wage. The equilibrium of the model is characterized by the following system of equations:

1. the definition of the price index, which determines inflation as a function of the relative reset price and the fraction of price changes

$$1 = n_t \left(\frac{p_t^{*}}{p_t}\right)^{1-\theta} + \left(1 - n_t\right) \pi_t^{\theta-1}, \tag{8}$$

2. the optimal reset price

$$\left(\frac{p_t^*}{p_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} = \frac{1}{\eta}\frac{b_{2t}}{b_{1t}}, \tag{9}$$

where b_{1t} and b_{2t} are determined by

$$b_{1t} = 1 + \beta\mathbb{E}_t\left(1-n_{t+1}\right)\pi_{t+1}^{\theta-1}b_{1t+1} \tag{10}$$

$$b_{2t} = p_t^{-\frac{1}{\eta}} + \beta\mathbb{E}_t\left(1-n_{t+1}\right)\pi_{t+1}^{\frac{\theta}{\eta}}b_{2t+1}, \tag{11}$$

3. the optimal choice of the fraction of price changes

$$\xi\left(n_t-\bar{n}\right) = b_{1t}\left(\left(\frac{p_t^*}{p_t}\right)^{1-\theta} - \pi_t^{\theta-1}\right) - \tau b_{2t}\left(\left(\frac{p_t^*}{p_t}\right)^{-\frac{\theta}{\eta}} - x_{t-1}^{-\frac{\theta}{\eta}}\pi_t^{\frac{\theta}{\eta}}\right), \tag{12}$$

4. the endogenous productivity term that captures the losses from misallocation

$$x_t^{-\frac{\theta}{\eta}} = n_t\left(\frac{p_t^*}{p_t}\right)^{-\frac{\theta}{\eta}} + \left(1-n_t\right)x_{t-1}^{-\frac{\theta}{\eta}}\pi_t^{\frac{\theta}{\eta}}.$$

Notice that inflation is equal to

$$\pi_t = \frac{P_t}{P_{t-1}} = \frac{p_t}{p_{t-1}}\frac{M_t}{M_{t-1}} = \frac{p_t}{s_t},$$

where

$$s_t = \frac{P_{t-1}}{M_t} = \frac{p_{t-1}}{\exp\left(\mu_t\right)}$$

is the previous period's price level scaled by current nominal spending. Because the growth rate of nominal spending is iid, the only two state variables in this economy are s_t and x_{t-1}, so the solution of the model is given by functions $p_t = \mathcal{P}\left(s_t,\, x_{t-1}\right)$, $x_t = \mathcal{X}\left(s_t,\, x_{t-1}\right)$ etc., that determine how output and inflation evolve over time in response to monetary policy shocks.

Relative to the Calvo model, the only new equation is equation (12) which characterizes how the fraction of price changes evolves over time. The Calvo model is a special case of our model that can be obtained by setting $\xi \to \infty$, in which case $n_t = \bar{n}$ is constant. In the Calvo model the previous period's losses from misallocation x_{t-1} do not affect the frequency of price changes, so the price level only depends on a single state variable, s_t.

We solve the system of functional equations that characterize the solution of the model using global projection methods, by approximating the equilibrium functions using Chebyshev polynomials. We found, however, that a third-order perturbation provides a very accurate approximation, suggesting that the model can be easily solved using readily-available solution techniques. See the Appendix for details.

3 Parameterization

We next explain how we parameterize the model. We first discuss the parameters we assign values to and then the parameters we calibrate endogenously.

3.1 Assigned Parameters

A period in the model is a quarter. We set three parameters to values conventional in the literature: a quarterly discount factor β of 0.99, a demand elasticity θ of 6 and a returns to scale parameter η of 2/3. In the robustness section below we show that our results are robust to perturbing θ and η.

3.2 Calibrated Parameters

The parameters we calibrate endogenously are those determining the average level and volatility of inflation, as well as the average frequency of price changes and its comovement with inflation. Specifically, we set the average growth rate of nominal spending μ, the standard deviation of nominal spending growth σ, the fraction of free price changes $\bar{n}$, and the price adjustment cost parameter ξ to reproduce the mean and standard deviation of inflation, the mean fraction of price changes, and the slope coefficient from regressing the fraction of price changes on the absolute value of inflation. This last statistic captures the extent to which the fraction of price changes and inflation comove in the time series.

3.2.1 The Data

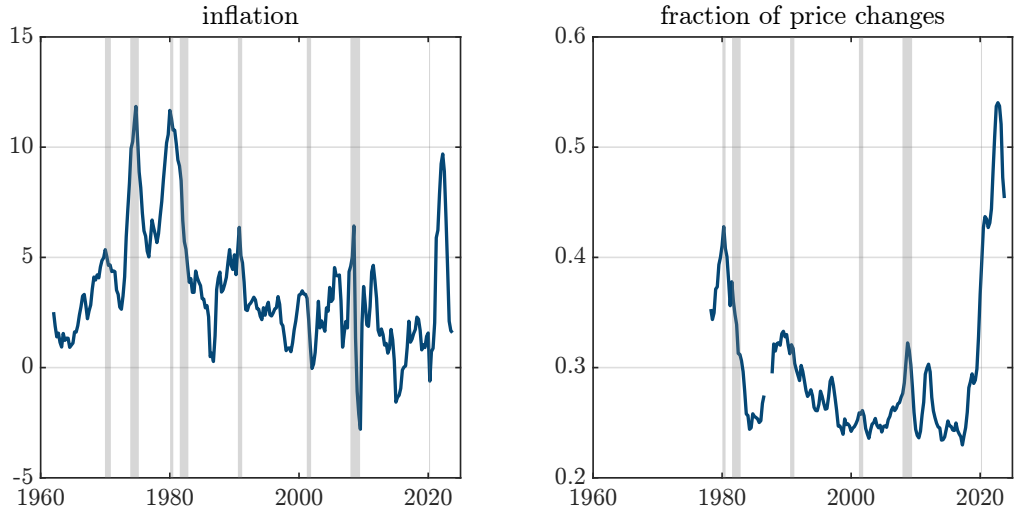
Our measure of inflation is the growth rate of the U.S. CPI, available from 1962:Q1 to 2023:Q4. We follow Nakamura et al. (2018) in using the series excluding shelter to ensure that the inflation data is compatible with the data on the fraction of price changes. The fraction of price changes is computed from the price quotes collected by the BLS that underlie the construction of the CPI.[10] Specifically, we use the monthly median fraction of price changes, excluding sales, available between 1978 and 2023. We convert the monthly series into a quarterly one, so the series we use is the fraction of prices that change in a quarter.

Figure 1 plots the year-to-year percent change in the price level and the average quarterly fraction of price changes in the preceding year. On average, approximately 25% of prices change in a given quarter in periods of low inflation. As documented by Nakamura et al.

[10] We are grateful to Hugh Montag and Daniel Villar for kindly sharing these data with us. See Nakamura et al. (2018) for a detailed description of how the data was constructed.

(2018), the fraction of price changes was relatively high, approximately 40% per quarter, in the high-inflation episode in the early 1980s. As documented by Montag and Villar (2023), the fraction of price changes spiked once again, to approximately 50%, during the post-Covid inflation episode. The fraction of price changes thus increases systematically in times of high inflation, a robust feature of the data documented, for example, by Gagnon (2009), Alvarez et al. (2018), Karadi and Reiff (2019) and Blanco et al. (2024).

Figure 1: Inflation and the Fraction of Price Changes



Notes: The gray bars indicate NBER recessions. The CPI series excludes shelter.

3.2.2 Parameter Values

For many of our subsequent exercises, we will contrast the predictions of our model with an endogenously-varying frequency of price changes to those of a standard Calvo model with a constant frequency. We therefore also calibrate the Calvo model using the same strategy, but discard the adjustment cost parameter and choose the fraction of price changes $\bar{n}$ to reproduce the average frequency of price changes.

Table 1 reports the results of the calibration. As Panel A of the table shows, both models reproduce the targeted moments perfectly. In both models and in the data, the average inflation rate is equal to 3.5%, the standard deviation of inflation is 2.7% and the average quarterly fraction of price change is 29.7%. In our model, the fraction of price changes also comoves systematically with inflation: the slope coefficient of a regression of the fraction of price changes on the absolute value of inflation is equal to 0.016, as in the data.

Panel B of the table reports the calibrated parameter values. The volatility of nominal spending growth is slightly lower in our model compared to the Calvo model. This is because in the Calvo model a sudden increase in inflation can only be explained by a large shock, whereas in our model the endogenous response of the frequency also contributes to inflation fluctuations. We also note that the fraction of free price changes is equal to 24.1% in our model, a number required to reproduce the fraction of price changes in low-inflation periods. The adjustment cost parameter ξ, though not interpretable on its own, implies that on average 0.65% of all labor is used in adjusting prices, a number in line with the evidence in Levy et al. (1997).

Table 1: Endogenously Calibrated Parameters

A. Targeted Moments

	Data	Our model	Calvo
mean inflation	3.517	3.517	3.517
s.d. inflation	2.739	2.739	2.739
mean frequency	0.297	0.297	0.297
slope of n_t on $\lvert\pi_t\rvert$	0.016	0.016	–

B. Calibrated Parameter Values

		Our model	Calvo
μ	mean spending growth rate	0.035	0.035
σ	s.d. monetary shocks	0.022	0.024
$\bar{n}$	fraction free price changes	0.241	0.297
ξ	adjustment cost	1.767	–

Notes: The mean nominal spending growth rate is annualized.

4 Steady State Analysis

Before characterizing the dynamics of U.S. inflation through the lens of our model, we first build intuition for the mechanism of the model by characterizing how the non-stochastic steady state of the model varies with trend inflation. We also use a first-order approximation around the non-stochastic steady state to provide intuition for how the economy responds to monetary policy shocks in environments with high and low trend inflation. Many of

the insights we derive below will carry through in the subsequent section which studies the responses to monetary policy shocks in the actual U.S. time series.

4.1 Steady State Fraction of Price Changes

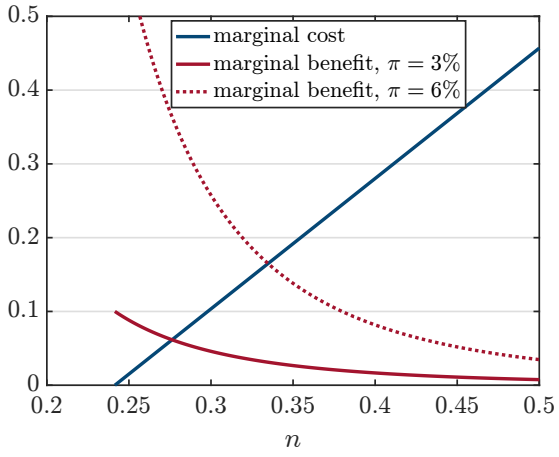
We first characterize how the steady state frequency of price changes varies with trend inflation. Letting $\pi = \exp(\mu)$ denote the trend level of inflation and variables without t subscripts denote the value of a variable in the non-stochastic steady state, we can show that the fraction of price changes is pinned down by

$$\xi\left(n-\bar{n}\right)=\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\frac{1}{n}\left(1-\pi^{\theta-1}-\tau\eta\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}\left(1-\pi^{\frac{\theta}{\eta}}\right)\right), \tag{13}$$

where the left-hand side of the equation is the marginal cost of increasing n and the right-hand side captures the marginal benefit to increasing n.[11]

The marginal cost is linearly increasing in n. Absent trend inflation, $\pi = 1$ and the marginal benefit of increasing n is equal to 0, implying that $n = \bar{n}$. Thus, absent trend inflation, the steady state of our model is identical to that of the Calvo model. More generally, with positive trend inflation, $\pi > 1$, the marginal benefit of changing prices is positive and decreases with n, as illustrated in Figure 2. The intersection of the marginal benefit and cost curves pins down the steady-state fraction of price changes. As the figure shows, higher trend inflation increases the marginal benefit of adjusting prices, thus increasing the fraction of price changes.

Figure 2: Steady-State Fraction of Price Changes



[11] See the Appendix for all the derivations.

We summarize how the steady-state fraction of price changes increases with trend inflation in the left panel of Figure C.2. The fraction of price changes increases from 24% per quarter at zero inflation to 41% when annual steady-state inflation is 10%.[12]

4.2 The Real Effects of Monetary Shocks

We next study the real effects of monetary shocks. We consider impulse responses to both small and large shocks and discuss how they depend on trend inflation.

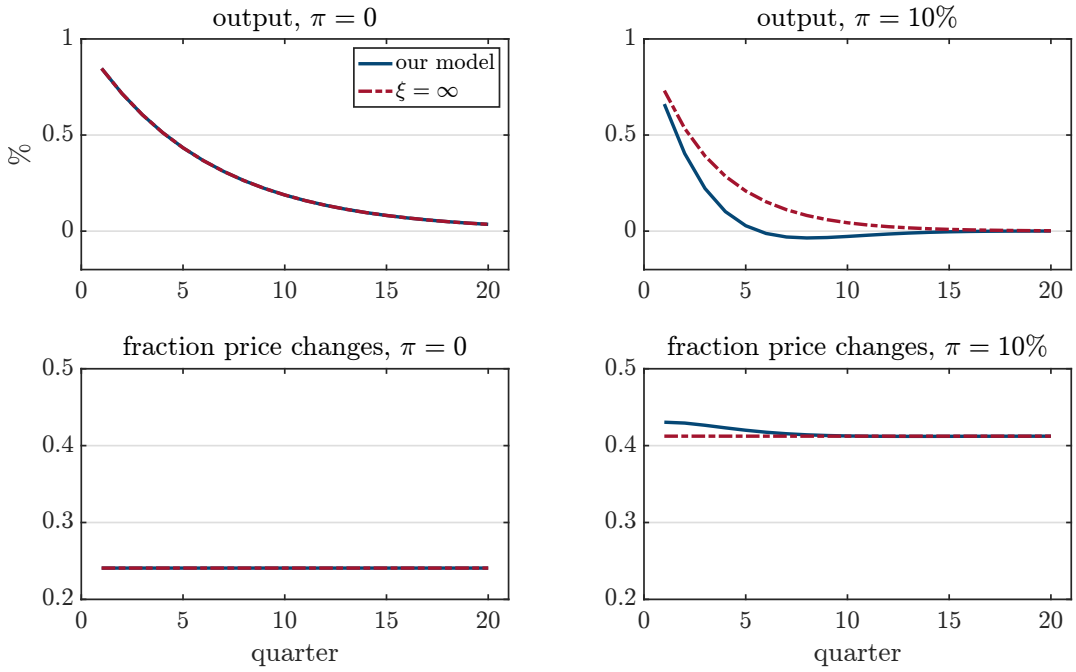
Impulse Responses to Small Shocks. We first study how our economy responds to small monetary policy shocks in environments with low and high trend inflation. We consider two economies, one with zero and another with 10% trend inflation, and report the impulse responses of output and the fraction of price changes to a 1% increase in nominal spending M_t. To build intuition, we consider a log-linear approximation of the model around the steady state of each economy. We contrast the responses in our model to those predicted by an otherwise identical model in which the fraction of price changes is equal to that in the steady state of our model, but $\xi = \infty$ so the fraction of price changes is constrained to not respond to shocks. As we discussed above, the steady-state fraction of price changes is approximately twice higher in the economy with 10% inflation compared to the economy with no inflation (0.41 vs. 0.24).

The left two panels of Figure 3 show the response of output (upper panel) and the fraction of price changes (lower panel) in the economy without trend inflation. Note that, up to a first-order approximation, the fraction of price changes does not respond to the monetary shock. Hence, our model with an endogenous adjustment frequency has identical responses to the economy with a time-invariant frequency.

The right two panels of Figure 3 depict the responses in an environment with 10% trend inflation. We make two observations. First, even when the fraction of price changes is constrained not to respond to the shock ($\xi = \infty$), the response of output is weaker in the economy with 10% trend inflation relative to the economy without trend inflation, owing to the larger steady-state frequency of price changes. Both the impact output response (0.73% vs. 0.85%), as well as the cumulative output response (2.72% vs. 5.50%) are lower. Second, the impulse response of output is smaller and more transient in our model relative to the $\xi = \infty$ economy with a constant frequency of price changes. Both the impact output response

[12] In Appendix C, we show that trend inflation reduces output and productivity by a lot less in our model relative of a model with a constant fraction of price changes.

Figure 3: Impulse Response to a 1% Monetary Shock



(0.66% vs. 0.73%) and the cumulative output response (1.22% vs 2.72%) fall considerably, owing to the increase in the fraction of price changes from 0.41 in steady state to 0.43 after the shock. Also note that in our model the output response falls below zero one year after the shock, owing to the overshooting of the price level caused by the increased frequency of price changes.

Although the increase in the fraction of price changes following a shock appears small, it leads to considerable aggregate price flexibility. Since, $P_t y_t = M_t$, the impact response of the price level to the monetary shock is 0.15% in the economy with zero trend inflation, 0.27% in the economy with 10% trend inflation and constant frequency, and 0.34% in our model with 10% trend inflation.

To see why such a small increase in the fraction of price changes has such a large effect on the aggregate price level, consider the log-linearized system of equilibrium conditions characterizing the evolution of inflation in our model. Letting hats denote log-deviations from the steady state, a first-order Taylor expansion of the expression for the aggregate price level in equation (8) implies that

$$\hat{\pi}_t = \underbrace{\frac{1}{(1-n)\,\pi^{\theta-1}} \frac{\pi^{\theta-1}-1}{\theta-1}}_{\mathcal{M}} \hat{n}_t + \underbrace{\frac{1-(1-n)\,\pi^{\theta-1}}{(1-n)\,\pi^{\theta-1}}}_{\mathcal{N}} \left(\hat{p}_t^* - \hat{p}_t\right). \tag{14}$$

The second term on the right-hand side of this expression is familiar from the standard Calvo model and describes how inflation responds to an increase in the relative reset price. The elasticity

$$\mathcal{N} = \frac{1 - (1-n)\,\pi^{\theta-1}}{(1-n)\,\pi^{\theta-1}}$$

increases with the fraction of price changes n and decreases with trend inflation π. As Coibion et al. (2012) point out, a higher trend inflation reduces the sensitivity of inflation to reset price changes because newly reset prices are larger and therefore have a smaller share in the consumption weights used to calculate the ideal price index.

