The Slope of the Phillips Curve and the Mandate of the Central Bank

Jean-Paul L'Huillier, Ernesto Pastén, and Gregory Phelan*

February 2026

Abstract

Under exogenous price stickiness, the divine coincidence suggests that the Central Bank can focus on inflation stabilization to maximize welfare. We show that endogenous price stickiness upsets this result. The pursuit of price stability may, in fact, increase price stickiness, flatten the Phillips curve, increase the distortions due to sticky prices, and lead to a welfare loss. Welfare can be improved if the Central Bank stabilizes the output gap directly (dual mandate). Our argument does not rely on markup shocks or nominal wage rigidities. Instead, the key to these insights is the endogeneity of the degree of price stickiness.

Keywords: Central bank design, policy-dependent price stickiness, Lucas critique.

JEL codes: E31, E52, E58.

*Brandeis University, Central Bank of Chile, and Williams College. We thank Olivier Blanchard, Anastasios Karantounias, Yuriy Gorodnichenko, Raphael Schoenle, Slavik Sheremirov, Alp Simsek, Sanjay R. Singh, and Michael Weber for useful conversations. All errors are our own.

1 Introduction

The design of the central bank is a fundamental topic in monetary policy, but there remains significant disagreement over what that mandate should be. In the European Union, the European Central Bank (ECB) operates under a single mandate focused solely on price stability.[1] In contrast, the Federal Reserve of the United States follows a dual mandate to stabilize both prices and output.[2] At first glance, the distinction between these mandates may seem to be of little consequence, particularly if the stabilization of prices and output naturally go hand in hand, a phenomenon often referred to as the "divine coincidence" (Blanchard and Gali, 2007).

However, we argue that the logic behind this divine coincidence is misleading under realistic characterizations of the dynamics of inflation. In fact, price stability can exacerbate output volatility by increasing price stickiness. Specifically, our analysis shows how central bank policies aimed at stabilizing prices can increase the stickiness of prices and create welfare losses. These findings suggest that a dual mandate, which seeks to stabilize both prices and output, may be more effective in promoting economic stability than a singular focus on price stability.

We use a microfounded model of firms' pricing decisions to show that the stickiness of prices can be affected by monetary policy. We consider a standard model with endogenous price stickiness following Blanco et al. (2024a). In this model, the volatility of shocks determines the misallocation losses due to sticky prices. When monetary policy is more aggressive in responding to inflation, inflation volatility drops and price stickiness increases. We show that this can flatten the Phillips Curve, widen the output gap, and exacerbate welfare losses.

In this context, the divine coincidence fails badly and in a unique way not currently considered in the literature. Reducing price volatility may *increase* the volatility of output, and so, surprisingly, the pursuit of price stabilization can actually decrease welfare, contrary to standard analyses in which price stickiness is exogenous (Woodford, 2003; Galí, 2015). Our results indicate that a dual central bank mandate with both price and output stabilization is desirable.

Our analysis identifies two key reasons why the divine coincidence fails under endogenous

[1] *Treaty on the Functioning of the European Union*, Article 127(1), Extracted from the Official Journal of the European Union, 26.10.2012.

[2] *Federal Reserve Reform Act of 1977*, §2A.

price stickiness.

First, with endogenous price stickiness, a more aggressive monetary policy pursuing price stability leads to stickier prices. As prices become stickier, the central bank needs to increase the intensity to which it reacts to the price level even more in order to achieve the same amount of stabilization. This feedback loop converges to a point where the price level may not react at all to movements in the output gap. In other words, the Phillips curve can become very, or completely, flat.[3] Hence, the inflation-targeting central bank is unable to close the output gaps.

Second, the the traditional failure of the divine coincidence occurs when closing an inflation gap does not necessarily close an output gap. We show that an even stronger failure is possible: targeting price stability can actually widen the output gap. If the central bank intensifies its focus on stabilizing prices, firms may respond by making prices even stickier. In this case, output gaps may become more volatile or persistent, as price setters no longer adjust prices in response to fluctuations in demand. Not only will the central bank fail to close output gaps, but output gaps can become even more pronounced due to the increased stickiness. The price level might become stable, not because the output gap has been closed, but because price stickiness with respect to aggregate demand shocks has increased to a point where price setters no longer react to output gap fluctuations.

Our findings challenge the conventional wisdom that price stability should be the central bank's sole objective. To understand how our results fit in the literature analyzing optimal monetary policy, it is useful to contrast them with an old and influential idea by Knut Wicksell. Following Wicksell (1898), business cycle stabilization can be aided by a central bank that uses the price level (or inflation) as a guide for setting interest rate policy. When the economy is running too hot, prices rise. This signals to the central bank that the interest rate should be increased. When the economy is running too cold, prices fall. This signals to the central bank that the nominal interest rate should be decreased. When the economy reaches equilibrium, the price level is stable, and the Central Bank should keep the interest rate unchanged, since it must be at its "natural" rate (also called the

[3] In our model the slope of the Phillips curve is a nonlinear function of the size of demand shocks, for given parameters of the interest rate rule. Hence, a related prediction is that for a given amount of stabilization (quantified by parameters of the interest rate rule), larger demand shocks lead to a steeper Phillips curve.

Wicksellian rate of interest).

While this analysis is elegant and intuitive, our model suggests that it is quite misleading when price stickiness is *endogenous* to monetary policy. If the central bank prioritizes price stability too aggressively, the price level may become stable not because the economy has reached a Wicksellian equilibrium, but because firms are choosing not to adjust prices at all. Indeed, in the model, prices can be sufficiently sticky with respect to demand shocks that the Phillips curve can flatten entirely, making the price level entirely unresponsive to the output gap. However, prices can still vary with respect to other shocks, such as cost-push shocks or idiosyncratic shocks, as explained in detail below. As a result, the Wicksellian rate may not be achievable through Wicksell's proposed mechanism, and long-lasting output gaps can persist even when inflation appears stable.

This insight has significant policy implications. A strict inflation-targeting regime may achieve price stability while undermining its ability to stabilize output and maximize welfare. Instead, we argue that central banks should adopt a dual mandate that explicitly considers both inflation and output stabilization.

