

Strike while the Iron is Hot: Optimal Monetary Policy with a Nonlinear Phillips Curve*

Peter Karadi^{1,2}, Anton Nakov^{1,2}, Galo Nuño^{3,2}, Ernesto Pastén⁴, and Dominik Thaler¹

¹European Central Bank

²CEPR

³Bank for International Settlements and Bank of Spain

⁴Central Bank of Chile

July 22, 2024

Abstract

We study the Ramsey optimal monetary policy within the [Goloso and Lucas \(2007\)](#) state-dependent pricing framework. The model provides microfoundations for a non-linear Phillips curve: the sensitivity of inflation to activity increases after large shocks due to an endogenous rise in the frequency of price changes, as observed during the recent inflation surge. In response to large cost-push shocks, optimal policy leverages the lower sacrifice ratio to reduce inflation and stabilize the frequency of price adjustments. At the same time, when facing efficient shocks, such as total factor productivity shocks, the optimal policy commits to strict price stability, similar to the approach in the standard [Calvo \(1983\)](#) model.

*We are grateful to Vladimir Asryan, Andres Blanco, Davide Debortoli, Eduardo Engel, Aurélien Eyquem, Jordi Galí, Mark Gertler, Mishel Ghassibe, Francesco Lippi, Albert Marcet, Virgiliu Midrigan, Giorgio Primiceri, Xavier Ragot, Morten Ravn, Tom Sargent, Edouard Schaal, and Jaume Ventura, as well as to participants at various conferences and seminars for their comments and suggestions. This manuscript was previously circulated as ‘Strike the Iron while it’s Hot: Optimal Monetary Policy with (S,s) Pricing’. The views expressed here are those of the authors only, and do not necessarily represent those of the BIS, Bank of Spain, the Central Bank of Chile, ECB, or the Eurosystem.

1 Introduction

The recent inflation surge has been accompanied by a significant increase in the frequency of price changes.¹ Concurrently, empirical evidence reveals marked time variations in the slope of the estimated Phillips curve, which characterizes the relationship between economic activity and inflation (Benigno and Eggertsson, 2023; Cerrato and Gitti, 2023). Traditional models of price setting (Calvo, 1983), which form the basis for optimal monetary policy analysis (Woodford, 2003; Galí, 2008), cannot explain these observations. In contrast, state-dependent pricing models are well suited to capturing both of them: firms’ price-adjustment decisions lead to endogenous variation in the repricing frequency and, thus, in the slope of the Phillips curve. Among the state-dependent models, the Golosov and Lucas (2007) menu-cost model has emerged as a benchmark for positive analysis. However, normative aspects of the model have received scant attention. The question of how the central bank should respond to large inflation surges remains unanswered. It is precisely this crucial gap that our paper aims to bridge.

Our analysis arrives at a novel insight: in response to large cost-push shocks, Ramsey optimal monetary policy in the Golosov and Lucas (2007) model should commit to quashing inflation and leaning against changes in the repricing frequency – a “strike while the iron is hot” policy. When shocks are small, policy prescriptions are similar to those in the Calvo model. However, for large shocks, a bolder anti-inflation policy is optimal. The reason is that the cost of such a policy in terms of output is smaller when the frequency of price changes increases in response to the shocks – as the Phillips curve becomes steeper, the sacrifice ratio falls.

Our state-dependent price-setting model closely follows Golosov and Lucas (2007). A representative household consumes a continuum of differentiated goods and provides labor in a centralized, frictionless market. Each good is produced by a single firm with only labor subject to aggregate productivity shocks, aggregate cost-push shocks, and idiosyncratic quality shocks.² Firms must incur a small, fixed, “menu cost” to adjust their prices. Thus,

¹See Montag and Villar (2023), Cavallo et al. (2023) and Blanco et al. (2024a)

²We depart from Golosov and Lucas (2007) model in this regard, which instead assumes idiosyncratic productivity shocks. We do so for ease of computation while its implications on our results are innocuous (see also Midrigan, 2011; Alvarez et al., 2021).

firms' pricing decisions are characterized by an (S,s) rule: when prices are within an endogenous band around the optimal reset price, firms keep them constant; otherwise, they pay the menu cost and update their price. The central bank sets the nominal interest rate. The model is calibrated to match the frequency and magnitude of price changes in the U.S. (Nakamura and Steinsson, 2008). We contrast the implications of our model to those of a Calvo model recalibrated to generate the same Phillips curve slope for small shocks (Auclert et al., 2024). This recalibration takes into account a special feature of the baseline model: the endogenous "selection" of large price changes, which raises the flexibility of the price level.

We start by exploring the implications of our menu-cost model assuming the central bank follows a Taylor rule. We show that the similarity of our baseline menu cost model and the (suitably recalibrated) Calvo model for small shocks does not generalize in the presence of large shocks (Blanco et al., 2024a; Cavallo et al., 2023). The reason is the repricing frequency: it increases endogenously in our baseline model as shocks become large, while it stays constant in the Calvo model. The repricing frequency raises price flexibility and generates a nonlinear relationship between inflation and the output gap: a nonlinear Phillips curve. When shocks are small, the repricing frequency remains close to its steady-state value and the slope of the Phillips curve stays equal to that of the Calvo model. As shocks become larger, the frequency increases, prices become more flexible and the slope of the Phillips curve steepens.

Next, we move to the core of the paper: the optimal design of monetary policy in a menu-cost model with a nonlinear Phillips curve. To this end, we solve the fully nonlinear Ramsey problem under commitment. We first analyze the Ramsey steady state. The model features a slightly positive steady-state inflation rate, at around 0.3%. This contrasts with the Calvo model, where optimal inflation is exactly zero. In our menu-cost model, positive inflation reduces the frequency and thus helps firms to economize on costly price adjustments. In particular, it counterbalances the impact of too frequent price increases relative to price decreases, a consequence of the asymmetry of the profit function: firms dislike more negative price misalignments when the demand for their product is high, relative to positive misalignments when the demand is low.

The well-known time-inconsistency problem of monetary policy is also present in our menu-cost model, although it is muted relative to Calvo. In both models, when the steady state is inefficient, monetary policy has the incentive to stimulate output via an unexpectedly easy policy (Galí, 2008). However, in the menu-cost model such a policy is less effective on output and more inflationary because the ensuing increase in the repricing rate raises the flexibility of the aggregate price level. The time-inconsistent motive to ease is thus considerably weaker.

We turn next to the optimal systematic response to shocks under a “timeless perspective” (Woodford, 2003). First, we show that the optimal response to efficient (technology, preference) shocks is characterized by the “divine coincidence” (Blanchard and Galí, 2007). In other words, optimal policy stabilizes both inflation and the output gap relative to the efficient level.

The optimal response to cost-push shocks in the Calvo model is to “lean against the wind”: the central bank temporarily drives output below its efficient level to contain the inflationary impact of a positive cost-push shock. The relationship between the inflation level and the change in the output gap is characterized by a near-linear target rule. In the menu cost model, the central bank also “leans against the wind” and, for small shocks, follows a linear target rule with a similar slope. The reason, however, is different. In the Calvo model, the central bank trades off the distortion associated with *average markup* volatility, linked to the variance in the output gap, to that due to higher *price dispersion*, proportional to the variance of inflation. In the menu-cost model, price dispersion actually can decrease with inflation as new adjusters are endogenously selected from those with the most misaligned prices. There is however a new distortion, namely the losses due to menu costs, which increase with inflation and tend to dominate the welfare effect of inflation on price dispersion.

While the balance between these distortions in the Calvo and our menu-cost model is similar for small shocks, the differences become relevant as the shock size increases. For large shocks, we find that monetary policy should be tighter in the menu cost model: the optimal prescription is to react more aggressively against inflation than in the case of small shocks or in a counterfactual fixed-frequency Calvo setting. The outcome is a nonlinear

target rule that significantly dampens the inflation surge for a unit decline in output gap as the shocks become large. Finally, we show that the source of the nonlinearity in the target rule is caused almost exclusively by the nonlinearity in the Phillips curve. We show that the planner’s preferences in our baseline model can be quite well approximated by the planner’s preferences in the Calvo model, which are near quadratic over the output gap and inflation. The Phillips curve, however, is different in the two frameworks: it is nonlinear in our baseline, while near-linear in Calvo. When we counterfactually insert the quadratic preferences of the Calvo model to our baseline, the target rule still remains nonlinear, confirming that it is the nonlinearity of the Phillips curve and, thus, the nonlinearity of the sacrifice ratio that drives the nonlinearity of the target rule: As inflation diverges more and more from its steady state value, the output cost of containing inflation diminishes, so the planner contains it more strongly.

We assess the optimal monetary policy response to the 2022-2023 inflation surge through the lens of our model. We construct a scenario that captures key features of the inflation surge in the US. We argue that a combination of both aggregate and relative-price shocks are needed to explain the evidence documented in the micro price data [Montag and Villar \(2023\)](#), which found a large increases both in the frequency and the dispersion of price changes. Relative-price shocks are especially relevant, as they generate a reason for an ‘efficient’ increase in frequency. The scenario generates realistic inflation dynamics when the monetary policy follows an inertial Taylor rule: inflation surges to around 9 percent temporarily and stays persistently above central bank’s 2 percent inflation target for a considerable amount of time. In the counterfactual optimal policy scenario, we find that monetary policy tightens aggressively and, at some output cost, keeps the inflation surge temporary with a peak below 4 percent.

Overall, our findings highlight the significance of an aggressive anti-inflationary stance by the central bank in the face of large shocks. By committing to policies that lean against inflation and stabilize the repricing frequency, the central bank can foster a more favorable macroeconomic outcome.

Related literature. Our paper builds on the seminal article by [Goloso and Lucas \(2007\)](#). These authors propose a menu cost model ([Barro, 1972](#); [Sheshinski and Weiss, 1977](#); Ca-

ballero and Engel, 1993) that has become the backbone of a positive literature studying the relationship between monetary non-neutrality and the distribution of price changes at the micro level (Midrigan, 2011; Costain and Nakov, 2011; Alvarez et al., 2016), as well as the impact of large aggregate shocks on inflation and activity (Karadi and Reiff, 2019; Alexandrov, 2020; Auer et al., 2021). The model describes well firms’ price setting behavior in diverse environments with both low and high inflation (Nakamura and Steinsson, 2008; Gagnon, 2009; Alvarez et al., 2019). This price-setting framework provides a microfounded state-dependent alternative to the canonical time-dependent Calvo (1983) model, with widely different implications in terms of both the extent of monetary non-neutrality and price-flexibility as a response to large shocks. Indeed, most familiar price-setting models, such as the random-menu-cost model of Dotsey et al. (1999); Alvarez et al. (2021), the Calvo-plus model of Nakamura and Steinsson (2010), the rational inattention framework by Woodford (2009), or the control cost framework by Costain and Nakov (2019), lie on a spectrum bracketed by these two polar cases. Normative prescriptions from the Golosov and Lucas (2007) model, therefore, can provide qualitative insights that generalize to a wide class of price-setting frameworks.

To the best of our knowledge, our paper is the first to solve for optimal monetary policy in the canonical Golosov and Lucas (2007) menu cost model. The main distinctive feature of the model, relative to the standard framework for monetary policy analysis based on Calvo (1983), for instance, as in Woodford (2003) and Galí (2008), is a state-dependent relationship between inflation and the output gap – a “non-linear Phillips curve” – which has received new empirical support following the recent inflation surge (Benigno and Eggertsson, 2023; Cerrato and Gitti, 2023; Blanco et al., 2024a). Our conclusion prescribing aggressive anti-inflationary policy after large shocks is a direct consequence of this non-linearity: a higher Phillips curve slope implies a favourable inflation-output trade-off that optimal policy should utilize. Crucially, the framework includes firms facing idiosyncratic shocks, which are essential to explain the large size of observed price changes (Golosov and Lucas, 2007), and raises relevant normative questions by providing an underlying cause for efficient relative price adjustments. Solving dynamic optimal policy in response to aggregate shocks in this environment complements previous research, which has restricted attention to settings without

firm-level shocks (Nakov and Thomas, 2014; Caratelli and Halperin, 2023)³, or to optimal *steady-state* inflation rate (Adam and Weber, 2019; Blanco, 2021; Nakov and Thomas, 2014).

This paper proposes a new algorithm to solve Ramsey optimal policy in heterogeneous-agent models, building on previous work in González et al. (2024). The algorithm (i) makes the infinite-dimensional planner’s problem finite-dimensional by approximating the infinite-dimensional value and distribution functions by piece-wise linear functions; (ii) accounts for the discrete price-adjustment choice using an endogenous grid; (iii) derives the FOCs of the planner’s problem by symbolic differentiation and (iv) solves the resulting set of equilibrium conditions non-linearly under perfect foresight over the sequence space. Our approach complements other methods to solve for Ramsey policy in heterogeneous agent models (Bhandari et al., 2021; Le Grand et al., 2022; Ragot, 2019; Nuño and Thomas, 2022; Smirnov, 2022).

2 Model

Our model is a variant of the Golosov and Lucas (2007) state-dependent pricing framework. The economy consists of a representative household, a unit mass of monopolistic producers facing fixed menu costs to update their prices, and a central bank that sets the nominal interest rate.

2.1 Households

There is a representative household that saves in one-period bonds whose nominal value is denoted by B_t . Bonds are in zero net supply. Workers supply labor hours N_t . The household maximizes

$$\max_{C_t, N_t, B_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(C_t, N_t), \quad (1)$$

subject to

$$P_t C_t + Q_t B_t + T_t = B_{t-1} + W_t N_t + D_t, \quad (2)$$

³Caratelli and Halperin (2023) focus on optimal policy in the face of *sector-specific* shocks and find that it can be characterized as *nominal wage targeting*.

where T_t are lump-sum taxes, W_t is the nominal wage, D_t are lump-sum dividends from firms, and Q_t is the price of the nominal bond. Aggregate consumption C_t is

$$C_t = \left\{ \int [A_t(j)C_t(j)]^{\frac{\epsilon-1}{\epsilon}} di \right\}^{\frac{\epsilon}{\epsilon-1}}, \quad (3)$$

where $C_t(j)$ is the quantity purchased of product $j \in [0, 1]$ and $A_t(j)$ is the quality of product j , following a random walk with stochastic volatility in logs:

$$\log A_t(j) = \log A_{t-1}(j) + \sigma_t \varepsilon_t(j),$$

and ε_t is an i.i.d Gaussian innovation. The demand for product j is,

$$C_t(j) = A_t(j)^{\epsilon-1} \left(\frac{P_t(j)}{P_t} \right)^{-\epsilon} C_t, \quad (4)$$

and the aggregate price index is

$$P_t = \left[\int_0^1 \left(\frac{P_t(j)}{A_t(j)} \right)^{1-\epsilon} di \right]^{\frac{1}{1-\epsilon}}. \quad (5)$$

We assume separable utility of the CRRA class, $u(C_t, N_t) = \frac{C_t^{1-\gamma}}{1-\gamma} - vN_t$. Solving for the FOCs, we obtain the labor supply condition,

$$w_t = vC_t^\gamma, \quad (6)$$

where $w_t = W_t/P_t$ is the real wage, and the Euler equation,

$$1 = \mathbb{E}_t [\Lambda_{t,t+1} e^{i_t - \pi_{t+1}}], \quad (7)$$

where $i_t \equiv \log(-Q_t)$ is the nominal interest rate,

$$\Lambda_{t,t+1} \equiv \beta \frac{u'(C_{t+1})}{u'(C_t)}. \quad (8)$$

2.2 Monopolistic producers

Production of good j is

$$Y_t(j) = A_t \frac{N_t(j)}{A_t(j)}, \quad (9)$$

where $N_t(j)$ is the labor input, and A_t is aggregate productivity.⁴

The nominal profit function is

$$\begin{aligned} D_t(j) &= P_t(j)Y_t(j) - (1 - \tau_t)W_tN_t(j) \\ &= P_t(j)^{1-\epsilon}A_t(j)^{\epsilon-1} \left(\frac{1}{P_t}\right)^{-\epsilon} C_t - (1 - \tau_t)\frac{W_t}{A_t}A_t(j)^\epsilon \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon} C_t \end{aligned} \quad (10)$$

where τ_t is an employment subsidy financed by lump-sum taxes. Notice that we have used the equilibrium condition $Y_t(j) = C_t(j)$. The real profit function thus is

$$\Pi_t(j) \equiv \frac{D_t(j)}{P_t} = C_t (\exp(p_t(j)))^{1-\epsilon} - C_t(1 - \tau_t)\frac{w_t}{A_t} (\exp(p_t(j)))^{-\epsilon} = \Pi(p_t(j), w_t, A_t), \quad (11)$$

where w_t is the real wage and

$$p_t(j) \equiv \log\left(\frac{P_t(j)}{A_t(j)P_t}\right)$$

is the quality-adjusted (log) *relative price*. When prices do not change in nominal terms, $p_t(j)$ evolves according to

$$p_t(j) = p_{t-1}(j) + \log\left(\frac{P_{t-1}(j)}{A_t(j)P_t}\right) - \log\left(\frac{P_{t-1}(j)}{A_{t-1}(j)P_{t-1}}\right) = p_{t-1}(j) - \sigma_t \varepsilon_t(j) - \pi_t.$$

From now on, we drop the index j for ease of notation. Without loss of generality, a firm resets its price with probability $\lambda_t(p)$. Price resetting involves the firm paying a fixed menu-cost η (in labor units). The optimal reset price p_t^* maximizes the firm's value, $p_t^* = \arg \max V_t(p)$, taking into account that this new price may not change for a random period of time. The firm's value is given by the equation

$$\begin{aligned} V_t(p) &= \Pi(p, w_t, A_t) + \mathbb{E}_t [(1 - \lambda_{t+1}(p - \sigma_{t+1}\varepsilon_{t+1} - \pi_{t+1})) \Lambda_{t,t+1} V_{t+1}(p - \sigma_{t+1}\varepsilon_{t+1} - \pi_{t+1})] \\ &\quad + \mathbb{E}_t \left[\lambda_{t+1}(p - \sigma_{t+1}\varepsilon_{t+1} - \pi_{t+1}) \Lambda_{t,t+1} \left(\max_{p'} V_{t+1}(p') - \eta w_{t+1} \right) \right]. \end{aligned}$$

which comprises the momentary profits $\Pi(\cdot)$ and the discounted continuation value $V_{t+1}(\cdot)$ evaluated when the price does not change at $t + 1$ with probability $1 - \lambda_{t+1}(\cdot)$, and when the firm sets a new price after paying the menu cost at $t + 1$, with probability $\lambda_{t+1}(\cdot)$. As in

⁴One advantage of introducing quality shocks as opposed to idiosyncratic productivity shocks is that the quality-adjusted relative price becomes the only state at the firm level, dramatically reducing the dimensionality of the optimal monetary policy problem.

Golosov and Lucas (2007) fixed menu-cost model, the adjustment probability is given by

$$\lambda_t(p) = \mathbb{1}[L_t(p) > 0]$$

where $\mathbb{1}[\cdot]$ is the indicator function, and

$$L_t(p) \equiv \max_{p'} V_t(p') - \eta w_t - V_t(p)$$

is the *gain from adjustment* (or loss from inaction), net of the menu cost.

2.3 Monetary policy rule

The central bank controls the short-term nominal interest rate i_t . In Section 4, we assume that the central bank follows a simple Taylor (1993) rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)(-\log \beta + \phi_\pi \pi_t + \phi_y (y_t - y_t^e)), \quad (12)$$

where ρ_i is the smoothing in the Taylor rule, y_t^e is the efficient-level of output, defined in Section 4, and $\phi_\pi > 1$ and ϕ_y are parameters. In Section 5, we assume instead that the central bank follows the optimal policy with commitment.

2.4 Aggregation

Firms' individual price-setting decisions give rise to a distribution of prices. Let the density of quality-adjusted log relative prices at the end of period t be $g_t(p)$. The definition of the aggregate price index can then be written as:

$$1 = \int e^{p(1-\epsilon)} g_t(p) dp, \quad (13)$$

Individual firms' labor demand aggregates up to

$$N_t = \frac{C_t}{A_t} \int e^{p(-\epsilon)} g_t(p) dp + \eta \int \lambda_t(p - \sigma_t \varepsilon - \pi_t) g_{t-1}(p) dp, \quad (14)$$

such that the total number of hours worked equals the total use of labor for production (the first term on the right-hand side) and the aggregation of labor allocated to price adjustment (the second term) – note that $\int \lambda_t(p - \sigma_t \varepsilon - \pi_t) g_{t-1}(p) dp$ is the *frequency* of price adjustments.

Consider next the law of motion of the price density function:

$$g_t(p) = (1 - \lambda_t(p)) \int g_{t-1}(p + \sigma_t \varepsilon + \pi_t) d\xi(\varepsilon) + \pi_t d\xi(\varepsilon) + \delta(p - p_t^*) \int \lambda_t(\tilde{p}) \left(\int g_{t-1}(\tilde{p} + \sigma_t \varepsilon + \pi_t) d\xi(\varepsilon) \right) d\tilde{p},$$

where $\delta(\cdot)$ is the Dirac delta function. This expression for $g_t(p)$ captures that, with probability $1 - \lambda_t$, the end of the period's density is just taken from the density of prices in period $t - 1$, after adjusting it for quality and inflation (where $\xi(\varepsilon)$ is the Gaussian distribution of the innovation ε). In turn, if the actual price equals the optimal reset price, $p = p_t^*$, the density $g_t(p)$ integrates the mass of all adjusting prices and, with probability λ_t , allocates it to the optimal reset price p_t^* .

2.5 Aggregate Shocks

The logarithm of aggregate productivity follows a first-order autoregressive process

$$\log A_t = \rho_A \log A_{t-1} + \varepsilon_{A,t},$$

where $\varepsilon_{A,t} \sim N(0, \sigma_A^2)$, and $\rho_A \in [0, 1]$ and σ_A are parameters. Likewise, we assume that the (lump-sum tax-financed) employment subsidy τ_t follows an autoregressive process which is interpretable as temporary *cost-push* shocks:

$$\tau_t - \tau = \rho_\tau (\tau_{t-1} - \tau) + \varepsilon_{\tau,t},$$

where τ is the steady-state employment subsidy, $\varepsilon_{\tau,t} \sim N(0, \sigma_\tau^2)$, and $\rho_\tau \in [0, 1]$ and σ_τ are parameters. We also consider autoregressive *dispersion* shocks to the idiosyncratic volatility of quality shocks:

$$\sigma_t - \sigma = \rho_\sigma (\sigma_{t-1} - \sigma) + \varepsilon_{\sigma,t},$$

where σ is the steady-state volatility, $\varepsilon_{\sigma,t} \sim N(0, \sigma_\sigma^2)$, and $\rho_\sigma \in [0, 1]$ and σ_σ are also parameters. We use these shocks to gauge the implications of shocks to price dispersion. Finally, in Section 4 we also assume i.i.d. shocks to the Taylor rule (12), $\varepsilon_{r,t} \sim N(0, \sigma_r^2)$, where σ_r is a parameter.