The first term on the right-hand side of (14) captures the impact of changes in the fraction of price changes on the inflation response. The elasticity

$$\mathcal{M} = \frac{1}{(1-n)\,\pi^{\theta-1}} \frac{\pi^{\theta-1} - 1}{\theta - 1}$$

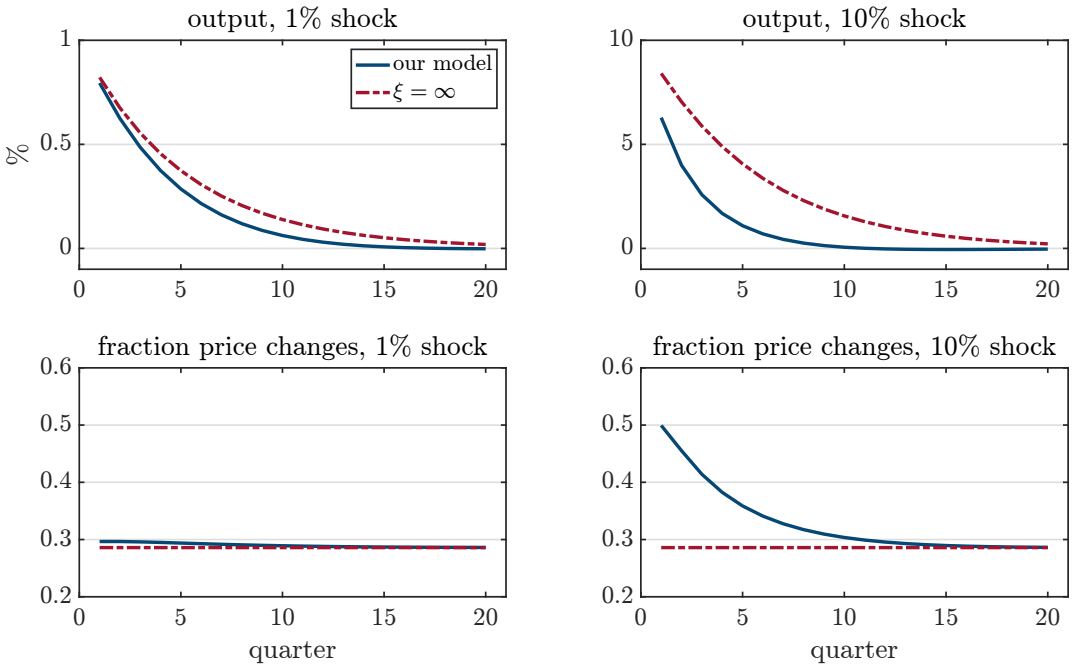
is equal to zero absent trend inflation and increases as π increases above 1. To understand why this is the case, note that inflation is approximately equal to the fraction of price changes n_t times the average price change conditional on adjustment. If the average price change is zero, as is the case absent trend inflation, an increase in the fraction of price changes does not affect inflation. In contrast, if the average price change is large, inflation greatly responds to changes in the fraction of price changes. For example, in our economy with 10% annual inflation the average size of a price change is approximately 4%. An increase in the frequency of price changes from 0.41 to 0.43 thus contributes an additional $(0.43 - 0.41) \times 4\% = 0.08\%$ to the increase in the price level in the period of the shock. Since the fraction of price changes mean-reverts gradually, these effects add up over time, implying that the response of the price level to a monetary shock is much more rapid relative to a setting with a constant fraction of price changes. This effect is reminiscent of the mechanism in the menu cost model of Caplin and Spulber (1987) in which small changes in the frequency of repricing render the aggregate price level flexible. Because newly-adjusting prices increase by a large amount, even small changes in the fraction of price changes add considerably to the response of inflation.

Impulse Responses to Large Shocks. So far we considered the responses to relatively small monetary policy shocks using a first-order approximation. To a first order, the increase in the frequency of price changes only contributes to aggregate price flexibility because adjusting firms respond to the underlying trend inflation, but not to the monetary policy shock. We next consider shocks of larger sizes and solve for the response of output non-linearly, thus

taking into account the interaction between the increase in the frequency of price changes and the increase in the average price change resulting from the shock. To compute these responses, we start from the non-stochastic steady state of the economy and consider a one-time, unanticipated, permanent increase in nominal spending M_t. We then calculate the resulting transition dynamics using a non-linear shooting method.

Figure 4 reports the responses to a 1% (left panels) and 10% (right panels) monetary shock starting from the steady state of our baseline model calibrated to match the U.S. data with 3.5% trend inflation. Notice first that even in response to small shocks the output responses in our model are weaker than in the model with a constant frequency. As explained above, the small increase in the fraction of price changes imparts considerable flexibility to the aggregate price level because these additional price changes incorporate the larger trend inflation into their price adjustment decisions. Moreover, the response of output to a 10% monetary policy shock is considerably smaller in our model, owing to the sharp increase in the fraction of price changes: 50% of prices change in response to this shock on impact. Consequently, the cumulative impulse response of output is one third of that in the model with a constant frequency of price changes.

Figure 4: Impulse Response to Small and Large Monetary Shocks



We thus conclude that our economy has many of the features of menu cost models, such

as the large responsiveness of the price level to movements in the fraction of price changes in periods of high inflation, as in Caplin and Spulber (1987) and considerable non-linearity in responses to shocks of different sizes, as in Blanco et al. (2024), but is considerably more tractable.[13]

4.3 The Phillips Curve and the Inflation Accelerator

We next derive the Phillips curve in our economy. We show that the slope of the Phillips curve increases rapidly with trend inflation due to a feedback loop between inflation and the frequency of price changes. On the one hand, an increase in the fraction of price changes increases inflation, more so in environments with higher trend inflation. On the other hand, an increase in inflation increases the firms' incentive to change prices, thus raising the frequency of price changes. We refer to this feedback loop as the *inflation accelerator*.

Log-linearizing the expression determining the optimal fraction of price changes (12) around the non-stochastic steady state, we have

$$\hat{n}_t = \mathcal{A}\hat{\pi}_t + \mathcal{B}\left(\hat{p}_t^* - \hat{p}_t\right) - \mathcal{C}\hat{x}_{t-1} + \frac{n-\bar{n}}{n}\hat{b}_{1t}, \tag{15}$$

where

$$\mathcal{A} = \frac{\theta-1}{\xi n}\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\frac{\pi^{\frac{\theta}{\eta}}-\pi^{\theta-1}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}$$

determines how sensitive the fraction of price changes is to inflation,

$$\mathcal{B} = \left(1-\tau\eta\right)\frac{\theta-1}{\xi n}\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\beta\left(1-n\right)\pi^{\theta-1}}\frac{1}{n}\frac{\pi^{\frac{\theta}{\eta}}-1}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}$$

determines how sensitive the fraction of price changes is to the relative reset price, and

$$\mathcal{C} = \frac{\theta-1}{\xi n}\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\beta\left(1-n\right)\pi^{\theta-1}}\frac{\pi^{\frac{\theta}{\eta}}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}$$

determines how sensitive the fraction of price changes is to past misallocation.

We note first that both $\mathcal{A}$ and $\mathcal{B}$ are equal to zero absent trend inflation. Thus, the frequency of price changes is, to a first-order, irresponsive to monetary policy shocks, as illustrated in Figure 3. In the presence of trend inflation these elasticities are positive and

[13] As Reiter and Wende (2024) show, by suitably modifying the adjustment cost function, a Rotemberg (1982)-style model–in which all firms adjust prices every period but face adjustment costs that depend on the size of price changes–can generate similar non-linearities.

decreasing in the adjustment cost parameter ξ. Combining the log-linearized expression for the price index (14) with (15) implies

$$\hat{\pi}_t = \frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}}\left(\hat{p}_t^* - \hat{p}_t\right) - \frac{\mathcal{MC}}{1-\mathcal{MA}}\hat{x}_{t-1} + \frac{\mathcal{M}}{1-\mathcal{MA}}\frac{n-\bar{n}}{n}\hat{b}_{1t}.$$

The elasticity of inflation to relative reset prices $\hat{p}_t^* - \hat{p}_t$ is equal to $\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}}$ and is amplified relative to the standard Calvo model whenever $\pi > 1$, so that $\mathcal{M}$, $\mathcal{A}$ and $\mathcal{B}$ are all positive. Intuitively, a higher desired reset price not only directly affects inflation with an elasticity $\mathcal{N}$, but also leads to more frequent price changes, which then increase inflation and further increase the incentives to reset prices. We refer to this feedback loop between the frequency of price changes, the optimal reset price and inflation as the *inflation accelerator*.

Finally, log-linearizing equations (9) – (11) which characterize the optimal reset price allows us to derive the Phillips curve

$$\begin{aligned}
\hat{\pi}_t &= \mathcal{K}\widehat{mc}_t + \beta(1-n)\left(\frac{\frac{\theta}{\eta}\pi^{\frac{\theta}{\eta}} - (\theta-1)\pi^{\theta-1}}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}} + \pi^{\frac{\theta}{\eta}}\right)\mathbb{E}_t\hat{\pi}_{t+1} \\
&+ \beta(1-n)\left(\frac{\pi^{\frac{\theta}{\eta}} - \pi^{\theta-1}}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}} - \pi^{\frac{\theta}{\eta}}\frac{\mathcal{M}}{1-\mathcal{MA}}\frac{n-\bar{n}}{n}\right)\mathbb{E}_t\hat{b}_{1t+1} \\
&- \beta n\frac{\pi^{\frac{\theta}{\eta}} - \pi^{\theta-1}}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}}\mathbb{E}_t\hat{n}_{t+1} \\
&+ \beta(1-n)\pi^{\frac{\theta}{\eta}}\frac{\mathcal{MC}}{1-\mathcal{MA}}\hat{x}_t - \frac{\mathcal{MC}}{1-\mathcal{MA}}\hat{x}_{t-1} + \frac{\mathcal{M}}{1-\mathcal{MA}}\frac{n-\bar{n}}{n}\hat{b}_{1t}. \qquad (16)
\end{aligned}$$

We note that all the terms in the last two rows of this expression drop out when the fraction of price changes is time-invariant.

The key elasticity is the slope of the Phillips curve: the elasticity of inflation with respect to real aggregate marginal cost, $mc_t = \frac{1}{\eta}\frac{W_t}{P_t}y_t^{\frac{1}{\eta}-1}$. As we show in the Appendix, this elasticity is equal to

$$\mathcal{K} = \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\left(1-\beta(1-n)\pi^{\frac{\theta}{\eta}}\right)\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}}. \qquad (17)$$

The first term of this expression captures the effect of strategic complementarities which are stronger the more elastic demand is, that is, the higher is θ and the stronger are decreasing returns to scale, that is, the lower is η. The second term captures the horizon effect: a transitory increase in marginal costs in period t only increases the optimal reset price by a

factor $1-\beta(1-n)\pi^{\frac{\theta}{\eta}}$ which reflects the discount factor and the probability that the current price will still be in effect in future periods. Finally, as discussed above, the last term captures the impact of higher reset prices on inflation. Note that absent the feedback effect of the frequency on inflation, that is, when $\mathcal{M}=0$, this expression reduces to the familiar slope of the Phillips curve in a Calvo model with trend inflation

$$\kappa=\frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\left(1-\beta(1-n)\pi^{\frac{\theta}{\eta}}\right)\frac{1-(1-n)\pi^{\theta-1}}{(1-n)\pi^{\theta-1}}.$$

The difference between these two slopes, $\mathcal{K}-\kappa$, reflects the inflation accelerator which is positive when $\pi>0$ and increases with π.

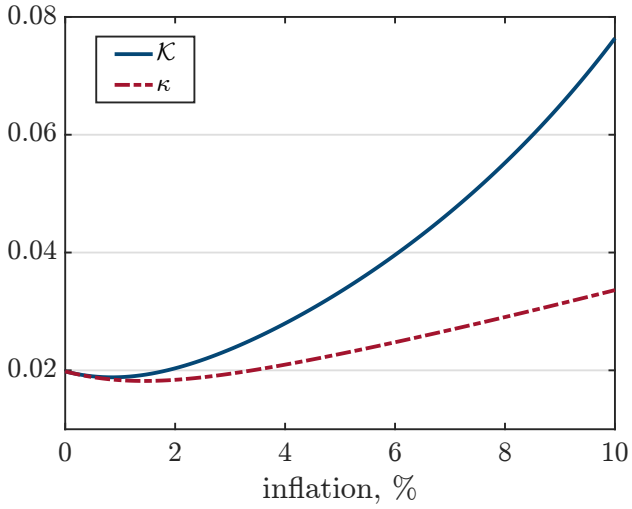
Figure 5 shows how the slope of the Phillips curve $\mathcal{K}$ varies with trend inflation in our baseline model. We gauge the importance of the inflation accelerator by contrasting the actual slope $\mathcal{K}$ with the slope κ that arises absent the inflation accelerator. In computing these two objects, we allow the fraction of price changes n to optimally increase with trend inflation according to equation (13).

At low rates of trend inflation the slope of the Phillips curve declines with inflation, owing to the negative relationship between $\mathcal{N}$ and π discussed above. Moreover, since the elasticity $\mathcal{M}$ governing the response of inflation to changes in the frequency of price changes is close to zero, the slope of the Phillips curve $\mathcal{K}$ is nearly the same as κ. The feedback effect is considerably amplified, however, at high rates of trend inflation. Though κ only increases from 0.020 to 0.034 as trend inflation increases from 0 to 10%, the overall slope $\mathcal{K}$ increases from 0.020 to 0.076. The inflation accelerator thus considerably magnifies the slope of the Phillips curve.

5 The Phillips Curve in the Time-Series

We next investigate how the slope of the Phillips curve evolves in the U.S. time series viewed through the lens of our model. To that end, we first identify the sequence of monetary policy shocks that allows our model to exactly reproduce the path of inflation in the post-war U.S. data. We then consider a log-linear approximation around the equilibrium point at each date and derive the slope of the Phillips curve. We show that the slope of the Phillips curve varies considerably, ranging from 0.02 in relatively low-inflation periods to 0.12 in the high-inflation periods of the 1970s and 1980s.

Figure 5: Trend Inflation and the Slope of Phillips Curve



5.1 Inflation and the Fraction of Price Changes

Recall that inflation in our model is a function of last period's price level p_{t-1}, the degree of misallocation, x_{t-1}, as well as the monetary policy shock ε_t

$$\pi_t = \pi\left(\frac{p_{t-1}}{\exp(\mu + \varepsilon_t)}, x_{t-1}\right). \tag{18}$$

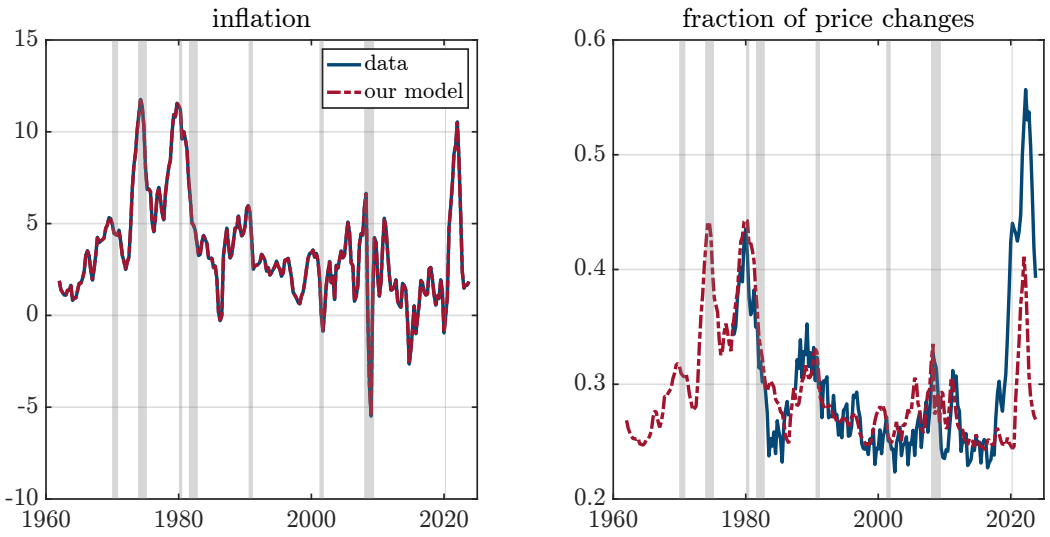
We initialize the economy in the stochastic steady state in 1962 and use the non-linear solution of our model to back out the monetary policy shocks that reproduce the observed U.S. inflation series. For visual clarity, we target an inflation series that removes high-frequency fluctuations using a 3-quarter centered moving average.[14]

Figure 6 shows the path of annualized quarterly inflation, which the model matches by construction, and the fraction of price changes in both the model and the data. The model reproduces well the relatively high fraction of price changes in the 1980s and its subsequent decline following the Volcker disinflation. Though the fraction of price changes also increases in our model during the post-Covid spike in inflation, the increase is not as large as in the data. Intuitively, our model predicts a stable relationship between inflation and the fraction of price changes. Since the post-Covid increase in inflation was not as large as that in the 1980s, the model predicts a smaller frequency response. Allowing for a reduction in the cost of changing prices to capture improvements in the technology for changing prices would

[14]In the Appendix we compare the raw inflation series and the smoother series we target and report the implied slope of the Phillips curve when we target the raw inflation data.

improve the model's fit and strengthen our conclusions that the slope of the Phillips curve considerably steepens in periods of high inflation.