It is worthwhile to highlight the following methodological aspect of our approach to the problem of optimal monetary policy. It strikes us as crucial to work with an endogenous model of nominal rigidities themselves to properly evaluate the implications of different long-term monetary policy *postures*. Indeed, such postures are adopted for long periods of times (more than one decade, at least), and hence one can reasonably expect that they can impact the way nominal rigidities arise in the economy. With these issues in mind, nominal rigidities can no longer be thought as an 'institutional' feature of the economy when considering the posture of monetary policy. As this paper shows, this modeling strategy brings new insights in terms of the long-term posture of the central bank.

Related Literature

There is a robust literature highlighting the presence, and failures, of the divine coincidence. The literature has highlighted several ways to break the divine coincidence. Our mechanism is distinct from those in the existing literature.

The most common strategy to break the divine coincidence is to introduce a cost-push (markup) shock, which creates a negative output gap and positive inflation. Alves (2014) shows that the divine coincidence only holds around zero inflation: if there is trend inflation, the divine coincidence fails because trend inflation endogenously acts like a cost-push shock. Another strand of the literature introduces additional wage or financial frictions or rigidities, as in Blanchard and Gali (2007) and Cúrdia and Woodford (2016). Another strand shows that with heterogeneous agents, decreasing inflation may not help all agents, e.g., Gornemann et al. (2016), Kaplan et al. (2018), and Auclert (2019). Limited attention and information constraints disrupt the divine coincidence by causing firms to adjust prices selectively and asymmetrically, making it harder for monetary policy to achieve both inflation and output stability simultaneously, see e.g., Woodford (2009), Mackowiak and Wiederholt (2009), and Mackowiak et al. (2021).

Our model highlights a novel and unique channel with distinct policy implications. In our setting, the problem is that price stickiness itself changes, and that should lead a central bank to put a larger weight on output gaps directly. In our paper, the divine coincidence fails because closing one gap changes the relationship between the gaps, and there is a critical asymmetry: prices are the things that firms are inclined to hold sticky, so the inflation gap is more likely to stop moving even as output keeps moving. Our novel result is not that targeting one gap affects the behavior of the other, but rather that responding to shocks asymmetrically weakens the price signal. We believe our novel mechanism is empirically relevant because the PC has been historically so flat.

We are not the first paper to consider how monetary policy affects the Phillips Curve or should operate in light of informational frictions. There is a long literature studying the optimality of inflation-targeting, Taylor, and Wicksellian monetary policy rules, including Svensson (1999, 2003); Giannoni and Woodford (2004); Boivin and Giannoni (2006); Reis (2013); Giannoni (2014). Benigno and Ricci (2011) shows that downward nominal wage rigidities produce nonlinear Phillips Curves that flatten at low inflation.

Methodologically, our paper is closest to Blanco et al. (2024b,b). These papers consider the time-varying choice of endogenous price setting frequencies and show that matching inflation dynamics requires matching variations in the frequency of price changes. These papers focus on

nonlinearities in the PC over the business cycle. In contrast, we are interested in differences in steady states caused by different monetary policy postures. Using an entirely different methodology, L'Huillier et al. (2022) consider how price targeting can limit the effectiveness of monetary policy in a game-theoretic model in which pricing rigidities are the result of a strategic friction between firms and consumers.

There is ample empirical evidence that the slope of the Phillips Curve responds to the economic environment. A burgeoning literature explores nonlinearities. We distinguish ourselves by the welfare analysis of central bank postures. Gitti (2024) and Gitti and Cerrato (2022) analyze U.S. metropolitan areas and find that the PC's slope more than tripled post-pandemic. Complementing this, Benigno and Eggertsson (2023) introduced a nonlinear New Keynesian PC model. In parallel important work, Blanco et al. (2024a,b) and Karadi et al. (2024) also explore nonlinear inflation dynamics in the presence of large shocks. Our story that the aggressiveness of monetary policy affects the slope of the Phillips curve is consistent with this evidence.

Organization. The remainder of the paper is organized as follows. Section 2 presents a simple NK model with endogenous price stickiness based on Blanco et al. (2024a) to show that monetary policy can flatten the Phillips Curve, widen the output gap, and lead to welfare losses. Section 3 presents our analytical results showing when the divine coincidence can fail with endogenous price stickiness in an environment with only demand shocks. Section 4 then presents a simple textbook calibration to show that with both demand and supply shocks, our results hold away from the conditions required in the analytical settings. Section 5 concludes. The Appendix provides proofs and additional materials.

2 Baseline Model

We describe the model, using the baseline ingredients from Blanco et al. (2024a) that are needed for our analysis. In this section we collect the minimal set of primitives, definitions, and equilibrium conditions that discipline the endogenous adjustment frequency decision, allowing us to derive analytical and quantitative results in later sections.

2.1 Setup

Time is discrete, $t = 0, 1, 2, \ldots$. There is a representative household with standard preferences over consumption C_t and labor N_t.

There is a continuum of monopolistically competitive firms indexed by $i \in [0,1]$. Each firm produces a continuum of differentiated products indexed by $k \in [0,1]$. Products are imperfect substitutes in final-good production. All products share the same technology and face the same demand elasticity. The final good is produced competitively using a CES aggregator over all products:

$$Y_t = \left(\int_0^1 \int_0^1 y_{i,k,t}^{\frac{\theta-1}{\theta}} \, di \, dk \right)^{\frac{\theta}{\theta-1}}, \qquad \theta > 1.$$

Let $p_{i,k,t}$ denote the nominal price of product k produced by firm i at date t, and let P_t denote the aggregate price index defined below. Cost minimization yields demand for each product:

$$y_{i,k,t} = \left(\frac{p_{i,k,t}}{P_t} \right)^{-\theta} Y_t, \tag{1}$$

and the aggregate price index

$$P_t = \left(\int_0^1 \int_0^1 p_{i,k,t}^{1-\theta} \, di \, dk \right)^{\frac{1}{1-\theta}}. \tag{2}$$

Each firm i chooses, in each period, a fraction $n_t \in (0,1)$ of its products whose prices are adjusted at date t. The remaining fraction $1 - n_t$ of products keep their previous nominal prices. By symmetry, the choice of n_t is symmetric across firms and products. When a product receives an adjustment opportunity, the firm resets its price to a common optimal reset price p_t^*, which is the same for all adjusting products by symmetry. Thus, nominal rigidity operates at the *product level*, while the adjustment decision n_t is made at the *firm level*.