2.6 Normalization and Equilibrium

In order to achieve high accuracy in the computations, we find it convenient to recast the problem in terms of the *distance* x between actual (log-) prices p and the optimal (log-) reset price p^* . Therefore, we define a new state variable $x_t \equiv p_t - p_t^*$. The dynamics of $x_t(j)$ are given by

$$x_t \equiv p_t - p_t^* = x_{t-1} + p_t - p_{t-1} - p_t^* + p_{t-1}^* = x_{t-1} - \sigma_t \varepsilon_t - \pi_t^*,$$

where $\pi_t^* \equiv p_t^* - p_{t-1}^* + \pi_t$ is the inflation rate of the (quality-adjusted) reset price. The advantage of this reformulation is that after a price reset, x_t always jumps back to zero.

Profits can then be expressed as

$$\Pi(x_t, p_t^*, w_t, A_t) = C_t (\exp(x_t(j) + p_t^*))^{1-\epsilon} - C_t (1 - \tau_t) \frac{w_t}{A_t} (\exp(x_t(j) + p_t^*))^{-\epsilon},$$

and the Bellman equation can thus be re-written as

$$\begin{aligned} V_t(x) &= \Pi(x, p_t^*, w_t, A_t) \\ &+ \mathbb{E}_t \left[(1 - \lambda_{t+1} (x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*)) \Lambda_{t,t+1} V_{t+1}(x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*) \right] \\ &+ \mathbb{E}_t \left[\lambda_{t+1} (x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*) \Lambda_{t,t+1} (V_{t+1}(0) - \eta w_{t+1}) \right]. \end{aligned} \quad (15)$$

The optimality of price updating implies the conditions

$$V_t(0) - \eta w_t = V_t(s_t) \quad (16)$$

$$V_t(0) - \eta w_t = V_t(S_t), \quad (17)$$

The optimality of the reset price requires $V_t'(0) = 0$. $V_t'(0)$ can be expressed as the sum of the marginal effect of x on current profits and on the continuation value conditional on not-updating and updating the price (where $\phi(\cdot)$ is the standard normal pdf):⁵

$$\begin{aligned} 0 = V_t'(0) &= \Pi_t'(0) + \frac{1}{\sigma} \Lambda_{t+1} \int_{s_{t+1}}^{S_{t+1}} V_{t+1}(x') \frac{\partial \phi \left(\frac{x-x'-\pi_{t+1}^*}{\sigma} \right)}{\partial x} \Bigg|_{x=0} dx' \\ &+ \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (V_{t+1}(0) - \eta w_{t+1}), \end{aligned} \quad (18)$$

⁵These three conditions are sufficient only if the firm's value function is convex in x_t . We check convexity ex-post.

Appendix A derives this expression. The law of motion of the density furthermore is

$$g_t(x) = (1 - \lambda_t(x)) \int g_{t-1}(x + \sigma_t \varepsilon + \pi_t^*) d\xi(\varepsilon) + \pi_t^* d\xi(\varepsilon) \quad (19)$$

$$+ \delta(x) \int \lambda_t(\tilde{x}) \left(\int g_{t-1}(\tilde{x} + \sigma_t \varepsilon + \pi_t^*) d\xi(\varepsilon) \right) d\tilde{x},$$

Finally, the aggregate price index and labor market clearing can be expressed as

$$1 = \int e^{(x+p_t^*)(1-\epsilon)} g_t(x) dx, \quad (20)$$

$$N_t = \frac{C_t}{A_t} \int e^{(x+p^*)(-\epsilon)} g_t(x) dx + \eta \int \lambda_t(x + p^* - \sigma_t \varepsilon - \pi_t) g_{t-1}(x) dx. \quad (21)$$

Equations (6), (7), (15)-(21), together with a policy such as (12), define an equilibrium in $g(\cdot)$, $V_t(\cdot)$, C_t , N_t , w_t , i_t , p_t^* , s_t , S_t , π_t^* .

3 Calibration

Our baseline calibration is presented in Table 1. There are four blocks: household's preferences, firms' price setting behavior, the Taylor rule used in our positive analysis in Section 4, and shocks processes.

This calibration relies as much as possible on the existing literature. Regarding preferences, the monthly discount factor is $0.96^{1/12}$. The elasticity of substitution across products is $\epsilon = 7$, so the frictionless net markup is $1/6$. We follow [Midrigan \(2011\)](#) in assuming log utility in consumption, so the relative risk aversion coefficient is $\gamma = 1$; and the weight on leisure in utility is $\nu = 1$. These assumptions yield that the real wage equals aggregate consumption, $w_t = C_t$ (and nominal wages equal nominal aggregate demand).

Regarding price setting, we set the menu cost $\eta = 0.045281$, and the standard deviation of idiosyncratic shocks in absence of dispersion shocks to be $\sigma = 0.025250$. These choices allow to match in the Ramsey steady state the 8.7% of average frequency of price changes as well as the 8.5% of average absolute size of non-zero price changes documented by [Nakamura and Steinsson \(2008\)](#) for the U.S. For the Taylor rule governing monetary policy in Section 4, we assume $\phi_\pi = 1.5$ and $\phi_y = 0.125$ as in [Taylor \(1993\)](#). A common assumption is a smoothing component in the Taylor rule, in our case, $\rho_i = 0.75^{1/3}$, to accommodate inertia in monetary

Table 1: Parameter values

Households			
β	$0.96^{1/12}$	Discount rate	Golosov and Lucas (2007)
ϵ	7	Elasticity of substitution	Golosov and Lucas (2007)
γ	1	Risk aversion parameter	Midrigan (2011)
ν	1	Utility weight on labor	Set to yield $w = C$
Price setting			
η	0.045	Menu cost	Set to match 8.7% of frequency of price changes and 8.5% of absolute size of price changes documented in Nakamura and Steinsson (2008)
σ	0.025	Std dev of quality shocks	
Monetary policy			
ϕ_π	1.5	Inflation coefficient in Taylor rule	Taylor (1993)
ϕ_y	0.125	Output gap coefficient in Taylor rule	Taylor (1993)
ρ_i	$0.75^{1/3}$	Smoothing coefficient	
Shocks			
ρ_A	$0.95^{1/3}$	Persistence of the TFP shock	Smets and Wouters (2007)
ρ_τ	$0.9^{1/3}$	Persistence of the cost-push shock	Smets and Wouters (2007)
ρ_σ	$0.75^{1/3}$	Persistence of the dispersion shock	

policy rates. Finally, the persistence of shocks is taken from Smets and Wouters (2007), once transformed from quarterly to monthly frequency: $0.95^{1/3}$ for aggregate productivity shocks, $0.9^{1/3}$ for aggregate subsidy shocks (interpreted as cost-push shocks). The standard deviation of idiosyncratic quality shocks at the product level is assumed to be $0.75^{1/3}$.

4 The non-linear Phillips curve

This section presents novel considerations for monetary policy design arising from the menu cost model in the presence of large shocks. The key distinguishing feature of the framework is the endogenous evolution of the frequency of price changes, which leads to a non-linear Phillips curve and, consequently, a state-dependent variation in the inflation-output trade-off. This differs from the Calvo (1983) model, where the frequency is exogenously given, the Phillips curve is approximately linear and the inflation-output trade-off is approximately state-invariant.

In order to provide some intuition about the working of this model, here we derive a

3-equation ‘macro-block’ of the model. This helps us to contrast it to the well-understood 3-equation New Keynesian model.

New Keynesian IS curve and the Taylor rule. The first equation of the macro block is the familiar new Keynesian IS curve. It can be derived from the Euler equation (7):

$$\tilde{y}_t^e = -\frac{1}{\gamma} (i_t - \mathbb{E}_t[\pi_{t+1}] - r^e) + \mathbb{E}_t[\tilde{y}_{t+1}^e] = -\frac{1}{\gamma} \sum_{j=0}^{\infty} \mathbb{E}_t[r_{t+j} - r_{t+j}^e]. \quad (22)$$

where the efficient output gap ($\tilde{y}_t^e \equiv \log Y_t/Y_t^e$) measures the (log) deviation of output from its efficient level, their difference expressed as its deviation from the steady state output gap, and r_t^e is the efficient real interest rate. The efficient output level and interest rates are those prevailing in a counterfactual efficient allocation, which is the solution to the social planning problem (see Appendix B.1). The expression also shows that output gap is determined by the cumulative sum of future real interest rate ($r_t = i_t - E_t[\pi_{t+1}]$) gaps. Under perfect foresight, the IS equation (22) is globally log-linear.

The second equation of the macro block, the Taylor rule (equation 12) is similarly globally log-linear, which we repeat here for convenience:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) (\phi_\pi \pi_t + \phi_y (y_t - y_t^e)),$$

The global log-linearity of these equations means that they are not responsible for the non-linear evolution of the variables of interest, like the output gap, or inflation, which are all expressed in log terms. The non-linearity, instead, comes from the third equation, the Phillips curve, which expresses the relationship between inflation and the output gap.

Linearized New Keynesian Phillips Curve. For small shocks, the menu cost model implies a relationship between inflation and output gap that is well approximated by a log-linear new Keynesian Phillips curve (NKPC) (Auclert et al., 2024) of the form:

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \tilde{y}_t^e + u_t. \quad (23)$$

where π_t and $\mathbb{E}_t[\pi_{t+1}]$ respectively are actual and expected inflation, \tilde{y}_t^e is the output gap, and u_t is a cost-push shock. The parameter κ captures the *slope* of the Phillips curve, that is, the impact on inflation of a unit change in the output gap. The slope is constant in the log-

linearized formula and therefore does not depend on the size of the shock or the state of the economy. Auclert et al. (2024) show that the slope κ in the Golosov and Lucas (2007) menu cost model is higher than in a Calvo model calibrated to the same steady-state frequency of price changes. The difference is the consequence of the so-called “selection effect”. While in the Calvo model it is random which firms can adjust their prices, in the menu cost model the prices that adjust are far from their optimum. Thereby, large price adjustments in the menu cost model are endogenously selected, which makes the price level more flexible and inflation more responsive to changes in the output gap. The log-linear relationship between inflation and the output gap, however, is valid only for small shocks in a neighborhood of the steady state.

Non-linear New Keynesian Phillips Curve. The relationship between inflation and output is non-linear in the menu cost model. To see why, consider a series of monetary policy shocks of an ever-increasing size. If the shocks are small, the responsiveness of the frequency of price changes is low and the economy behaves similarly to a (suitably recalibrated) standard sticky-price economy *à la* Calvo. When the shocks are large, however, the frequency of price changes increases: firms update prices more often thus reducing the degree of price stickiness. For a sufficiently large shock, the frequency reaches 100%, that is, prices become completely flexible, and monetary policy loses its ability to influence activity. The output gap stays somewhat over its steady state level, as fully flexible prices help eliminate relative-price distortions caused by the interaction of idiosyncratic shocks and sticky prices.

The non-linear relationship between inflation and the output gap is state-dependent because it depends potentially on all state variables of the model, including the initial relative price distribution, as well as the history of past aggregate shocks. Despite the state-dependence of the Phillips curve, we can characterize numerically some of its relevant features. Panel (a) in Figure 1 plots the relationship between π_t on the y-axis and the cumulative output gap $\sum \beta^i \mathbb{E}_t \tilde{y}_{t+i}^e$ on the x-axis for varying shock sizes. Specifically, we simulate random realizations of i.i.d monetary policy shocks of varying size initiated from the stationary distribution of quality-adjusted relative prices at the deterministic steady state. We store the impulse responses of inflation and the output gap.⁶ Each point on the curve shows

⁶These impulse responses are computed under perfect foresight. For small shocks, this is equivalent to the first-

a particular realization of inflation and cumulative discounted output gap for a particular shock size.⁷ We are thus plotting κ_t in the expression (24) below, which is a generalization of the log-linear New Keynesian Phillips Curve (23) after repeated forward substitutions and assuming a non-linear coefficient κ_t :

$$\pi_t = \kappa_t \sum_{i=0}^{\infty} \beta^i \mathbb{E}_t \tilde{y}_{t+i}^e + \sum_{i=0}^{\infty} \beta^i \mathbb{E}_t u_{t+i} \quad (24)$$

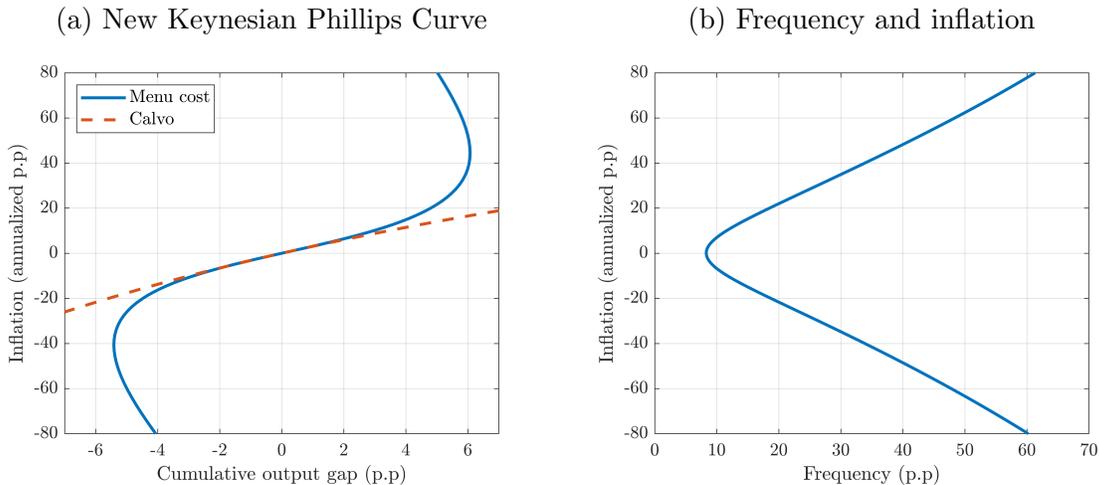


Figure 1: Inflation-output trade-offs in the menu cost and Calvo models.

Note: The figure is produced by computing the impulse responses to monetary policy shocks of different magnitudes and signs assuming no persistence of the Taylor rule ($\rho_i = 0$). The solid blue line is the baseline menu cost model and the dashed-red line is the Calvo model.

The figure shows that the relationship between activity and prices is strongly nonlinear, as large shocks have a larger impact on inflation relative to the output gap. The relationship akin to the slope of a New Keynesian Phillips curve is non-linear in menu cost models. Panel (b) displays the frequency of price changes at different inflation levels. The frequency is exactly 8.7% when the inflation π_t is zero. This is the deterministic steady state, when there are no aggregate shocks but idiosyncratic shocks yield a non-degenerate distribution of price changes. For higher magnitudes of shocks, the frequency of price adjustments increases gradually. At frequencies over 40%, the New Keynesian Phillips curve becomes backward

order approximation to the stochastic problem, as discussed by [Boppart et al. \(2018\)](#). For large shocks, its interpretation is similar to that in [Cavallo et al. \(2023\)](#): an unexpected once-and-for-all large shock that hits the economy in the deterministic steady state.

⁷Due to the certainty equivalence of first-order perturbation approximations, expected output gaps are equal to the actual impulse-response realizations.

bending. At this level, monetary policy impulses reach their maximum effectiveness in stimulating activity, and any larger policy easing raises inflation with smaller output effects. At some point (not shown), as frequency reaches 100%, monetary policy becomes ineffective in influencing the output gap.⁸ However, such shock sizes are fairly extreme.

While the location of the inflation-output gap pairs on the Phillips curve is informative about the *size dependence* of the shocks, the slope of the Phillips curve at each point $d\kappa(\Theta_t)/dy$, is informative about the *state dependence* of marginal shocks. It reflects the relative impact on inflation versus the output gap of a small shock, provided the economy is in a particular aggregate state Θ_t . Analogously, the slope reflects the state dependence of monetary policy and the inflation-output trade-off of a marginal monetary policy shock: how much cumulative output gap needs to decline to reduce inflation by a unit, also known as the *sacrifice ratio* of monetary policy. Figure 11 in Appendix C.1 shows the evolution of the slope of the Phillips curve for positive shocks (the relationship is analogous for negative shocks) for a range of realistic frequency values. The slope almost doubles by the time the frequency reaches 20%, a magnitude documented during the 2022-2023 inflation surge (Montag and Villar, 2023). While in a low-frequency/low-inflation environment, the sacrifice ratio is high, it becomes much lower once frequency/inflation increases.⁹

Finally, the *sign-dependence or asymmetry* of the non-linear Phillips curve is quantitatively small in the case of shocks with moderate persistence as shown by Figure 1. However, it can become quantitatively sizable when the aggregate shocks are highly persistent. The asymmetry of the Phillips curve is the consequence of the asymmetry of the firms' profit function: low relative prices lead to high demand, while high relative prices lead to low demand. Therefore, a negative deviation of the relative price from the optimum causes higher losses than a similar-sized positive deviation. An aggregate shock, therefore, which generates positive inflation and reduces relative prices can be expected to generate stronger endogenous responses and more flexible prices, than those causing negative inflation and higher relative prices. Persistent shocks increase the expected cumulative impact on the relative price if firms do not adjust, amplifying the asymmetry.

⁸At 100% frequency, output gap stabilizes at a level above its steady state value, as full price flexibility eliminates distortions due to idiosyncratic quality shocks.

⁹Blanco et al. (2024b) also discuss how the sacrifice ratio changes with the level of inflation

Shock-size dependence in response to a cost-push shock. Figure 2 displays instead the response to cost-push shocks of different magnitudes. The shock is implemented as a persistent decline in the firms' employment subsidy τ_t . Inflation, and the repricing frequency are evaluated on the period after the shock arrival; output gap is evaluated at its peak response usually around 5-6 months after impact; whereas the real interest rate is the cumulative over the time-span of the shock span. The figure displays both the baseline menu cost model and the standard Calvo model. The Calvo parameter is set to account for the selection effects as in Auclert et al. (2024).¹⁰ The value of the frequency in this case is implausibly large, 40%.

Interestingly, for a certain range of shock sizes (frequency increases below 10 p.p.) the non-linearity is quantitatively small. In this region, a linearized framework can be expected to approximate the true model quite well, and the responses coincide with those in the Calvo model.¹¹ For larger cost-push shocks, however, the menu cost model deviates from its Calvo counterpart. A key reason behind this development is the increase in the frequency of price changes, which makes the price level endogenously more flexible. As the frequency approaches 20% (increases more than 11%), the non-linear model is significantly different from its linearized counterpart. The non-linearity in the frequency response is a robust feature of menu cost models: the new price increases that are triggered by the small shock are almost fully offset by the new decreases that are canceled. For large shocks, however, price decreases disappear and price increases generate a large frequency response (Gagnon, 2009; Karadi and Reiff, 2019; Alvarez and Neumeyer, 2020; Alexandrov, 2020). The more flexible inflation response after large shocks in the menu cost framework leads to tighter systematic policy relative to Calvo as shown by the real interest rate on panel (d). This leads to a marginally amplified downturn in the menu cost model (see panel b).

¹⁰It delivers the same slope of the Phillips curve for small inflation levels, as illustrated by the dashed red line in Figure 1.

¹¹Figure 14 in Appendix C.3 shows how the full impulse responses to a small cost-push shock are similar to those of the recalibrated Calvo model.

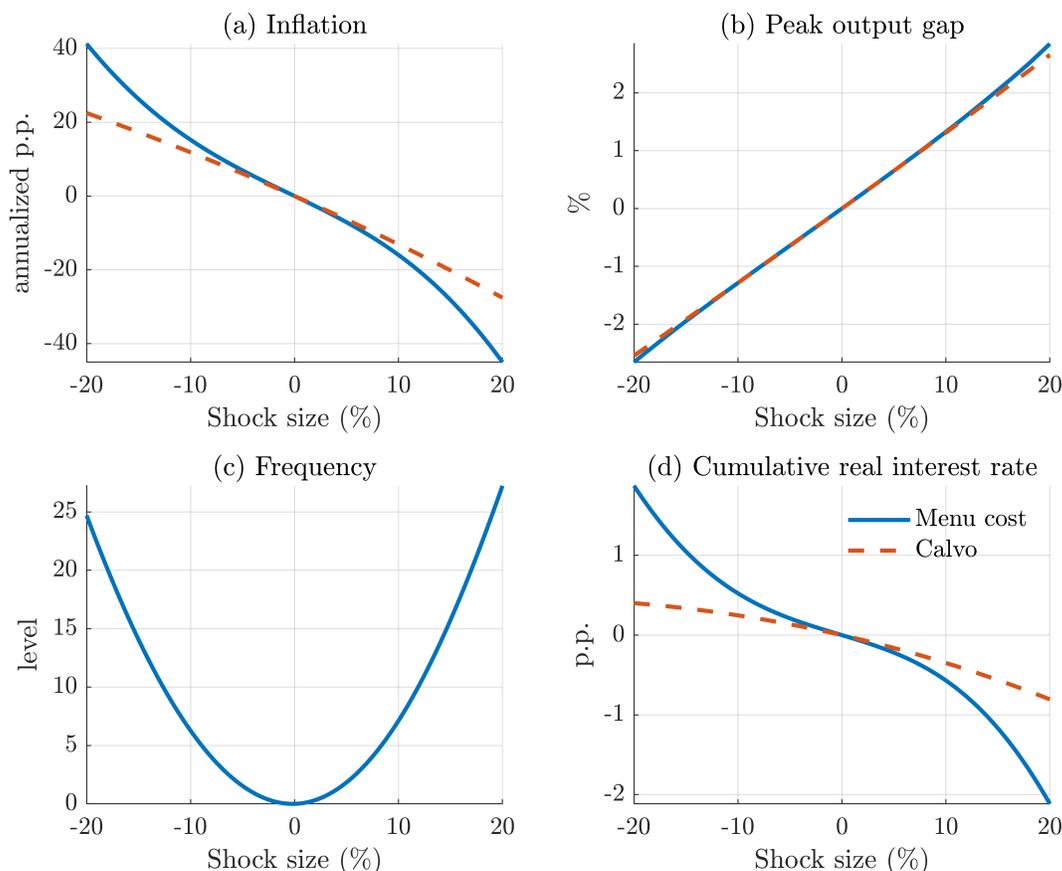


Figure 2: Response to a cost-push shock for different shock magnitudes.

Note: The figure displays the difference in the value of inflation, and the frequency of price changes between the period of the shock arrival and the value in the deterministic steady state. The output gap is evaluated at the peak of the impact relative to the deterministic steady state. The real interest rate is evaluated over the entire life of the shock.

5 Optimal monetary policy problem and computational approach

We turn next to the analysis of optimal monetary policy. We consider optimal monetary policy under commitment. In this section we introduce the central bank's problem and present a new computational method to deal with the complexities associated with the high-dimensionality of this problem.