Figure 6: Inflation and the Fraction of Price Changes



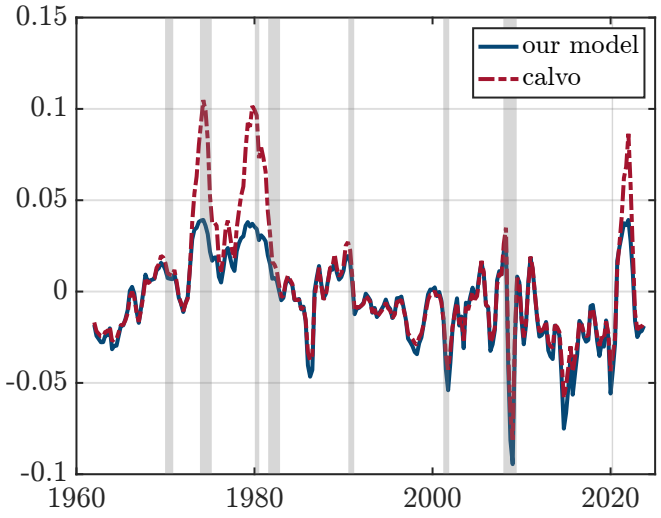
Notes: The gray bars indicate NBER recessions. The left panel plots the annualized quarterly inflation and the right panel plots the quarterly fraction of price changes. Both data series are smoothed with a 3-quarter centered moving average.

5.2 The Slope of the Phillips Curve

We first illustrate the non-linear relationship between inflation and the output gap implied by our model by contrasting the time series of the output gap with that predicted by the Calvo model. Recall that both models exactly match the time series of inflation in the data and have the same average frequency of price changes. As Figure 7 shows, the output gap in the Calvo model is considerably larger than in our model in the high-inflation years of the 1970s and 1980s, occasionally reaching 10%. In contrast, in our model the output gap is never above 4%. Our model therefore implies that the monetary authority is limited in its ability to use expansionary monetary policy to achieve an increase in the output gap.

We next discuss in more detail the determinants of the slope of the Phillips curve at each point in time in our model. To this end, we consider the impact of an additional monetary policy shock $\tilde{\varepsilon}_t$ that changes the growth rate of nominal spending to $\tilde{\mu}_t = \mu_t + \tilde{\varepsilon}_t$ in period t. We use a tilde to denote the value of a variable following this additional shock and hats denote the log-deviation of a variable from the original equilibrium. For example, $\hat{\pi}_t = \log \tilde{\pi}_t - \log \pi_t$ denotes the response of inflation to the shock.

Figure 7: Output Gap



Notes: The gray bars indicate NBER recessions.

As we show in the Appendix, the expression relating inflation to the fraction of price changes and the relative reset price is now

$$\hat{\pi}_t = \underbrace{\frac{1}{(1-n_t)\,\pi_t^{\theta-1}}\frac{\pi_t^{\theta-1}-1}{\theta-1}}_{\mathcal{M}_t}\hat{n}_t + \underbrace{\frac{1-(1-n_t)\,\pi_t^{\theta-1}}{(1-n_t)\,\pi_t^{\theta-1}}}_{\mathcal{N}_t}(\hat{p}_t^* - \hat{p}_t)\,.$$

This expression is similar to that derived in equation (14) which perturbed the economy around the non-stochastic steady state, except that now the actual values of inflation π_t and fraction of price changes n_t determine how inflation reacts to an increase in the optimal reset price and the fraction of price changes. Once again, if inflation is high in a given period, the elasticity $\mathcal{M}_t$ that determines how inflation responds to an additional increase in the fraction of price changes is high as well: in times of elevated inflation the desired price change is high, so even a small increase in the frequency greatly increases aggregate price flexibility.

Consider next the expression describing how the fraction of price changes responds to shocks. Up to a first-order approximation,

$$\hat{n}_t = \mathcal{A}_t\hat{\pi}_t + \mathcal{B}_t\,(\hat{p}_t^* - \hat{p}_t) - \mathcal{C}_t\hat{x}_{t-1} + \frac{n_t - \bar{n}}{n_t}\hat{b}_{1t},$$

where once again the elasticities vary over time as a function of inflation and the optimal reset price. For example,

$$\mathcal{A}_t = \frac{\theta-1}{\xi n_t}b_{1t}\left(\left(\frac{p_t^*}{p_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)}\pi_t^{\frac{\theta}{\eta}}\left(x_{t-1}\right)^{-\frac{\theta}{\eta}} - \pi_t^{\theta-1}\right),$$

and

$$\mathcal{B}_t = (1-\tau\eta)(\theta-1)\frac{b_{1t}}{\xi n_t}\left(\left(\frac{p_t^*}{p_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)}\pi_t^{\frac{\theta}{\eta}}(x_{t-1})^{-\frac{\theta}{\eta}} - \left(\frac{p_t^*}{p_t}\right)^{1-\theta}\right).$$

Finally, the slope of the Phillips curve is equal to

$$\mathcal{K}_t = \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{p_t^{-\frac{1}{\eta}}}{b_{2t}}\frac{\mathcal{M}_t\mathcal{B}_t+\mathcal{N}_t}{1-\mathcal{M}_t\mathcal{A}_t}.$$

We once again find it useful to compare the slope of the Phillips curve in our model to that in a model with a time-varying frequency of price changes that is constrained not to respond to the additional shocks $\tilde{\varepsilon}_t$. In this case the slope is equal to

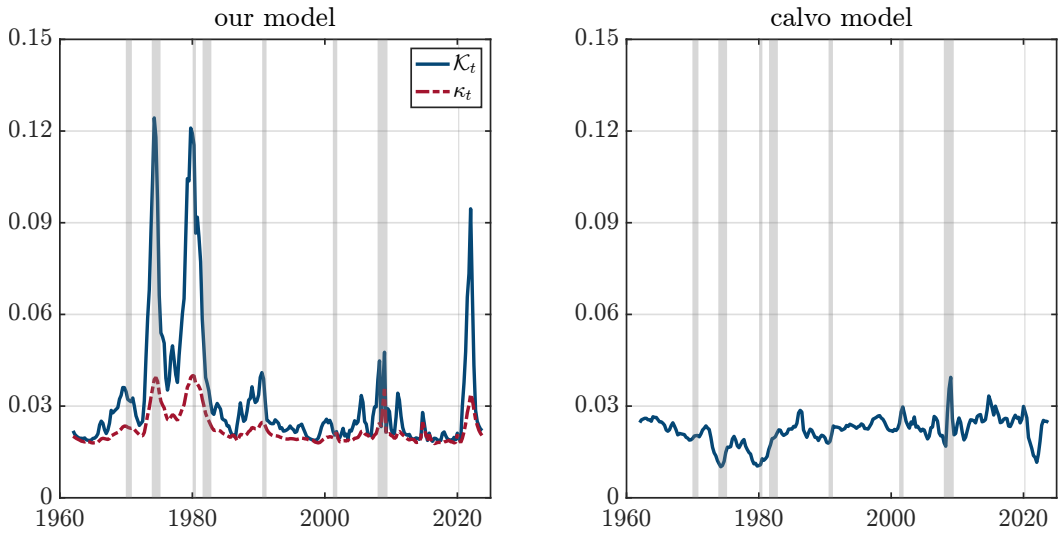
$$\kappa_t = \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{p_t^{-\frac{1}{\eta}}}{b_{2t}}\mathcal{N}_t$$

and captures that an elevated fraction of price changes affects the slope of the Phillips curve. As before, the difference between the two slopes, $\mathcal{K}_t - \kappa_t$, captures the inflation accelerator which now varies over time and reflects the endogenous response of the frequency of price changes and its disproportionately larger contribution to aggregate price flexibility in periods of high inflation.

The left panel of Figure 8 depicts the slope of the Phillips curve in our model and contrasts it to κ_t. The slope of the Phillips curve $\mathcal{K}_t$ fluctuates significantly over time, reaching a low of 0.02 in low-inflation periods and increasing to as high as 0.12 in times of high inflation. Critically, the inflation accelerator is largely responsible for the steeper slope in high-inflation periods: absent the inflation accelerator, the slope κ_t peaks at only 0.04, a third of the overall effect. Even though our model does not fully reproduce the sharp increase in the frequency of price changes post-Covid, the slope of the Phillips curve increased by a factor of five, from 0.019 in the first quarter of 2019 to 0.095 in the first quarter of 2022, an increase once again largely accounted for by the inflation accelerator: κ_t only increased from 0.018 to 0.034 in this period.

For comparison, the right panel of Figure 8 reports the slope predicted by the Calvo model discussed in Section 3. This slope fluctuates much less than in our model and, in fact, decreases in periods of high inflation. This is because, as discussed in Section 4.2, the price level is less responsive to changes in the reset price when inflation is high, that is, $\mathcal{N}_t$ falls, a mechanism that reduces the slope of the Phillips curve.

Figure 8: The Slope of the Phillips Curve



Notes: The gray bars indicate NBER recessions.

5.3 Time-Varying Responses to Monetary Policy Shocks

We next study the consequences of the elevated slope of the Phillips curve in periods of high inflation for how the economy responds to monetary policy shocks. We isolate the role of the inflation accelerator by considering a log-linear approximation of the model's equilibrium conditions around each date t and expressing the system as

$$\mathbf{A}_t \mathbf{z}_t = \mathbf{B}_t \mathbf{z}_{t-1} + \mathbf{C}_t \mathbf{z}_{t+1}, \tag{19}$$

where $\mathbf{z}_t$ is a vector that collects all the equilibrium variables, expressed in log-deviations from the initial equilibrium, and the matrices $\mathbf{A}_t$ to $\mathbf{C}_t$ collect the time-varying elasticities describing the log-linearized equilibrium conditions, including the elasticities $\mathcal{M}_t$, $\mathcal{N}_t$ and $\mathcal{K}_t$ defined above.[15] In calculating the matrices $\mathbf{A}_t$ to $\mathbf{C}_t$ we use the history of monetary shocks that reproduces the inflation data up to the period of the shock, and then set the path of all future shocks to zero.[16]

We use the representation in equation (19) to recover the solution of the model,

$$\mathbf{z}_t = \mathbf{Q}_t \mathbf{z}_{t-1}, \tag{20}$$

where

$$\mathbf{Q}_t = \left(\mathbf{A}_t - \mathbf{C}_t \mathbf{Q}_{t+1}\right)^{-1} \mathbf{B}_t.$$

[15] See the Appendix for a full list of the log-linearized equilibrium equations.

[16] We also considered an alternative approach in which we simulated a large number of histories of shocks going forward and found that the average response is similar to that described here.

We do this for periods $t = 1 \dots T$, where period 1 corresponds to the period in the sample for which we calculate the impulse response, and T is sufficiently large so that the impact of initial conditions dies out. Equation (20) thus allows us to compute a conditional forecast of how the economy would respond to an additional change in monetary policy at any point in time. The log-linearized solution (20) produces impulse responses to relatively small shocks, say of 1%, that are very similar to those obtained from the non-linear solution. For larger shocks, the non-linearity in our model is stronger, further reinforcing our conclusions below.

We find the log-linear approximation above useful because it allows us to isolate the role of the inflation accelerator in determining how output responds to monetary policy shocks. To this end, we recompute the solution of the model by setting $\mathcal{M}_t = 0$ at every date and leaving all other elasticities unchanged. This alternative solution captures what the responses would be in the absence of the inflation accelerator.

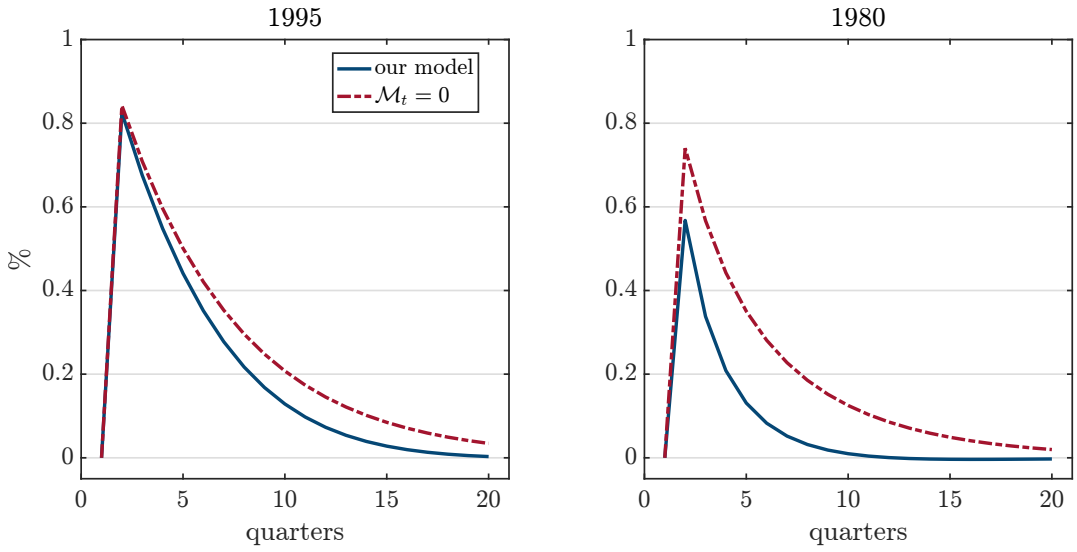
Figure 9 compares the impulse response of output to a 1% monetary shock in the the first quarter of 1995, when inflation was relatively low, 2.8% year-on-year, and in the first quarter of 1980, when inflation was much higher, 11.7%. As the figure shows, the real effects of the monetary shock are much smaller in 1980: our model predicts that the cumulative response of output is equal to only 1.4%, much smaller than the 4.0% in 1995. Part of this difference is mechanically accounted for by the higher frequency of price changes in 1980: 44% vs. 27%. The bulk of the difference, however, is accounted for by the inflation accelerator: setting $\mathcal{M}_t = 0$ increases the cumulative impulse response to a shock in 1980 significantly, to 3.7%. Thus, the endogenous increase in the fraction of price changes significantly increases the flexibility of the aggregate price level.

5.4 The Sacrifice Ratio

The time-varying nature of the slope of the Phillips curve in our model has important implications for the tradeoffs policymakers face in stabilizing prices and real activity. We illustrate how these tradeoffs change over time by calculating a measure of the *sacrifice ratio*. Specifically, we ask: what is the drop in output required to reduce inflation by one percentage point during the course of one year? We use the non-linear solution of the model to back out the change in nominal spending necessary to achieve this reduction in inflation and then calculate the average decline in output during the course of the four quarters of that year. We repeat this experiment for every date and report the results in Figure 10.

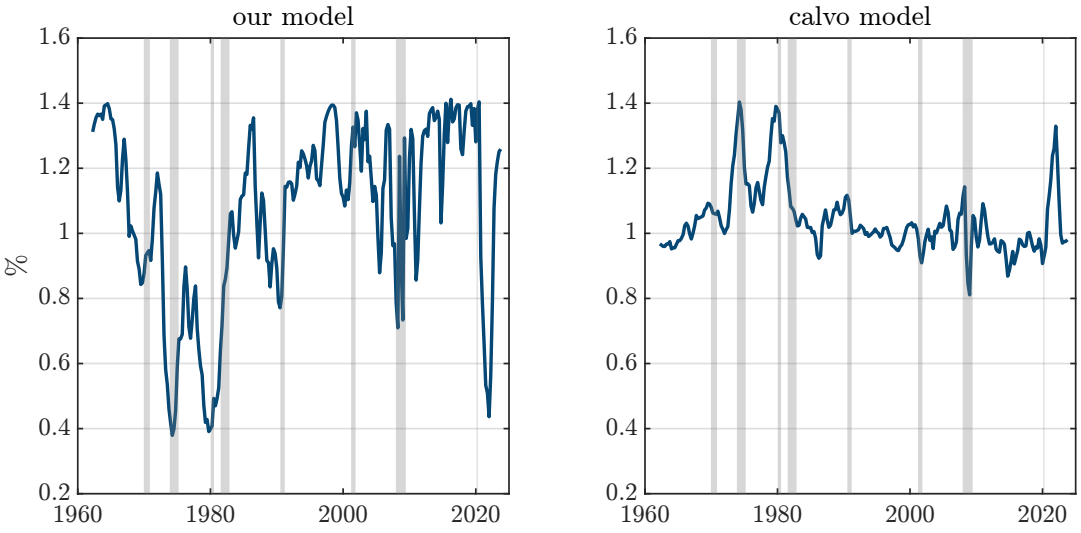
As the left panel of the figure shows, in periods of low inflation the sacrifice ratio is

Figure 9: Output Responses to Monetary Shocks In Different Periods



Notes: The year-on-year inflation rate was equal to 2.8% in the first quarter of 1995 and 11.7% in the first quarter of 1980.

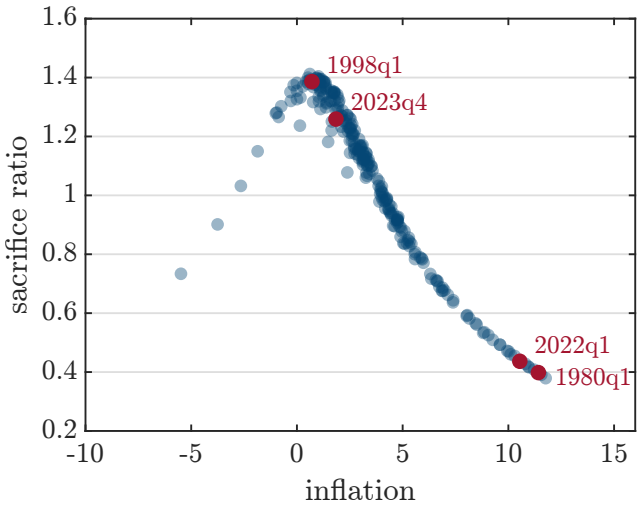
Figure 10: Sacrifice Ratio



Notes: The gray bars indicate NBER recessions.

approximately 1.4%. That is, output would have to fall by 1.4% on average over the course of the year in order for the monetary authority to reduce inflation by one percentage point. When inflation is at its peak, in the 1970s and 1980s, the sacrifice ratio is only 0.4%. Thus, reducing inflation by one percentage point in that period would have been a lot less costly. Interestingly, even though our model does not fully match the increase in the fraction of price changes post Covid, it predicts a sharp decline in the sacrifice ratio, from 1.4% prior to the pandemic to approximately 0.45% when inflation was at its peak in 2022.[17] In contrast, the sacrifice ratio fluctuates much less over time in the Calvo model and, in fact, increases in times of high inflation.

Figure 11: Inflation and the Sacrifice Ratio



We therefore conclude that our model implies that if inflation is high to begin with, bringing it down requires a smaller drop in output than if inflation is low. We illustrate this point in Figure 11, which shows a scatterplot of the sacrifice ratio against inflation and highlights the high-inflation period in the 1980s and the low-inflation period at the end of the 1990s. The figure also highlights the post-Covid period and shows that the sacrifice ratio was 0.45% at the beginning of 2022, when inflation was high, and increased to 1.25% as inflation fell by the end of 2023.