The object n_t should be interpreted as the *intensive margin of price adjustment within firms*: it is the fraction of products whose prices are updated by each firm in period t. This interpretation is central to the Blanco et al. (2024a) framework and allows firms to smooth price adjustment costs by adjusting subsets of products rather than all prices simultaneously.

Adjusting a larger fraction of products is costly. Firms incur a convex cost $\mathscr{C}(n_t)$, specified below, which captures managerial, informational, or operational costs associated with coordinating price changes across many products within the firm. Blanco et al. (2024a) specify an adjustment cost that penalizes deviations of the adjustment probability from a benchmark $\bar{n}$.

Nominal aggregate demand is driven by an exogenous money growth process. We specify shocks in greater detail below.

2.2 Pricing block

The pricing block can be written in terms of: (i) the reset price p_t^*, (ii) the aggregate price level P_t (or inflation π_t), (iii) auxiliary present-value objects that summarize discounted profits of adjusting firms, and (iv) a dispersion (misallocation) term.

Given that a fraction n_t of products reset to p_t^* and a fraction $1-n_t$ keep their previous prices, the aggregate price index satisfies

$$1 = n_t \left(\frac{p_t^*}{P_t}\right)^{1-\theta} + (1-n_t)\pi_t^{\theta-1}, \tag{3}$$

where $\pi_t \equiv P_t/P_{t-1}$ is gross inflation. Equation (3) is the key aggregation identity linking inflation to the reset price. The reset price equates discounted marginal revenue to discounted marginal cost.

Let x_t denote a price-dispersion (or misallocation) term, whose law of motion can be written as

$$x_t^{-\theta/\eta} = n_t \left(\frac{p_t^*}{P_t}\right)^{-\theta/\eta} + (1-n_t)\, x_{t-1}^{-\theta/\eta}\, \pi_t^{\theta/\eta}. \tag{4}$$

This recursion captures the idea that dispersion (and hence real misallocation) is reduced by price adjustment and increased by inflation-induced drift in relative prices among non-adjusters.

The subsequent sections specialize this baseline to derive a closed-form solution for n and the implied Phillips-curve slope.

3 Analytical Results

The key insight of the Blanco et al. (2024a) setup is that the fraction of price adjustments n_t is endogenous, and thus the slope of the Phillips Curve is endogenous. We now analytically characterize how variations in monetary policy can influence the slope of the Phillips Curve, perhaps with detrimental welfare consequences.

To motivate ideas, we first consider a reduced-form relationship between the slope of the Phillips Curve (PC) and monetary policy (MP) in a standard New Keynesian (NK) model: the more the central bank (CB) targets inflation, the flatter is the Phillips Curve. If that relationship is sufficiently strong, then we have two results. First, the divine coincidence (DC) fails: a more aggressive monetary policy rule may fail to decrease the output gap. Second, more aggressive policy can increase welfare losses. After considering a simple NK setting, we return to the model with endogenous price adjustment frequency to micro-found these reduced-form results.

Thus, the analysis proceeds in three steps: (i) characterize equilibrium responses as a function of the Phillips-curve slope κ; (ii) derive how κ responds endogenously to inflation volatility; and (iii) close the model by determining inflation volatility in equilibrium.

3.1 NK Setup

We start from the simplified 3-equation New Keynesian model:

$$x_t = \mathbb{E}_t[x_{t+1}] - \frac{1}{\sigma}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^n), \tag{IS}$$

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa x_t, \tag{PC}$$

$$i_t = \phi_\pi \pi_t + \phi_x x_t \tag{MP}$$

where σ is relative risk aversion, β is the discount rate, κ is the slope of the PC, and ϕ_π and ϕ_x are the coefficients in the CB's Taylor rule. The baseline inflation accelerator models nests this setting when adjustment costs are infinite and therefore the adjust frequency is fixed at $\bar{n}$.

Assume a one-time natural rate shock to r_0^n, and $r_t^n = 0$ for $t \geq 1$. The model is essentially static:

since the shock is i.i.d., all variables for $t \geq 1$ are at steady state (no deviations). It is immediate that the equilibrium values of inflation π_0 and output gap x_0 are given by

$$\pi_0 = \frac{\kappa}{\sigma + \kappa\phi_\pi + \phi_x} r_0^n, \quad x_0 = \frac{1}{\sigma + \kappa\phi_\pi + \phi_x} r_0^n, \quad i_0 = \frac{\kappa\phi_\pi + \phi_x}{\sigma + \kappa\phi_\pi + \phi_x} r_0^n.$$

Naturally, all else equal, increasing ϕ_π or ϕ_x decreases inflation and the output gap. Write the denominator as $D \equiv \sigma + \kappa\phi_\pi + \phi_x$. The equilibrium welfare losses can be expressed as

$$\mathscr{L} = \hat{x}_t^2 + \frac{\varepsilon}{\kappa}\pi_t^2 = \frac{1 + \varepsilon\kappa}{D^2}. \tag{5}$$

If κ decreases with policy, and if policy fails to close the output gap ($\hat{x}$ increases), then stronger policy can increase the loss function. This is what we wish to examine.

Endogenous Kappa Suppose that κ is a function of inflation volatility. This could capture endogenous choices about price flexibility, and firms will naturally have more of an incentive to use flexible price policies when inflation is more volatile. For now we just posit such a relationship.

Suppose that κ is a linear function of inflation volatility. From above we have $\pi_0 = \frac{\kappa}{D} r_0^n$, and therefore a linear endogenous PC would satisfy

$$\kappa = \frac{\kappa}{D} C,$$

for some constant C. This requires

$$\sigma + \kappa\phi_\pi + \phi_x = C.$$

The following result is immediate:

Proposition 1. *If κ is a linear function of inflation volatility, then (i) the DC fails: the output gap is a constant; but (ii) increasing ϕ_π decreases welfare losses.*

There are three main observations:

1. The PC slope κ is decreasing in ϕ_π.

2. The DC fails. In equilibrium the output gap is a constant, $x_0 = r_0^n/C$ independent of MP.

3. Even though the DC fails, aggressive MP will decrease the welfare losses: $\mathscr{L} = \frac{1+\varepsilon\kappa}{D^2} = \frac{1+\varepsilon\kappa}{C^2}$, and since κ decreases with ϕ_π, welfare losses decrease if ϕ_π increases.