5.1 Ramsey problem

The central bank maximizes households' welfare under commitment (Galí, 2008). The problem implies that the central bank selects the path for all equilibrium variables subject to all

the competitive equilibrium conditions:

$$\max_{\{g_t^c(\cdot), g_t^0, V_t(\cdot), C_t, w_t, p_t^*, s_t, S_t, \pi_t^*\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\gamma}}{1-\gamma} - v \frac{C_t}{A_t} \left(\int e^{(x+p_t^*)(-\epsilon_t)} g_t^c(p) dx + g_t^0 e^{(p_t^*)(-\epsilon)} \right) - v \eta g_t^0 \right)$$

subject to

$$\begin{aligned} w_t &= v C_t^\gamma, \\ V_t(x) &= \Pi(x, p_t^*, w_t, A_t) + \Lambda_{t,t+1} \frac{1}{\sigma} \int_{s_t}^{S_t} \left[V_{t+1}(x') \phi \left(\frac{(x-x') - \pi_{t+1}^*}{\sigma} \right) \right] dx' + \\ &\quad \Lambda_{t,t+1} \left(1 - \frac{1}{\sigma} \int_{s_t}^{S_t} \left[\phi \left(\frac{(x-x') - \pi_{t+1}^*}{\sigma} \right) \right] dx' \right) [(V_{t+1}(0) - \eta w_{t+1})], \\ V_t(s_t) &= V_t(0) - \eta w_t, \\ V_t(S_t) &= V_t(0) - \eta w_t, \\ V_t'(0) &= 0, \\ g_t^c(x) &= \frac{1}{\sigma} \int_{s_{t-1}}^{S_{t-1}} g_{t-1}^c(x_{-1}) \phi \left(\frac{x_{-1} - x - \pi_t^*}{\sigma} \right) dx_{-1} + g_{t-1}^0 \phi \left(\frac{-x - \pi_t^*}{\sigma} \right), \\ g_t^0 &= 1 - \int_{s_t}^{S_t} g_t^c(x) dx, \\ 1 &= \int e^{(x+p_t^*)(1-\epsilon)} g_t^c(x) dx + g_t^0 e^{(p_t^*)(1-\epsilon)}. \end{aligned}$$

where $\phi(\cdot)$ is the probability density function of a normal random variable, g_t^c is the continuous part of the price-gap density, g_t^0 is the frequency of repricing, that is, the probability mass of updated prices, and s and S are the endogenous boundaries of the inaction region for x . The nominal interest rates consistent with this path of nominal and real variables can then be recovered from the household's Euler equation (7). Importantly, the constraint set of the planner's problem is continuous and differentiable despite the fact that the individual firm's price policy function is not. This is so because each firm has zero mass, and thus the discontinuity in a single firm's behavior does not lead to a discontinuity in aggregates. Furthermore, note that both $V(x)$ and $g^c(x)$ are continuously differentiable over the relevant range (s, S) .

5.2 Computational method

We propose a new algorithm similar to that in [González et al. \(2024\)](#) but applied to discrete time. The idea is to represent the problem as a high-dimensional dynamic programming problem in which the Bellman equation and the law of motion of the price-gap distribution are constraints.

The problem of the Ramsey planner is complicated by the fact that the value function $V_t(\cdot)$ and the distribution $g_t(\cdot)$ are infinite-dimensional variables. This poses a challenge when solving the optimal monetary policy problem, as we need to compute the first-order conditions (FOCs) with respect to these infinite-dimensional variables.¹² Our algorithm first discretizes the planner’s objective and constraints (the private equilibrium conditions) and then determines the FOCs, instead of first determining the FOCs for the planner’s continuous space problem, and then discretizing them. Thus we transform the original infinite-dimensional problem into a high-dimensional problem, in which the value function and the state density are replaced by large vectors with a dimensionality equal to the number of grid points used to approximate the individual state space. This approximation needs to be smooth and good enough to capture the higher-order effects of policy.

An additional challenge in our particular problem is that a simple discrete-state approximation may fail, as the private equilibrium conditions include discrete choices. Therefore, we approximate the distribution and value function not by discrete functions on a predetermined grid, but by *piece-wise linear* functions over an *endogenous* grid. The endogenous grid is selected to always include the two boundaries of the inaction region (points $x = s$ and $x = S$) and the optimal price gap ($x = 0$). Furthermore, we explicitly take the mass point at 0 into account in the distribution. Integrals to compute expectations are evaluated algebraically, conditional on those piecewise linear functions.

¹²There are a number of proposals in the literature to deal with this problem. [Nuño and Thomas \(2022\)](#), [Smirnov \(2022\)](#), and [Dávila and Schaab \(2022\)](#) deal with the full infinite-dimensional planner’s problem in continuous time. This implies that the Kolmogorov forward (KF) and the Hamilton-Jacobi-Bellman (HJB) equations are constraints faced by the central bank. They derive the planner’s FOCs using calculus of variations, thus expanding the original problem to also include the Lagrange multipliers, which in this case are also infinite-dimensional. These papers solve the resulting differential equation system using the upwind finite-difference method of [Achdou et al. \(2021\)](#). [Bhandari et al. \(2021\)](#) make the continuous cross-sectional distribution finite-dimensional by assuming that there are N agents instead of a continuum. They then derive standard FOCs for the planner. In order to cope with the large dimensionality of their problem, they employ a perturbation technique. [Le Grand et al. \(2022\)](#) employ the finite-memory algorithm proposed by [Ragot \(2019\)](#). It requires changing the original problem such that, after K periods, the state of each agent is reset. This way the cross-sectional distribution becomes finite-dimensional.

As we show in Appendix D, the Bellman equation can thus be approximated over a grid of price gaps x as

$$\mathbf{V}_t = \mathbf{\Pi}_t + [\mathbf{A}_t \mathbf{V}_{t+1} - \mathbf{b}_{t+1} \eta w_{t+1}]$$

where \mathbf{V}_t and \mathbf{b}_t are vectors with the value function and the expected adjustment probability evaluated at different grid points, respectively, and \mathbf{A}_t is a matrix that captures the idiosyncratic transitions due to firm-level quality shocks and aggregate inflation. Similarly, the law of motion of the density for $x \neq 0$ is

$$\mathbf{g}_t^c = \mathbf{F}_t \mathbf{g}_{t-1}^c + \mathbf{f}_t^\top g_{t-1}^0.$$

where \mathbf{g}_t^c and \mathbf{f}_t are vectors representing the probability distribution function and the scaled and shifted normal distribution, respectively, \mathbf{F}_t is a matrix that captures the evolution of the price distribution due to firm-level quality shocks and aggregate inflation, and

$$g_t^0 = 1 - \mathbf{e}_t^\top \mathbf{g}_t.$$

is the mass point at $x = 0$. Here \mathbf{e}_t is a vector of weights corresponding to the trapezoid rule. The labor market clearing condition and the definition of the price index can be written in a similar form. The computational appendix D provides further details.

Once we have the discretized version of the problem, we find the planner's FOCs by symbolic differentiation. This delivers a large-dimensional system of difference equations, as we have Lagrange multipliers associated with each gridpoint of the value function or the probability function.

Next, we find the Ramsey steady state. To do so, we construct a nonlinear multidimensional function mapping one variable, in our case inflation, to the rest of the steady-state equilibrium variables. We then combine this function with the planner's FOCs. As the system is linear in the Lagrange multipliers, the solution is to finding the zero of a nonlinear function of the initial variable (i.e., inflation), for instance, using the Newton method. Finally, to compute the dynamics of the Ramsey problem, we solve the system of difference equations non-linearly in the sequence space also using the Newton method.

The symbolic differentiation and the two applications of the Newton algorithm can be

conveniently automated using several available software packages. In our case, we employ Dynare (Adjemian et al., 2023), but the approach is also compatible with the nonlinear sequence-space Jacobian toolboxes. This algorithm can be employed to compute optimal policies in a large class of heterogeneous-agent models. Compared to other techniques, it stands out for being easy to implement. González et al. (2024) show that this algorithm delivers the same results as computing the FOCs by hand using calculus of variations and then discretizing the model. Our algorithm runs in a few minutes on a normal laptop.

6 Optimal monetary policy: results

We now proceed to investigate the model’s normative prescriptions.

6.1 The steady state under the optimal policy

The solution of the Ramsey planner’s problem has a steady state state featuring a slightly positive inflation of 0.3%.¹³ This is different from the standard New Keynesian model with Calvo pricing (Galí, 2008), where the optimal inflation in the Ramsey steady state is zero. This value of optimal inflation in the menu-cost model is very close to the value of steady-state inflation that maximizes steady-state welfare, which in turn is also very close to the value of inflation that minimizes the frequency of price adjustments.

What explains the positive optimal inflation? The key is the asymmetry of the profit function (11). For a firm, a negative price gap is more undesirable than a positive price gap of the same size because a negative price gap $-x$ leads to much larger sales at a markup-loss of $-x$, while the positive price gap x leads to only somewhat smaller sales at a markup gain of x . This implies that the Ss bands are asymmetric: the lower Ss band is closer to the optimal price gap than the upper one (see Figure 3). Thus, in the zero-inflation steady state, there is more mass of firms close to the lower threshold of the inaction band than to the upper threshold. As a result, there are more upward than downward price adjustments. Small positive inflation raises the optimal reset price p^* and shifts the Ss bands leftwards and thus reduces the number of upward price movements by more than it increases the number

¹³In our numerical exploration, we have only found a single steady state.

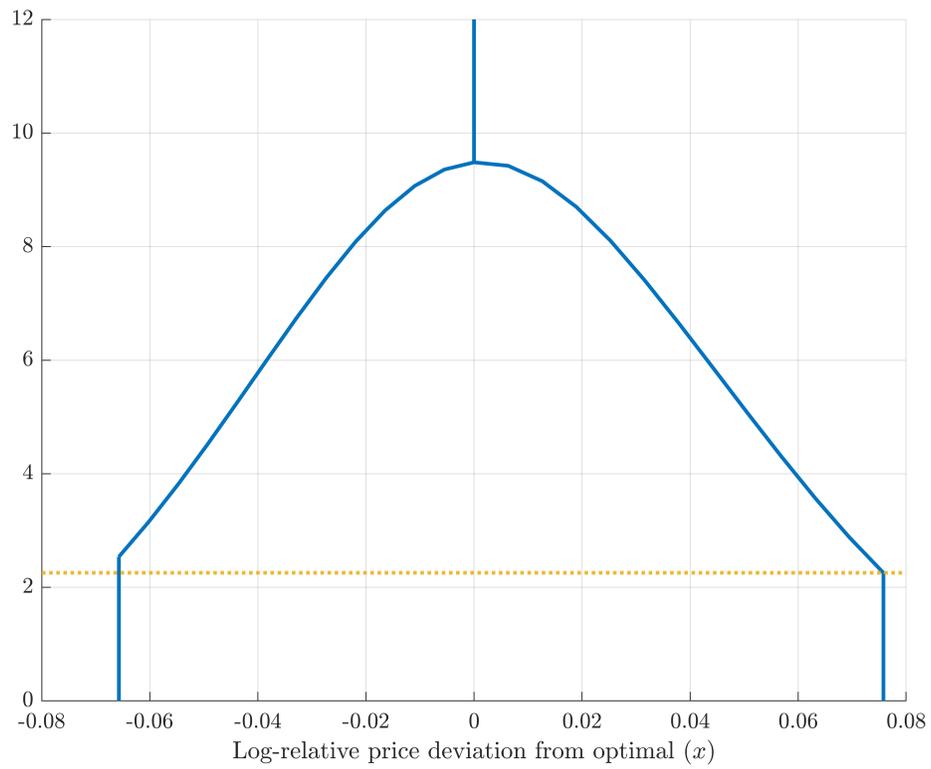


Figure 3: Steady-state price-gap density.

Note: The figure displays the steady-state price-gap density $g(x)$ with zero inflation. The dashed yellow line indicates the mass of firms at the upper Ss band.

of downward price movements. The frequency of price adjustments decreases and with it the distortions caused by menu costs. Quantitatively this effect is small, but not negligible.

6.2 Time-0 problem

We now turn to investigating the time inconsistency of optimal policy. To assess its magnitude, we solve the optimal policy problem, starting from the price distribution in the Ramsey steady state, assuming that the central bank faces no previous pre-commitment. In this case, the Lagrange multipliers associated with forward-looking equations are initially set to zero. This problem is often referred to as the “time-0 problem” (Woodford, 2003).

The solid blue lines in Figure 4 show the time path under the optimal policy. The labor subsidy is set to zero in this exercise, which, therefore, ceases to offset any markup distortions caused by the firms’ market power. The optimal policy is time-inconsistent: without pre-commitment, the central bank engineers a temporary expansion. Thereby, it raises welfare by bringing output closer to its efficient level at a cost of elevated price distortions arising from the higher inflation.

The dashed red line on Figure 4 shows the equivalent time-0 response in the Calvo model. The figure shows that the incentives to surprise is substantially weaker in the menu-cost model: both the inflation and output gap increases are smaller relative to the Calvo model. The main reason is that the price level is more flexible in the state dependent model. The unexpected easing causes a sizable inflation hike but only a small and temporary output increase. The sizable increase in the frequency of price changes contributes to the flexibility of the price level, which is already elevated due to endogenous selection of large price changes. Overall, the magnitude of time inconsistency is lower, because the welfare gain from the higher output is smaller, but the costs from the price distortions caused by the higher inflation are higher, albeit less persistent.

There is a countervailing force that raises the time inconsistency in our baseline model relative to the Calvo model. Namely, the markup distortions are higher, as discussed below, and a labor subsidy of $\tau = 1/\epsilon$ is insufficient to bring the distortions caused by the average markup to zero, as it is the case in the Calvo model. A time-0 optimal policy, therefore, stays time inconsistent even with a $\tau = 1/\epsilon$ labor subsidy (not shown). The optimal policy easing

in this scenario, however, is small, two orders of magnitude smaller than those under no labor subsidy. Therefore, this channel is too weak to counteract the opposite effect caused by the more flexible price level detailed in the previous paragraphs. A corollary to the negligibility of the time inconsistency with an appropriate labor subsidy is that the analysis in the next section, where we focus on adopting a timeless perspective, would go through also if we adopted a time-0 perspective.

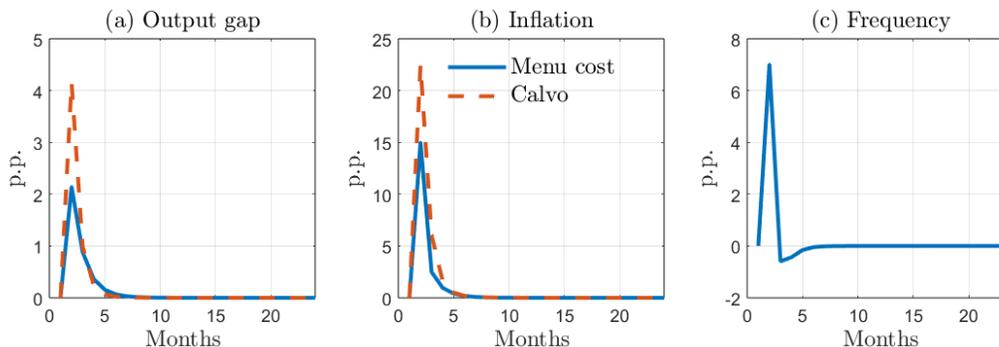


Figure 4: Time inconsistency of the optimal policy.

Note: The figure compares the time-0 optimal policies in the menu-cost model and in the Calvo model.

6.3 Timeless optimal monetary response to efficient shocks

While the time-0 policy is interesting to analyze time inconsistency, our main focus is on the systematic monetary policy response to shocks. We then analyze *timeless* optimal monetary policy (Woodford, 2003; Galí, 2008). This corresponds to the optimal monetary policy starting from the Ramsey steady state when a shock hits the economy and all of the Lagrange multipliers are initialized at their steady-state values.

First, we consider shocks that affect the efficient allocation, such as TFP shocks. In the standard New Keynesian model with Calvo prices, the response to such shocks is characterized by strict price stability: the central bank steers real interest rates to replicate the path of natural interest rates, which leads to inflation and the output gap remaining at zero. This is commonly known as the “*divine coincidence*” (Blanchard and Galí, 2007).

A version of such divine coincidence also holds in our economy. As we have shown in Section 6.1, strict price stability would be suboptimal in our menu-cost model, and optimal trend inflation is also positive. As a response to an efficient shock, optimal timeless com-

mitment policy keeps inflation and the output gap at their steady-state levels and makes the real interest rate follow its natural level. As inflation remains constant, the frequency of repricing also stays constant. In other words, the optimal policy offsets the dynamic impact of the efficient shocks in a form of “dynamic divine coincidence”.

The conclusion is that strict targeting of the optimal steady-state inflation rate simultaneously minimizes inefficient output fluctuations and the costs of nominal rigidities. Notice that the shape of the Phillips curve plays no role in this result and thus the prescription is the same for small and large shocks. The next subsection turns to a case when the Phillips curve matters.

6.4 Timeless optimal monetary response to cost-push shocks

Impulse response to a cost-push shock. Cost-push shocks are well known to break the divine coincidence in the standard New Keynesian model. For such shocks, a policy of “leaning against the wind” is optimal, that is the central bank tolerates a temporary increase in inflation to cushion the decline in output below its efficient level. The central bank trades off this decline in the output gap with the initial bout of inflation under the commitment that inflation will be negative in the future so that the price level is eventually restored. Such a policy exploits the Phillips curve relationship and is thus the relevant case to explore the implications of a nonlinear Phillips curve.

The optimal monetary policy response to a cost-push shock in our model is also characterized by leaning against the wind. Figure 5 shows the response of the economy to an inflationary large cost-push shock under the optimal policy, which pushes there repricing frequency above 20%. Inflation goes up (panel b) and the efficient output gap drops (panel a). The price level initially increases, but then decreases and eventually returns to the steady-state trend.

However, due to the nonlinearity of the Phillips curve, the policy is now size-dependent. The figure compares the response to a large shock (solid green line) with that of a small shock that is scaled linearly by the relative shock size (black dashed line). The optimal monetary policy is tighter for the large shock: this prevents inflation from increasing as much as in the small-shock case (panel b). The tighter policy is also successful in preventing a decline

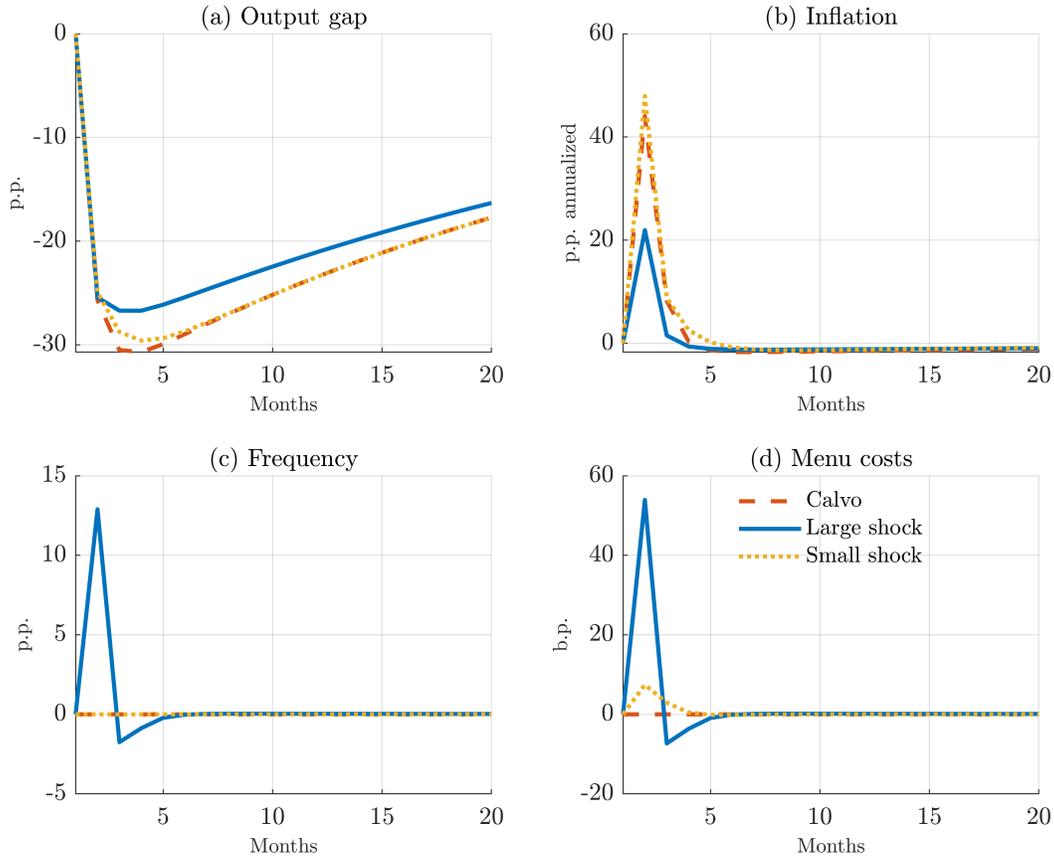


Figure 5: Impulse responses to a cost-push shock in the Calvo model, and to a large and a small cost-push shock in the menu-cost model under the optimal monetary policy.

Note: all displayed variables except frequency are linearly scaled in the small shock according to the ratio between the large and small shocks.

in the output gap as large as with small shocks (panel a).

Finally, the figure also displays the optimal policy response in the case of a Calvo model calibrated to replicate the slope of the Phillips curve. The policy approximately coincides with that of the menu-cost model for small shocks.

Non-linear targeting rule. To interpret these results, we first need to understand the impact of higher frequency on the conventional New Keynesian model. In the linearized Calvo model, optimal policy is characterized by a targeting rule linking inflation to the change in the output gap. The relationship is constant and depends only on the elasticity of substitution parameter ϵ . Strikingly, the relationship is independent of the frequency of price changes, proportional to the Calvo parameter θ . The reason behind this result is that

two opposing effects perfectly offset each other. On the one hand, frequency raises the slope of the Phillips curve $\kappa = (1 - \theta)(1 - \beta\theta)/\theta$, and therefore reduces the sacrifice ratio making the cost of reducing inflation lower. On the other hand, higher frequency reduces the relative weight of inflation in the welfare function, as this is given by ϵ/κ . The reason is that higher frequency reduces the relative price distortions caused by inflation, so the planner is willing to tolerate more inflation.

The frequency of price changes also influences the magnitude of the Phillips-curve wedge (u_t) that a cost-push shock generates. The shock $u_t = \kappa(y_t^e - y_t^n)$ is equal to the product of the slope of the Phillips curve (κ), which increases with the frequency of price changes, and the difference between the efficient output (y_t^e) and the natural level of output (y_t^n) (see Galí, 2008, pp. 97). The cost-push shock (τ_t) drives the difference between the efficient and the natural output. An increase in frequency thus amplifies the impact of the shock on the Phillips curve.¹⁴

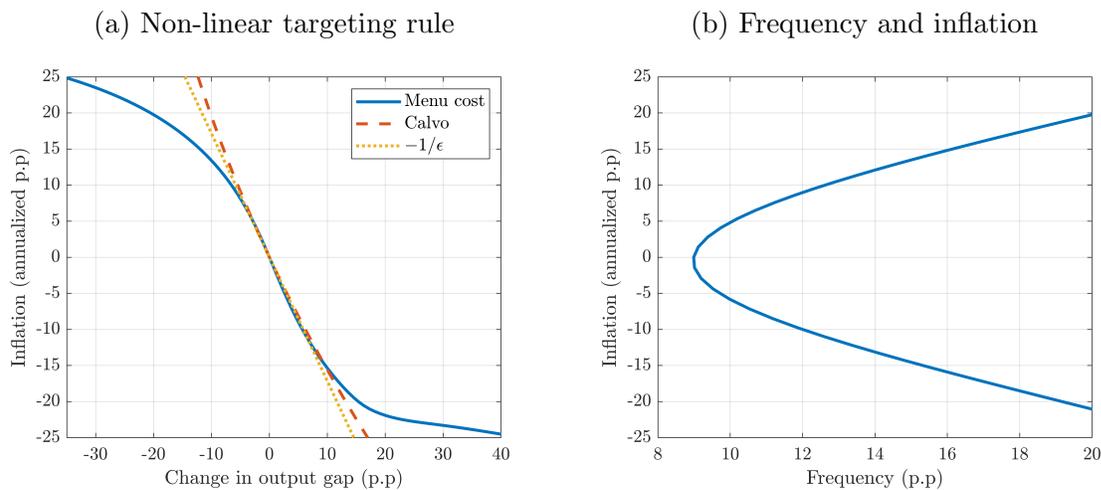


Figure 6: Optimal target rule under optimal policy.