[17] Hobijn et al. (2023) also argue that the sacrifice ratio fell after the onset of the pandemic due to the steepening of the Phillips curve.

6 Robustness

In this section, we show that our results are robust to eliminating strategic complementarities in price setting, to adding idiosyncratic shocks to match the dispersion in price changes, to considering a more conventional interest rate rule and to adding multiple sources of aggregate uncertainty.

6.1 The Role of Strategic Complementarities

In our baseline model we assumed a moderate degree of strategic complementarities in pricing by setting $\eta = 2/3$ and $\theta = 6$. Here we gauge the robustness of our results to eliminating strategic complementarities by setting $\eta = 1$. We consider two economies, one in which $\theta = 6$, as in our baseline model, and one in which $\theta = 3$. In both of these, we keep $\beta = 0.99$. We re-calibrate each of these economies to match the same moments as in the baseline calibration.

Table 2 shows that both economies match the targeted moments perfectly. Eliminating strategic complementarities reduces the curvature of the profit function and thus the firms' incentives to adjust prices so, as Panel B of the table shows, the model requires smaller adjustment costs to match the extent to which the fraction of price changes comoves with inflation, as in menu cost models (see Blanco et al., 2024). When $\theta = 6$, adjustment costs amount to 0.12% of all labor costs and when $\theta = 3$, they amount to 0.04% of the total labor costs, lower than the 0.65% implied by our baseline calibration.

Figure 12 shows that eliminating strategic complementarities increases the slope of the Phillips curve considerably, more so when θ is lower. Intuitively, in our baseline model in which $\eta = 2/3$, the slope of the Phillips curve is dampened by a factor of $1 + \theta\left(\frac{1}{\eta} - 1\right) = 4$, which is no longer present when $\eta = 1$. Additionally, reducing θ increases the impact of newly reset prices on inflation, as captured by the elasticity $\mathcal{N}_t$, further increasing the slope of the Phillips curve. Our earlier conclusion stands, however: the slope of the Phillips curve greatly increases in times of high inflation, primarily due to the inflation accelerator. For example, as the left panel of the figure shows, the slope $\mathcal{K}_t$ reaches a low of 0.07 and increases to as much as 0.55 in times of high inflation. Absent the inflation accelerator, κ_t increases from 0.07 to only 0.17.

Table 2: Calibration: Alternative Parameterization

A. Targeted Moments

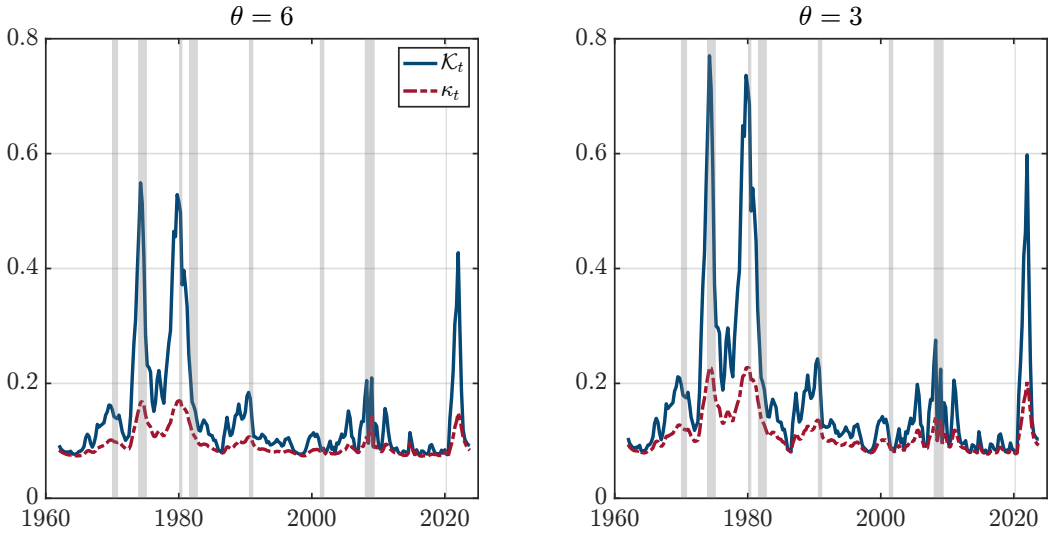
	Data	$\theta = 6$	$\theta = 3$
mean inflation	3.517	3.517	3.517
s.d. inflation	2.739	2.739	2.739
mean frequency	0.297	0.297	0.297
slope of n_t on $\lvert\pi_t\rvert$	0.016	0.016	0.016

B. Calibrated Parameter Values

		$\theta = 6$	$\theta = 3$
μ	mean spending growth rate	0.035	0.035
σ	s.d. monetary shocks	0.019	0.018
$\bar{n}$	fraction free price changes	0.232	0.227
ξ	adjustment cost	0.365	0.109

Note: The mean nominal spending growth rate is annualized.

Figure 12: Slope of the Phillips Curve, No Strategic Complementarities



Notes: The gray bars indicate NBER recessions.

6.2 Adding Idiosyncratic Shocks

For clarity of exposition, in our baseline model we abstracted from idiosyncratic shocks, so that model cannot reproduce the large dispersion in price changes observed in the data. We

show next that our model can easily be extended to allow for such idiosyncratic shocks, as in the literature on menu cost models inspired by Golosov and Lucas (2007), and that our results are robust to this. To continue maintaining tractability, we assume that these shocks are at the product level.[18] Since idiosyncratic now generate a motive for price changes even when inflation is zero, we assume that $\bar{n} = 0$, so there are no free price changes.

Environment. We assume that intermediate good firms produce a continuum of products k with technology

$$y_{ikt} = z_{ikt} \left(l_{ikt}\right)^{\eta},$$

where z_{ikt} is a product-level shock that evolves according to

$$z_{ikt} = z_{ikt-1} e^{\sigma_z \epsilon_{ikt}}$$

and ϵ_{ikt} is a iid standard normal random variable. Intermediate goods are aggregated into a final good with a CES aggregator

$$y_t = \left(\int_0^1 \int_0^1 \left(\frac{y_{ikt}}{z_{ikt-1}} \right)^{\frac{\theta-1}{\theta}} \mathrm{d}k\mathrm{d}i \right)^{\frac{\theta}{\theta-1}},$$

so the shock z_{ikt} has the interpretation of a quality shock (Midrigan, 2011, Blanco et al., 2024). This implies that the demand for an individual product,

$$y_{ikt} = z_{ikt} \left(\frac{z_{ikt} P_{ikt}}{P_t} \right)^{-\theta} y_t,$$

is a function of its quality-adjusted price and the corresponding aggregate price index is

$$P_t = \left(\int_0^1 \int_0^1 \left(z_{ikt} P_{ikt} \right)^{1-\theta} \mathrm{d}k\mathrm{d}i \right)^{\frac{1}{1-\theta}}.$$

Price Setting and Equilibrium. Adding idiosyncratic shocks requires that we evaluate the price index and misallocation in equations (3) and (4) using the quality-adjusted product prices, $z_{ikt} P_{ikt}$. The firms' price setting problem can then be characterized as in the baseline model. In a symmetric equilibrium, the price index is

$$1 = n_t \left(\frac{p_t^*}{p_t} \right)^{1-\theta} + (1 - n_t) \exp\left(\frac{\sigma_z^2}{2} (1-\theta)^2 \right) \pi_t^{\theta-1}.$$

[18] Assuming, instead, that shocks are to individual firms would require keeping track of the distribution of firm prices, as in menu cost models.

The only difference relative to the definition of the price index in equation (8) is the term $\exp\left(\frac{\sigma_z^2}{2}(1-\theta)^2\right)$, which effectively decreases inflation. To see this, notice that if the relative reset price, $p_t^*/p_t = 1$, inflation is equal to $\pi_t = -\exp\left(\frac{\sigma_z^2}{2}(\theta-1)\right) < 1$, reflecting that consumers reallocate spending towards cheaper quality-adjusted varieties. A similar adjustment appears in the law of motion of misallocation,[19]

$$x_t^{-\frac{\theta}{\eta}} = n_t\left(\frac{p_t^*}{p_t}\right)^{-\frac{\theta}{\eta}} + (1-n_t)\exp\left(\frac{\sigma_z^2}{2}\left(\frac{\theta}{\eta}\right)^2\right)\left(\frac{x_{t-1}}{\pi_t}\right)^{-\frac{\theta}{\eta}},$$

the optimal fraction of price changes,

$$\xi n_t = b_{1t}\left(\left(\frac{p_t^*}{p_t}\right)^{1-\theta} - \exp\left(\frac{\sigma_z^2}{2}(1-\theta)^2\right)\pi_t^{\theta-1}\right) - \tau b_{2t}\left(\left(\frac{p_t^*}{p_t}\right)^{-\frac{\theta}{\eta}} - \exp\left(\frac{\sigma_z^2}{2}\left(\frac{\theta}{\eta}\right)^2\right)x_{t-1}^{-\frac{\theta}{\eta}}\pi_t^{\frac{\theta}{\eta}}\right),$$

and the expressions for b_{1t} and b_{2t},

$$\begin{aligned} b_{1t} &= 1 + \beta\exp\left(\frac{\sigma_z^2}{2}(1-\theta)^2\right)\mathbb{E}_t\left((1-n_{t+1})\pi_{t+1}^{\theta-1}b_{1t+1}\right) \\ b_{2t} &= p_t^{-\frac{1}{\eta}} + \beta\exp\left(\frac{\sigma_z^2}{2}\left(\frac{\theta}{\eta}\right)^2\right)\mathbb{E}_t\left((1-n_{t+1})\pi_{t+1}^{\frac{\theta}{\eta}}b_{2t+1}\right). \end{aligned}$$

Phillips Curve. The expression for the slope of the Phillips curve is identical to that in the baseline model discussed in Section 5.2, except that the terms involving $\pi_t^{\theta-1}$ in the elasticities $\mathcal{A}_t$, $\mathcal{B}_t$, $\mathcal{M}_t$ and $\mathcal{N}_t$ are scaled by $\exp\left(\frac{\sigma_z^2}{2}(1-\theta)^2\right)$ and those involving $\pi_t^{\frac{\theta}{\eta}}$ are scaled by $\exp\left(\frac{\sigma_z^2}{2}\left(\frac{\theta}{\eta}\right)^2\right)$.

Quantification. We assign the same values to the discount factor, returns to scale and demand elasticity. In addition to the moments targeted in the baseline model, we also match the standard deviation of individual price changes, Δp_{ikt}, which is 0.129 in the BLS data (Morales-Jimenez and Stevens, 2024), by appropriately setting the standard deviation of idiosyncratic shocks, σ_z.[20] Table 3 reports the results of the calibration. We note that the the model now requires a larger adjustment cost, ξ, since idiosyncratic shocks increase the marginal benefit of adjusting and also make it more sensitive to inflation.

Figure 13 shows how the slope of the Phillips curve, $\mathcal{K}_t$, evolves over time in this model. As in the baseline, this slope fluctuates considerably, ranging from 0.01 in the 1990s to 0.05 in the

[19] In the presence of idiosyncratic shocks, the losses from misallocation are large. As shown by Blanco et al. (2024), this problem can be remedied by assuming that the elasticity of substitution inside the firm is lower than θ. In this case, misallocation would be a function of the within-firm demand elasticity.

[20] In Appendix D, we derive closed form expressions for the moments of the distribution of price changes.

Table 3: Endogenously Calibrated Parameters, Model with Idiosyncratic Shocks

A. Targeted Moments

	Data	Model
mean inflation	3.517	3.517
s.d. inflation	2.739	2.739
mean frequency	0.297	0.297
slope of n_t on $\lvert\pi_t\rvert$	0.016	0.015
s.d. price changes	0.129	0.129

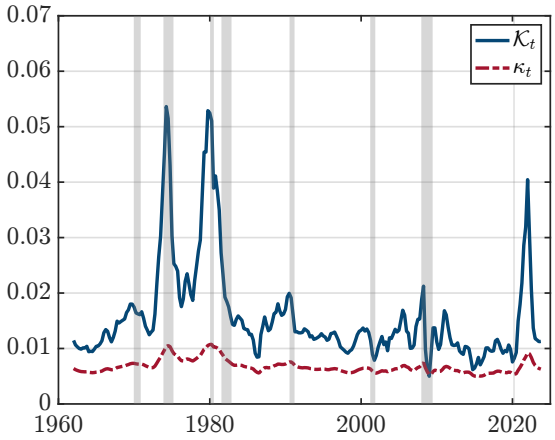
B. Calibrated Parameter Values

		Model
μ	mean spending growth rate	0.035
σ	s.d. monetary shocks	0.023
ξ	adjustment cost	17.00
σ_z	s.d. idiosyncratic shocks	0.068

Notes: The mean nominal spending growth rate is annualized.

1970s. While idiosyncratic shocks dampen the slope of the Phillips curve because consumers reallocate towards products whose prices have not adjusted, the inflation accelerator is as powerful as in the baseline. For example, absent the inflation accelerator, the slope κ_t peaks at only 0.01 in the 1970s, one-fifth of the overall effect.

Figure 13: Slope of the Phillips Curve, Model with Idiosyncratic Shocks



Notes: The gray bars indicate NBER recessions.

6.3 A Taylor Rule Monetary Policy

Until now, we have assumed that monetary policy follows a nominal spending rule. We show next that our results are robust to assuming instead that monetary policy follows an interest rate rule. Specifically, we follow Justiniano and Primiceri (2008) and assume that

$$\frac{1+i_t}{1+i} = \left(\frac{1+i_{t-1}}{1+i}\right)^{\phi_i} \left(\left(\frac{\pi_t}{\pi}\right)^{\phi_\pi} \left(\frac{y_t}{y_{t-1}}\right)^{\phi_y}\right)^{1-\phi_i} u_t, \tag{21}$$

where π is the inflation target, $1+i = \pi/\beta$ is the steady-state nominal interest rate, the parameters ϕ_i, ϕ_π and ϕ_y determine the inertia in the interest rate rule and the sensitivity of monetary policy to fluctuations in inflation and output growth, and where the monetary policy shock, u_t, evolves according to

$$\log u_t = \rho_u \log u_{t-1} + \sigma_u \varepsilon_t,$$

where the innovations $\varepsilon_t \sim N(0,1)$. We solve the model using a third-order perturbation.

We assign the same values to the discount factor, returns to scale and demand elasticity as in our baseline model. We set $\phi_i = 0.65$, $\phi_\pi = 2.35$ and $\phi_y = 0.51$, the median estimates reported by Justiniano and Primiceri (2008). We consider two economies, one in which $\rho_u = 0$, in which we target the same moments as in the baseline, and one in which $\rho_u > 0$, in which we also target the autocorrelation of year-on-year inflation, that is, of the series reported in Figure 1. Table 4 reports the results of the calibration.

Figure 14 reports how the slope of the Phillips curve varies in the U.S. time series viewed through the lens of these two models. As in the baseline, the slope of the Phillips curve varies substantially over time, from approximately 0.02 in the 1990s to as high as 0.12–0.14 in the 1970s and 1980s. Once again, the inflation accelerator in responsible for the bulk of the steepening of the Phillips curve in times of high inflation.

6.4 Adding Multiple Aggregate Shocks

To illustrate the mechanism as transparently as possible, in the baseline model we assume that the only source of aggregate uncertainty comes from monetary shocks. In this section, we incorporate other sources of aggregate fluctuations and estimate the model using likelihood methods. Specifically, we study an economy in which, in addition to shocks to the interest rate rule, firms are also subject to aggregate shocks to productivity and a time-varying tax on labor, which we refer to as a "cost-push" shock.

Table 4: Endogenously Calibrated Parameters, Taylor Rule

A. Targeted Moments

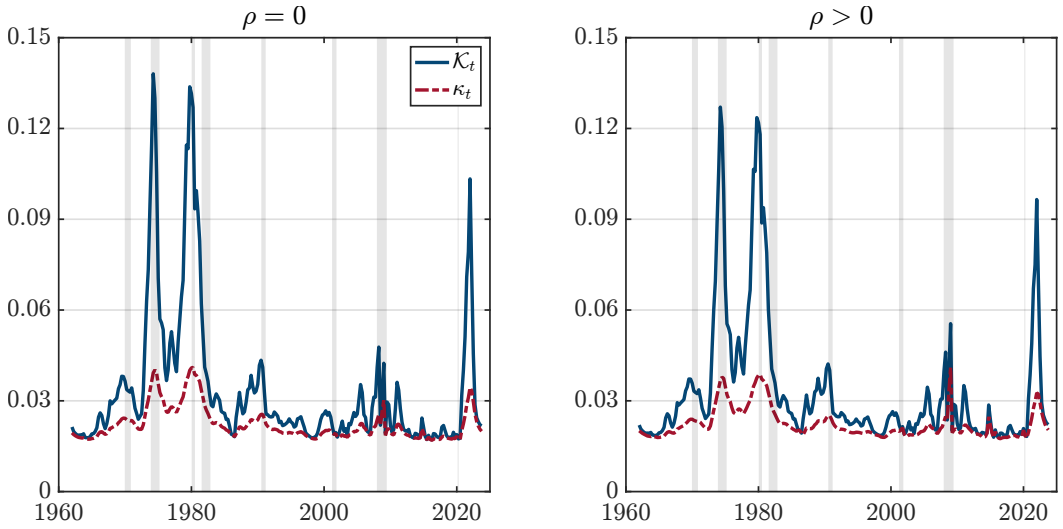
	Data	$\rho_u = 0$	$\rho_u > 0$
mean inflation	3.517	3.517	3.517
s.d. inflation	2.739	2.739	2.739
mean frequency	0.297	0.297	0.297
slope of n_t on $\lvert\pi_t\rvert$	0.016	0.016	0.016
autocorr. inflation	0.942	*0.913*	0.942

B. Calibrated Parameter Values

		$\rho_u = 0$	$\rho_u > 0$
$\log \pi$	inflation target	0.040	0.037
σ_u	s.d. monetary shocks $\times 100$	2.626	0.551
ρ_u	persistence monetary shocks	–	0.685
$\bar{n}$	fraction free price changes	0.241	0.241
ξ	adjustment cost	1.671	1.688

Note: The inflation target is annualized. We italicize the autocorrelation of inflation implied by the economy with $\rho_u = 0$, which is not a target in the calibration.