We assumed a very particular relationship between κ and inflation volatility, which implies a particular relationship with ϕ_π, namely a unit elasticity. More generally, suppose that κ is some function of ϕ_π. To fix ideas, let $\kappa = K\phi_\pi^{-\gamma}$ so that γ is the elasticity with respect to ϕ_π. We have the following result

Proposition 2. *For some $\bar{\gamma} > 1$, if κ is a super-elastic function of ϕ_π, i.e., $\gamma > \bar{\gamma} > 1$, then increasing ϕ_π increases welfare losses.*

For welfare losses to increase, the PC needs to be sufficiently elastic. In this case, aggressive MP not only fails to close the output gap, but the output gap actually widens as ϕ_π increases.

We can characterize this result more sharply when inflation targeting policy is very aggressive. Suppose monetary policy is sufficiently aggressive that the policy term dominates the IS denominator, so that

$$\sigma + \kappa\phi_\pi + \phi_x \approx \kappa\phi_\pi,$$

because ϕ_π is large. We call this the ϕ_π policy-dominant region, defined by $\kappa(\phi_\pi)\phi_\pi \gg \sigma + \phi_x$.

Proposition 3. *Consider the static NK model with demand shocks only and let the PC slope depend on policy according to $\kappa(\phi_\pi) = K\phi_\pi^{-\gamma}$ with $\gamma > 0$, and consider the ϕ_π policy-dominant region. Then the following results hold:*

1. *The equilibrium output gap satisfies $x_0 \propto \phi_\pi^{\gamma-1}$. Hence the output gap widens with more aggressive inflation stabilization if and only if $\gamma > 1$.*

2. *The welfare loss,*

$$\mathscr{L}_0^{NK} = \frac{1+\varepsilon\kappa}{D^2}(r_0^n)^2,$$

is increasing in ϕ_π if and only if $\gamma > 1$.

3.2 Microfoundation for Endogenous PC

The last section assumed a reduced-form relationship between the PC slope and MP. We now use a microfounded model to derive the precise relationship between the PC slope and MP.

To derive analytical results, we consider a tractable variant of Blanco et al. (2024a) in which the PC slope κ is endogenously determined at the stochastic steady state. First, we assume firms choose a *time-invariant* adjustment rate $n_t = n \; \forall t$. In the original formulation, firms choose n_t each period as part of a dynamic decision problem. Second, we replace the quadratic cost $\frac{\xi}{2}(n-\bar{n})^2$ in Blanco et al. (2024a) with a power function anchored at zero ($\bar{n} = 0$):

$$C_{\text{adj}}(n) = \frac{\xi}{\nu} n^{\nu}, \qquad \xi > 0, \; \nu > 0.$$

This normalization implies that in a deterministic environment ($\sigma_\pi^2 = 0$), n is optimally zero. The presence of shocks will lead firms to choose $n > 0$ even when trend inflation is zero.

We study small fluctuations around a symmetric, zero-inflation reference point—the stochastic steady state. In the original model, Blanco et al. (2024a) solve for a deterministic steady state in the absence of shocks. In contrast, we seek an analytically tractable *stochastic steady state* in which firms choose an adjustment frequency n, taking into account that aggregate shocks generate price dispersion and misallocation. We show explicitly how the optimal n varies with the volatility of aggregate shocks and how this feeds into the slope of the New Keynesian Phillips curve.

Price dynamics Let $p_{i,t}$ denote firm i's price for a given product and P_t the aggregate price index. Define the log relative price

$$r_{i,t} \equiv \log\left(\frac{p_{i,t}}{P_t}\right).$$

When a firm receives a price-adjustment opportunity for that good, it resets to the frictionless optimal price. When a firm adjusts at date t, it chooses a reset price p_t^* that solves the standard optimal pricing problem. The corresponding reset relative price is

$$r_t^* \equiv \log\left(\frac{p_t^*}{P_t}\right).$$

Aggregate inflation is

$$\hat{\pi}_t \equiv \log P_t - \log P_{t-1},$$

and shocks are assumed to have zero mean. Under a first-order (log-linear) approximation, the reset relative price is linear in aggregate shocks and state variables. Hence, $E[r_t^*] = 0$ since expected inflation is zero.

Log-linearizing the equilibrium system around the non-stochastic steady state with constant n implies that inflation responds linearly to the aggregate shock,

$$\hat{\pi}_t = \psi \varepsilon_t,$$

for some coefficient $\psi > 0$ determined by the underlying model. Since our dispersion and n-choice formulas depend only on the variance of inflation innovations, it is convenient to define

$$\sigma_\pi^2 \equiv \psi^2 \sigma_M^2$$

and work directly with the reduced-form inflation process

$$\hat{\pi}_t = u_t, \qquad u_t \sim \mathcal{N}(0, \sigma_\pi^2), \quad \text{i.i.d.}, \tag{6}$$

where $\hat{\pi}_t \equiv \log P_t - \log P_{t-1}$ denotes log-deviations of inflation from its deterministic trend. Firms take $\sigma_\pi^2 = \text{Var}(\hat{\pi}_t)$ as given.

Misallocation loss and Endogenous n Each firm chooses a time-invariant adjustment probability $n \in [0,1]$. With probability n, the firm adjusts a given product's price at date t and sets $r_{i,t} = 0$. With probability $1-n$, its price remains unchanged, so

$$r_{i,t} = r_{i,t-1} - \hat{\pi}_t = r_{i,t-1} - u_t.$$

Under CES demand, second-order approximations imply that the static misallocation loss from dispersion in log relative prices is proportional to the variance of $r_{i,t}$. We therefore write the

reduced-form misallocation loss as follows:

Lemma 1. *The misallocation costs from sticky prices and price dispersion can be written*

$$\textit{Misallocation}(n;\sigma_\pi^2) = \Lambda\,\sigma_\pi^2\frac{1-n}{n}, \tag{7}$$

where $\Lambda > 0$ depends on θ and other parameters (e.g., output scale)

In a stationary environment, the firm's per-period expected loss is

$$L(n;\sigma_\pi^2) = \frac{\xi}{\nu}n^\nu + \Lambda\sigma_\pi^2\frac{1-n}{n}.$$

The firm chooses n to minimize L. Lower volatility implies less price dispersion and therefore weaker incentives to adjust. We have the following result.