Note: Panel (a) shows the relationship between annualized inflation and the change in output gap on impact as a function of a series of cost-push shocks under the optimal policy (blue solid line); and contrasts it to the optimal target rule in the Calvo model (dashed red line) and in the linearized Calvo model (dotted yellow line). Panel (b) shows the frequency at each inflation level.

Figure 6 shows the relationship between annualized impact inflation and impact change

¹⁴This feature influences the optimal policy: by leaning against frequency, optimal policy reduces the Phillips curve slope (κ) and, consequently, mitigates the distorting effects of the shock. However, as panel c of Figure 7 shows, the frequency is strictly increasing in shock size τ_t even under optimal policy, so this channel is not strong enough to offset the direct impact of the shock size on frequency. Therefore, the slope of the Phillips curve is increasing (and the sacrifice ratio is decreasing) in the shock size even under optimal policy, qualitatively similarly to the case under Taylor rule as outlined in Section C.1

in output gap as a response to cost-push shocks of different sizes. The slope of the target rule at zero output gap is very close to $-1/\epsilon$ (dotted yellow line), the slope of the target rule in the linear Calvo model. In the case of small shocks, the change in frequency is relatively modest, and thus the logic of the Calvo framework still applies. The logic generalizes in the Calvo model also for large shocks: the target rule stays approximately linear (dashed red line) with a slope close to $-1/\epsilon$.

In the menu cost model the relationship is non-linear (blue solid line). At a frequency of around 12%, the slope of the optimal target rule becomes significantly lower than its slope at the steady-state frequency of 8.5%. The relationship means that after a large shock that increases frequency substantially, the planner is stabilizing inflation relative to the output gap on the margin more than after small shocks. The central bank thus “leans against frequency”, tightening policy more aggressively in the case of a large shock that increases frequency. The targeting rule is not independent of frequency: as frequency increases the central bank has more incentives to quench inflation to minimize price distortions and is more effective in doing so due to a lower sacrifice ratio. These two effects trump the amplification of the cost-push wedge and explain our results.

Impact of shock size. Figure 7 illustrates the optimal response to a cost-push shock for a whole range of shock sizes. The impact response is plotted as a function of the shock size (solid blue line) and is contrasted to the Calvo model (dashed red line), which roughly coincides with the linear extrapolation of the slope in the menu-cost model around small shocks. The deviation between the two lines illustrates again the nonlinearity of the economy under optimal policy for large shocks. The larger the shock the more inflation and frequency are contained as the central bank tightens monetary policy more aggressively. The central bank thus ‘strikes while the iron is hot’.

Second, there is a certain sign-dependence, as the optimal policy response to positive and negative shocks differ. This is related to the asymmetries discussed in Section 4.

Welfare decomposition. To understand the rationale for the optimal response of the model, we decompose welfare into three components. The welfare gap (W_0) relative to the welfare in the first-best efficient equilibrium (W_0^e) comes from three underlying sources. First, it comes from the *misallocation* due to distortions caused by non-zero markups ($\mu_t(i) =$

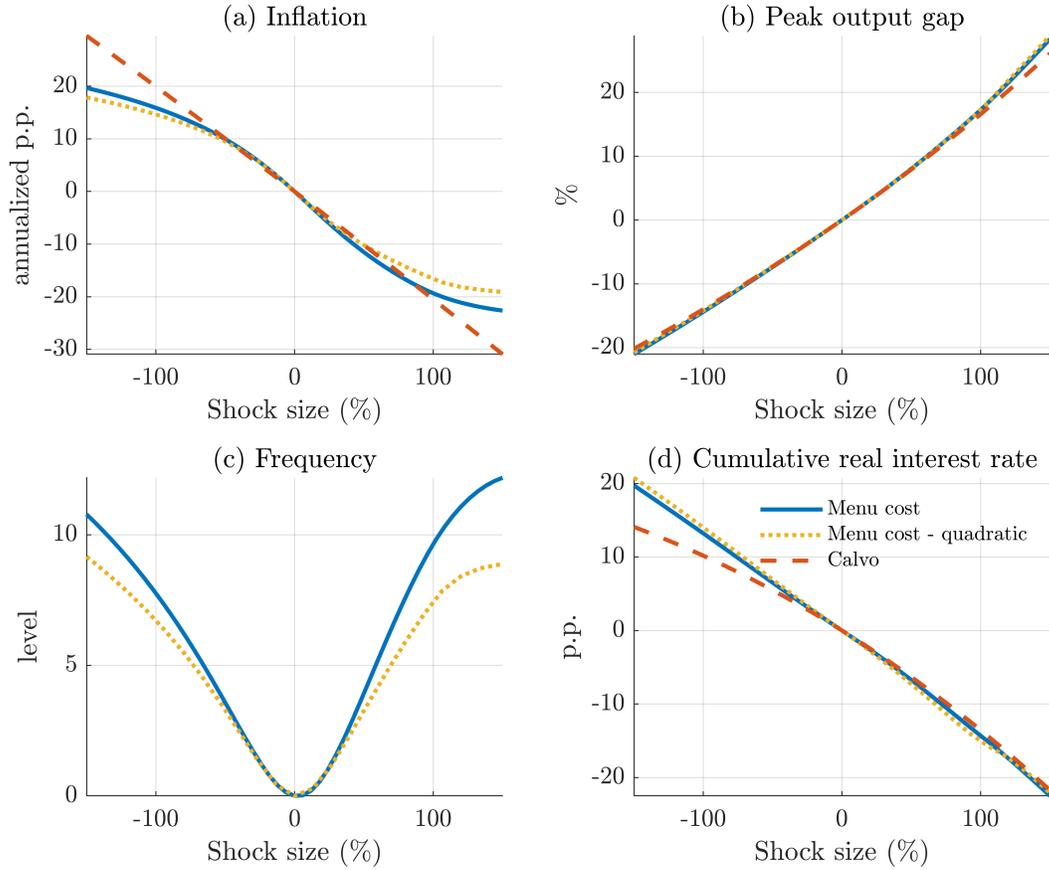


Figure 7: Optimal response to a cost-push shock for different shock magnitudes.

Note: The figure displays the difference in the value of inflation, output gap, and repricing frequency between the period after the shock arrival and the value in the deterministic steady state. The real interest rate is evaluated over the entire life of the shock.

$p_t(i) - mc_t(i)$) as in the first-best all markups are zero. This in, turn, depends on the *average markup gap* linked to the average markup $\bar{\mu}_t$, and to the *price dispersion*, which affects the dispersion of the demeaned markups $\zeta_t^{\mu - \bar{\mu}}$. Finally, the welfare gap comes from the resource loss from the labor allocated to price adjustment (*menu costs*). In Appendix B.2, we derive that the welfare loss, in the baseline case of $\gamma = 1$, equals

$$W_0 - W_0^e = \sum_{t=0}^{\infty} \beta^t \left(\underbrace{-\log(1 - \tau_t) - \bar{\mu}_t - \left(\frac{1}{e^{\bar{\mu}_t(1 - \tau_t)} - 1} \right)}_{\text{Average markup gap}} - \underbrace{\left(\frac{1}{e^{\bar{\mu}_t(1 - \tau_t)} (\zeta_t^{\mu - \bar{\mu}} - 1)} \right)}_{\text{Price dispersion}} \underbrace{- \eta g_t^0}_{\text{Menu costs}} \right). \quad (25)$$

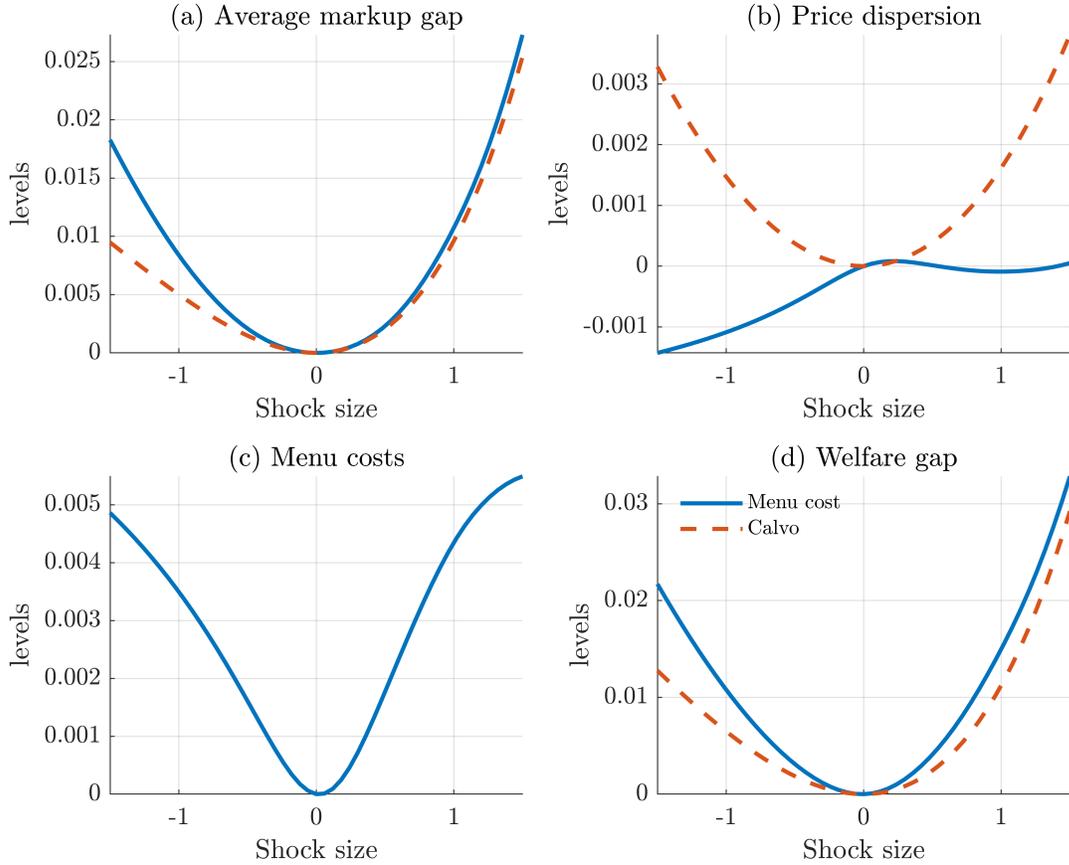


Figure 8: Welfare decomposition of a cost-push shock for different shock magnitudes.

Note: The figure displays the difference in the different components of the welfare gap in equation (25).

We have set the steady-state labor subsidy τ to offset the average markup distortions. In particular, we make sure that $\bar{\mu} = -\log(1 - \tau)$ and thus the steady-state average markup gap is zero. In this case, the steady-state welfare gap comes exclusively from the markup

dispersion ($\zeta^{\mu-\bar{\mu}}$) and the menu costs. In our baseline calibration, these terms are similar in magnitude and equal to 0.4 percent. Menu costs with such a magnitude are reasonable: Zbaracki et al. (2004) estimates menu costs including managerial and marketing costs around 1.2 percent of firm revenues, which is higher, but of similar magnitude.

Figure 8 displays the decomposition for the menu-cost and the Calvo models. In the Calvo model there are only two components. First, the average markup gap, which is directly related to the output gap, and second, the price dispersion. Both terms increase with shock size. In the menu-cost model, however, price dispersion does not increase, but *decreases*, with shock size. This is due to the selection effect. Quantitatively, price dispersion plays a minor role in the welfare gap. Instead, in the menu cost model the price adjustment costs are important. Like price dispersion in Calvo, they increase with shock size.

The average markup gap equals the square of the output gap up to a second-order approximation. In Calvo, price dispersion depends on the square of inflation. This gives rise to the standard output gap-inflation trade-off that characterizes New Keynesian models. In menu-cost models, the menu costs depend linearly on frequency, which in turn depends on inflation, as discussed in Section 4. This relationship can be approximated by a quadratic function. Therefore we have a similar inflation-output gap trade-off, but the central bank dislikes inflation not because it increases misallocation due to price dispersion, but because it increases the losses associated with menu costs.

Alternative central bank objectives. While the previous explanation addresses the motivation of the central banker, it does not answer the question of why optimal policy in the menu-cost model displays such a strong degree of non-linearity. To answer that question, and following the previous discussion, we consider a purely quadratic objective

$$\frac{1}{2} \left[\mathbb{E} (\pi_t - \pi_{ss})^2 + \frac{\kappa}{\epsilon} \mathbb{E} (\tilde{y}_t)^2 \right]$$

In the Calvo model, this expression is exact up to a second-order approximation (with $\pi_{ss} = 0$). The variance of the inflation term is the welfare cost of price dispersion, and the variance of the output gap term is the welfare cost of changes in the average markup. In the menu cost model, we impose that quadratic objective. In this case, the first term is still a second-order approximation of the welfare cost of the average markup, and the second

term serves as an approximation of the relationship between the inflation and the sum of price dispersion and menu costs. This approximation neglects the state dependency of this relationship.

The dotted yellow line in Figure 7 above displays the optimal response in the menu-cost model as a function of shock size in the case of a quadratic objective. The result is strikingly similar to the result in the baseline case. In particular, the model with a quadratic objective yields a nonlinear target rule. This confirms that the main source of the non-linearity in the menu cost model is not to be found in the non-linear shape of the welfare function, but in the non-linearity of the Phillips curve relationship.

7 Implications for the 2022-2023 US inflation surge

Scenario description. We use our model to assess its optimal policy prescriptions in a stylized scenario akin to the 2022-2023 post-Covid US inflation surge. Besides the large increase in frequency, a further notable feature of the inflationary episode was the significant *increase* in the dispersion of the price changes (Montag and Villar, 2023). This is relevant because large aggregate shocks with uniform impact on firms' optimal reset prices, including demand, supply, and cost-push shocks in our framework, would *decrease* the price-change dispersion. The reason is that the new price changes, resulting from price gaps pushed over one of the inaction thresholds by the shock, are necessarily closer to the mean of price changes than the canceled price adjustments, which are around the other inaction threshold. The contradictory evidence, therefore, suggests that concurrent forces raised the dispersion of reset prices, like an increase in the volatility of idiosyncratic demand or supply shocks.¹⁵ In this section, we construct a scenario as a combination of a cost-push shock, a dispersion shock, and a monetary policy that follows an inertial Taylor rule. The scenario broadly captures key features of the 2022-2023 US inflation surge. Then we use our model to characterize the counterfactual optimal Ramsey commitment policy.

Figure 9 shows the impulse responses to a combination of a cost-push and a dispersion shock, alongside the impulse responses to the two shock components. The magnitudes of

¹⁵Similar idiosyncratic volatility shock and its interaction with aggregate shocks is analyzed in Vavra (2014). Berger and Vavra (2019) explore the possibility of increased dispersion caused by heterogeneous response to aggregate shocks.

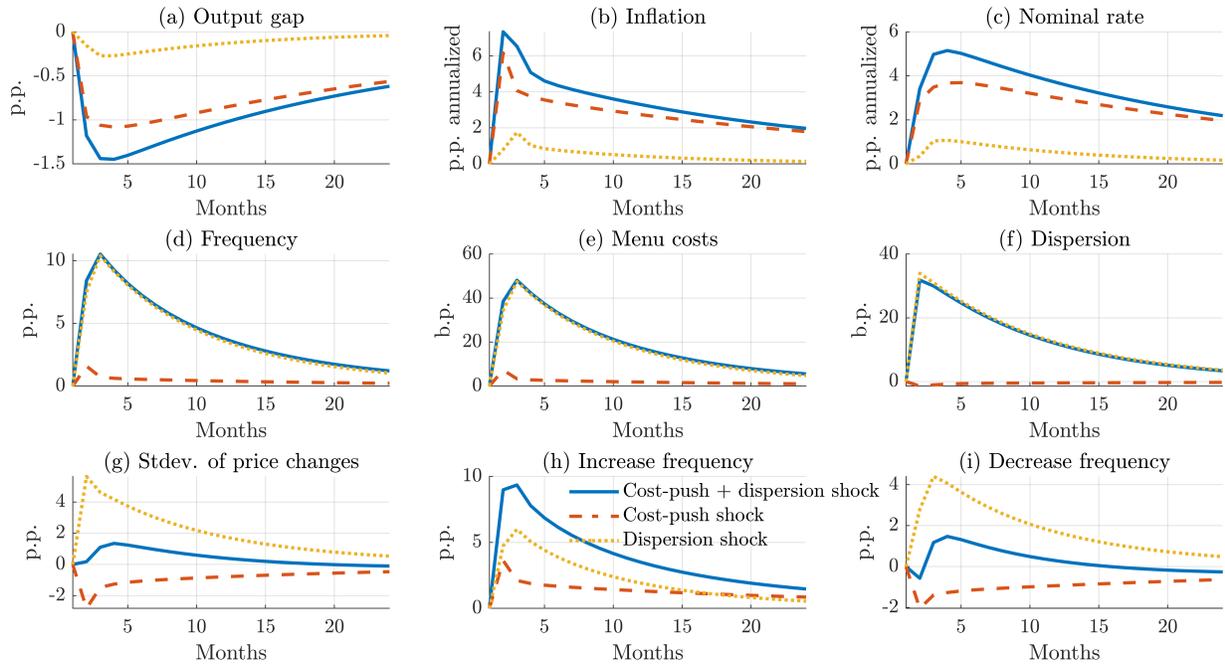


Figure 9: Inflation surge caused by a cost-push and dispersion shocks.

Note: The figure shows a scenario (blue solid line) as a response to a combination of a large cost push τ_t and a dispersion σ_t shock, which captures some key features of the 2022-2023 inflation surge. The figure also shows the individual responses to the cost-push shock (red dashed line) and the dispersion shock (yellow dotted line) components.

the shocks were targeted to capture both the increase in the frequency of price changes to around 20%, and the increase in the standard deviation of the price changes by around 1.5 percentage points, as observed during the 2022-2023 inflation surge (Montag and Villar, 2023).

The scenario captures broad features of the 2022-2023 US inflation surge. The shocks raise the inflation rate from the steady-state 2 percent to around 9 percent.¹⁶ Around one-third of the inflation surge is temporary and recedes quickly, while the remaining two-thirds are persistent and decline slowly. The nominal interest rate increases quickly by around 5 percent, broadly in line with the magnitude of the Fed’s 2022 policy tightening. The output gap declines by 1.5 percentage points.¹⁷ The figure also shows that the scenario captures the

¹⁶We have reduced the interest-rate smoothing parameter of the inertial Taylor rule from 0.75 quarterly rate to 0.25 quarterly rate to capture the unusually fast pace of interest-rate increase observed in 2022.

¹⁷The decline in output is contrary to the evidence. A more realistic scenario should also include a positive demand shock raising output, caused, for example, by heightened consumption demand from excess savings built up during the Covid-19 lock-downs and from fiscal transfers. Excluding such a demand shock is a conservative choice: a scenario with demand shock would prescribe an even more aggressive optimal policy tightening than our baseline. The reason is that monetary policy would optimally fully offset an efficient demand shock, while it leans only partially against

sizable and persistent increase in the frequency of price changes.

The figure also confirms that the standard deviation of price changes helps to identify the magnitude of the dispersion shock accompanying the aggregate shock. The standard deviation of price changes declines as a response to the cost-push shock, while it increases significantly after the dispersion shock. Their combination in the baseline matches the targeted 1.5 percent increase. Notably, the exercise also captures the around 2 percentage-point *increase* in the frequency of price decreases, even though it is not a moment we explicitly target. An aggregate shock alone would have reduced the frequency of price decreases, while a sole dispersion shock would have increased it by more than 4 percentage points. All in all, the large increase in the overall frequency is driven mostly by the increase in price increases reaching around 14% from around 6%, similar to the data (Montag and Villar, 2023).

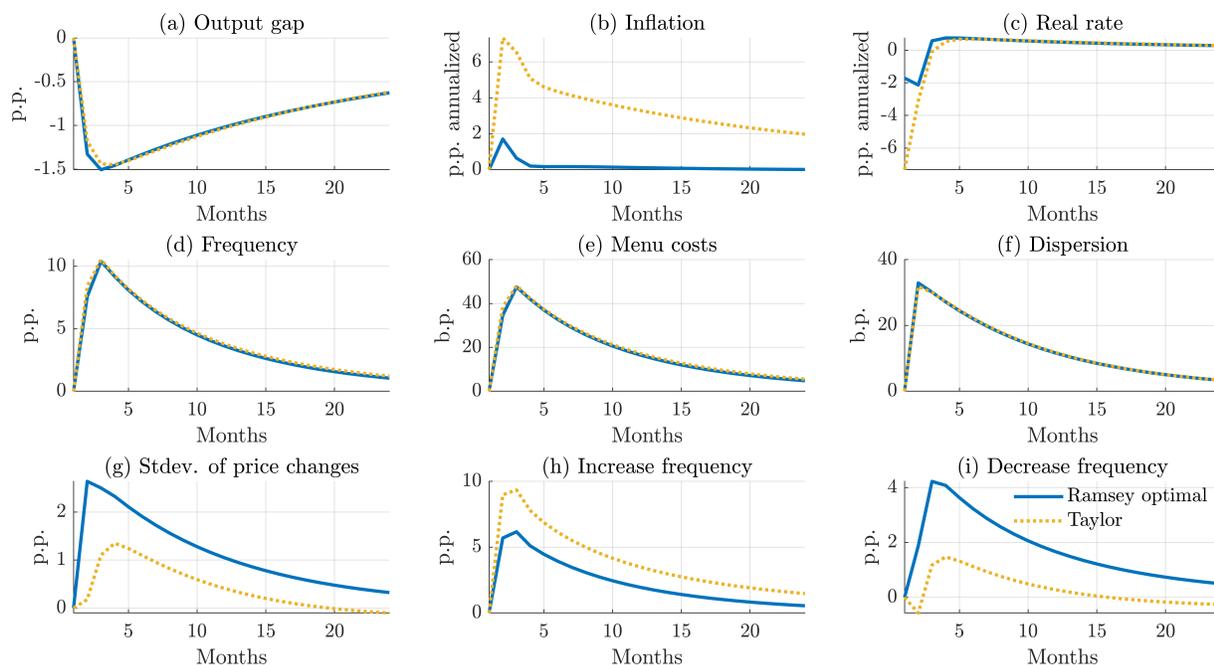


Figure 10: Optimal monetary policy response to a joint cost-push and dispersion shock.

Note: The figure shows the optimal Ramsey commitment response (blue solid line) to a combination of a large cost push τ_t and a dispersion σ_t shocks. It contrasts it to the policy following an inertial Taylor rule (yellow dotted line).