Figure 14: Slope of the Phillips Curve, Taylor Rule



Notes: The gray bars indicate NBER recessions.

The setup of the model and the derivations are, to a very large extent, as in the baseline. We highlight here the new elements. First, we assume a more general utility function where

consumers' life-time utility is

$$\mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \left(\log c_t - \frac{h_t^{1+\gamma}}{1+\gamma} \right),$$

where $1/\gamma$ is the Frisch elasticity of labor supply. Second, the technology of intermediate goods producers is

$$y_{ikt} = z_t \left(l_{ikt} \right)^{\eta},$$

where z_t is aggregate productivity and follows an AR(1) process. Third, the per period profits from producing a good k are

$$P_{ikt} y_{ikt} - \tau_t W_t l_{ikt},$$

where τ_t is a time-varying tax on labor that also follows an AR(1) process. Lastly, the central bank follows the Taylor rule described in Section 6.3 and the monetary shocks also follow an AR(1) process.

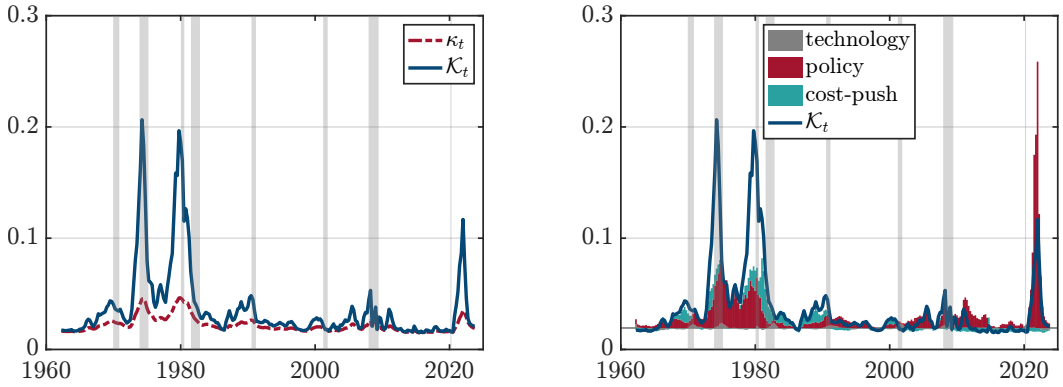
We estimate this model using Bayesian likelihood methods. We discuss the details of the estimation in Appendix D. To compute the slope of the Phillips curve, we use the solution of the model to back out the sequences of technology, cost-push and monetary shocks that allow the model to exactly reproduce the time-series paths of inflation, output growth and the interest rate. Even though the shocks were not chosen to match the path of the fraction of price changes, the model reproduces it well (see Figure D.5).

The left panel of Figure 15 plots the path for the slope of the Phillips curve, $\mathcal{K}_t$, predicted by the model. As in the baseline model, the slope fluctuates considerably over time: from 0.02 in the 1990s to 0.2 in the 1970s and 1980s. Once again, the inflation accelerator is responsible for the bulk of these fluctuations.

In the right panel of Figure 15, we isolate the contribution of each shock to the slope of the Phillips curve by subjecting the economy to one shock at a time. The figure shows that monetary policy shocks are primarily responsible for the increase in the slope of the Phillips curve in periods of high inflation. Notably, the model features important non-linearities from the endogenously increasing fraction of price changes in the 1970s and 1980s: the overall increase in the slope of the Phillips curve predicted by shocks in isolation is much smaller than when the economy is subject to all three shocks.[21]

[21] In Appendix D, we provide a similar decomposition for inflation.

Figure 15: Slope of the Phillips Curve, Model with Multiple Aggregate Shocks



Notes: The gray bars indicate NBER recessions.

7 Conclusions

A widely documented fact is that the frequency of price changes increases in periods of high inflation. We developed a tractable sticky price model in which the frequency of price changes varies endogenously over time. Tractability stems from assuming that firms sell a continuum of products and choose how many, but not which, prices to adjust each period. This eliminates the need to keep track of the price distribution, so our model admits exact aggregation and reduces to a one-equation extension of the Calvo model. The model predicts that the frequency of price changes increases in times of high inflation. The endogenous response of the frequency of price changes to shocks implies a powerful feedback loop between inflation and the frequency of price changes, which we refer to as the *inflation accelerator*. One one hand, an increase in the frequency increases inflation, more so the higher is inflation to begin with. On the other hand, a increase in inflation increases the benefits to adjusting prices and thus further increases the frequency.

When applied to the post-war U.S. time-series data, the model predicts that the slope of the Phillips curve fluctuates considerably over time, ranging from 0.02 in the 1990s to 0.12 in the 1970s and 1980s. The inflation accelerator is responsible for the bulk of this increase. Our findings imply that the tradeoff between inflation and output stabilization is also time-varying: reducing inflation from 10% to 9% is a lot less costly than reducing it from 3% to 2%. Because our model is highly tractable, it can be relatively easily extended to incorporate richer sources of aggregate uncertainty, other frictions and used to conduct empirical and policy analysis. We leave these exercises for future work.

References

Afrouzi, Hassan and Choongryul Yang, "Dynamic Rational Innatention and the Phillips Curve," 2021.

Alvarez, Fernando and Francesco Lippi, "Price Setting With Menu Cost for Multiproduct Firms," *Econometrica*, 2014, *82* (1), 89–135.

_ , _ , and Aleksei Oskolkov, "The Macroeconomics of Sticky Prices with Generalized Hazard Functions," *The Quarterly Journal of Economics*, 11 2021, *137* (2), 989–1038.

_ , Hervé Le Bihan, and Francesco Lippi, "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach," *American Economic Review*, October 2016, *106* (10), 2817–51.

_ , Martin Beraja, Martín Gonzalez-Rozada, and Pablo Andres Neumeyer, "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models," *The Quarterly Journal of Economics*, 09 2018, *134* (1), 451–505.

Ascari, Guido and Tiziano Ropele, "Trend Inflation, Taylor Principle, and Indeterminacy," *Journal of Money, Credit and Banking*, 2009, *41* (8), 1557–1584.

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.

Bakhshi, Hasan, Hashmat Khan, Pablo Burriel-Llombart, and Barbara Rudolf, "The New Keynesian Phillips Curve Under Trend Inflation and Strategic Complementarity," *Journal of Macroeconomics*, 2007, *29* (1), 37–59.

Ball, Laurence M and Sandeep Mazumder, "Inflation Dynamics and the Great Recession," Working Paper 17044, National Bureau of Economic Research May 2011.

Barro, Robert J., "A Theory of Monopolistic Price Adjustment," *The Review of Economic Studies*, 01 1972, *39* (1), 17–26.

Blanco, Andres, "Optimal Inflation Target in an Economy with Menu Cost and Zero Lower Bound," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 108–141.

_ , Corina Boar, Callum Jones, and Virgiliu Midrigan, "Non-linear Inflation Dynamics in Menu Cost Economies," 2024.

Burstein, Ariel and Christian Hellwig, "Welfare Costs of Inflation in a Menu Cost Model," *American Economic Review*, May 2008, *98* (2), 438–43.

Calvo, Guillermo A., "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–398.

Caplin, Andrew S. and Daniel F. Spulber, "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, 1987, *102* (4), 703–726.

Cavallo, Alberto, Francesco Lippi, and Ken Miyahara, "Large Shocks Travel Fast," *Forthcoming in American Economic Review: Insights*, 2024.

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland, "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?," *The Review of Economic Studies*, 03 2012, *79* (4), 1371–1406.

Devereux, Michael B. and James Yetman, "Menu Costs and the Long-Run Output–Inflation Trade-off," *Economics Letters*, 2002, *76* (1), 95–100.

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.

Fitzgerald, Terry, Callum Jones, Mariano Kulish, and Juan Pablo Nicolini, "Is There a Stable Relationship between Unemployment and Future Inflation?," *American Economic Journal: Macroeconomics (Forthcoming)*, 2024.

Flynn, Joel P., George Nikolakoudis, and Karthik A. Sastry, "A Theory of Supply Function Choice and Aggregate Supply," 2024.

Gagnon, Etienne, "Price Setting during Low and High Inflation: Evidence from Mexico," *The Quarterly Journal of Economics*, 08 2009, *124* (3), 1221–1263.

Gertler, Mark and John Leahy, "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, 2008, *116* (3), 533–572.

Golosov, Mikhail and Robert Lucas, "Menu Costs and Phillips Curves," *Journal of Political Economy*, 2007, *115* (2), 171–199.

Gordon, N.J., D.J. Salmond, and A.F.M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings F (Radar and Signal Processing)*, 1993, *140*, 107–113.

Greenwald, Daniel, "The Mortgage Credit Channel of Macroeconomic Transmission," 2018.

Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jon Steinsson, "The Slope of the Phillips Curve: Evidence from U.S. States," *The Quarterly Journal of Economics*, 02 2022, *137* (3), 1299–1344.

Hobijn, Bart, Russell Miles, James Royal, and Jing Zhang, "The Recent Steepening of Phillips Curves," *Chicago Fed Letter*, January 2023.

Justiniano, Alejandro and Giorgio E. Primiceri, "Potential and Natural Output," 2008.

Karadi, Peter and Adam Reiff, "Menu Costs, Aggregate Fluctuations, and Large Shocks," *American Economic Journal: Macroeconomics*, July 2019, *11* (3), 111–46.

Kiley, Michael T., "Endogenous Price Stickiness and Business Cycle Persistence," *Journal of Money, Credit and Banking*, 2000, *32* (1), 28–53.

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.

Maliar, Lilia and Serguei Maliar, "Merging Simulation and Projection Approaches to Solve High-Dimensional Problems with an Application to a New Keynesian Model," *Quantitative Economics*, 2015, *6* (1), 1–47.

Midrigan, Virgiliu, "Menu Costs, Multi-product Firms, and Aggregate Fluctuations," *Econometrica*, 2011, *79* (4), 1139–1180.

Montag, Hugh and Daniel Villar, "Price-Setting During the Covid Era," Technical Report, Board of Governors of the Federal Reserve System August 2023.

Morales-Jimenez, Camilo and Luminita Stevens, "Price Rigidities in U.S. Business Cycles," *Working paper*, 2024.

Nakamura, Emi, Jón Steinsson, 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.

Reiter, Michael and Adrian Wende, "Generalized Rotemberg Price-Setting," CESifo Working Paper Series 11297 2024.

Romer, David, "Staggered Price Setting with Endogenous Frequency of Adjustment," *Economics Letters*, 1990, *32* (3), 205–210.

Rotemberg, Julio J., "Sticky Prices in the United States," *Journal of Political Economy*, 1982, *90* (6), 1187–1211.

Sheshinski, Eytan and Yoram Weiss, "Inflation and Costs of Price Adjustment," *The Review of Economic Studies*, 06 1977, *44* (2), 287–303.

Taylor, John B., "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 1980, *88* (1), 1–23.

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit, and Banking*, 2016, *48*, 253–291.

Appendix

For Online Publication

A Detailed Derivations

Here we provide detailed derivations of the main results discussed in text.

A.1 Steady State

We start by characterizing how the key equilibrium variables depend on $\pi = \exp(\mu)$ in the non-stochastic steady state. In steady state, the equilibrium conditions are

$$\left(\frac{p^*}{p}\right)^{1-\theta} = \frac{1-(1-n)\,\pi^{\theta-1}}{n}, \tag{22}$$

$$x^{-\frac{\theta}{\eta}} = \frac{n\left(\frac{p^*}{p}\right)^{-\frac{\theta}{\eta}}}{1-(1-n)\,\pi^{\frac{\theta}{\eta}}}, \tag{23}$$

$$\left(\frac{p^*}{p}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} = \frac{1}{\eta}\frac{b_2}{b_1}, \tag{24}$$

$$b_1 = \frac{1}{1-\beta\,(1-n)\,\pi^{\theta-1}}, \tag{25}$$

$$b_2 = \frac{p^{-\frac{1}{\eta}}}{1-\beta\,(1-n)\,\pi^{\frac{\theta}{\eta}}}, \tag{26}$$

$$\xi\,(n-\bar{n}) = b_1\left(\left(\frac{p^*}{p}\right)^{1-\theta} - \pi^{\theta-1} - \tau\frac{b_2}{b_1}\left(\left(\frac{p^*}{p}\right)^{-\frac{\theta}{\eta}} - x^{-\frac{\theta}{\eta}}\pi^{\frac{\theta}{\eta}}\right)\right). \tag{27}$$

We first derive an expression for p and x as a function of n and π. Combining equations (24), (25) and (26) implies that

$$\left(\frac{p^*}{p}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} = \frac{1}{\eta}\frac{1-\beta\,(1-n)\,\pi^{\theta-1}}{1-\beta\,(1-n)\,\pi^{\frac{\theta}{\eta}}}p^{-\frac{1}{\eta}}. \tag{28}$$

Using equation (22) and $p = 1/y$, we have that the price level and output satisfy

$$p^{-\frac{1}{\eta}} = y^{\frac{1}{\eta}} = \eta\frac{1-\beta\,(1-n)\,\pi^{\frac{\theta}{\eta}}}{1-\beta\,(1-n)\,\pi^{\theta-1}}\left(\frac{n}{1-(1-n)\,\pi^{\theta-1}}\right)^{\frac{1+\theta\left(\frac{1}{\eta}-1\right)}{\theta-1}}. \tag{29}$$

To find the losses from misallocation, combine equations (22) and (23) and write

$$x^{-\frac{\theta}{\eta}} = \frac{n}{1-(1-n)\,\pi^{\frac{\theta}{\eta}}}\left(\frac{1-(1-n)\,\pi^{\theta-1}}{n}\right)^{\frac{1}{\eta}\frac{\theta}{\theta-1}},$$

which implies

$$x^{\theta}=\left(\frac{1-(1-n)\,\pi^{\frac{\theta}{\eta}}}{n}\right)^{\eta}\left(\frac{1-(1-n)\,\pi^{\theta-1}}{n}\right)^{-\frac{\theta}{\theta-1}}.$$

To find n we use equation (28), which can be rearranged as

$$\xi\left(n-\bar{n}\right)=\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\left(\left(1-\tau\eta\right)\left(\frac{p^{*}}{p}\right)^{1-\theta}-\pi^{\theta-1}+\tau\eta\left(\frac{p^{*}}{p}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)}x^{-\frac{\theta}{\eta}}\pi^{\frac{\theta}{\eta}}\right),$$

or, using equations (22) and (23),

$$\xi\left(n-\bar{n}\right)=\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\left(\left(1-\tau\eta\right)\frac{1-(1-n)\,\pi^{\theta-1}}{n}-\pi^{\theta-1}+\tau\eta\frac{1-(1-n)\,\pi^{\theta-1}}{n}\frac{n\pi^{\frac{\theta}{\eta}}}{1-(1-n)\,\pi^{\frac{\theta}{\eta}}}\right).$$

Since

$$\frac{1-(1-n)\,\pi^{\theta-1}}{n}-\pi^{\theta-1}=\frac{1-\pi^{\theta-1}}{n},$$

this expression simplifies to

$$\xi\left(n-\bar{n}\right)=\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\frac{1}{n}\left(1-\pi^{\theta-1}-\tau\eta\frac{1-(1-n)\,\pi^{\theta-1}}{1-(1-n)\,\pi^{\frac{\theta}{\eta}}}\left(1-\pi^{\frac{\theta}{\eta}}\right)\right).$$

A.2 Log-Linear Approximation Around the Steady State

Recall that the system is

$$1=n_{t}\left(\frac{p_{t}^{*}}{p_{t}}\right)^{1-\theta}+\left(1-n_{t}\right)\pi_{t}^{\theta-1} \tag{30}$$

$$x_{t}^{-\frac{\theta}{\eta}}=n_{t}\left(\frac{p_{t}^{*}}{p_{t}}\right)^{-\frac{\theta}{\eta}}+\left(1-n_{t}\right)\left(x_{t-1}\right)^{-\frac{\theta}{\eta}}\pi_{t}^{\frac{\theta}{\eta}} \tag{31}$$

$$\xi\left(n_{t}-\bar{n}\right)=b_{1t}\left(\left(1-\tau\eta\right)\left(\frac{p_{t}^{*}}{p_{t}}\right)^{1-\theta}-\left(\pi_{t}\right)^{\theta-1}+\tau\eta\left(\frac{p_{t}^{*}}{p_{t}}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)}\left(x_{t-1}\right)^{-\frac{\theta}{\eta}}\pi_{t}^{\frac{\theta}{\eta}}\right) \tag{32}$$

$$b_{1t}=1+\beta\mathbb{E}_{t}\left(1-n_{t+1}\right)\left(\pi_{t+1}\right)^{\theta-1}b_{1t+1} \tag{33}$$

$$b_{2t}=p_{t}^{-\frac{1}{\eta}}+\beta\mathbb{E}_{t}\left(1-n_{t+1}\right)\left(\pi_{t+1}\right)^{\frac{\theta}{\eta}}b_{2t+1} \tag{34}$$

$$\left(\frac{p_{t}^{*}}{p_{t}}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)}=\frac{1}{\eta}\frac{b_{2t}}{b_{1t}} \tag{35}$$

$$\pi_{t}=\pi\frac{p_{t}}{p_{t-1}}\exp\left(\varepsilon_{t}\right) \tag{36}$$

We next log-linearize all these equations. We use hats to denote the log-deviation of variables from their non-stochastic steady state levels.