Proposition 4. *In the stochastic steady-state, lower inflation volatility leads to a lower adjustment rate n. The optimal constant adjustment probability is*

$$n(\sigma_\pi^2) = \left(\frac{\Lambda}{\xi}\right)^{\frac{1}{\nu+1}}(\sigma_\pi^2)^{\frac{1}{\nu+1}}. \tag{8}$$

Hence, $\frac{\partial n}{\partial \sigma_\pi^2} > 0$, $\lim_{\sigma_\pi^2\to 0} n(\sigma_\pi^2) = 0$, and the elasticity between n and inflation variance is $\frac{1}{\nu+1}$, where ϕ is the elasticity of the adjustment cost function.

Embedding into the New Keynesian Phillips curve We now embed the pricing block to derive ψ in equilibrium. Let θ_C be the Calvo probability of *not* adjusting, so that $\theta_C = 1-n$. Log-linearizing around zero-inflation, the standard NK PC is

$$\pi_t = \beta E_t\pi_{t+1} + \kappa(n)\,mc_t,$$

where mc_t is real marginal cost and

$$\kappa(n) = \Xi\,\frac{n}{1-n}\big(1-\beta(1-n)\big), \tag{9}$$

with $\Xi > 0$ a collection of preference and technology parameters. When volatility is low, we can derive the following result:

Proposition 5. *For low* σ_π^2, $\kappa(\sigma_\pi^2) \approx \Xi(1-\beta)\left(\frac{\Lambda}{\xi}\right)^{\frac{1}{\nu+1}}(\sigma_\pi^2)^{\frac{1}{\nu+1}}$. *In other words,* $\kappa(\sigma_\pi^2) \propto (\sigma_\pi^2)^{\frac{1}{\nu+1}}$.

In the stochastic steady-state, lower inflation volatility leads to a flatter PC through the adjustment rate n. For small volatility, the elasticity between κ and inflation variance is $\frac{1}{\nu+1}$, where ϕ is the elasticity of the adjustment cost function. As aggregate volatility falls, the optimal adjustment frequency $n(\sigma_\pi^2)$ shrinks, which in turn flattens the Phillips curve by driving $\kappa(\sigma_\pi^2)$ toward zero. The resulting $\kappa(\sigma_\pi^2)$ can be used directly in the standard three-equation New Keynesian model (IS curve, Phillips curve, and monetary policy rule) to study how reduced macro volatility endogenously flattens the Phillips curve.

Notice that if the cost function has $\nu = 1$ (linear), then the Phillips curve slope is proportional to the volatility of shocks σ, since $\kappa(\sigma_\pi^2) \propto (\sigma_\pi^2)^{1/2} = \sigma_\pi$ in the small-volatility region. This is the condition already considered in the previous section.

3.3 Endogenous determination of inflation volatility

In the pricing block above, firms take as given the variance of the innovation to inflation,

$$\hat{\pi}_t = u_t, \qquad u_t \sim \mathcal{N}(0, \sigma_\pi^2),$$

and the optimal adjustment frequency $n(\sigma_\pi^2)$ is derived under the assumption that σ_π^2 is an exogenous primitive. In a New Keynesian environment, σ_π^2 is itself an equilibrium object. We now show how the inflation elasticity ψ—the mapping from fundamental shocks to inflation innovations—is endogenously determined.

Solving yields the impact response of inflation,

$$\pi_0 = \psi(\kappa)\, r_0^n, \qquad \psi(\kappa) \equiv \frac{\kappa}{\sigma + \kappa\phi_\pi + \phi_x}. \tag{10}$$

If $\mathrm{Var}(r_0^n) = \sigma_r^2$, then the implied variance of inflation is

$$\sigma_\pi^2(\kappa) \equiv \mathrm{Var}(\pi_0) = \psi(\kappa)^2 \sigma_r^2 = \left(\frac{\kappa}{\sigma + \kappa\phi_\pi + \phi_x} \right)^2 \sigma_r^2. \tag{11}$$

In the pricing block, κ is not exogenous. Firms choose their adjustment frequency n as a function of the variance of inflation innovations: $n = n(\sigma_\pi^2)$, and $\kappa = \kappa\big(n(\sigma_\pi^2)\big)$. Combining (10)–(11) with the pricing-block mapping $n(\sigma_\pi^2)$ and the PC mapping $\kappa(n)$ delivers the fixed-point condition

$$\sigma_\pi^2 = \left(\frac{\kappa\big(n(\sigma_\pi^2)\big)}{\sigma + \phi_\pi\, \kappa\big(n(\sigma_\pi^2)\big) + \phi_x} \right)^2 \sigma_r^2. \tag{12}$$

Equation (12) shows that the elasticity ψ is itself an equilibrium object in the closed model:

$$\psi = \frac{\kappa\big(n(\sigma_\pi^2)\big)}{\sigma + \phi_\pi\, \kappa\big(n(\sigma_\pi^2)\big) + \phi_x}.$$

Thus the variance of inflation shocks, the optimal adjustment frequency, the slope of the Phillips curve, and the elasticity ψ are jointly determined. Lower volatility in fundamentals (e.g., as a result of more aggressive monetary policy) leads to a lower induced inflation volatility, which in turn reduces n, flattens the Phillips curve, and feeds back into a smaller ψ—a mechanism that becomes important when embedding our pricing block into a general-equilibrium NK model.

3.4 Inflation and Output Responses to a Demand Shock

We now characterize the interaction between monetary policy and the endogenous PC slope κ in an environment where inflation volatility is generated exclusively by demand shocks. This case illustrates the possibility of a perverse widening of the output gap for sufficiently concave adjustment costs. By "perverse" we mean that stronger inflation stabilization widens the output gap. Hence, endogenously we can get a strong failure of the divine coincidence.