How would optimal Ramsey policy have reacted to such a combination of shocks? The

the inefficient cost-push shock in our scenario.

answer is not obvious, as the dispersion shock generates variation in optimal relative prices and, consequently, a welfare-enhancing increase in the frequency of price changes that monetary policy might not want to counteract. Figure 10 shows the counterfactual optimal policy response. It shows that the optimal policy would prescribe a substantially tighter policy in the initial periods than the Taylor rule. The optimal policy would have led to an increase in inflation by only around 1 percentage point at the cost of a somewhat lower output gap. This is substantially lower than the 7 percentage points increase in inflation under the Taylor rule. The policy would have reduced the frequency of price changes only marginally: it would have stayed persistently high, reaching close to 20%. The low optimal inflation, furthermore, would be accompanied by a substantial shift from price increases towards price decreases, which are not costlier than price increases in our framework.¹⁸

8 Conclusion

This paper analyzes the Ramsey problem within a menu-cost model à la Golosov and Lucas (2007). A key contribution is the identification of a new incentive for the central bank: In the presence of large *cost-push* shocks, optimal monetary policy should commit to quelling inflation and stabilizing the repricing frequency. Along the trajectory of optimal commitment, the central bank utilizes the reduction in the sacrifice ratio, leading to lower inflation at a cost of a more moderate decline in output. This policy prescription diverges markedly from those of the standard New Keynesian model with exogenous timing of price adjustment, which fails to capture such non-linear dynamics. When confronted with demand or efficiency shocks (e.g., TFP shocks), our findings indicate that the optimal policy in the menu-cost model involves a commitment to full price stability, akin to the standard Calvo model with exogenous timing of price changes.

In sum, our research underscores the importance of an aggressive anti-inflationary policy by the central bank in the face of large shocks. By committing to policies that curb inflation and stabilize the repricing frequency, the central bank can deliver a more favorable macroe-

¹⁸There is no reason for downward nominal rigidity in our framework, up to first order. Price increases are more frequent than price decreases even under zero inflation only due to the asymmetry of the profit function, a higher-order feature. It raises the incentives of firms to adjust their relative prices if it is below the optimal reset price than when it is above it.

conomic outcome. Our analysis is confined to the case of nominal price rigidities, we leave the interaction with wage rigidities for future research.

References

- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll (2021) “Income and wealth distribution in macroeconomics: A continuous-time approach,” *The Review of Economic Studies*, Vol. 89, pp. 45–86.
- Adam, Klaus and Henning Weber (2019) “Optimal Trend Inflation,” *American Economic Review*, Vol. 109, pp. 702–737.
- Adjemian, Stéphane, Houtan Bastani, Michel Juillard, Frédéric Karamé, Ferhat Mihoubi, Willi Mutschler, Johannes Pfeifer, Marco Ratto, Sébastien Villemot, and Normann Rion (2023) “Dynare: Reference Manual Version 5,” working papers, HAL.
- Alexandrov, Andrey (2020) “The Effects of Trend Inflation on Aggregate Dynamics and Monetary Stabilization,” crc tr 224 discussion paper series, University of Bonn and University of Mannheim, Germany.
- Alvarez, Fernando, Martin Beraja, Martín Gonzalez-Rozada, and Pablo Andrés Neumeyer (2019) “From Hyperinflation to Stable Prices: Argentina’s Evidence on Menu Cost Models,” *The Quarterly Journal of Economics*, Vol. 134, pp. 451–505.
- Alvarez, Fernando, Hervé Le Bihan, and Francesco Lippi (2016) “The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach,” *American Economic Review*, Vol. 106, pp. 2817–51.
- Alvarez, Fernando, Francesco Lippi, and Aleksei Oskolkov (2021) “The Macroeconomics of Sticky Prices with Generalized Hazard Functions*,” *The Quarterly Journal of Economics*, Vol. 137, pp. 989–1038.
- Alvarez, Fernando and Pablo Andres Neumeyer (2020) “The Passthrough of Large Cost Shocks in an Inflationary Economy,” in Gonzalo Castex, Jordi Galí, and Diego Saravia eds. *Changing Inflation Dynamics, Evolving Monetary Policy*, Vol. 27 of Central Banking, Analysis, and Economic Policies Book Series: Central Bank of Chile, Chap. 2, pp. 007–048.
- Auclert, Adrien, Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub (2024) “New Pricing Models, Same Old Phillips Curves?” *The Quarterly Journal of Economics*, Vol. 139, pp. 121–186.
- Auer, Raphael, Ariel Burstein, and Sarah M Lein (2021) “Exchange Rates and Prices: Evidence from the 2015 Swiss Franc Appreciation,” *American Economic Review*, Vol. 111, pp. 652–686.
- Barro, Robert J. (1972) “A Theory of Monopolistic Price Adjustment,” *Review of Economic Studies*, Vol. 39, pp. 17–26.
- Benigno, Pierpaolo and Gauti Eggertsson (2023) “It’s Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve,” NBER Working Papers 31197, National Bureau of Economic Research, Inc.
- Berger, David and Joseph Vavra (2019) “Shocks versus Responsiveness: What Drives Time-Varying Dispersion?” *Journal of Political Economy*, Vol. 127, pp. 2104–2142.
- Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J Sargent (2021) “Inequality, Business Cycles, and Monetary-Fiscal Policy,” *Econometrica*, Vol. 89, pp. 2559–2599.

- Blanchard, Olivier and Jordi Galí (2007) “The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s so Different from the 1970s?” in *International Dimensions of Monetary Policy*: National Bureau of Economic Research, Inc, pp. 373–421.
- Blanco, Andres, Corina Boar, Callum J Jones, and Virgiliu Midrigan (2024a) “Nonlinear Inflation Dynamics in Menu Cost Economies,” NBER Working Papers 32094, National Bureau of Economic Research.
- Blanco, Andrés (2021) “Optimal Inflation Target in an Economy with Menu Costs and a Zero Lower Bound,” *American Economic Journal: Macroeconomics*, Vol. 13, pp. 108–141.
- Blanco, Andrés, Corina Boar, Callum J. Jones, and Virgiliu Midrigan (2024b) “The Inflation Accelerator,” NBER Working Papers 32531, National Bureau of Economic Research, Inc.
- Boppart, Timo, Per Krusell, and Kurt Mitman (2018) “Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative,” *Journal of Economic Dynamics and Control*, Vol. 89, pp. 68–92.
- Caballero, Ricardo and Eduardo Engel (1993) “Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy,” *The Review of Economic Studies*, Vol. 60, pp. 95–119.
- Calvo, Guillermo A. (1983) “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, Vol. 12, pp. 383 – 398.
- Caratelli, Daniele and Basil Halperin (2023) “Optimal monetary policy under menu costs,” unpublished manuscript.
- Cavallo, Alberto, Francesco Lippi, and Ken Miyahara (2023) “Large Shocks Travel Fast,” NBER Working Papers 31659, National Bureau of Economic Research, Inc.
- Cerrato, Andrea and Giulia Gitti (2023) “Inflation Since COVID: Demand or Supply,” unpublished manuscript.
- Costain, James and Anton Nakov (2011) “Distributional Dynamics under Smoothly State-Dependent Pricing,” *Journal of Monetary Economics*, Vol. 58, pp. 646 – 665.
- (2019) “Logit Price Dynamics,” *Journal of Money, Credit and Banking*, Vol. 51, pp. 43–78.
- Dávila, Eduardo and Andreas Schaab (2022) “Optimal Monetary Policy with Heterogeneous Agents: A Timeless Ramsey Approach,” *Working Paper*.
- Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999) “State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output,” *The Quarterly Journal of Economics*, Vol. 114, pp. 655–690.
- Gagnon, Etienne (2009) “Price Setting during Low and High Inflation: Evidence from Mexico,” *The Quarterly Journal of Economics*, Vol. 124, pp. 1221–1263.
- Galí, Jordi (2008) *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*: Princeton University Press.
- Golosov, Mikhail and Robert E. Lucas (2007) “Menu Costs and Phillips Curves,” *Journal of Political Economy*, Vol. 115, pp. 171–199.

- González, Beatriz, Galo Nuño, Dominik Thaler, and Silvia Albrizio (2024) “Firm Heterogeneity, Capital Misallocation and Optimal Monetary Policy,” Working Paper Series 2890, European Central Bank.
- Karadi, Peter and Adam Reiff (2019) “Menu Costs, Aggregate Fluctuations, and Large Shocks,” *American Economic Journal: Macroeconomics*, Vol. 11, pp. 111–146.
- Le Grand, Francois, Alais Martin-Baillon, and Xavier Ragot (2022) “Should monetary policy care about redistribution? Optimal monetary and fiscal policy with heterogeneous agents,” Technical report, Paris School of Economics.
- Midrigan, Virgiliu (2011) “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, Vol. 79, pp. 1139–1180.
- Montag, Hugh and Daniel Villar (2023) “Price-Setting During the Covid Era,” *FEDS Notes*.
- Nakamura, Emi and Jón Steinsson (2008) “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *The Quarterly Journal of Economics*, Vol. 123, pp. 1415–1464.
- (2010) “Monetary Non-neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, Vol. 125, pp. 961–1013.
- Nakov, Anton and Carlos Thomas (2014) “Optimal Monetary Policy with State-Dependent Pricing,” *International Journal of Central Banking*, Vol. 36.
- Nuño, Galo and Carlos Thomas (2022) “Optimal Redistributive Inflation,” *Annals of Economics and Statistics*, pp. 3–63.
- Ragot, Xavier (2019) “Managing Inequality over the Business Cycles: Optimal Policies with Heterogeneous Agents and Aggregate Shocks,” 2019 Meeting Papers 1090, Society for Economic Dynamics.
- Sheshinski, Eytan and Yoram Weiss (1977) “Inflation and Costs of Price Adjustment,” *Review of Economic Studies*, Vol. 44, pp. 287–303.
- Smets, Frank and Rafael Wouters (2007) “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *The American Economic Review*, Vol. 97, pp. 586–606.
- Smirnov, Danila (2022) “Optimal Monetary Policy in HANK,” Technical report.
- Taylor, John B (1993) “Discretion versus Policy Rules in Practice,” in *Carnegie-Rochester Conference Series on Public Policy*, Vol. 39, pp. 195–214, Elsevier.
- Vavra, Joseph (2014) “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *The Quarterly Journal of Economics*, Vol. 129, pp. 215–258.
- Woodford, Michael (2003) *Interest and Prices: Foundations of a Theory of Monetary Policy*: Princeton University Press.
- (2009) “Information-Constrained State-Dependent Pricing,” *Journal of Monetary Economics*, Vol. 56, pp. S100–S124.
- Zbaracki, Mark J., Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen (2004) “Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets,” *The Review of Economics and Statistics*, Vol. 86, pp. 514–533.

A Optimality condition of the reset price

If the value function is convex, the optimal reset price is fully characterized by the first order condition.¹⁹ Here we reformulate the value function conveniently and derive the first order condition reported in the text. In the main text the value function was

$$\begin{aligned} V_t(x) &= \Pi(x, p_t^*, w_t, A_t) \\ &\quad + \mathbb{E}_t \left[\Lambda_{t,t+1} (1 - \lambda_{t+1} (x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*)) V_{t+1}(x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*) \right] \\ &\quad + \mathbb{E}_t \left[\Lambda_{t,t+1} \lambda_{t+1} (x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*) (V_{t+1}(0) - \eta w_{t+1}) \right] \end{aligned}$$

Since the only source of uncertainty is the idiosyncratic shocks we can write:

$$\begin{aligned} V_t(x) &= \Pi(x, p_t^*, w_t, A_t) \\ &\quad + \Lambda_{t,t+1} \int \left[(1 - \lambda_{t+1} (x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*)) V_{t+1}(x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*) \phi(\varepsilon) \right] d\varepsilon \\ &\quad + \Lambda_{t,t+1} (V_{t+1}(0) - \eta w_{t+1}) \int \left[\lambda_{t+1} (x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*) \phi(\varepsilon) \right] d\varepsilon \end{aligned}$$

where $\phi(\cdot)$ denotes the standard normal pdf. The expression $x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*$ denotes the state of the firm tomorrow, conditional on the state x today and tomorrows shock $\sigma_{t+1} \varepsilon_{t+1}$. We can thus replace it by $x' \equiv x - \sigma_{t+1} \varepsilon_{t+1} - \pi_{t+1}^*$ and $\varepsilon_{t+1} \equiv \frac{x - x' - \pi_{t+1}^*}{\sigma_{t+1}}$ and apply a change of variable to the integral:

$$\begin{aligned} V_t(x) &= \Pi(x, p_t^*, w_t, A_t) + \Lambda_{t,t+1} \frac{1}{\sigma_{t+1}} \int \left[(1 - \lambda_{t+1} (x')) V_{t+1}(x') \phi \left(\frac{x - x' - \pi_{t+1}^*}{\sigma_{t+1}} \right) \right] dx' \\ &\quad + \Lambda_{t,t+1} (V_{t+1}(0) - \eta w_{t+1}) \frac{1}{\sigma_{t+1}} \int \left[\lambda_{t+1} (x') \phi \left(\frac{x - x' - \pi_{t+1}^*}{\sigma_{t+1}} \right) \right] dx' \end{aligned}$$

In Golosov Lucas the probability of updating a price λ_{t+1} given any state of the nature is either 1 or 0. In particular it is 0 between the (s,S) bands and 1 otherwise. Therefore we can restrict the bounds of the first integral to the (s,S) bands. Furthermore we can replace the second integral, which is the probability mass of updating the price, by 1 minus the

¹⁹We verify convexity ex post.

probability mass of not updating the price.

$$\begin{aligned}
V_t(x) &= \Pi(x, p_t^*, w_t, A_t) + \frac{\Lambda_{t,t+1}}{\sigma_{t+1}} \int_{s_{t+1}}^{S_{t+1}} \left[V_{t+1}(x') \phi \left(\frac{x - x' - \pi_{t+1}^*}{\sigma_{t+1}} \right) \right] dx' \\
&\quad + \Lambda_{t,t+1} (V_{t+1}(0) - \eta w_{t+1}) \left\{ 1 - \frac{1}{\sigma_{t+1}} \int_{s_{t+1}}^{S_{t+1}} \phi \left(\frac{x - x' - \pi_{t+1}^*}{\sigma_{t+1}} \right) dx' \right\} \\
&= \Pi(x) + \frac{\Lambda_{t,t+1}}{\sigma_{t+1}} \int_{s_{t+1}}^{S_{t+1}} \left[V_{t+1}(x') \phi \left(\frac{x - x' - \pi_{t+1}^*}{\sigma_{t+1}} \right) \right] dx' \\
&\quad + \Lambda_{t,t+1} (V_{t+1}(0) - \eta w_{t+1}) \left\{ 1 - \frac{1}{\sigma_{t+1}} \left[\Phi \left(\frac{x - S_{t+1} - \pi_{t+1}^*}{\sigma_{t+1}} \right) - \Phi \left(\frac{x - s_{t+1} - \pi_{t+1}^*}{\sigma_{t+1}} \right) \right] \right\}
\end{aligned}$$

where in the last line we denote the standard normal cdf by $\Phi(\cdot)$ and suppressed the dependence of Π on aggregate variables in the notation.

Now take the derivative of $V_t(x)$ and reformulate, to get to expression (18) in the main text:

$$\begin{aligned}
V_t'(x) &= \Pi_t'(x) + \frac{1}{\sigma_{t+1}} \Lambda_{t+1} \frac{\partial \int_{s_{t+1}}^{S_{t+1}} V_{t+1}(x') \phi \left(\frac{x - x' - \pi_{t+1}^*}{\sigma_{t+1}} \right) dx'}{\partial x} \\
&\quad + \Lambda_{t+1} \left(\phi \left(\frac{x - S_{t+1} - \pi_{t+1}^*}{\sigma_{t+1}} \right) - \phi \left(\frac{x - s_{t+1} - \pi_{t+1}^*}{\sigma_{t+1}} \right) \right) \left(V_{t+1}(0) - \kappa \frac{w_{t+1}}{A_{t+1}} \right) \\
&= \Pi_t'(x) + \frac{1}{\sigma_{t+1}} \Lambda_{t+1} \int_{s_{t+1}}^{S_{t+1}} V_{t+1}(x') \frac{\partial \phi \left(\frac{x - x' - \pi_{t+1}^*}{\sigma_{t+1}} \right)}{\partial x} dx' \\
&\quad + \Lambda_{t+1} \left(\phi \left(\frac{x - S_{t+1} - \pi_{t+1}^*}{\sigma_{t+1}} \right) - \phi \left(\frac{x - s_{t+1} - \pi_{t+1}^*}{\sigma_{t+1}} \right) \right) \left(V_{t+1}(0) - \kappa \frac{w_{t+1}}{A_{t+1}} \right)
\end{aligned}$$

B Efficiency and welfare analysis

B.1 Efficient and natural level of output

For the welfare analysis, it is instructive to derive the efficient and natural levels of aggregate and product-level output.

Efficient output. The efficient level of output is the solution to a social planning problem. The problem maximizes household welfare in equation (1) subject to (i) the aggregate consumption equation (3), (ii) aggregate labor supply in ($N_t = \int_i N_t(j)$) and (iii) product-level production functions in (9) with respect to product-level consumption and labor ($C_t(j), N_t(j), j \in [0, 1], t = 0, 1, 2, \dots$).

After some algebra, the optimization problem simplifies to

$$\max_{N_t(j)} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{\left[A_t \left(\int N_t(j) \frac{\varepsilon-1}{\varepsilon} di \right)^{\frac{\varepsilon}{\varepsilon-1}} \right]^{1-\gamma}}{1-\gamma} - v \int N_t(j) di,$$

subject to $\int N_t(j)di = N_t$. The solution implies that the efficient output fluctuates with aggregate productivity, but is independent of demand shocks as well as of cost-push shocks. In particular, the efficient level of output is

$$Y_t^e = C_t^e = A_t N_t^e = v^{-1/\gamma} A_t^{1/\gamma}. \quad (26)$$

For our parametrization, $v = 1$ and $\gamma = 1$, we thus have that

$$\begin{aligned} N_t^e &= 1 \\ C_t^e &= A_t \end{aligned} \quad (27)$$

The efficient real interest rate is implicitly defined by the Euler equation after substituting in efficient consumption:

$$r_t^e = -\log \beta - \gamma(1 - \rho_A) \log A_t$$

The efficient labor supply is equal across products and the efficient product-level consumption varies across products j inversely proportional to the product-level quality, in particular

$$\begin{aligned} N_t^e(j) &= N_t^e \\ C_t^e(j) &= \frac{A_t N_t^e}{A_t(j)}. \end{aligned}$$

Natural output. The natural level of output is the counterfactual output with flexible prices. Under flexible prices, firms maximize their real profit function (11) in each period t by choosing

$$\frac{P_t^n(j)}{A_t(j)P_t^n} = \frac{\varepsilon}{\varepsilon - 1} (1 - \tau_t) \frac{w_t^n}{A_t}.$$

The expression implies that the quality-adjusted relative price is homogeneous across products j . Together with the definition of the price-level (5), this implies that the natural level of the quality-adjusted log relative price is zero ($p_t(j) = 0$). Or equivalently, the natural level of relative price is equal to the quality: $P_t^n(j)/P_t^n = A_t(j)$. The product-level demand function and the unit quality-adjusted relative price implies that product-level natural consumption is inversely proportional to the quality of product j :

$$C_t^n(j) = \frac{1}{A_t(j)} C_t^n.$$

Furthermore, the natural real wage, output and labor are given by the following closed-form

expressions:

$$\begin{aligned} w_t^n &= A_t \frac{\varepsilon - 1}{\varepsilon} \frac{1}{1 - \tau_t} \\ Y_t^n &= C_t^n = \left(\frac{w_t^n}{v} \right)^{1/\gamma} \\ N_t^n &= \frac{Y_t^n}{A_t}. \end{aligned}$$

Notably, the productivity shock affects the natural and the efficient output similarly, but the cost-push (labor-tax) shocks only affect the natural level of output.

B.2 Welfare decomposition

Markups. The real marginal cost of firm j is

$$MC_t(j) = \frac{\partial ((1 - \tau_t)w_t N_t(j))}{\partial Y_t(j)} = \frac{(1 - \tau_t)w_t A_t(j)}{A_t},$$

where we have used that $N_t(j) = A_t(j)Y_t(j)/A_t$.

The (log-) markup $\mu_t(j)$ is the (log-) difference between the relative price and the real marginal cost:

$$\mu_t(j) = \log \frac{P_t(j)}{P_t} - \log \frac{(1 - \tau_t)w_t A_t(j)}{A_t} = \log \frac{P_t(j)}{A_t(j)P_t} - \log \frac{(1 - \tau_t)w_t}{A_t} = p_t(j) - mc_t, \quad (28)$$

where $p_t(j)$ is the quality-adjusted relative price and

$$mc_t \equiv \log (MC_t(j)/A_t(j)) = \log ((1 - \tau_t)w_t/A_t) = \log(1 - \tau_t) - \log A_t + \log v + \gamma \log Y_t,$$

is the ‘aggregate component’ of the marginal cost. Notice that we have employed eq. (6).

Aggregate markup and output. The aggregate (log-) markup $\bar{\mu}_t$ is

$$\bar{\mu}_t = \log \left(\int_i e^{\mu_t(j)(1-\varepsilon)} \right)^{\frac{1}{1-\varepsilon}} = \log \left(\int_i \frac{e^{p_t(j)(1-\varepsilon)}}{e^{mc_t(1-\varepsilon)}} \right)^{\frac{1}{1-\varepsilon}} = \log \frac{1}{e^{mc_t}} \left(\int_i e^{p_t(j)(1-\varepsilon)} \right)^{\frac{1}{1-\varepsilon}} = -mc_t,$$

where we used the observation that the average quality-adjusted relative price is one (eq. 13). Therefore

$$\bar{\mu}_t = -\log(1 - \tau_t) + \log A_t - \log v - \gamma \log Y_t \quad (29)$$

or equivalently

$$e^{\bar{\mu}_t} = \frac{A_t}{v(1 - \tau_t)Y_t^\gamma} \quad (30)$$

expressing the tight relationship between average markup and the output.

Taking into account eq. (26), the efficient output gap can be expressed

$$\log Y_t - \log Y_t^e = \frac{1}{\gamma} (-\log(1 - \tau_t) - \bar{\mu}_t), \quad (31)$$

which is proportional to the negative average markup.

Markup dispersion. The dispersion of the quality-adjusted relative prices (ζ_t^p) is

$$\begin{aligned} \zeta_t^p &= \int_j e^{p_t(j)(-\varepsilon)} g(p_t(j)) = \int_j e^{(\mu_t(j) + mc_t)(-\varepsilon)} g(\mu_t(j) + mc_t) = \\ & \int_j e^{(\mu_t(j) - \bar{\mu}_t)(-\varepsilon)} g(\mu_t(j) - \bar{\mu}_t) \equiv \zeta_t^{\mu - \bar{\mu}} \end{aligned}$$

where $\zeta_t^{\mu - \bar{\mu}}$ is the dispersion of the demeaned markups which equals price dispersion.