A.2.1 Price Index

Log-linearizing equation (30) gives

$$\hat{\pi}_t = \underbrace{\frac{1}{(1-n)\,\pi^{\theta-1}}\frac{\pi^{\theta-1}-1}{\theta-1}}_{\mathcal{M}}\hat{n}_t + \underbrace{\frac{1-(1-n)\,\pi^{\theta-1}}{(1-n)\,\pi^{\theta-1}}}_{\mathcal{N}}\left(\hat{p}_t^* - \hat{p}_t\right) \tag{37}$$

A.2.2 Frequency of Price Changes

Log-linearizing equation (32) we have

$$\xi n\hat{n}_t = \xi\left(n-\bar{n}\right)\hat{b}_{1t} + \frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\left(\left(1-\tau\eta\right)\frac{1-\left(1-n\right)\pi^{\theta-1}}{n}\left(1-\theta\right)\left(\hat{p}_t^*-\hat{p}_t\right) - \pi^{\theta-1}\left(\theta-1\right)\hat{\pi}_t\right)$$
$$+\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\tau\eta\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}\pi^{\frac{\theta}{\eta}}\left(\left(1+\theta\left(\frac{1}{\eta}-1\right)\right)\left(\hat{p}_t^*-p_t\right) - \frac{\theta}{\eta}\hat{x}_{t-1} + \frac{\theta}{\eta}\hat{\pi}_t\right)$$

or, equivalently,

$$\begin{aligned}\hat{n}_t &= \frac{n-\bar{n}}{n}\hat{b}_{1t} + \underbrace{\frac{1}{\xi n}\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\left(\tau\theta\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}\pi^{\frac{\theta}{\eta}} + \pi^{\theta-1}\left(1-\theta\right)\right)}_{\mathcal{A}}\hat{\pi}_t \\ &+ \underbrace{\frac{1}{\xi n}\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\beta\left(1-n\right)\pi^{\theta-1}}\left(\left(1-\tau\eta\right)\frac{1}{n}\left(1-\theta\right) + \tau\eta\frac{\pi^{\frac{\theta}{\eta}}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}\left(1+\theta\left(\frac{1}{\eta}-1\right)\right)\right)}_{\mathcal{B}}\left(\hat{p}_t^*-\hat{p}_t\right) + \\ &- \underbrace{\frac{1}{\xi n}\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\tau\theta\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}\pi^{\frac{\theta}{\eta}}}_{\mathcal{C}}\hat{x}_{t-1}\end{aligned} \tag{38}$$

Since $\tau\theta = \theta - 1$, we can simply the expression for $\mathcal{A}$ to

$$\mathcal{A} = \frac{\theta-1}{\xi n}\frac{1}{1-\beta\left(1-n\right)\pi^{\theta-1}}\frac{\pi^{\frac{\theta}{\eta}}-\pi^{\theta-1}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}},$$

which is zero when $\pi = 1$ and is increasing in π. Thus, not only is inflation more responsive to changes in the frequency in economies with higher trend inflation, but the frequency is itself more sensitive to inflation when trend inflation is higher.

Similarly, we can simplify the expression for $\mathcal{B}$ to

$$\mathcal{B} = \left(1-\tau\eta\right)\frac{\theta-1}{\xi n}\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\beta\left(1-n\right)\pi^{\theta-1}}\frac{1}{n}\frac{\pi^{\frac{\theta}{\eta}}-1}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}},$$

which is also zero when $\pi = 1$ and is increasing with π. Finally,

$$\mathcal{C} = \frac{\theta-1}{\xi n}\frac{1-\left(1-n\right)\pi^{\theta-1}}{1-\beta\left(1-n\right)\pi^{\theta-1}}\frac{\pi^{\frac{\theta}{\eta}}}{1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}}.$$

A.2.3 Optimal Reset Price

Log-linearizing equations (33)-(35) and rearranging implies

$$\begin{aligned}\hat{p}_t^* - \hat{p}_t \;&= -\tfrac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\tfrac{1}{\eta}\left(1-\beta\left(1-n\right)\pi^{\frac{\theta}{\eta}}\right)\hat{p}_t + \tfrac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\beta\left(1-n\right)\left(\tfrac{\theta}{\eta}\pi^{\frac{\theta}{\eta}} - \left(\theta-1\right)\pi^{\theta-1}\right)\mathbb{E}_t\hat{\pi}_{t+1} \\ &+\beta\left(1-n\right)\pi^{\frac{\theta}{\eta}}\mathbb{E}_t\left(\hat{p}_{t+1}^* - \hat{p}_{t+1}\right) + \tfrac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\beta\left(\pi^{\frac{\theta}{\eta}} - \pi^{\theta-1}\right)\mathbb{E}_t\left(\left(1-n\right)\hat{b}_{1t+1} - n\hat{n}_{t+1}\right) \quad (39)\end{aligned}$$

A.2.4 Losses from Misallocation

Log-linearizing equation (31) we have

$$\hat{x}_t = \left(1-\left(1-n\right)\pi^{\frac{\theta}{\eta}}\right)\left(\hat{p}_t^* - \hat{p}_t\right) - \frac{\eta}{\theta}\left(1-\pi^{\frac{\theta}{\eta}}\right)\hat{n}_t + \left(1-n\right)\pi^{\frac{\theta}{\eta}}\left(\hat{x}_{t-1} - \hat{\pi}_t\right)$$

A.2.5 Equation for b_{1t}

Log-linearizing equation (33) we have

$$\hat{b}_{1t} = \beta\left(1-n\right)\pi^{\theta-1}\left(\theta-1\right)\mathbb{E}_t\hat{\pi}_{t+1} + \beta\left(1-n\right)\pi^{\theta-1}\mathbb{E}_t\hat{b}_{1t+1} - \beta n\pi^{\theta-1}\mathbb{E}_t\hat{n}_{t+1}$$

A.2.6 Slope of Phillips Curve

Combining equations (37) and (38) implies

$$\hat{\pi}_t = \frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}}\left(\hat{p}_t^* - \hat{p}_t\right) - \frac{\mathcal{MC}}{1-\mathcal{MA}}\hat{x}_{t-1} + \frac{\mathcal{M}}{1-\mathcal{MA}}\frac{n-\bar{n}}{n}\hat{b}_{1t}$$

To derive an expression for inflation, we multiply both sides of equation (39) by $\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}}$ and add $-\frac{\mathcal{MC}}{1-\mathcal{MA}}x_{t-1} + \frac{\mathcal{M}}{1-\mathcal{MA}}\frac{n-\bar{n}}{n}\hat{b}_{1t}$. Then, the LHS of equation (39) is equal to $\hat{\pi}_t$. Adding and subtracting $\beta\left(1-n\right)\pi^{\frac{\theta}{\eta}}\left(-\frac{\mathcal{MC}}{1-\mathcal{MA}}\hat{x}_t + \frac{\mathcal{M}}{1-\mathcal{MA}}\frac{n-\bar{n}}{n}\mathbb{E}_t\hat{b}_{1t+1}\right)$ to the RHS of (39) to express $\mathbb{E}_t\left(\hat{p}_{t+1}^* - \hat{p}_{t+1}\right)$ as a function of expected inflation and rearranging, implies that

$$\begin{aligned}\hat{\pi}_t = \mathcal{K}\widehat{mc}_t &+ \beta\left(1-n\right)\left(\frac{\frac{\theta}{\eta}\pi^{\frac{\theta}{\eta}} - \left(\theta-1\right)\pi^{\theta-1}}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}} + \pi^{\frac{\theta}{\eta}}\right)\mathbb{E}_t\hat{\pi}_{t+1} + \\ &+\beta\left(1-n\right)\left(\frac{\pi^{\frac{\theta}{\eta}} - \pi^{\theta-1}}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}} - \pi^{\frac{\theta}{\eta}}\frac{\mathcal{M}}{1-\mathcal{MA}}\frac{n-\bar{n}}{n}\right)\mathbb{E}_t\hat{b}_{1t+1} \\ &-\beta n\frac{\pi^{\frac{\theta}{\eta}} - \pi^{\theta-1}}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}}\mathbb{E}_t\hat{n}_{t+1} \\ &+\beta\left(1-n\right)\pi^{\frac{\theta}{\eta}}\frac{\mathcal{MC}}{1-\mathcal{MA}}\hat{x}_t - \frac{\mathcal{MC}}{1-\mathcal{MA}}\hat{x}_{t-1} + \frac{\mathcal{M}}{1-\mathcal{MA}}\frac{n-\bar{n}}{n}\hat{b}_{1t},\end{aligned}$$

where we used that $\widehat{mc}_t = \left(-\frac{1}{\eta}\hat{p}_t\right)$.

The slope of the Phillips curve is

$$\mathcal{K} = \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\left(1-\beta\left(1-n\right)\pi^{\frac{\theta}{\eta}}\right)\frac{\mathcal{MB}+\mathcal{N}}{1-\mathcal{MA}}.$$

A.3 Log-Linear Approximation Around Each Point in Time

We next log-linearized the model around each point in time. To do this we consider the impact of an additional monetary shock $\tilde{\varepsilon}_t$ which changes the money growth rate to

$$\tilde{\mu}_t = \mu_t + \tilde{\varepsilon}_t.$$

We let tildes denote the value of the equilibrium variables following this additional shock and hat denote the log-difference between the tilde equilbrium variable and the original one, e.g. $\hat{\pi}_t = \log\tilde{\pi}_t - \log\pi_t$. The equilibrium of the model can then be described by the system of equations (30)-(36), where each equilibrium variable is replaced by its tilde counterpart. In what follows, we log-linearize this new system, but refer to the original equations for brevity.

A.3.1 Price Index

Log-linearizing equation (30) and using

$$n_t\left(\frac{p_t^*}{p_t}\right)^{1-\theta} = 1-\left(1-n_t\right)\pi_t^{\theta-1}$$

implies

$$\hat{\pi}_t = \underbrace{\frac{1}{\left(1-n_t\right)\pi_t^{\theta-1}}\frac{\pi_t^{\theta-1}-1}{\theta-1}}_{\mathcal{M}_t}\hat{n}_t + \underbrace{\frac{1-\left(1-n_t\right)\pi_t^{\theta-1}}{\left(1-n_t\right)\pi_t^{\theta-1}}}_{\mathcal{N}_t}\left(\hat{p}_t^*-\hat{p}_t\right).$$

A.3.2 Frequency of Price Changes

Log-linearizing (32) implies

$$\begin{aligned}\xi n_t\hat{n}_t = {} & \xi\left(n_t-\bar{n}\right)\hat{b}_{1t} + b_{1t}\left(1-\tau\eta\right)\left(\frac{p_t^*}{p_t}\right)^{1-\theta}\left(1-\theta\right)\left(\hat{p}_t^*-\hat{p}_t\right) - b_{1t}\pi_t^{\theta-1}\left(\theta-1\right)\hat{\pi}_t + \\ & + b_{1t}\tau\eta\left(\frac{p_t^*}{p_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)}\pi_t^{\frac{\theta}{\eta}}\left(x_{t-1}\right)^{-\frac{\theta}{\eta}}\left(\left(1+\theta\left(\frac{1}{\eta}-1\right)\right)\left(\hat{p}_t^*-\hat{p}_t\right)+\frac{\theta}{\eta}\hat{\pi}_t-\frac{\theta}{\eta}\hat{x}_{t-1}\right),\end{aligned}$$

which can be rearranged as

$$\hat{n}_t = \frac{n_t - \bar{n}}{n_t}\hat{b}_{1t} + \underbrace{\frac{\theta-1}{\xi n_t} b_{1t}\left(\left(\frac{p_t^*}{p_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} \pi_t^{\frac{\theta}{\eta}} (x_{t-1})^{-\frac{\theta}{\eta}} - \pi_t^{\theta-1}\right)}_{\mathcal{A}_t}\hat{\pi}_t +$$

$$+ \underbrace{\frac{b_{1t}}{\xi n_t}\left((1-\tau\eta)\left(\frac{p_t^*}{p_t}\right)^{1-\theta}(1-\theta) + \tau\eta\left(1+\theta\left(\frac{1}{\eta}-1\right)\right)\left(\frac{p_t^*}{p_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} \pi_t^{\frac{\theta}{\eta}} (x_{t-1})^{-\frac{\theta}{\eta}}\right)}_{\mathcal{B}_t}(\hat{p}_t^* - \hat{p}_t)$$

$$- \underbrace{(\theta-1)\frac{b_{1t}}{\xi n_t}\left(\frac{p_t^*}{p_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} \pi_t^{\frac{\theta}{\eta}} (x_{t-1})^{-\frac{\theta}{\eta}}}_{\mathcal{C}_t}\hat{x}_{t-1}.$$

Using

$$\tau\eta\left(1+\theta\left(\frac{1}{\eta}-1\right)\right) = (\theta-1)(1-\tau\eta)$$

allows us to write

$$\mathcal{B}_t = (1-\tau\eta)(\theta-1)\frac{b_{1t}}{\xi n_t}\left(\left(\frac{p_t^*}{p_t}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} \pi_t^{\frac{\theta}{\eta}} (x_{t-1})^{-\frac{\theta}{\eta}} - \left(\frac{p_t^*}{p_t}\right)^{1-\theta}\right).$$

A.3.3 Optimal Reset Price

The log-linearized versions of equations (33)-(35) are

$$\begin{aligned}
b_{1t}\hat{b}_{1t} &= \beta\mathbb{E}_t(1-n_{t+1})(\pi_{t+1})^{\theta-1} b_{1t+1}\left((\theta-1)\hat{\pi}_{t+1} + \hat{b}_{1t+1}\right) - \beta\mathbb{E}_t n_{t+1}(\pi_{t+1})^{\theta-1} b_{1t+1}\hat{n}_{t+1} \\
b_{2t}\hat{b}_{2t} &= -(p_t)^{-\frac{1}{\eta}}\frac{1}{\eta}\hat{p}_t + \beta\mathbb{E}_t(1-n_{t+1})(\pi_{t+1})^{\frac{\theta}{\eta}} b_{2t+1}\left(\frac{\theta}{\eta}\hat{\pi}_{t+1} + \hat{b}_{2t+1}\right) - \beta\mathbb{E}_t n_{t+1}(\pi_{t+1})^{\frac{\theta}{\eta}} b_{2t+1}\hat{n}_{t+1}
\end{aligned}$$

and

$$\hat{p}_t^* - \hat{p}_t = \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\left(\hat{b}_{2t} - \hat{b}_{1t}\right).$$

Because we consider a one-time unanticipated shock $\tilde{\varepsilon}_t$, the resulting transition, that is, the evolution of variables involving hats is deterministic, so we can write

$$\begin{aligned}
\hat{b}_{1t} &= \beta\left(\mathbb{E}_t(1-n_{t+1})(\pi_{t+1})^{\theta-1}\frac{b_{1t+1}}{b_{1t}}\right)\left((\theta-1)\hat{\pi}_{t+1} + \hat{b}_{1t+1}\right) - \beta\left(\mathbb{E}_t n_{t+1}(\pi_{t+1})^{\theta-1}\frac{b_{1t+1}}{b_{1t}}\right)\hat{n}_{t+1} \\
\hat{b}_{2t} &= \frac{(p_t)^{-\frac{1}{\eta}}}{b_{2t}}\left(-\frac{1}{\eta}\hat{p}_t\right) + \beta\left(\mathbb{E}_t(1-n_{t+1})(\pi_{t+1})^{\frac{\theta}{\eta}}\frac{b_{2t+1}}{b_{2t}}\right)\left(\frac{\theta}{\eta}\hat{\pi}_{t+1} + \hat{b}_{2t+1}\right) \\
&- \beta\left(\mathbb{E}_t n_{t+1}(\pi_{t+1})^{\frac{\theta}{\eta}}\frac{b_{2t+1}}{b_{2t}}\right)\hat{n}_{t+1}.
\end{aligned}$$

Subtracting the first expression from the second and multiplying by $\frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}$ gives

$$\begin{aligned}
\hat{p}_t^* - \hat{p}_t &= \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)} \frac{(p_t)^{-\frac{1}{\eta}}}{b_{2t}} \widehat{mc}_t + \\
&+ \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)} \beta \underbrace{\mathbb{E}_t \left(1-n_{t+1}\right) \left(\frac{\theta}{\eta} \left(\pi_{t+1}\right)^{\frac{\theta}{\eta}} \frac{b_{2t+1}}{b_{2t}} - \left(\theta-1\right) \left(\pi_{t+1}\right)^{\theta-1} \frac{b_{1t+1}}{b_{1t}}\right)}_{\mathcal{F}_t} \hat{\pi}_{t+1} + \\
&+ \underbrace{\beta \mathbb{E}_t \left(1-n_{t+1}\right) \left(\pi_{t+1}\right)^{\frac{\theta}{\eta}} \frac{b_{2t+1}}{b_{2t}}}_{\mathcal{G}_t} \left(\hat{p}_{t+1}^* - \hat{p}_{t+1}\right) \\
&+ \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)} \beta \underbrace{\mathbb{E}_t \left(1-n_{t+1}\right) \left(\left(\pi_{t+1}\right)^{\frac{\theta}{\eta}} \frac{b_{2t+1}}{b_{2t}} - \left(\pi_{t+1}\right)^{\theta-1} \frac{b_{1t+1}}{b_{1t}}\right)}_{\mathcal{H}_t} \hat{b}_{1t+1} \\
&- \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)} \beta \underbrace{\mathbb{E}_t n_{t+1} \left(\left(\pi_{t+1}\right)^{\frac{\theta}{\eta}} \frac{b_{2t+1}}{b_{2t}} - \left(\pi_{t+1}\right)^{\theta-1} \frac{b_{1t+1}}{b_{1t}}\right)}_{\mathcal{I}_t} \hat{n}_{t+1}
\end{aligned}$$

A.3.4 Losses from Misallocation

Log-linearizing equation (31) implies

$$-\frac{\theta}{\eta} x_t^{-\frac{\theta}{\eta}} \hat{x}_t = n_t \left(\frac{p_t^*}{p_t}\right)^{-\frac{\theta}{\eta}} \left(\hat{n}_t - \frac{\theta}{\eta}\left(\hat{p}_t^* - \hat{p}_t\right)\right) + \left(1-n_t\right) \left(x_{t-1}\right)^{-\frac{\theta}{\eta}} \pi_t^{\frac{\theta}{\eta}} \left(-\frac{\theta}{\eta}\hat{x}_{t-1} + \frac{\theta}{\eta}\hat{\pi}_t\right) - n_t \left(x_{t-1}\right)^{-\frac{\theta}{\eta}} \pi_t^{\frac{\theta}{\eta}} \hat{n}_t,$$

which can be rearranged as

$$\hat{x}_t = n_t x_t^{\frac{\theta}{\eta}} \left(\frac{p_t^*}{p_t}\right)^{-\frac{\theta}{\eta}} \left(\hat{p}_t^* - \hat{p}_t\right) + \frac{\eta}{\theta} n_t \left(\left(\frac{x_t}{x_{t-1}}\right)^{\frac{\theta}{\eta}} \pi_t^{\frac{\theta}{\eta}} - x_t^{\frac{\theta}{\eta}} \left(\frac{p_t^*}{p_t}\right)^{-\frac{\theta}{\eta}}\right) \hat{n}_t + \left(1-n_t\right) \left(\frac{x_t}{x_{t-1}}\right)^{\frac{\theta}{\eta}} \pi_t^{\frac{\theta}{\eta}} \left(\hat{x}_{t-1} - \hat{\pi}_t\right).$$