For analytical tractability, we suppose MP responds aggressively to inflation (i.e., ϕ_π is large). The following proposition summarizes the results:

Proposition 6. *Adjustment costs affect the divine coincidence as follows:*

1. *With convex costs ($\nu > 1$), $x_0(\phi_\pi) \to 0$ as $\phi_\pi \to \infty$. A stronger rule improves stabilization because the PC does not flatten very quickly.*
2. *With linear costs ($\nu = 1$), $x_0(\phi_\pi) \approx$ constant. The output gap is invariant to ϕ_π. The DC fails, but welfare increases as ϕ_π increases because a lower κ decreases misallocation losses.*
3. *With concave costs ($0 < \nu < 1$), $x_0(\phi_\pi) \to \frac{1}{\sigma+\phi_x} r_0^n$. A stronger rule* worsens *the output gap because the PC flattens sufficiently fast.*

Thus concavity in adjustment costs generates the possibility that the incentive to increase the adjustment frequency n is strong enough that κ collapses as ϕ_π rises, driving x_0 upward. The case $\nu = 1$ is a knife-edge in which the output gap is invariant to policy but welfare still improves.

If adjustment costs are concave, then the PC slope can flattens sufficiently fast so that the output gap widens and, per our analysis above, the standard reckoning of welfare losses increases with more aggressive MP. Intuitively, concave adjustment costs make firms disproportionately sensitive to reductions in inflation volatility, causing price flexibility to collapse as policy becomes more aggressive. The inflation accelerator model also includes welfare losses because of adjustment costs when choosing n. This creates a third term in the welfare loss expression in addition to the normal terms from output gaps and price dispersion (inflation).

The Output-Gap Coefficient ϕ_x We now study how the MP response to the output gap, governed by ϕ_x, influences the endogenous PC slope κ and the initial output-gap response x_0. The analysis parallels the previous section. When varying ϕ_x, the divine coincidence does not fail. As before, we consider when ϕ_x is large to get asymptotic results. For all $\nu > 0$, the output gap shrinks as ϕ_x increases. No perverse region arises in which the DC fails; increasing ϕ_x always improves stabilization.

Proposition 7. *When ϕ_x is large, $x_0(\phi_x) \sim \frac{r_0^n}{\phi_x}$ as $\phi_x \to \infty$.*

With demand shocks only, a stronger real-side response (ϕ_x) always improves stabilization. An increase in ϕ_x never produces a wider output gap: the perverse region possible under variations in ϕ_π does not arise for variations in ϕ_x. Thus, while more assertive MP in the form of higher ϕ_π can

lead to a failure of the DC and decrease average welfare, more assertive MP in the form of higher ϕ_x does not have these potential tradeoffs.

3.5 Welfare losses with endogenous price flexibility

In the static NK environment, the associated welfare loss is

$$\mathscr{L}_0^{NK} \equiv x_0^2 + \frac{\varepsilon}{\kappa}\pi_0^2 = \frac{1+\varepsilon\kappa}{D^2}(r_0^n)^2. \tag{13}$$

We now incorporate real resource adjustment costs from endogenous price flexibility. The per-period adjustment cost is

$$\mathscr{L}_0^{adj} \equiv \frac{\xi}{\nu} n^{\nu}, \qquad \xi > 0,\ \nu > 0. \tag{14}$$

In the small-n region we use the standard approximation $\kappa \approx Kn$, and $K \equiv \Xi(1-\beta)$, so that $n \approx \kappa/K$ and hence

$$\mathscr{L}_0^{adj} \approx \frac{\xi}{\phi}\left(\frac{\kappa}{K}\right)^{\phi}. \tag{15}$$

The total per-period loss is therefore

$$\mathscr{L}_0 \equiv \mathscr{L}_0^{NK} + \mathscr{L}_0^{adj},$$

which can be written

$$\frac{1+\varepsilon\kappa}{D^2}(r_0^n)^2 + \frac{\xi}{\nu}n^{\nu} \;\approx\; \frac{1+\varepsilon\kappa}{D^2}(r_0^n)^2 + \frac{\xi}{\nu}\left(\frac{\kappa}{K}\right)^{\nu}. \tag{16}$$

Suppose the firm's endogenous adjustment choice implies a reduced-form relationship

$$\kappa(\phi_\pi) = \bar{K}\,\phi_\pi^{-\gamma}, \qquad \gamma > 0, \tag{17}$$

in the relevant region. Proposition 3 shows that in the policy-dominant regime, the New Keynesian component of losses increases with ϕ_π if and only if $\gamma > 1$.

In the demand-shock-only equilibrium, the elasticity γ is pinned down by the adjustment-cost

exponent ϕ. In particular, the firm's optimality implies

$$n \propto (\sigma_\pi^2)^{\frac{1}{\nu+1}}, \qquad \kappa \approx Kn \;\Rightarrow\; \kappa \propto (\sigma_\pi^2)^{\frac{1}{\nu+1}}.$$

In the static NK equilibrium, inflation is proportional to the demand shock with coefficient κ/D, so inflation variance scales as $\sigma_\pi^2 \propto (\kappa/D)^2$. Combining these relations yields

$$\kappa \propto \left(\frac{\kappa}{D}\right)^{\frac{2}{\nu+1}}.$$

In the policy-dominant regime $D \approx \kappa\phi_\pi$, this implies

$$\kappa(\phi_\pi) \propto \phi_\pi^{-\frac{2}{\nu+1}},$$

so that the policy elasticity is

$$\gamma = \frac{2}{\nu+1}. \tag{18}$$

Hence $\gamma > 1$ if and only if $\nu < 1$.

Equation (18) implies that concave adjustment costs ($\nu < 1$) generate a super–elastic Phillips curve slope $\kappa(\phi_\pi)$ in the policy–dominant regime. In that case, Proposition 3 implies that the output gap term and the New Keynesian component of losses can rise as ϕ_π increases.