Welfare. Finally, consider the case with $\gamma = 1$. We can express the difference between welfare (W_t) from the welfare in the efficient equilibrium (W_t^e) subject to an MIT shock in period 0 as

$$\begin{aligned} W_0 - W_0^e &= \sum_{t=0}^{\infty} \beta^t (U_t - U_t^e) = \sum_{t=0}^{\infty} \beta^t ((\log C_t - N_t) - (\log C_t^e - N_t^e)) \\ &= \sum_{t=0}^{\infty} \beta^t ((\log Y_t - N_t) - (\log Y_t^e - 1)) \\ &= \sum_{t=0}^{\infty} \beta^t \left(-\log(1 - \tau_t) - \bar{\mu}_t - \frac{C_t}{A_t} \int e^{p(-\varepsilon)} g_t(p) dp - \eta g_t^0 + 1 \right) \\ &= \sum_{t=0}^{\infty} \beta^t \left(-\log(1 - \tau_t) - \bar{\mu}_t - \left(\frac{1}{e^{\bar{\mu}_t}(1 - \tau_t)} \zeta_t^{\mu - \bar{\mu}} - 1 \right) - \eta g_t^0 \right) \\ &= \sum_{t=0}^{\infty} \beta^t \left(-\log(1 - \tau_t) - \bar{\mu}_t - \left(\frac{1}{e^{\bar{\mu}_t}(1 - \tau_t)} - 1 \right) - \left(\frac{1}{e^{\bar{\mu}_t}(1 - \tau_t)} (\zeta_t^{\mu - \bar{\mu}} - 1) \right) - \eta g_t^0 \right) \end{aligned}$$

where U_t^e is the utility in the efficient equilibrium and where we have used (27) and $C_t = Y_t$ in line 1 and (31) and (14) in line 2 and (30) and $C_t = Y_t$ and $\gamma = 1$ in line 3. The final expression decomposes welfare into terms related to average markup, markup dispersion, and adjustment costs.

C Additional figures

C.1 Slope of the nonlinear New Keynesian Phillips curve

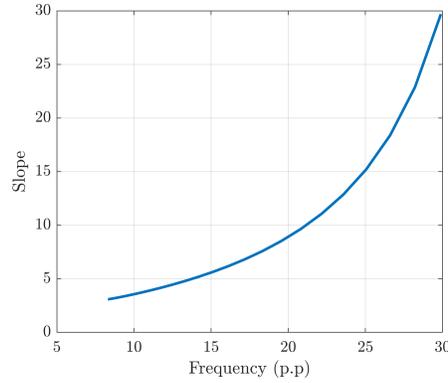


Figure 11: Slope of the nonlinear New Keynesian Phillips curve in the Golosov-Lucas model. *Note:* The figure is produced by computing the impulse responses to monetary policy shocks of different magnitudes and signs and then computing the slope.

C.2 Nonlinear New Keynesian Phillips curve

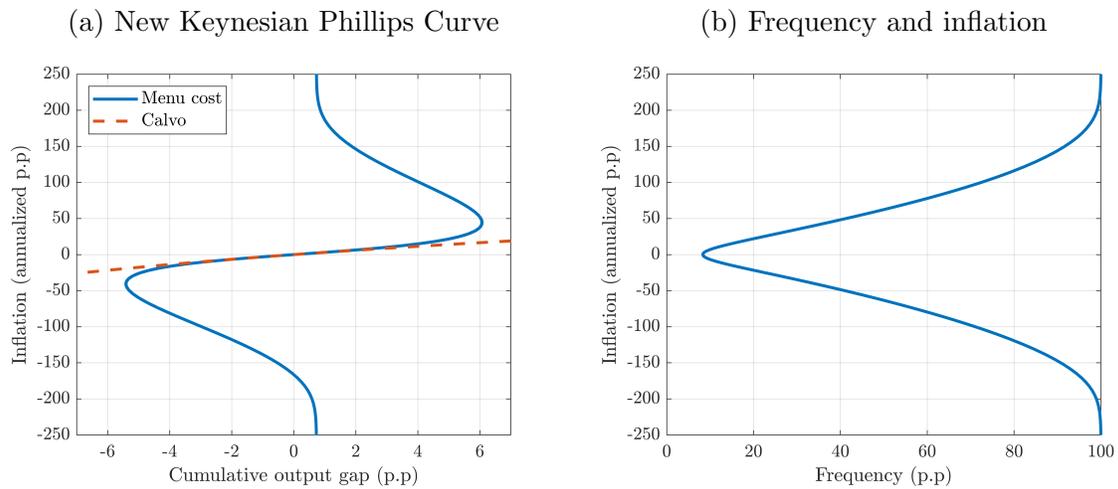


Figure 12: Inflation-output trade-offs in the menu-cost model.

Note: The figure is produced by computing the impulse responses to monetary policy shocks of different magnitudes and signs.

C.3 Impulse responses under a Taylor rule

Cost-push shocks. The solid green lines on the first row of Figure 13 show the responses to a large cost-push shock. The shock is implemented as a persistent decline in the firms' employment subsidy τ_t . The shock size is calibrated to generate a 20% frequency on impact (a frequency increase of $20\% - 8.7\% = 11.3\%$, see panel d), a magnitude that has been documented during the 2022-2023 inflation surge (Montag and Villar, 2023). The exercise assumes that monetary policy follows an inertial Taylor rule (eq. 12). The shock captures some realistic features of the recent inflation surge. First, it generates a large increase in inflation (panel b), characterized by an initial temporary spike followed by a period of persistent inflation. Second, the shock causes a mild downturn as the output gap decreases (panel a).

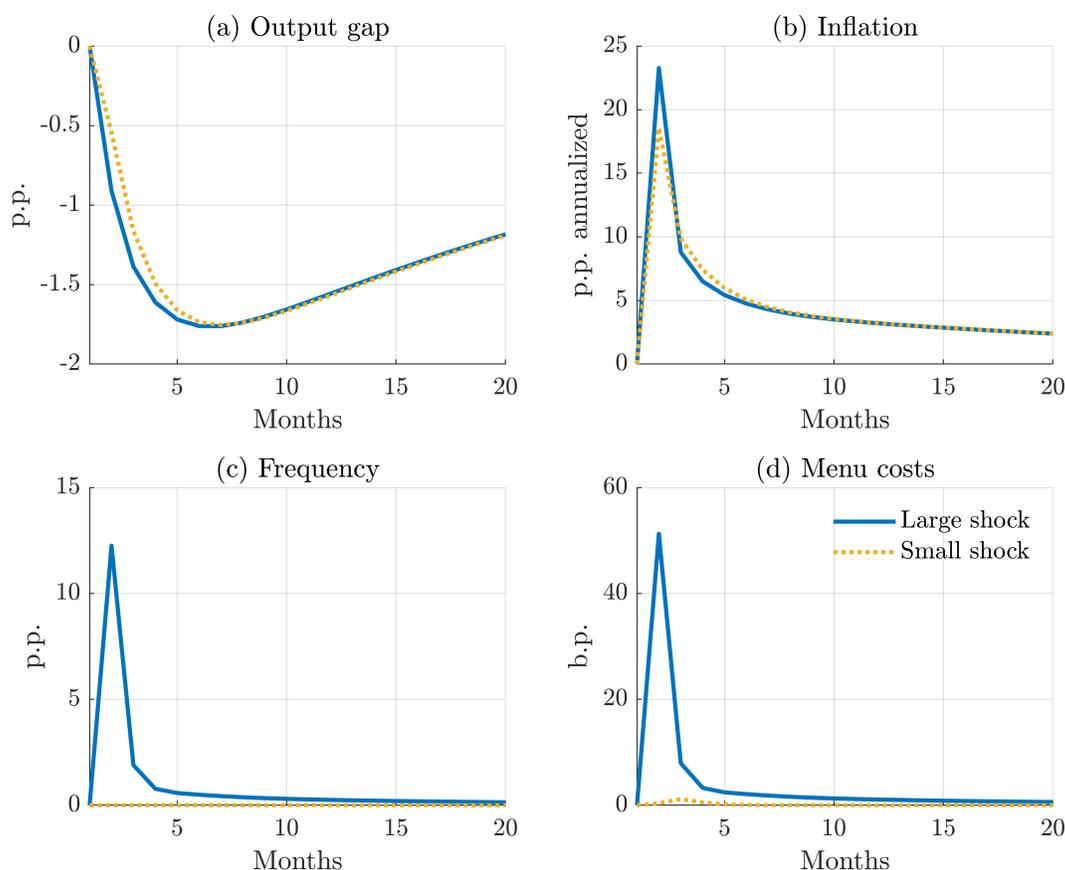


Figure 13: Impulse responses to a large and a small cost-push shock in the menu-cost model. *Note:* all displayed variables except frequency are linearly scaled in the small shock according to the ratio between the large and small shocks.

In contrast, the black dashed lines show the responses to a *small* cost-push shock, scaled up linearly to the size of the large shock. The small shock is a 25 basis point decrease in the annualized employment subsidy. The difference between the figures illustrates the non-linearity of the model. The inflation response is roughly 25% larger after the large shock

than after the linearly scaled small shock. The key reason behind this is the sizable difference between the frequency responses: while the frequency jumps after the large shock, it stays almost unchanged in response to the small shock.²⁰

Comparison with the Calvo model. We compare the results above, for the small shock re-scaled, with the standard Calvo model in Figure 14. We consider two calibrations for the Calvo parameter. The dotted-dashed green line corresponds to a calibration in which the frequency of price adjustment, which in Calvo is constant, matches the frequency of price adjustments of 8.7% in the Golosov-Lucas model. In this case, the response of both inflation and output gap is much attenuated compared to the menu-cost model (solid gray line).

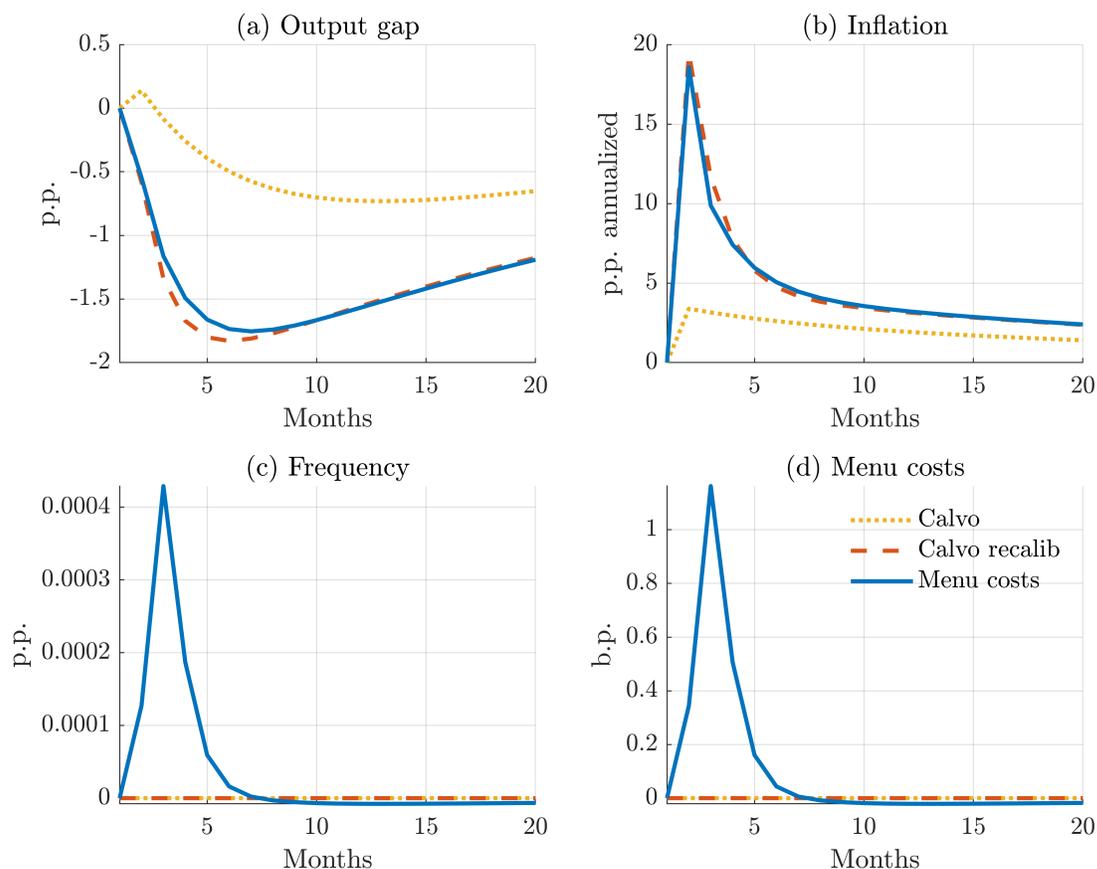


Figure 14: Impulse responses to a small cost-push shock in the menu-cost and in the Calvo model. *Note:* all displayed variables except frequency are linearly-scaled in the small shock according to the ratio between the large and small shocks. The Calvo recalibrated model replicates the slope of the Phillips curve

The dashed blue line in Figure 14 shows the recalibrated Calvo model. The Calvo parameter in that case is set to match the slope of the Phillips curve in the Golosov-Lucas model (discussed below). The recalibrated Calvo model approximately replicates the dynamics of the Golosov-Lucas model in response to small shocks. The fact that the Golosov-Lucas

²⁰The larger initial inflation bout in the case of large shocks is then compensated by a lower inflation path, relative to the small shock counterfactual, from month 3 onwards, and both paths coincide after 9 months.

model can be locally approximated for small inflation levels by a re-calibrated Calvo model with a large frequency was studied by [Auclert et al. \(2024\)](#).

Price dispersion shocks. Figure 15 shows the impulse responses to a shock to the volatility of the idiosyncratic quality shocks (σ_t). The shock raises the dispersion of optimal reset prices and generates a persistent increase in the frequency of price changes ([Vavra, 2014](#)). The observed increase in the dispersion of price changes during the recent inflation surge ([Montag and Villar, 2023](#)) indicates the presence of a similar idiosyncratic dispersion shock either parallel or as a result of the aggregate shocks ([Berger and Vavra, 2019](#)). The size of the large shock is calibrated to generate a 11.3% increase in frequency, which coincides with the peak of the cost-push shock. Notably, the shock is inflationary. This is primarily due to the asymmetry of the value function: firms with too low quality-adjusted prices face high demand, so they are relatively more motivated to *increase* their prices than firms with too high quality-adjusted prices are motivated to *decrease* theirs due to their low demand. The output gap declines primarily as a result of the high price distortions. The differences relative to the linearly scaled small volatility shock (black dashed lines) reveal a sizable non-linearity of the model with respect to this shock. Nonetheless, notice how the impact of the shock on inflation and output is one order of magnitude lower than that of the cost-push shock presented in Figure 13.

TFP and monetary policy shocks. Figures 16 and 17 show analogous impulse responses to a TFP and monetary policy shocks, respectively. These shocks are calibrated to replicate the 11.3% increase in frequency. The model again displays significant nonlinearities in both cases, as the inflation surge is larger, and the output gap surge is smaller, in the case of a large shock.

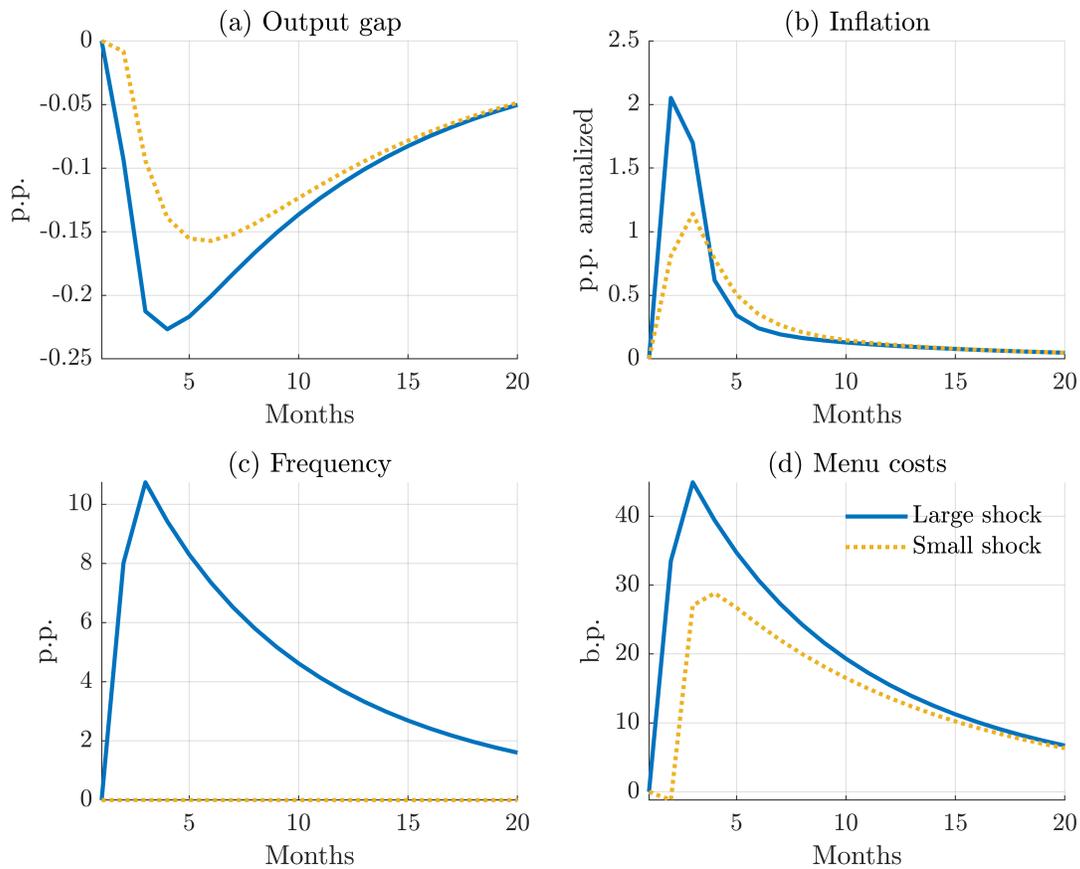


Figure 15: Impulse responses to a large and a small dispersion shock in the menu-cost model. *Note:* all displayed variables except frequency are linearly-scaled in the small shock according to the ratio between the large and small shocks.

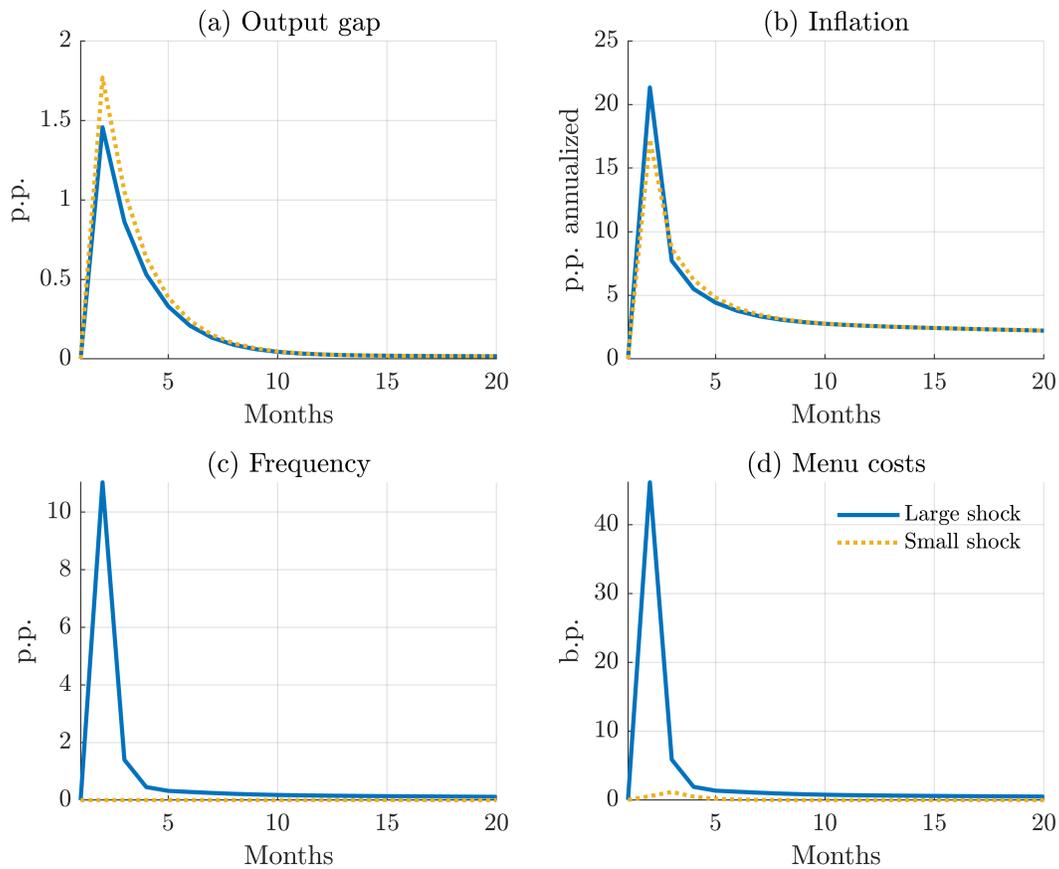


Figure 16: Impulse responses to a large and a small TFP shock in the menu-cost model. *Note:* all displayed variables except frequency are linearly-scaled in the small shock according to the ratio between the large and small shocks.

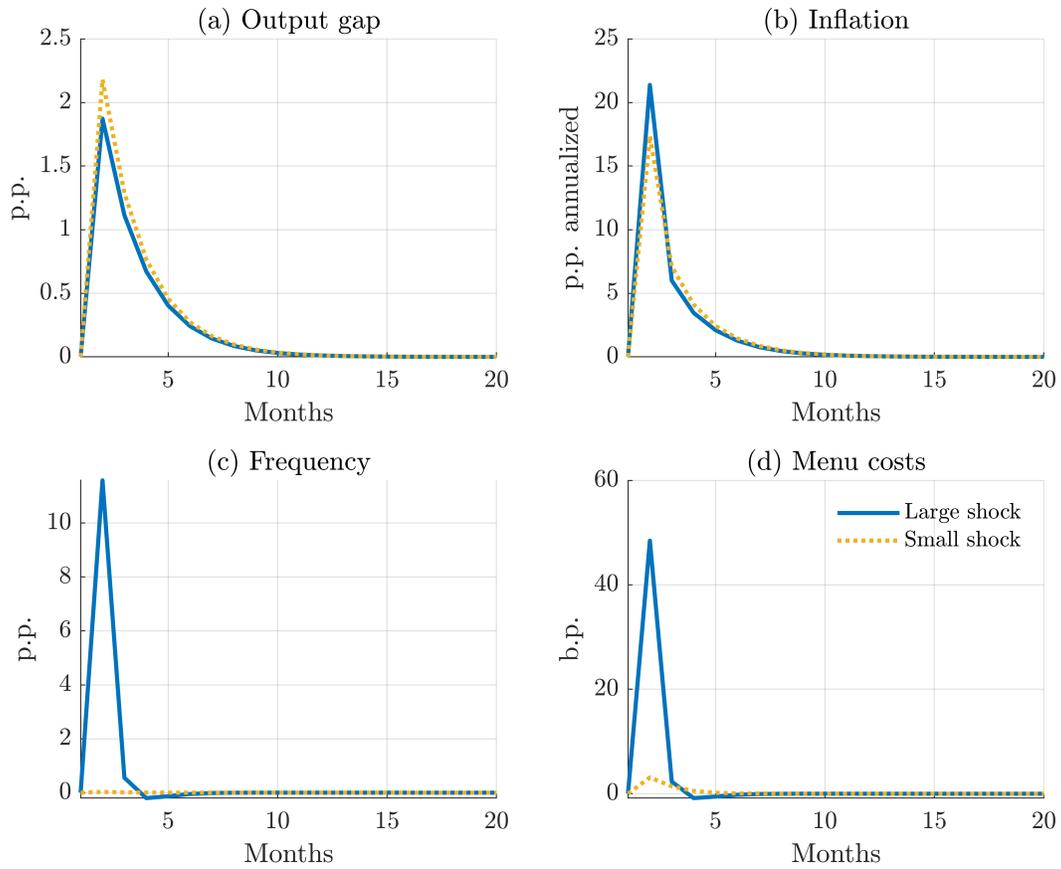


Figure 17: Impulse responses to a large and a small monetary policy shock in the menu-cost model. *Note:* all displayed variables except frequency are linearly-scaled in the small shock according to the ratio between the large and small shocks.