A.3.5 Equation for b_{1t}

Lastly, log-linearizing equation (33) gives

$$\hat{b}_{1t} = \underbrace{\beta \mathbb{E}_t \left(1-n_{t+1}\right) \left(\pi_{t+1}\right)^{\theta-1} \frac{b_{1t+1}}{b_{1t}}}_{\mathcal{D}_t} \left(\left(\theta-1\right)\hat{\pi}_{t+1} + \hat{b}_{1t+1}\right) - \underbrace{\beta \mathbb{E}_t n_{t+1} \left(\pi_{t+1}\right)^{\theta-1} \frac{b_{1t+1}}{b_{1t}}}_{\mathcal{E}_t} \hat{n}_{t+1}.$$

A.3.6 Phillips Curve

Following the same steps as in Section A.2.6 allows us to write the Phillips curve

$$
\begin{aligned}
\hat{\pi}_t &= \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{(p_t)^{-\frac{1}{\eta}}}{b_{2t}}\frac{\mathcal{M}_t\mathcal{B}_t+\mathcal{N}_t}{1-\mathcal{M}_t\mathcal{A}_t}\widehat{mc}_t+\frac{\frac{\mathcal{M}_t\mathcal{B}_t+\mathcal{N}_t}{1-\mathcal{M}_t\mathcal{A}_t}}{\frac{\mathcal{M}_{t+1}\mathcal{B}_{t+1}+\mathcal{N}_{t+1}}{1-\mathcal{M}_{t+1}\mathcal{A}_{t+1}}}\beta\mathbb{E}_t\left(1-n_{t+1}\right)\left(\pi_{t+1}\right)^{\frac{\theta}{\eta}}\frac{b_{2t+1}}{b_{2t}}\hat{\pi}_{t+1}\\
&+ \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{M}_t\mathcal{B}_t+\mathcal{N}_t}{1-\mathcal{M}_t\mathcal{A}_t}\beta\mathbb{E}_t\left(1-n_{t+1}\right)\left(\frac{\theta}{\eta}\left(\pi_{t+1}\right)^{\frac{\theta}{\eta}}\frac{b_{2t+1}}{b_{2t}}-\left(\theta-1\right)\left(\pi_{t+1}\right)^{\theta-1}\frac{b_{1t+1}}{b_{1t}}\right)\hat{\pi}_{t+1}+\\
&+ \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{M}_t\mathcal{B}_t+\mathcal{N}_t}{1-\mathcal{M}_t\mathcal{A}_t}\beta\mathbb{E}_t\left(1-n_{t+1}\right)\left(\left(\pi_{t+1}\right)^{\frac{\theta}{\eta}}\frac{b_{2t+1}}{b_{2t}}-\left(\pi_{t+1}\right)^{\theta-1}\frac{b_{1t+1}}{b_{1t}}\right)\hat{b}_{1t+1}\\
&- \frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{\mathcal{M}_t\mathcal{B}_t+\mathcal{N}_t}{1-\mathcal{M}_t\mathcal{A}_t}\beta\mathbb{E}_t n_{t+1}\left(\left(\pi_{t+1}\right)^{\frac{\theta}{\eta}}\frac{b_{2t+1}}{b_{2t}}-\left(\pi_{t+1}\right)^{\theta-1}\frac{b_{1t+1}}{b_{1t}}\right)\hat{n}_{t+1}\\
&- \frac{\mathcal{M}_t\mathcal{C}_t}{1-\mathcal{M}_t\mathcal{A}_t}\hat{x}_{t-1}+\frac{\mathcal{M}_t}{1-\mathcal{M}_t\mathcal{A}_t}\frac{n_t-\bar{n}}{n_t}\hat{b}_{1t}\\
&- \frac{\frac{\mathcal{M}_t\mathcal{B}_t+\mathcal{N}_t}{1-\mathcal{M}_t\mathcal{A}_t}}{\frac{\mathcal{M}_{t+1}\mathcal{B}_{t+1}+\mathcal{N}_{t+1}}{1-\mathcal{M}_{t+1}\mathcal{A}_{t+1}}}\beta\mathbb{E}_t\left(1-n_{t+1}\right)\left(\pi_{t+1}\right)^{\frac{\theta}{\eta}}\frac{b_{2t+1}}{b_{2t}}\left(-\frac{\mathcal{M}_{t+1}\mathcal{C}_{t+1}}{1-\mathcal{M}_{t+1}\mathcal{A}_{t+1}}\hat{x}_t+\frac{\mathcal{M}_{t+1}}{1-\mathcal{M}_{t+1}\mathcal{A}_{t+1}}\frac{n_{t+1}-\bar{n}}{n_{t+1}}\hat{b}_{1t+1}\right),
\end{aligned}
$$

so the slope of the Phillips curve is

$$
\mathcal{K}_t=\frac{1}{1+\theta\left(\frac{1}{\eta}-1\right)}\frac{(p_t)^{-\frac{1}{\eta}}}{b_{2t}}\frac{\mathcal{M}_t\mathcal{B}_t+\mathcal{N}_t}{1-\mathcal{M}_t\mathcal{A}_t}.
$$

B Solution Method

We describe our global solution method and show that it produces dynamics that are similar to those obtained by solving the model using a third-order perturbation.

To solve the model globally, we use Chebyshev polynomials to approximate all equilibrium objects. Our baseline model has two state variables, last period's misallocation, x_{t-1}, as well as last period's nominal price level deflated by this period's nominal spending,

$$s_t = \frac{P_{t-1}}{M_t},$$

which evolves according to

$$s_{t+1} = \frac{p_t}{\exp(\mu + \varepsilon_t)},$$

where recall that $p_t = P_t/M_t$ is the nominal price level detrended by nominal spending. Letting $\Phi(s_t, x_{t-1})$ denote a row vector collecting the basis functions (tensor product of univariate Chebyshev polynomials) and γ^i a column vector of coefficients characterizing a particular variable, say, π, we have

$$\pi(s_t, x_{t-1}) = \Phi(s_t, x_{t-1}) \times \gamma^\pi.$$

We used 7 polynomials in the s_t dimension and 6 in the x_{t-1} dimension, so γ^i is 42×1 and $\Phi(s_t, x_{t-1})$ is 1×42.

We use a simulation-based approach to ensure accuracy in the region of the state-space most often visited in equilibrium and pin down the coefficients γ^i by minimizing the errors in the equilibrium conditions at all points that the economy visits in response to a history of 10,000 monetary shocks.[22] We use a time-iteration algorithm. For a given guess of the coefficients γ^i we calculate all the equilibrium variables for every date using a 5-node Gaussian quadrature to compute expectations and update the coefficients using least-squares projection.

Table B.1 reports several statistics that describe the accuracy of the solution method. In Panel A we report the mean and maximum absolute error in the equilibrium conditions, expressed relative to the value of each respective variable. The first column shows that the projection method produces small errors in the equilibrium conditions. The second column gauges the accuracy of a third-order perturbation: the average error produced by this alternative approach is relatively small, 0.11% of the value of the respective variable, but occasionally the perturbation approach produces large errors, with a maximum value of 35%. Importantly, as Panel B of the table illustrates, both approaches imply similar values for the moments we targeted in calibration, suggesting that the occasionally large errors implied by the perturbation approach do not significantly affect the model's key predictions.

Figure B.1 plots a time-series of inflation and the frequency of price changes produced by the projection and perturbation-based solution methods for the same history of monetary policy shocks. This figure includes the period in which the perturbation method produces the largest error of 35%, which occurs during the disinflation episode in period 37. The inflation

[22] Because ours is a relatively simple problem which converges fast, we do not use the clustering approach suggested by Maliar and Maliar (2015).

Table B.1: Accuracy of Solution

A. Errors in Equilibrium Conditions

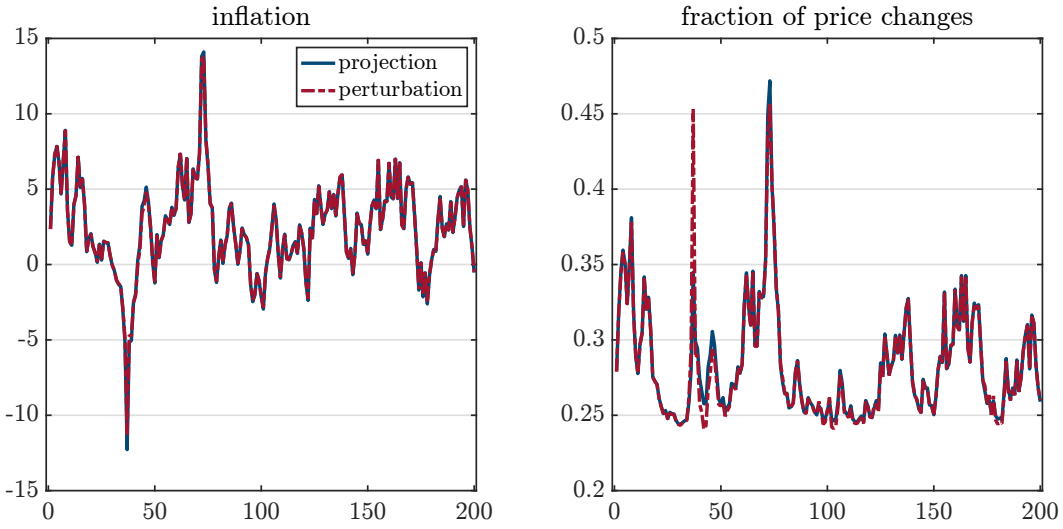
	projection	perturbation
mean abs. error	3.7×10^{-6}	1.1×10^{-3}
max abs. error	2.7×10^{-3}	3.5×10^{-1}

B. Targeted Moments

	projection	perturbation
mean inflation	3.517	3.517
s.d. inflation	2.739	2.727
mean frequency	0.297	0.296
slope of n_t on $\lvert\pi_t\rvert$	0.016	0.016

rate implied by the projection method is equal to -12.3% at this date, while that implied by the perturbation method is equal to -11.3%. As the right panel of the figure shows, the perturbation method overstates the fraction of firms that change prices on this date: 0.45 vs. 0.34. With the exception of this episode, the two approaches produce very similar values for inflation and the frequency of price changes, suggesting that the perturbation method provides a fairly accurate approximation.

Figure B.1: Simulated Time-Series Paths



C Additional Figures

Figure C.2 shows the implications of endogenizing the fraction of price changes for output and productivity. The steady-state values of output and productivity can be expressed solely as a function of trend inflation and the fraction of price changes. Specifically, output is determined by

$$y^{\frac{1}{\eta}} = \eta \frac{1-\beta(1-n)\pi^{\frac{\theta}{\eta}}}{1-\beta(1-n)\pi^{\theta-1}} \left(\frac{n}{1-(1-n)\pi^{\theta-1}} \right)^{\frac{1+\theta\left(\frac{1}{\eta}-1\right)}{\theta-1}},$$

and productivity x^θ is given by

$$x^\theta = \left(\frac{1-(1-n)\pi^{\frac{\theta}{\eta}}}{n} \right)^{\eta} \left(\frac{1-(1-n)\pi^{\theta-1}}{n} \right)^{-\frac{\theta}{\theta-1}}.$$

Absent trend inflation, output is equal to $y = \eta^\eta$ and productivity is equal to $x^\theta = 1$ and are therefore both equal to their flexible-price counterparts. The middle and right panels of Figure C.2 plot output and productivity as a function of trend inflation. For reference, we also plot steady-state outcomes in an otherwise identical economy in which the fraction of price changes is constant and equal to $\bar{n}$. The figure shows that both output and productivity generally decrease with trend inflation[23], but much less in our model with an endogenous fraction of price changes, a result reminiscent of Devereux and Yetman (2002), Bakhshi et al. (2007) and the menu cost model of Blanco (2021).

Figure C.3 contrasts the raw quarterly annualized inflation series with its counterpart smoothed using a centered 3-quarter moving average. We target the smoothed series in Section 5. Figure C.4 reports the slope of the Phillips curve we obtain when targeting the raw inflation data. Notice that the slope spikes to a level above 0.2 for one quarter in 1980, but is otherwise comparable to the magnitudes we report in the main text.

D Robustness

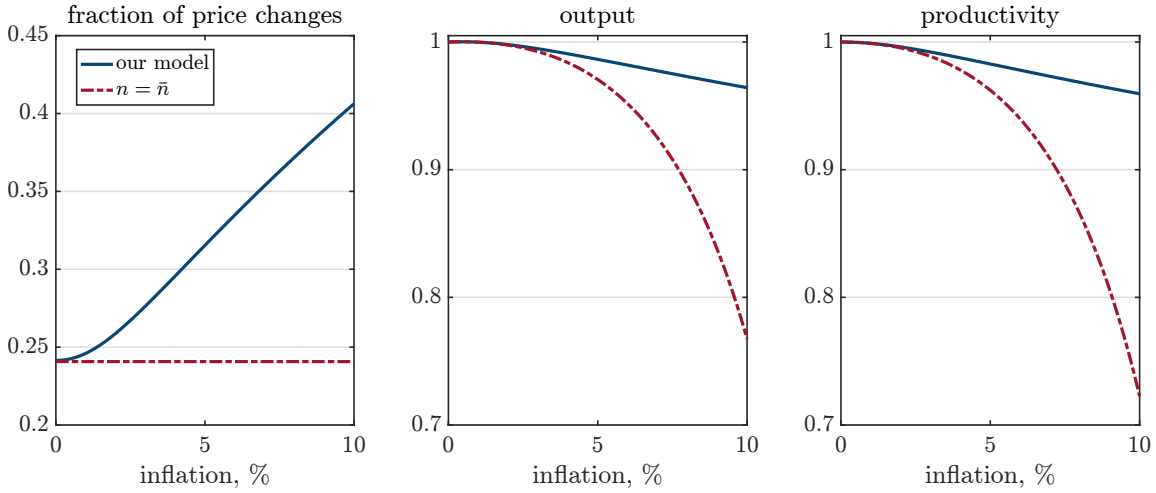
D.1 A Model with Idiosyncratic Shocks

In this section, we derive expressions that characterize moments of the distribution of price changes. Let $\hat{p}_{ikt} := \log\left(\frac{P_t}{z_{ikt}P_{ikt}}\right)$ be the inverse of the log-price gap and $M_t(q) = \int_0^1 \int_0^1 \hat{p}_{ikt}^q \mathrm{d}k\mathrm{d}i$ be the q-th moment of the distribution in period t. We can then recursively express $M_t(1)$ and $M_t(2)$ as

$$\begin{aligned} M_t(1) &= -n_t \log(p_t^*/p_t) + (1-n_t)[\log(\pi_t) + M_{t-1}(1)], \\ M_t(2) &= n_t \log(p_t^*/p_t)^2 + (1-n_t)\left[\log(\pi_t)^2 + \sigma_z^2 + 2\log(\pi_t)M_{t-1}(1) + M_{t-1}(2)\right]. \end{aligned}$$

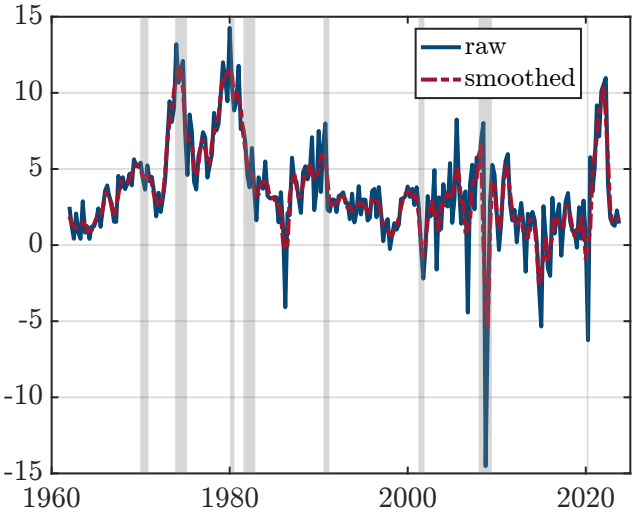
[23] As Ascari and Ropele (2009) point out, in the Calvo model the relationship between output and trend inflation is hump-shaped at low rates of inflation. This is also the case in our model: output peaks at a level of approximately 0.02% above its flexible-price (zero trend inflation) level when inflation is equal to 0.5%, but this effect is too small to be visible in the figure.

Figure C.2: Steady State Outcomes



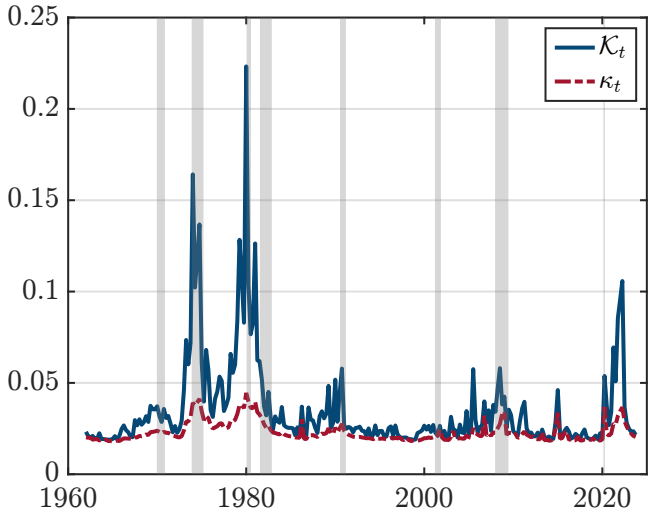
Notes: The figure traces out how variables adjust in steady state in response to changes in the parameter μ that determines steady-state inflation. The x-axis reports the annualized value of steady-state inflation π. We report the quarterly fraction of price changes. The level of output is normalized to 1 at zero steady state inflation.