At the same time, adjustment costs (14) fall with ϕ_π whenever $\kappa(\phi_\pi)$ falls with ϕ_π. Thus the sign of $d\mathscr{L}_0/d\phi_\pi$ depends on the relative quantitative importance of $\mathscr{L}_0^{NK}$ and $\mathscr{L}_0^{adj}$. In particular, when $\nu < 1$ and ξ is sufficiently small so that adjustment costs are quantitatively negligible in (16), total losses inherit the comparative statics of $\mathscr{L}_0^{NK}$. The following result is therefore immediate:

Corollary 1. *In the policy-dominant, low-volatility region, welfare losses increase with ϕ_π if $\nu < 1$ and ξ is small so that*

$$\frac{d\mathscr{L}_0}{d\phi_\pi} > 0$$

in the policy-dominant regime.

In sum, allowing firms to endogenously choose their price-adjustment intensity may fundamen-

tally alter the transmission of monetary policy. When price flexibility responds sufficiently elastically to inflation volatility—an outcome that arises naturally when adjustment costs are concave—more aggressive inflation stabilization can flatten the Phillips curve so rapidly that the output gap fails to close, and may even widen, following a demand shock. In this case, the divine coincidence breaks down in a strong sense: tighter inflation targeting lowers inflation volatility but amplifies real fluctuations. Whether this mechanism raises or lowers welfare depends on the relative magnitude of standard NK losses and the resource costs associated with price adjustment. The analytical results characterize sharp thresholds under which this tradeoff emerges and clarify how they depend on structural features of the adjustment technology. In the next section, we embed this pricing block into a quantitative NK model to assess the empirical relevance of these mechanisms and to evaluate their implications for optimal monetary policy.

4 Textbook Calibration

We have shown analytically with demand shocks alone when endogenous price-setting frequency can undermine the divine coincidence. We now consider a standard textbook model with both supply and demand shocks to consider how quantitatively robust are these analytical results. It turns out that the presence of supply shocks strengthens, rather than weakens, our main results. We show that by incorporating endogenous price rigidities, the optimal Taylor rule coefficients change: the central banks should respond less aggressively to inflation and should respond more aggressively to the output gap. We first include the minimal set of additional equations to solve the model numerically in a stochastic steady state.

4.1 Setup

As before, we extend the pricing block of Blanco et al. (2024a) to allow for stochastic inflation volatility, while remaining as close as possible to their original structure. We impose a time–invariant adjustment frequency $n_t \equiv n$ and characterize a stochastic steady state in which firms choose n taking stationary expectations over future inflation. We first describe how to calculate the

endogenous choice of n generally. Let gross inflation follow

$$\pi_{t+1} = \bar{\pi}\exp(u_{t+1}), \qquad \bar{\pi} \equiv e^{\mu} \geq 1,$$

where u_{t+1} is i.i.d. normal with $E[\exp(u_{t+1})] = 1$. A convenient parameterization is

$$u_{t+1} \sim \mathcal{N}\left(-\tfrac{1}{2}\sigma_u^2,\, \sigma_u^2\right),$$

which implies $E[\pi_{t+1}] = \bar{\pi}$. In a linearized equilibrium, these will be determined endogenously in equilibrium.

For any exponent a, lognormality implies $E[\pi_{t+1}^a] = \bar{\pi}^a \exp\left(\frac{1}{2}a(a-1)\sigma_u^2\right)$. Define the two moment objects that enter the pricing block:

$$m_1 \equiv E[\pi^{\theta-1}] = \bar{\pi}^{\theta-1}\exp\left(\tfrac{1}{2}(\theta-1)(\theta-2)\sigma_u^2\right), \tag{19}$$

$$m_2 \equiv E[\pi^{\theta/\eta}] = \bar{\pi}^{\theta/\eta}\exp\left(\tfrac{1}{2}\tfrac{\theta}{\eta}\left(\tfrac{\theta}{\eta}-1\right)\sigma_u^2\right). \tag{20}$$

Endogenous adjustment frequency With $n_t \equiv n$, the auxiliary present–value objects in Blanco et al. (2024a) become constants. The recursions for b_{1t} and b_{2t} imply

$$b_1 = 1 + \beta(1-n)E[\pi^{\theta-1}]\, b_1 = \frac{1}{1-\beta(1-n)m_1}, \tag{21}$$

$$b_2 = p^{-1/\eta} + \beta(1-n)E[\pi^{\theta/\eta}]\, b_2 = \frac{p^{-1/\eta}}{1-\beta(1-n)m_2}. \tag{22}$$

The price–dispersion (misallocation) term satisfies

$$x^{-\theta/\eta} = n\left(\frac{p^*}{p}\right)^{-\theta/\eta} + (1-n)x^{-\theta/\eta}E[\pi^{\theta/\eta}],$$

so that in the stochastic steady state

$$x^{-\theta/\eta} = \frac{n\left(\frac{p^*}{p}\right)^{-\theta/\eta}}{1-(1-n)m_2}. \tag{23}$$

Thus dispersion is strictly positive whenever $m_2 > 1$, which occurs with either positive trend inflation or positive inflation volatility.

The optimal reset price condition implies

$$\left(\frac{p^*}{p}\right)^{1+\theta\left(\frac{1}{\eta}-1\right)} = \frac{1}{\eta}\frac{b_2}{b_1} = \frac{1}{\eta}\,p^{-1/\eta}\,\frac{1-\beta(1-n)m_1}{1-\beta(1-n)m_2}. \tag{24}$$

With $\bar{\pi} > 1$ and/or $\sigma_u^2 > 0$, the optimal reset relative price generally differs from unity.

The firm's optimality condition for the adjustment frequency n is identical in structure to equation (13) in Blanco et al. (2024a), except that deterministic powers of inflation are replaced by their stationary expectations. Replacing their adjusment cost function, which is $\xi/2(n-\bar{n})^2$, with

$$C(n) = \frac{\xi}{\nu}n^{\nu},$$

the stochastic steady–state analog of their equation (13) is

$$\xi n^{\nu-1} = \frac{1}{1-\beta(1-n)m_1}\cdot\frac{1}{n}\left[1-m_1-\tau\eta\,\frac{1-(1-n)m_1}{1-(1-n)m_2}\,(1-m_2)\right]. \tag{25}$$

Equation (25) jointly captures two forces shaping the optimal adjustment frequency. First, positive trend inflation ($\bar{\pi} > 1$) raises m_1 and m_2, generating an incentive to adjust prices more frequently even in the absence of uncertainty, as in the deterministic analysis of Blanco et al. (2024a). Second, inflation volatility ($\sigma_u^2 > 0$) raises m_1 and m_2 through lognormal corrections, increasing expected misallocation losses and further strengthening the incentive to raise n.