D Computational algorithm

This appendix explains how we convert the infinite dimensional planner’s problem into a finite dimensional one. To do so we employ three distinct computational methods. First, we approximate the distribution and value functions by piece-wise linear functions over a set of nodes. Second, we use endogenous nodes, such that both the (s,S) bands and the optimal reset price are “on the grid”. Third, given the approximation from step 1, we evaluate integrals analytically. Method one makes the problem finite dimensional. Methods two and three ensure that the approximation is highly accurate, smooth and computationally efficient.

To explain this procedure, the rest of the appendix proceeds in four steps. First, we define some auxiliary functions that will be useful later on. Second, we transform the private equilibrium conditions that contain the value and distribution functions to prepare them for the endogenous grid. Third, we approximate the value and distribution functions in the such transformed equations by a piece-wise linear function. Fourth, we evaluate the integrals analytically. Together with the private equilibrium conditions not containing value and distribution functions, the result is a discrete set of equations that can conveniently be written down in matrix notation.

Once we have converted the planner’s infinite-dimensional problem into a finite-dimension problem in this way, we can then derive the planner’s first order conditions. For this step we use symbolic differentiation, and in particular dynare’s Ramsey comand. The resulting set of first order conditions is then solved in the sequence space under perfect foresight. Here we employ a standard Newton method, and in particular we use dynare’s perfect foresight solver command. To determine the appropriate initial and terminal conditions, and an initial guess for the transition paths, we need the non-stochastic steady state of the model. For this step we determine the steady state of the private equilibrium conditional on a particular value of the policy instrument π using a standard Newton based solution method. We then use this function and exploit the linearity of the first order conditions to convert the high-dimensional problem of solving for the SS into a one dimensional problem, which is easily solved with a newton solver. This last step is performed by dynare’s steady command. That is, we have to manually convert the problem into a finite-dimension problem and find the steady state conditional on a policy, the rest is conveniently automated in dynare. We now explain how we convert the problem.

D.1 Preliminaries

To begin with, let us normalize the variable x_t as

$$\mathbf{x}_t = \begin{cases} \frac{x_t}{s_t} & \text{if } x_t < 0 \\ \frac{x_t}{S_t} & \text{else} \end{cases} \quad (32)$$

Under this normalization, the optimal price is at $\mathbf{x}_t = 0$, the upper (s, S) band at $\mathbf{x}_t = 1$ and the lower (s, S) band at $\mathbf{x}_t = -1$. This will later allow us to have all critical points (s, S, p^*) on the grid.

The law of motion of \mathbf{x}_t conditional on not updating can be derived from $x_t = x_{t-1} -$

$\sigma\varepsilon - \pi_t^*$:

$$\mathbf{x}_t = \begin{cases} \frac{x_t}{S_t} = \frac{x_{t-1} - \sigma\varepsilon - \pi_t^*}{S_t} = \begin{cases} \frac{x_{t-1} - \sigma\varepsilon - \pi_t^*}{S_{t-1}} \frac{S_{t-1}}{S_t} = \mathbf{x}_{t-1} \frac{S_{t-1}}{S_t} - \frac{\sigma\varepsilon + \pi_t^*}{S_t} & \text{if } \mathbf{x}_t > 0, \text{ if } \mathbf{x}_{t-1} > 0 \\ \frac{x_{t-1} - \sigma\varepsilon - \pi_t^*}{s_{t-1}} \frac{s_{t-1}}{S_t} = \mathbf{x}_{t-1} \frac{s_{t-1}}{S_t} - \frac{\sigma\varepsilon + \pi_t^*}{S_t} & \text{if } \mathbf{x}_t > 0, \text{ if } \mathbf{x}_{t-1} < 0 \end{cases} \\ \frac{x_t}{s_t} = \frac{x_{t-1} - \sigma\varepsilon - \pi_t^*}{s_t} = \begin{cases} \frac{x_{t-1} - \sigma\varepsilon - \pi_t^*}{S_{t-1}} \frac{S_{t-1}}{-s_t} = \mathbf{x}_{t-1} \frac{S_{t-1}}{s_t} - \frac{\sigma\varepsilon + \pi_t^*}{s_t} & \text{if } \mathbf{x}_t < 0, \text{ if } \mathbf{x}_{t-1} > 0 \\ \frac{x_{t-1} - \sigma\varepsilon - \pi_t^*}{s_{t-1}} \frac{s_{t-1}}{s_t} = \mathbf{x}_{t-1} \frac{s_{t-1}}{s_t} - \frac{\sigma\varepsilon + \pi_t^*}{s_t} & \text{if } \mathbf{x}_t < 0, \text{ if } \mathbf{x}_{t-1} < 0 \end{cases} \end{cases} \quad (33)$$

Now let us define functions which we will use in the next sections to redefine the value and distribution functions. For compactness, let us adopt the notation where $\hat{s}_t(\mathbf{x}_t)$ picks the respective extremes (s, S) depending on the value of \mathbf{x}_t following 32. For brevity, at times we will drop the dependence on \mathbf{x}_t and just write \hat{s}_t .

Solving 33 for \mathbf{x}_t , \mathbf{x}_{t-1} and ε , we can then define the following functions, which we use later on to shorten notation.

$$\mathbf{x}_t = \mathbf{x}_{t-1} \frac{\hat{S}_{t-1}}{\hat{S}_t} - \frac{\sigma\varepsilon + \pi_t^*}{\hat{S}_t} \equiv \hat{f}(\mathbf{x}_{t-1}) \quad (34)$$

$$\mathbf{x}_{t-1} = \mathbf{x}_t \frac{\hat{S}_t}{\hat{S}_{t-1}} + \frac{\sigma\varepsilon + \pi_t^*}{\hat{S}_{t-1}} \equiv f(\mathbf{x}_t) \quad (35)$$

$$\varepsilon = \frac{(\hat{S}_{t-1}\mathbf{x}_{t-1} - \hat{S}_t\mathbf{x}_t) - \pi_t^*}{\sigma} \equiv h(\mathbf{x}_{t-1}, \mathbf{x}_t) \quad (36)$$

D.2 Approximating the distribution and value functions by piecewise linear functions on an endogenous grid

Now we redefine the value and distribution functions over the variable \mathbf{x} and approximate them by piece-wise linear functions. Our starting point infinite dimensional problem of the planner as laid out in section 5.1. We consider each of the equations containing the distribution and value functions in turn.

D.2.1 Distribution

Recall the distribution function

$$g_t(x) = (1 - \lambda_t(x)) \int g_{t-1}(x + \sigma_t\varepsilon + \pi_t^*) d\xi(\varepsilon) + \delta(x) \int \lambda_t(\tilde{x}) \left(\int g_{t-1}(\tilde{x} + \sigma_t\varepsilon + \pi_t^*) d\xi(\varepsilon) \right) d\tilde{x}$$

with

$$\int_{s_t}^{S_t} g_t(x) dx = 1 \quad (37)$$

where $\delta(x)$ is the Diract Delta function that captures the mass point of those firms who update their prices.

We split the distribution into the continuous distribution of agents who do not update their prices plus a mass point of updaters at $x = 0$ (this is already reflected in section 5.1):

$$g_t^c(x) = (1 - \lambda_t(x)) \int g_{t-1}^c(x + \sigma\varepsilon + \pi_t^*) d\xi(\varepsilon)$$

$$g_t^0 = \int \lambda_t(\tilde{x}) \int g_{t-1}^c(\tilde{x} + \sigma\varepsilon - \pi_t^*) d\xi(\varepsilon) d\tilde{x}$$

Furthermore, by 37 we can re-express the second equation as:

$$g_t^0 = 1 - \int_{s_t}^{S_t} g_t^c(x) dx$$

Now rewrite it using the newly defined re-normalized \mathbf{x} where $x = \mathbf{x}\hat{s}_t$ as in equation 32: define $g_t(\mathbf{x}\hat{s}_t) \equiv \mathbf{g}_t(\mathbf{x})$ and $g_t^c(\mathbf{x}\hat{s}_t) \equiv \mathbf{g}_t^c(\mathbf{x})$ and - with a slight abuse of notation - $\lambda_t(\mathbf{x}\hat{s}_t) \equiv \lambda_t(\mathbf{x})$ and write

$$\mathbf{g}_t^c(\mathbf{x}) = (1 - \lambda_t(\mathbf{x})) \int \mathbf{g}_{t-1}(f(\mathbf{x})) d\xi(\varepsilon)$$

$$\mathbf{g}_t^0 = 1 - \int_{-1}^1 \mathbf{g}_t^c(\mathbf{x}) \hat{s}_t(\mathbf{x}) d\mathbf{x}$$

Note that in the second equation we have applied a change of variable to the integral. In particular, we have used the following substitution.

$$\int_{s_t}^{S_t} g_t^c(x) dx = \int_{s_t}^{S_t} g_t^c(\mathbf{x}\hat{s}_t(\mathbf{x})) d\mathbf{x}\hat{s}_t(\mathbf{x})$$

$$= \int_{s_t}^{S_t} \mathbf{g}_t^c(\mathbf{x}) d\mathbf{x}\hat{s}_t(\mathbf{x}) = \int_{s_t/\hat{s}_t(\mathbf{x})}^{S_t/\hat{s}_t(\mathbf{x})} \hat{s}_t(\mathbf{x}) \mathbf{g}_t^c(\mathbf{x}) d\mathbf{x} = \int_{-1}^1 \hat{s}_t(\mathbf{x}) \mathbf{g}_t^c(\mathbf{x}) d\mathbf{x}$$

Applying a change of variable to the integral using the definition $\mathbf{x}' \equiv f(\mathbf{x})$, we re-write the first equation as follows:

$$\mathbf{g}_t^c(\mathbf{x}) = (1 - \lambda_t(\mathbf{x})) \int \frac{\hat{s}_{t-1}(\mathbf{x}')}{\sigma} \mathbf{g}_{t-1}(\mathbf{x}') \phi(h(\mathbf{x}', \mathbf{x})) d\mathbf{x}'$$

where $\phi(\cdot)$ is the standard normal pdf. The end of period distribution has mass only between the (s,S) bands, i.e. in the range $x \in [-1, 1]$. We can thus restrict the boundaries of the integral accordingly.

$$\mathbf{g}_t^c(\mathbf{x}) = (1 - \lambda_t(\mathbf{x})) \int_{-1}^1 \frac{\hat{s}_{t-1}(\mathbf{x}')}{\sigma} \mathbf{g}_{t-1}(\mathbf{x}') \phi(h(\mathbf{x}', \mathbf{x})) d\mathbf{x}'$$

So far we have only rewritten the law of motion of the firm distribution g_t . Now we introduce the only approximation we rely on for g_t .

We approximate \mathbf{g}^c by a piece-wise linear function with equally spaced nodes $\mathbf{x}_1, \dots, \mathbf{x}_I = -1, \dots, 0, \dots, 1$ with $\mathbf{g}_t^c(\mathbf{x} | \mathbf{x}_i < \mathbf{x} < \mathbf{x}_{i+1}) \approx \mathbf{g}_t^c(\mathbf{x}_i) + \frac{\mathbf{x} - \mathbf{x}_i}{\mathbf{x}_{i+1} - \mathbf{x}_i} \frac{g_{t-1}^c(\mathbf{x}_{i+1}) - g_{t-1}^c(\mathbf{x}_i)}{\mathbf{x}_{i+1} - \mathbf{x}_i}$.

Note that the auxiliary grid for \mathbf{x} is exogenous. However, this exogenous auxiliary grid defines an *endogenous grid* for $x = \hat{s}_t \mathbf{x}$, which, at each t , exactly spans the area between the (s,S) bands and has a node at 0.

From now on, with a slight abuse of notation, \mathbf{g}_t^c denotes the piece-wise linear approximated function and $\mathbf{g}_t^c(\mathbf{x}_i < \mathbf{x} < \mathbf{x}_{i+1})$ denotes a linear piece of it. Thus, the functions are approximated as

$$\mathbf{g}_t^c(\mathbf{x}) = (1 - \lambda_t(\mathbf{x})) \left[\sum_{i=1}^{I-1} \int_{\mathbf{x}_i}^{\mathbf{x}_{i+1}} \frac{\hat{s}_{t-1}(\mathbf{x}')}{\sigma} \mathbf{g}_{t-1}^c(\mathbf{x}_i < \mathbf{x}' < \mathbf{x}_{i+1}) \phi(h(\mathbf{x}', \mathbf{x})) d\mathbf{x}' + \mathbf{g}_{t-1}^0 \phi(h(0, \mathbf{x})) \right]$$

$$\mathbf{g}_t^0 = 1 - \sum_{i=1}^{I-1} \int_{\mathbf{x}_i}^{\mathbf{x}_{i+1}} \mathbf{g}_t^c(\mathbf{x}_i < \mathbf{x} < \mathbf{x}_{i+1}) \hat{s}_t(\mathbf{x}) d\mathbf{x}$$

Notice that in the expressions above, in each interval $\mathbf{x}_i < \mathbf{x} < \mathbf{x}_{i+1}$ the integrand is continuous since \mathbf{x} and \mathbf{x}' are of constant sign.

D.2.2 Other Aggregation Equations

The private equilibrium conditions contain two further aggregation equations that contain the function $\mathbf{g}(\cdot)$, for which we use the piecewise linear approximation of $\mathbf{g}^c(\cdot)$.

Let us recall the aggregate price index and the labor market clearing condition

$$e^{p_t^*(\epsilon-1)} = \int e^{x(1-\epsilon)} g_t(x) dx$$

$$N_t = \frac{C_t}{A_t} e^{p_t^*(-\epsilon)} \int e^{x(-\epsilon)} g_t(x) d(x) + \eta \int \lambda_t(x + p_t^* - \sigma_t \epsilon - \pi_t) g_{t-1}(x) d(x)$$

which we approximate as follows, after the change of variable to \mathbf{x}

$$e^{p_t^*(\epsilon-1)} = \sum_{i=1}^{I-1} \int_{\mathbf{x}_i}^{\mathbf{x}_{i+1}} e^{x(1-\epsilon)} \mathbf{g}_t^c(\mathbf{x}_{i-1} < \mathbf{x} < \mathbf{x}_{i+1}) \hat{s}_t(\mathbf{x}) d\mathbf{x} + \mathbf{g}_t^0$$

$$N_t = \frac{C_t}{A_t} e^{p_t^*(-\epsilon)} \sum_{i=1}^{I-1} \int_{\mathbf{x}_i}^{\mathbf{x}_{i+1}} (e^{x(-\epsilon)} \mathbf{g}_t^c(\mathbf{x}_{i-1} < \mathbf{x} < \mathbf{x}_{i+1}) \hat{s}_t(\mathbf{x}) d\mathbf{x} + \mathbf{g}_{t-1}^0) + \eta \mathbf{g}_{t-1}^0$$

D.2.3 Value Function

Recall the value function

$$\begin{aligned} V_t(x) = & \Pi_t(x) + \Lambda_{t+1} \int (1 - \lambda_{t+1}(x - \sigma\varepsilon - \pi_t^*)) V_{t+1}(x - \sigma\varepsilon - \pi_t^*) d\xi(\varepsilon) \\ & + \Lambda_{t+1} (V_{t+1}(0) - \eta w_{t+1}) \int \lambda_{t+1}(x - \sigma\varepsilon - \pi_t^*) d\xi(\varepsilon) \end{aligned}$$

Now write it using the newly defined re-normalized \mathbf{x} , with $\mathbf{V}(\mathbf{x}) \equiv V(x/\hat{s})$

$$\begin{aligned} \mathbf{V}_t(\mathbf{x}) = & \Pi_t(\mathbf{x}) + \Lambda_{t+1} \int \left(1 - \lambda_{t+1}(\hat{f}(\mathbf{x}))\right) \mathbf{V}_{t+1}(\hat{f}(\mathbf{x})) d\xi(\varepsilon) \\ & + \Lambda_{t+1} (\mathbf{V}_{t+1}(0) - \eta w_{t+1}) \int \lambda_{t+1}(\hat{f}(\mathbf{x})) d\xi(\varepsilon) \end{aligned}$$

Note that by definition we have that $\mathbf{V}_t(0) - \eta \frac{w_{t+1}}{A_{t+1}} = \mathbf{V}_t(-1) = \mathbf{V}_t(1)$ and $\mathbf{V}'_t(0) = 0$. The first 2 equalities are straightforward. The latter is discussed in the next subsection.

After the change of variable $\mathbf{x}' = \hat{f}(\mathbf{x})$, we can rewrite it as

$$\begin{aligned} \mathbf{V}_t(\mathbf{x}) = & \Pi_t(\mathbf{x}) + \Lambda_{t+1} \frac{1}{\sigma} \int \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) \mathbf{V}_{t+1}(\mathbf{x}') \phi(h(\mathbf{x}, \mathbf{x}')) d\mathbf{x}' \\ & + \Lambda_{t+1} (\mathbf{V}_{t+1}(0) - \eta w_{t+1}) \frac{1}{\sigma} \int \hat{s}_{t+1}(\mathbf{x}') \lambda_{t+1}(\mathbf{x}') \phi(h(\mathbf{x}, \mathbf{x}')) d\mathbf{x}' \end{aligned}$$

Since the price updating probability $\lambda_{t+1}(\mathbf{x})$ is 1 for any \mathbf{x} outside the (s,S) bands, we can restrict the first integral to the range $[-1, 1]$. The last term in the second line (which captures the probability of updating a price tomorrow, given the current state) can be replaced by 1 minus the probability of not updating the price tomorrow. The latter is given by an integral over the range $[-1, 1]$. So we write:

$$\begin{aligned} \mathbf{V}_t(\mathbf{x}) = & \Pi_t(\mathbf{x}) + \Lambda_{t+1} \frac{1}{\sigma} \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) \mathbf{V}_{t+1}(\mathbf{x}') \phi(h(\mathbf{x}, \mathbf{x}')) d\mathbf{x}' \\ & + \Lambda_{t+1} (\mathbf{V}_{t+1}(0) - \eta w_{t+1}) \left(1 - \frac{1}{\sigma} \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) \phi(h(\mathbf{x}, \mathbf{x}')) d\mathbf{x}'\right) \end{aligned}$$

So far we have normalized the support of the value function. Additionally, it is convenient to normalize further the value function itself. We normalize the value function by its maximal value $\mathbf{V}_t(0)$, and denote the normalized value function by $\mathbf{v}_t(\mathbf{x})$: $\mathbf{v}_t(\mathbf{x}) \equiv \mathbf{V}_t(\mathbf{x}) - \mathbf{V}_t(0)$. The expression above can be re-written as:

$$\begin{aligned}
v_t(\mathbf{x}) &\equiv V_t(\mathbf{x}) - V_t(0) = \\
&= \Pi_t(\mathbf{x}) - \Pi_t(0) \\
&+ \frac{1}{\sigma} \Lambda_{t+1} \left(\int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) \left[V_{t+1}(\mathbf{x}') \phi \left(\frac{\mathbf{x} - \mathbf{x}' - \pi_t^*}{\sigma} \right) - V_{t+1}(\mathbf{x}') \phi \left(\frac{0 - \mathbf{x}' - \pi_t^*}{\sigma} \right) \right] d\mathbf{x}' \right) \\
&+ \Lambda_{t+1} \frac{1}{\sigma} \left(- \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) \left[\phi \left(\frac{\mathbf{x} - \mathbf{x}' - \pi_t^*}{\sigma} \right) - \phi \left(\frac{0 - \mathbf{x}' - \pi_t^*}{\sigma} \right) \right] d\mathbf{x}' \right) (V_{t+1}(0) - \eta w_{t+1}) \\
&= \Pi_t(\mathbf{x}) - \Pi_t(0) \\
&+ \frac{1}{\sigma} \Lambda_{t+1} \left(\int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) \left[v_{t+1}(\mathbf{x}') \left(\phi \left(\frac{\mathbf{x} - \mathbf{x}' - \pi_t^*}{\sigma} \right) - \phi \left(\frac{0 - \mathbf{x}' - \pi_t^*}{\sigma} \right) \right) \right] d\mathbf{x}' \right) \\
&+ \Lambda_{t+1} \frac{1}{\sigma} \left(- \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) \left[\phi \left(\frac{\mathbf{x} - \mathbf{x}' - \pi_t^*}{\sigma} \right) - \phi \left(\frac{0 - \mathbf{x}' - \pi_t^*}{\sigma} \right) \right] d\mathbf{x}' \right) (-\eta w_{t+1})
\end{aligned}$$

Following our approach for $\mathbf{g}^c(\cdot)$, we approximate $\mathbf{v}(\cdot)$ by a piece-wise linear function with nodes $\mathbf{x}_1, \dots, \mathbf{x}_I = -1, \dots, 0, \dots, 1$ with $v_t(\mathbf{x} | \mathbf{x}_i < \mathbf{x} < \mathbf{x}_{i+1}) \approx v_t(\mathbf{x}_i) + \frac{\mathbf{x} - \mathbf{x}_i}{\mathbf{x}_{i+1} - \mathbf{x}_i} \frac{v_t(\mathbf{x}_{i+1}) - v_t(\mathbf{x}_i)}{\mathbf{x}_{i+1} - \mathbf{x}_i}$.