Figure C.3: Inflation Data



Notes: The gray bars indicate NBER recessions.

Figure C.4: Slope of the Phillips Curve, Target Raw Inflation Data



Notes: The gray bars indicate NBER recessions.

Let $\mathbb{E}_t[\Delta p^m]$ denote the m-th moment of the log of price changes in period t. By definition, and under the assumption of random adjustment, it satisfies

$$\mathbb{E}_t[\Delta p^m] = \int_0^1 \int_0^1 \left[\log\left(\frac{P^*_{ikt}}{P_{ikt-1}}\right)\right]^m \mathrm{d}k\mathrm{d}i,$$

where P^*_{ikt} is the optimal nominal reset price and P_{ikt-1} is the nominal price in the previous period. It is straightforward to show that

$$\begin{aligned}
\mathbb{E}_t[\Delta p] &= \log(p^*_t/p_t) + \log(\pi_t) + M_{t-1}(1), \\
\mathbb{E}_t[\Delta p^2] &= \log(p^*_t/p_t\pi_t)^2 + \sigma_z^2 + 2\log(p^*_t/p_t\pi_t)M_{t-1}(1) + M_{t-1}(2), \\
\mathbb{S}td_t[\Delta p] &= \sqrt{\mathbb{E}_t[\Delta p^2] - \mathbb{E}_t[\Delta p]^2}.
\end{aligned}$$

To provide intuition for the micro-price statistics, notice that in the steady state, $M(1) = -\log(p^*/p) + \frac{1-n}{n}\log(\pi)$. Since $\mathbb{E}[\Delta p] = \log(p^*/p) + \log(\pi) + M(1)$, we obtain

$$\begin{aligned}
\mathbb{E}[\Delta p] &= \log(p^*/p) + \log(\pi) + M(1) \\
&= \log(p^*/p) + \log(\pi) - \log(p^*/p) + \frac{1-n}{n}\log(\pi) \\
&= \frac{\log(\pi)}{n}.
\end{aligned}$$

Thus, inflation is equal to the frequency of price changes times the size of price changes. Similarly, assuming zero inflation in the steady state ($\log(\pi) = 0$), we find that $M(2) =$

$-\log(p^*/p)^2 + \frac{1-n}{n}\sigma_z^2$. Since $\mathbb{E}[\Delta p^2] = -\log(p^*/p)^2 + \sigma_z^2 + M_{t-1}(2)$, we obtain

$$\begin{aligned}\mathbb{E}[\Delta p^2] &= \log(p^*/p)^2 + \sigma_z^2 + 2\log(p^*/p)^2 - \log(p^*/p)^2 + \frac{1-n}{n}\sigma_z^2 \\ &= \frac{\sigma_z^2}{n} + 2\log(p^*/p)^2 \\ &\approx \frac{\sigma_z^2}{n}.\end{aligned}$$

Thus, the variance of idiosyncratic shocks is equal to the frequency of price changes times the variance of price changes, plus a small correction term related to consumer substitution effects.

D.2 A Model with Multiple Aggregate Shocks

In this section, we describe how we estimate the model with multiple aggregate shocks in Section 6.4. We assume that each of the three aggregate shocks, z_t, τ_t and u_t, follow an AR(1) process with persistence ρ_z, ρ_τ and ρ_u, and standard deviation of innovations σ_z, σ_τ and σ_u

$$\begin{aligned}\log z_t &= (1-\rho_z)\log z + \rho_z \log z_{t-1} + \sigma_z \epsilon_{zt} \\ \log \tau_t &= (1-\rho_\tau)\log \tau + \rho_\tau \log \tau_{t-1} + \sigma_\tau \epsilon_{\tau t} \\ \log u_t &= (1-\rho_u)\log u + \rho_u \log u_{t-1} + \sigma_u \epsilon_{ut}\end{aligned}$$

where $z = u = 1$, and $\tau = 1 - 1/\theta$.

To parameterize this model, we first set a number of parameters to values typical in the literature. We then conduct a Bayesian likelihood estimation to estimate the variance and standard deviations of the shock processes, as well as the parameter ξ determining the cost of changing prices and the fraction of free price changes $\bar{n}$.

D.2.1 Data

We use quarterly data on inflation, output and the nominal federal funds rate from 1962:Q1 to 2023:Q4. As in the baseline, we use the CPI price index excluding shelter for inflation. We use real gross domestic product for output growth, and the fed funds rate for the nominal interest rate. For periods when the zero lower bound binds, we substitute the fed funds rate with the shadow interest rate calculated by Wu and Xia (2016). Since there is no trend growth in output in our model, we de-trend output in the data using a linear trend. We also use data on the frequency of price changes from Montag and Villar (2023) and Nakamura et al. (2018) that we discussed in the main text.

D.2.2 Solution

Let y_t be a vector of observables, s_t a vector of hidden state variables and ϑ a vector of parameters. The state and measurement equations are

$$s_t = \Phi(s_{t-1}, \epsilon_t, \vartheta) \tag{40}$$

$$y_t = H_t s_t + \nu_t \tag{41}$$

where s_t is the vector of state variables, y_t is the vector of observed variables, ϑ is the vector of parameters, and ϵ_t and ν_t are shocks, with ϵ_t the state innovation with density $N(0, \Omega)$, and where ν_t is a measurement innovation with density $N(0, R)$.

We use the third-order approximation for equation (40) that we solve using Dynare. Recall that in Appendix B we showed the third-order approximation is accurate when compared to the global solution. Because we solve the model with a third-order approximation, we use a particle filter to compute the likelihood of the model. Finally, because the fraction of price changes is not available for all time periods of our sample, we conduct a mixed-frequency estimation, in which the set of observables in our measurement equation (41) varies over time, reflected in the time-varying matrix H_t.

We discuss the steps of the estimation below. Our goal is to characterize the posterior distribution of ϑ. The prior is $p(\vartheta)$. By Bayes rule, the posterior is proportional to the likelihood and the prior

$$p(\vartheta \mid y^T) \propto p(y^T \mid \vartheta) p(\vartheta)$$

where $p(y^T \mid \vartheta)$ is the likelihood. By the chain rule, we can write

$$\begin{aligned} p(y^T \mid \vartheta) &= \prod_{s=0}^{T} p(y_s \mid y_{s-1}, \ldots, y_0, \vartheta) \\ &= p(y_0, \vartheta) \prod_{s=1}^{T} p(y_s \mid y^{s-1}, \vartheta) \end{aligned} \tag{42}$$

where $p(y_s \mid y^{s-1}, \vartheta)$ is the one-step ahead predictive density, and where the last line follows from our assumption of i.i.d. normal errors.

Equation (42) is the prediction error decomposition. We want the terms in the prediction error decomposition. The filter we use to compute these terms depends on the functional forms of Φ and Ψ in the state-space representation. Because Φ and Ψ are non-linear functions, we cannot use the Kalman filter, so we use a particle filter instead.

Particle Filter. The particle filter simulates many paths (particles) of the state of the system through the nonlinear state-space (40), and weights those particles according to how well they fit the data. This weight is the likelihood of the observed measurement given the state represented by the particle. Each period, we resample the particles to generate a new set of particles; particles with higher likelihood weights are more likely to be chosen. This focuses the particle set on regions of higher likelihood.

To compute the likelihood of the model we want the elements of the one-step ahead predictive density

$$p(y_t \mid y^{t-1}, \vartheta) = \int p(y_t \mid s_t, \vartheta) p(s_t \mid y^{t-1}, \vartheta) \, \mathrm{d}s_t.$$

We know the distribution of the measuerment error $r(\varepsilon)$ and $H(\bullet)$, so we can evaluate

$$p(y_t \mid s_t, \vartheta) = |J_H| \bullet r\left[H^{-1}(y_t, s_t, \vartheta)\right],$$

where J_H is the Jacobian.

We need to evaluate the posterior predictive density

$$p(s_t \mid y^{t-1}, \vartheta) = \int p(s_t \mid s_{t-1}, \vartheta) p(s_{t-1} \mid y^{t-1}, \vartheta) \, \mathrm{d}s_{t-1}.$$

Then we update using the data y_t to get $p(s_t \mid y^t, \vartheta)$

$$p(s_t \mid y^t, \vartheta) = \frac{p(y_t \mid s_t, \vartheta) p(s_t \mid y^{t-1}, \vartheta)}{p(y_t \mid y^{t-1}, \vartheta)}.$$

This is computationally challenging, because we map probability densities into probabilities densities, so we use the Monte Carlo bootstrap filter of Gordon et al. (1993).

Monte Carlo Bootstrap Particle Filter Algorithm. The algorithm recursively produces discrete approximations to the distribution of states s_t conditional on time $t-1$ information (forecasting distribution) as well as time t information (updated distribution). The algorithm takes existing particles s_{t-1}^j, simulating forward using the state equation to obtain particles s_t^j, and then reweights those particles based on the new data. We can compute the likelihood as a byproduct of computing those weights.

The algorithm is as follows.

1. At $t = 0$ initialize particles by sampling from the prior $p(x_0)$.
2. For each period $t = 1, \ldots, T$:
   (a) Map sample from $p(s_{t-1} \mid y^{t-1}, \vartheta)$ into sample from $p(s_t \mid y^{t-1}, \vartheta)$.

$$p(s_t \mid y^{t-1}) = \int p(s_t \mid s_{t-1}, \vartheta) p(s_{t-1} \mid y^{t-1}, \vartheta) \, \mathrm{d}s_{t-1}$$

   Simulate $p(s_t \mid y^{t-1}, \vartheta)$ by sampling in parts:
   i. Draw $x_{i,t-1}$ from $p(s_{t-1} \mid y^{t-1}, \vartheta)$.
   ii. Draw x_{it} from $p(s_t \mid x_{i,t-1}, \vartheta)$. In this step, we iterate on the state equation.
   iii. Marginalize with respect to s_{t-1} by discarding $x_{i,t-1}$.

   (b) For each draw i, evaluate $p(y_t \mid x_{it}, \vartheta)$ using the measurement equation:

$$p(y_t \mid s_t, \vartheta) = |J_H| \bullet r\left[H^{-1}(y_t, s_t, \vartheta)\right],$$

   where J_H is the Jacobian.

   (c) Use Monte-Carlo integration to evaluate $p(y_t \mid y^{t-1}, \vartheta)$:

$$p(y_t \mid y^{t-1}, \vartheta) \approx \frac{1}{N} \sum_i p(y_t \mid x_{it}, \vartheta).$$

   (d) *Resample.* Map sample from $p(s_t \mid y^{t-1}, \vartheta)$ into sample from $p(s_t \mid y^t, \vartheta)$.

i. With draws $p(y_t \mid x_{it}, \vartheta)$ from $p(s_t \mid y^{t-1}, \vartheta)$, normalize so that probability weights sum to 1.

ii. Resample the particles x_{it} using the new probability weights, using multinomial resampling.

iii. This gives an evenly weighted sample from $p(s_t \mid y^t, \vartheta)$.

3. The approximated log-likelihood is given by

$$\log p(y^T \mid \vartheta) = \sum_{t=1}^{T} \log p(y \mid y^{t-1}, \vartheta).$$

Degeneracy. In practice, an important consideration is that the particles do not degenerate, in that the distribution of likelihood weights across the particles is highly uneven. This occurs when the forecast distribution may be highly mismatched with the updated distribution. A consequence of this is that the likelihood function can become imprecisely estimated. In our application in this paper, to handle cases like these, we assign zero likelihood weight to period in which one particle forms at least 50%. In pratice, this occurs in quarters in which there are large swings in the data: 2009:Q1 and 2020:Q2.

Parallelization. The particle filter can be computationally intensive, since we may need to work with a very large set of particles to obtain a low-variance estimate of the likelihood (an estimate of the likelihood that does not vary much with different draws of particles). In Matlab, we use single program multiple data (spmd) blocks over the particles within the computation of the likelihood, i.e. in step 2 of the algorithm described above. This yields substantial speed gains.

Metropolis Hastings Algorithm. The final step is to summarize the procedure that we use to generate the Markov Chain that yields the posterior distribution of ϑ. It has a single block, corresponding to parameters ϑ. The algorithm is as follows.

Let N be the length of the chain. For $j = 1, \ldots, N$,

1. Draw a new proposal of ϑ, denoted by $\hat{\vartheta}_j$, from a thick-tailed proposal density centered at $\hat{\vartheta}_{j-1}$ to ensure sufficient coverage of the parameter space. In practice, we draw random values from the Student t-distribution, and scale the draws by covariance matrix of the Hessian implied by the posterior maximization of the linearized version of the model.

2. The proposal $\hat{\vartheta}_j$ is accepted with probability $\frac{p(\hat{\vartheta}_j|y^T)}{p(\hat{\vartheta}_{j-1}|y^T)}$. If $\hat{\vartheta}_j$ is accepted, then set $\hat{\vartheta}_{j-1} = \hat{\vartheta}_j$. In practice, it is important to compute both $p(\hat{\vartheta}_j \mid y^T)$ and $p(\hat{\vartheta}_{j-1} \mid y^T)$ within in each step j, due to the variance of the estimate of the likelihood using the particle filter.

We use four chains of length 15,000 and discard the first quarter of the chain as a burn-in. We choose a scaling parameter for the proposal density for the Metropolis Hastings algorithm; that yields an acceptance probability of about 35% on average across the two chains.

D.3 Parameter Estimates

We assign the same values as in the baseline to the discount factor, demand elasticity, returns to scale. We parameterize the Taylor rule as in Section 6.3. We set γ, the inverse of the labor supply elasticity, equal to 2.

Table D.2 reports the parameters estimates from the Bayesian estimation. We choose standard priors for the model's parameters. Importantly, we use relatively diffuse priors for the key parameter governing the cost of adjusting prices, ξ, as well as the fraction of free price changes, $\bar{n}$. We find that, at the median of the posterior distribution, the value of ξ and $\bar{n}$ are similar to those estimated in the baseline model.

Table D.2: Parameter Estimates, Model with Multiple Aggregate Shocks

	Prior			Posterior			
	Type	Mean	SD	Mode	Median	5%	95%
σ_z	IG	1.00	2.0	1.16	1.26	0.85	1.82
σ_τ	IG	1.00	2.0	7.23	8.84	5.69	14.29
σ_u	IG	0.25	5.0	0.34	0.35	0.24	0.48
ρ_z	B	0.50	0.1	0.99	0.99	0.96	1.00
ρ_τ	B	0.50	0.1	0.74	0.70	0.53	0.83
ρ_u	B	0.50	0.1	0.38	0.41	0.16	0.64
ξ	N	0.75	0.5	1.37	1.70	1.06	2.82
$\bar{n}$	N	0.23	0.1	0.23	0.23	0.21	0.25

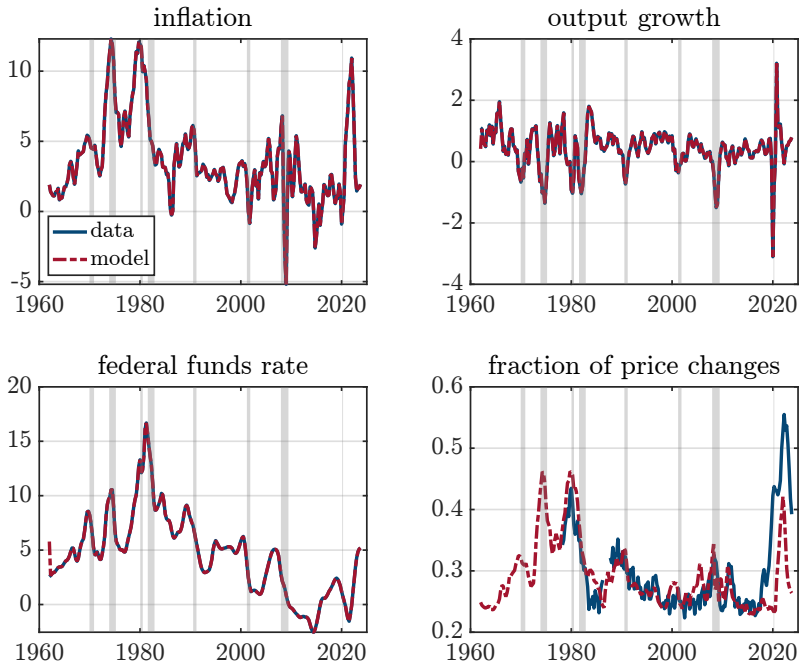
Notes: IG = Inverse Gamma, B = Beta, N = Normal distribution.

D.4 Slope of the Phillips Curve

As in the baseline, to gauge how the Phillips curve evolves over time, we initialize the economy in the stochastic steady state in 1962 and use the non-linear solution of the model to back out the sequences of productivity, cost-push and monetary shocks that reproduce the observed time series for inflation, output growth and the nominal interest rate. Figure D.5 shows these time series. Importantly, as the bottom-right panel of the figure shows, the model reproduces relatively well the time series for the fraction of price changes.

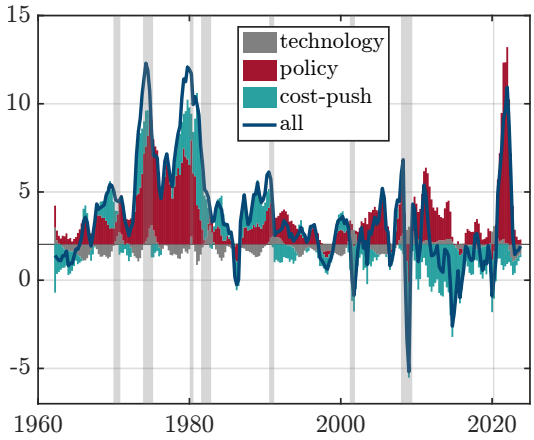
Figure D.6 illustrates the role of each aggregate shock for the path of inflation.

Figure D.5: Inflation, Output Growth, Interest Rate and the Fraction of Price Changes



Notes: The gray bars indicate NBER recessions.

Figure D.6: Decomposition of Inflation



Notes: The gray bars indicate NBER recessions.