From now on, with a slight abuse of notation, \mathbf{v}_t denotes the piece-wise linear approximated function and $v_t(\mathbf{x}_i < \mathbf{x} < \mathbf{x}_{i+1})$ denotes a linear piece of it. Thus, the function is approximated as

$$\begin{aligned}
v_t(\mathbf{x}) &= \Pi_t(\mathbf{x}) - \Pi_t(0) \\
&+ \frac{1}{\sigma} \Lambda_{t+1} \sum_{i=1}^{I-1} \int_{\mathbf{x}_i}^{\mathbf{x}_{i+1}} \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) v_{t+1}(\mathbf{x}_i < \mathbf{x}' < \mathbf{x}_{i+1}) (\phi(h(\mathbf{x}, \mathbf{x}')) - \phi(h(0, \mathbf{x}'))) d\mathbf{x}' \\
&+ \frac{1}{\sigma} \Lambda_{t+1} (-\eta w_{t+1}) \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') (1 - \lambda_{t+1}(\mathbf{x}')) (\phi(h(\mathbf{x}, \mathbf{x}')) - \phi(h(0, \mathbf{x}'))) d\mathbf{x}'
\end{aligned}$$

D.2.4 Optimality condition for reset price

We proceed in the same way for the derivative of the value function. We start with

$$\begin{aligned}
0 = V_t'(0) &= \Pi_t'(0) + \frac{1}{\sigma} \Lambda_{t+1} \int_{s_{t+1}}^{S_{t+1}} V_{t+1}(x') \frac{\partial \phi \left(\frac{x - x' - \pi_{t+1}^*}{\sigma} \right)}{\partial x} \Bigg|_{x=0} dx' \\
&+ \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (V_{t+1}(0) - \eta w_{t+1}),
\end{aligned}$$

where

$$\left. \frac{\partial \phi \left(\frac{x-x'-\pi_{t+1}^*}{\sigma} \right)}{\partial x} \right|_{x=0} = \frac{1}{\sqrt{2\pi}\sigma} \frac{-\pi_{t+1}^* + x'}{\sigma} e^{-\frac{1}{2} \left(\frac{-\pi_{t+1}^* + x'}{\sigma} \right)^2}$$

Change of variable to \mathbf{x}

$$\begin{aligned} 0 &= \Pi_t'(0) + \Lambda_{t+1} \frac{1}{\sigma} \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') \mathbf{v}_{t+1}(\mathbf{x}') \frac{1}{\sqrt{2\pi}\sigma} h(0, \mathbf{x}') e^{-\frac{1}{2}(h(0, \mathbf{x}'))^2} d\mathbf{x}' \\ &+ \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (\mathbf{v}_{t+1}(0) - \eta w_{t+1}), \end{aligned}$$

Now we re-express this in terms of $\mathbf{v}(\mathbf{x})$

$$\begin{aligned} 0 &= \Pi_t'(0) + \Lambda_{t+1} \frac{1}{\sigma} \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') \mathbf{v}_{t+1}(\mathbf{x}') \frac{1}{\sqrt{2\pi}\sigma} h(0, \mathbf{x}') e^{-\frac{1}{2}(h(0, \mathbf{x}'))^2} d\mathbf{x}' \\ &+ \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (-\eta w_{t+1}), \end{aligned}$$

and apply the piece-wise linear approximation of $\mathbf{v}(\mathbf{x})$

$$\begin{aligned} 0 &= \Pi_t'(0) + \Lambda_{t+1} \frac{1}{\sigma} \sum_{i=1}^{I-1} \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') \mathbf{v}_{t+1}(\mathbf{x}_i < \mathbf{x}' < \mathbf{x}_{i+1}) \frac{1}{\sqrt{2\pi}\sigma} h(0, \mathbf{x}') e^{-\frac{1}{2}(h(0, \mathbf{x}'))^2} d\mathbf{x}' \\ &+ \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (-\eta w_{t+1}). \end{aligned}$$

D.3 Solving for Integrals

Let us collect the approximated functions we defined so far. So far we carried around the probability of price updating $\lambda_t(\mathbf{x})$ in our notation. In Golosov Lucas this probability is 0 between the (s,S) bands. Applying this, we get:

$$\begin{aligned} \mathbf{v}_t(\mathbf{x}) &= \Pi_t(\mathbf{x}) - \Pi_t(0) \\ &+ \frac{1}{\sigma} \Lambda_{t+1} \sum_{i=1}^{I-1} \int_{\mathbf{x}_i}^{\mathbf{x}_{i+1}} \hat{s}_{t+1}(\mathbf{x}') \mathbf{v}_{t+1}(\mathbf{x}_i < \mathbf{x}' < \mathbf{x}_{i+1}) (\phi(h(\mathbf{x}, \mathbf{x}')) - \phi(h(0, \mathbf{x}'))) d\mathbf{x}' \\ &+ \frac{1}{\sigma} \Lambda_{t+1} (-\eta w_{t+1}) \int_{-1}^1 \hat{s}_{t+1}(\mathbf{x}') (\phi(h(\mathbf{x}, \mathbf{x}')) - \phi(h(0, \mathbf{x}'))) d\mathbf{x}' \end{aligned} \quad (38)$$

$$\begin{aligned}
0 &= \Pi'_t(0) + \Lambda_{t+1} \frac{1}{\sigma} \int_{-1}^1 \hat{s}_{t+1} \mathbf{v}_{t+1}(\mathbf{x}') \frac{1}{\sqrt{2\pi}\sigma} h(0, \mathbf{x}') e^{-\frac{1}{2}(h(0, \mathbf{x}'))^2} d\mathbf{x}' \\
&+ \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (-\eta w_{t+1})
\end{aligned} \tag{39}$$

$$\mathbf{g}_t^c(\mathbf{x}) = \sum_{i=1}^{I-1} \int_{x_i}^{x_{i+1}} \frac{\hat{s}_{t-1}(\mathbf{x}')}{\sigma} \mathbf{g}_{t-1}^c(x_i < \mathbf{x}' < x_{i+1}) \phi(h(\mathbf{x}', \mathbf{x})) d\mathbf{x}' + \mathbf{g}_{t-1}^0 \phi(h(0, \mathbf{x})) \tag{40}$$

$$\mathbf{g}_t^0 = 1 - \sum_{i=1}^{I-1} \int_{x_i}^{x_{i+1}} \mathbf{g}_t^c(x_i < \mathbf{x}' < x_{i+1}) \hat{s}_t(\mathbf{x}) d\mathbf{x} \tag{41}$$

$$e^{p_i^*(\epsilon-1)} = \sum_{i=1}^{I-1} \int_{x_i}^{x_{i+1}} e^{x(1-\epsilon)} \mathbf{g}_t^c(x_i < \mathbf{x}' < x_{i+1}) \hat{s}_t(\mathbf{x}) d\mathbf{x} + \mathbf{g}_t^0 \tag{42}$$

$$N_t = \frac{C_t}{A_t} e^{p^*(-\epsilon)} \left(\sum_{i=1}^{I-1} \int_{x_i}^{x_{i+1}} e^{x(-\epsilon)} \mathbf{g}_t^c(x_{i-1} < \mathbf{x} < x_{i+1}) \hat{s}_t(\mathbf{x}) d\mathbf{x} + \mathbf{g}_{t-1}^0 \right) + \eta \mathbf{g}_{t-1}^0 \tag{43}$$

The individual integrals in all of these expressions can be computed analytically, since the integrands consist of affine functions multiplied by expressions that have closed form antiderivatives. We now determine the solution of those integrals, equation by equation. Given the coefficients of the affine functions, which depend on the values of $\mathbf{v}_{t+1}(\mathbf{g}_{t-1})$ at the grid points x_i , we can then write the solutions as a function that is linear in the elements of the vector $\mathbf{v}_{t+1}(x_i)$ ($\mathbf{g}_{t-1}(x_i)$). We now explain this for the simple case of the integral in equation 41. The other equations require some more tedious algebra, which we omit for brevity, but are conceptually equivalent.

D.3.1 Mass Point

The integral over an affine function $f(x)$ from x_1 to x_2 is given by

$$\int_{x_1}^{x_2} f(x) dx = \frac{(f(x_1) + f(x_2))}{2} (x_2 - x_1)$$

Thus

$$\sum_{i=1}^{I-1} \int_{x_i}^{x_{i+1}} f(x) dx = \sum_{i=1}^{I-1} \frac{(f(x_i) + f(x_{i+1}))}{2} (x_{i+1} - x_i)$$

Collecting the common terms on the right-hand side we get

$$\sum_{i=1}^{I-1} \int_{x_i}^{x_{i+1}} f(x) dx = \frac{\Delta x}{2} \left(f(x_1) + 2 \sum_{i=2}^{I-1} f(x_i) + f(x_I) \right)$$

Applying this formula to equation 41, which defines the mass point at $x = 0$, and rearranging terms we get

$$\mathbf{g}_t^0 = 1 - \mathbf{e}_t^T \mathbf{g}_t^c \quad (44)$$

where $\mathbf{e}_t^T = [0.5 \ 1 \dots 1 \ 0.5] \Delta x$. Note that this formula corresponds to the trapezoid rule.

D.3.2 Aggregate Price Index

By the same logic, the aggregate price index 42 is computed as

$$e^{p_t^*(\epsilon-1)} = \sum_{i=1}^I (\mathbf{g}_t^c(\mathbf{x}_i) \mathbb{1}_{i \neq 1} h_{t,i,i-1,1-\epsilon} + \mathbf{g}_t^c(\mathbf{x}_i) \mathbb{1}_{i \neq I} h_{t,i,i+1,1-\epsilon}) + \mathbf{g}_t^0 \quad (45)$$

where

$$h_{t,i,j,\epsilon} = \frac{(e^{(\epsilon)\mathbf{x}_i s_{t,i,j}} ((\epsilon) (\mathbf{x}_i - \mathbf{x}_j) s_{t,i,j} - 1) + e^{(\epsilon)\mathbf{x}_j s_{t,i,j}})}{(\epsilon)^2 (\mathbf{x}_i - \mathbf{x}_j) s_{t,i,j}}$$

and where $\mathbb{1}_{i \neq 1}, \mathbb{1}_{i \neq I}$ is an indicator function equal to 1 when i is different than 1 or I . The indicator ensures that the integral is appropriately evaluated between the (s, S) bands.

Hence, we can re-write equation 45 in matrix form as

$$e^{p_t^*(\epsilon-1)} = \mathbf{h}_{t,1-\epsilon}^T \mathbf{g}_t^c + \mathbf{g}_t^0 \quad (46)$$

where \mathbf{g}_t^c is the matrix form of the distribution function \mathbf{g}_t^c and where the vector $\mathbf{h}_{t,1-\epsilon}$ is

$$\mathbf{h}_{t,1-\epsilon} = \sum_{i=1}^I [\mathbb{1}_{i \neq 1} h_{t,i,i-1,1-\epsilon} + \mathbb{1}_{i \neq I} h_{t,i,i+1,1-\epsilon}]_{i=1}^I$$

D.3.3 Labor Market

Following closely the previous subsection, the labor market condition 43 is computed as

$$N_t = \frac{C_t}{A_t} e^{p^*(-\epsilon)} \left(\sum_{i=1}^I (\mathbf{g}_t^c(\mathbf{x}_i) \mathbb{1}_{i \neq 1} h_{t,i,i-1,-\epsilon} + \mathbf{g}_t^c(\mathbf{x}_i) \mathbb{1}_{i \neq I} h_{t,i,i+1,-\epsilon}) + \mathbf{g}_{t-1}^0 \right) + \eta \mathbf{g}_{t-1}^0$$

which we re-write in matrix form as

$$N_t = \frac{C_t}{A_t} e^{p^*(-\epsilon)} (\mathbf{h}_{t,-\epsilon}^T \mathbf{g}_t^c + \mathbf{g}_{t-1}^0) + \eta \mathbf{g}_{t-1}^0 \quad (47)$$

D.3.4 Distribution

Once we have evaluated the integrals, the distribution function 40 can be written as:

$$\mathbf{g}_t^c(\mathbf{x}_j) = \sum_{i=1}^{I-1} \frac{1}{2\sqrt{2\pi}} \mathbf{g}_{t-1}^c(\mathbf{x}_i) [\mathbb{1}_{i \neq 1} f_{t,i,i-1,j} + \mathbb{1}_{i \neq I} f_{t,i,i+1,j}] + \mathbf{g}_{t-1}^0 \phi\left(\frac{-\hat{s}_{t,j} \mathbf{x}_j - \pi_t^*}{\sigma}\right) \quad (48)$$

where $f_{t,i,\bar{i},j}$ and $\mathcal{P}_{t,i,j,l}$ are defined below.

$$f_{t,i,\bar{i},j} = \frac{\sqrt{2\pi} (\mathcal{P}_{t,\bar{i},j}) \left(\operatorname{erf}\left(\frac{\mathcal{P}_{t,\bar{i},j}}{\sqrt{2\sigma}}\right) - \operatorname{erf}\left(\frac{\mathcal{P}_{t,i,j}}{\sqrt{2\sigma}}\right) \right) + 2\sigma \left(\exp\left(-\frac{\mathcal{P}_{t,\bar{i},j}^2}{2\sigma^2}\right) - \exp\left(-\frac{\mathcal{P}_{t,i,j}^2}{2\sigma^2}\right) \right)}{|x_i \hat{s}_{t-1,i} - x_{\bar{i}} \hat{s}_{t-1,\bar{i}}|}$$

$$\mathcal{P}_{t,i,j} = -x_i \hat{s}_{t-1,i} + x_j \hat{s}_{t,j} + \pi_t^*$$

For compactness, let us define

$$\begin{aligned} \mathbf{g}_t^c &\equiv [\mathbf{g}_t^c(\mathbf{x}_j)]_{j=1}^I \\ \mathbf{F}_t &\equiv \left[\sum_{i=1}^{I-1} \frac{1}{2\sqrt{2\pi}} [\mathbb{1}_{i \neq 1} f_{t,i,i-1,j} + \mathbb{1}_{i \neq I} f_{t,i,i+1,j}] \right]_{i=1,j=1}^{I,I} \\ \mathbf{f}_t &\equiv \left[\phi\left(\frac{-\hat{s}_{t,j} \mathbf{x}_j - \pi_t^*}{\sigma}\right) \right]_{j=1}^I \end{aligned}$$

where \mathbf{g}_t^c and \mathbf{f}_t vectors with the probability mass function and the scaled and shifted normal distribution, respectively, \mathbf{F}_t is a matrix that captures the idiosyncratic transitions due to firm-level quality shocks. Thus, equation 48 can be represented in matrix form as

$$\mathbf{g}_t^c = \mathbf{F}_t \mathbf{g}_{t-1}^c + \mathbf{f}_t^T \mathbf{g}_{t-1}^0 \quad (49)$$

D.3.5 Value function

Once we have evaluated the integrals, and denoting the standard normal cdf by $\Phi(\cdot)$ and the central grid point by i_0 (i.e. for $\mathbf{x}_{i_0} = 0$), the value function 38 can be written as

$$\begin{aligned} \mathbf{v}_t(\mathbf{x}_j) &= \Pi_{j,t} - \Pi_{j,t}(0) \\ &+ \Lambda_{t+1} \sum_{i=1}^I \frac{1}{2\sqrt{2\pi}} \mathbf{v}_{t+1}(\mathbf{x}_i) (\mathbb{1}_{i \neq 1} (a_{t,i,i-1,j} - a_{t,i_0,i_0-1,j}) + \mathbb{1}_{i \neq I} (a_{t,i,i+1,j} - a_{t,i_0,i_0+1,j})) \\ &+ \Lambda_{t+1} (-\eta w_{t+1}) \left(\Phi\left(\frac{\mathcal{P}_{t+1,j,I}}{\sigma}\right) - \Phi\left(\frac{\mathcal{P}_{t+1,j,1}}{\sigma}\right) - \Phi\left(\frac{\mathcal{P}_{t+1,i_0,I}}{\sigma}\right) + \Phi\left(\frac{\mathcal{P}_{t+1,i_0,1}}{\sigma}\right) \right) \end{aligned} \quad (50)$$

where

$$a_{t,i,\bar{i},j} = \frac{\sqrt{2\pi} (\mathcal{P}_{t+1,j,\bar{i}}) \left(\operatorname{erf} \left(\frac{\mathcal{P}_{t+1,j,\bar{i}}}{\sqrt{2\sigma}} \right) - \operatorname{erf} \left(\frac{\mathcal{P}_{t+1,j,i}}{\sqrt{2\sigma}} \right) \right) + 2\sigma \left(\exp \left(-\frac{(\mathcal{P}_{t+1,j,\bar{i}})^2}{2\sigma^2} \right) - \exp \left(-\frac{(\mathcal{P}_{t+1,j,i})^2}{2\sigma^2} \right) \right)}{|x_i \hat{s}_{t+1,i} - x_{\bar{i}} \hat{s}_{t+1,\bar{i}}|} \quad (51)$$

For compactness, let us define

$$\begin{aligned} \mathbf{v}_t &\equiv [\mathbf{v}_t(x_j)]_{j=1}^I \\ \mathbf{\Pi}_t &\equiv [\Pi_{j,t} - \Pi_{j,t}(0)]_{j=1}^I \\ \mathbf{A}_t &\equiv \left[\Lambda_{t+1} \sum_{i=1}^I \frac{1}{2\sqrt{2\pi}} (\mathbf{1}_{i \neq 1} (a_{t,i,i-1,j} - a_{t,i_0,i_0-1,j}) + \mathbf{1}_{i \neq I} (a_{t,i,i+1,j} - a_{t,i_0,i_0+1,j})) \right]_{i=1,j=1}^{I,I} \\ \mathbf{b}_{t+1} &\equiv \left[\Lambda_{t+1} \left(\Phi \left(\frac{\mathcal{P}_{t+1,j,I}}{\sigma} \right) - \Phi \left(\frac{\mathcal{P}_{t+1,j,1}}{\sigma} \right) - \Phi \left(\frac{\mathcal{P}_{t+1,i_0,I}}{\sigma} \right) + \Phi \left(\frac{\mathcal{P}_{t+1,i_0,1}}{\sigma} \right) \right) \right]_{j=1}^I \end{aligned}$$

where \mathbf{v}_t and \mathbf{b}_{t+1} are vectors that evaluate the value function and the adjustment probability at different grid points, $\mathbf{\Pi}_t$ is the vector of profit differences, while \mathbf{A}_t is a matrix that represents the idiosyncratic transition due to firm-level quality shocks and price updating. Thus, equation 50 can be represented in matrix form as

$$\mathbf{v}_t = \mathbf{\Pi}_t + [\mathbf{A}_t \mathbf{v}_{t+1} - \mathbf{b}_{t+1} \eta w_{t+1}] \quad (52)$$

D.3.6 Optimality condition for reset price

After evaluating the integral, we can write the optimality condition 39 as

$$\begin{aligned} 0 &= \Pi'_t(0) + \Lambda_{t+1} \sum_{i=1}^I \mathbf{v}_{t+1}(x_i) \frac{1}{2} (\mathbf{1}_{i \neq 1} c_{t,i,i-1,i_0} + \mathbf{1}_{i \neq I} c_{t,i,i+1,i_0}) \\ &\quad + \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (-\eta w_{t+1}) \end{aligned} \quad (53)$$

where

$$c_{t,i,\bar{i},j} = \frac{\operatorname{erf} \left(\frac{\mathcal{P}_{t+1,j,i}}{\sqrt{2\sigma}} \right) - \operatorname{erf} \left(\frac{\mathcal{P}_{t+1,j,\bar{i}}}{\sqrt{2\sigma}} \right)}{x_i s_{t+1,i} - x_{\bar{i}} s_{t+1,\bar{i}}} - \frac{\sqrt{\frac{2}{\pi}} \exp \left(-\frac{(\mathcal{P}_{t+1,j,i})^2}{2\sigma^2} \right)}{\sigma} \quad (54)$$

We can write this equation using matrix notation:

$$\begin{aligned} 0 &= \Pi'_t(0) + \mathbf{c}_{t+1}^T \mathbf{v}_{t+1} \\ &\quad + \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (-\eta w_{t+1}) \end{aligned} \quad (55)$$

where

$$\mathbf{c}_{t+1} = \left[\Lambda_{t+1} \frac{1}{2} (\mathbf{1}_{i \neq 1} c_{t,i,i-1,i_0} + \mathbf{1}_{i \neq I} c_{t,i,i+1,i_0}) \right]_{i=1}^I \quad (56)$$

D.4 Final equation system

Collecting the thus derived equations, and combining them with the remainder of the private equilibrium conditions (which contain no infinite dimensional objects) and the objective, we can approximate the infinite dimensional planner's problem by the following finite dimensional planner's problem

$$\max_{\{\mathbf{g}_t^c, \mathbf{g}_t^0, \mathbf{v}_t, C_t, w_t, p_t^*, s_t, S_t, \pi_t^*\}_{t=0}^\infty} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\gamma}}{1-\gamma} - v \left(\frac{C_t}{A_t} e^{p^*(-\epsilon)} (\mathbf{h}_{t,-\epsilon}^T \mathbf{g}_t^c + \mathbf{g}_{t-1}^0) + \eta \mathbf{g}_{t-1}^0 \right) \right)$$

subject to

$$\begin{aligned} w_t &= v C_t^\gamma \\ \mathbf{v}_t &= \mathbf{\Pi}_t + \mathbf{A}_t \mathbf{v}_{t+1} - \mathbf{b}_{t+1} \eta w_{t+1} \\ \mathbf{v}_{t,1} &= -\eta w_t \\ \mathbf{v}_{t,I} &= -\eta w_t \\ 0 &= \Pi'_t(0) + \mathbf{c}_{t+1}^T \mathbf{v}_{t+1} + \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (-\eta w_{t+1}) \\ \mathbf{g}_t^c &= \mathbf{F}_t \mathbf{g}_{t-1}^c + \mathbf{f}_t^T \mathbf{g}_{t-1}^0 \\ \mathbf{g}_t^0 &= 1 - \mathbf{e}_t^T \mathbf{g}_t^c \\ e^{p_t^*(-\epsilon-1)} &= \mathbf{h}_{t,1-\epsilon}^T \mathbf{g}_t^c + \mathbf{g}_t^0 \end{aligned}$$

Here the choice variables \mathbf{v}_t and \mathbf{g}_t^c are vectors of length I . The rest of the choice variables are scalars. Note that the choice variables p_t^*, s_t, S_t, π_t^* appear in the problem implicitly inside the vectors and matrices.

As already explained at the beginning, we solve for the FOCs of this system by symbolic differentiation. The resulting system of FOCs is then solved in the sequence space. The last subsection explains how we find the steady state, which serves as initial and terminal condition for these dynamic simulations.

D.5 Steady state

To solve for the steady state of the private equilibrium conditions, given a policy π , the algorithm is as follows. In steady state we have $w = v C^\gamma$, and $R = (1 + \pi)/\beta$. Start with a guess for the real wage w , the optimal rest price p^* , and the (s,S) bands s and S then:

1. Compute consumption $C = \left(\frac{w}{v}\right)^{1/\gamma}$.
2. Using $\pi = \bar{\pi}$, C and the 4 initial guesses, solve for that stationary value function using the Bellman equation and the stationary distribution using the law of motion of the distribution. Both have closed form solutions given the guesses.

3. Compute the residuals of the 4 unused equations

$$\begin{aligned}
 \mathbf{v}_{t,1} &= -\eta w_t \\
 \mathbf{v}_{t,I} &= -\eta w_t \\
 0 &= \Pi'_t(0) + \mathbf{c}_{t+1}^T \mathbf{v}_{t+1} + \Lambda_{t+1} \left(\phi \left(\frac{S_{t+1} - \pi_{t+1}^*}{\sigma} \right) - \phi \left(\frac{s_{t+1} - \pi_{t+1}^*}{\sigma} \right) \right) (-\eta w_{t+1}) \\
 e^{p_t^*(\epsilon-1)} &= \mathbf{h}_{t,1-\epsilon}^T \mathbf{g}_t^c + \mathbf{g}_t^0
 \end{aligned}$$

4. Use a newton method to update the 4 guesses (w , p^* , s , S) and return to step 1, until convergence of the residuals.