When $\sigma_u^2 \to 0$, equation (25) collapses to the non-stochastic steady-state condition in Blanco et al. (2024a). When $\bar{\pi} = 1$, it reduces to the zero-trend stochastic steady-state case analyzed earlier.

Welfare In a standard NK model with exogenous price rigidity, the welfare loss function is

$$\mathscr{L} = \frac{1}{2}\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left(\Xi\cdot\hat{x}_t^2 + \frac{\varepsilon}{\lambda}\pi_t^2\right)\right] \tag{26}$$

which assumes no distortion in steady state (we can use $\hat{x}$), and where $\Xi > 0$ is a coefficient that depends on structural parameters of the model.[4] We can write $\kappa \equiv \lambda \cdot \Xi$.[5] Then dropping the $1/2$, we can write the welfare losses in flow terms as

$$\mathscr{L} = \Xi \operatorname{Var}(\hat{x}_t) + \frac{\Xi \cdot \varepsilon}{\kappa} \operatorname{Var}(\pi_t) \tag{27}$$

and we can divide out Ξ to write

$$\mathscr{L} = \operatorname{Var}(\hat{x}_t) + \frac{\varepsilon}{\kappa} \operatorname{Var}(\pi_t). \tag{28}$$

However, we now need to include the adjustment costs to n in our welfare losses. The cost $C(n) = \frac{\xi}{\nu} n^{\nu}$ is a flow cost paid each period. Since firms make a once-and-for-all choice of n, the welfare losses are given by

$$\mathscr{L} = \frac{1}{2} \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \left(\Xi \cdot \hat{x}_t^2 + \frac{\varepsilon}{\lambda} \pi_t^2 + \frac{\xi}{\nu} n^{\nu} \right) \right]$$

As before, we can right the losses in flow terms as

$$\mathscr{L} = \Xi \left(\operatorname{Var}(\hat{x}_t) + \frac{\varepsilon}{\kappa} \operatorname{Var}(\pi_t) \right) + \frac{\xi}{\nu} n^{\nu}. \tag{29}$$

4.2 Calibration

To fix ideas, I solve the model using the standard textbook parameters around zero trend inflation. I follow Blanco et al. (2024a) and I calibrate ξ so that at the baseline calibration $n^* = 0.297$ in steady-state. I then solve out the standard model with fixed n and the model with endogenous n while varying ϕ_π. I calculate volatilities of output gap and inflation, and the welfare losses, as a function of ϕ_π. I then do the same exercise for ϕ_x. I first set $\nu = 2$ as in Blanco et al. (2024a), and I show how variations in this parameter affects outcomes.

[4] p. 111 in Gali. In Gali, $\Xi = \left(\sigma + \frac{\varphi+\alpha}{1-\alpha}\right)$.

[5] Again, in the textbook we have $\kappa \equiv \lambda \left(\sigma + \frac{\varphi+\alpha}{1-\alpha}\right)$.

Table 1 lists the parameter values in the model. The top panel are taken from the textbook values. The next panel follow Blanco et al. (2024a), and there are a final set that are calibrated.

Table 1: Baseline Parameter Values and Calibration

Texbook Parameter	Value	Description
β	0.99	Household discount factor
σ	1.0	Intertemporal elasticity of substitution^{-1}
ϕ_x	0.50	Output gap coefficient in Taylor rule
ϕ_π	1.5	Inflation coefficient in Taylor rule (policy grid)
ρ_{rn}	0.80	Persistence of natural rate (demand) shock
ρ_u	0.50	Persistence of cost-push shock
σ_{rn}	0.01	Std. dev. of demand shock
σ_u	0.01	Std. dev. of cost-push shock
Parameters from Blanco et al		
n^*	0.297	Target price adjustment frequency (baseline)
ε	6	Elasticity of substitution across goods
η	$2/3$	Labor share parameter
τ	$1-1/\varepsilon$	Steady-state markup wedge
To calibrate/normalized		
ν	2.0 or 1.0	Curvature of price-adjustment cost
ξ	calibrated	Adjustment cost scale (targets n^* at $\phi_\pi = 1.5$)
Ξ	1.0	Phillips curve slope normalization

Disciplining the pricing–flexibility elasticity It remains to discipline the elasticity ν. This section illustrates how existing empirical estimates of monetary policy rules and Phillips–curve slopes can be used to discipline the reduced–form elasticity governing endogenous price flexibility in our model. The goal is not to provide a full structural estimation, but rather to assess whether empirically plausible magnitudes are consistent with the region of the model in which aggressive inflation stabilization can flatten the Phillips curve sufficiently fast to undermine the divine coincidence.

Empirical moments. We focus on two sets of moments: (i) estimates of the inflation response coefficient in the Taylor rule, ϕ_π, across monetary policy regimes; and (ii) estimates of the slope of the Phillips curve, κ, across time.

We consider empirical estimates in which the Taylor coefficient changes by a lot and the PC flattens only minimally. This biases our results toward a more elastic adjustment cost.

For monetary policy, a standard reference is Clarida et al. (2000), who estimate regime–specific Taylor rules for the United States. In their baseline specification, the estimated inflation coefficient is below unity in the pre–Volcker period and well above unity in the post–Volcker period. As a representative calibration, we take

$$\phi_{\pi,\text{pre}} \approx 0.83, \qquad \phi_{\pi,\text{post}} \approx 2.15,$$

consistent with their reported estimates.

For the Phillips curve slope, we use estimates from Hazell et al. (2022), who exploit cross–state variation in unemployment and inflation. Their preferred specification includes time fixed effects to purge movements in long–run inflation expectations. Using their estimates with time fixed effects, a representative pre/post split implies

$$\kappa_{\text{pre}} \approx 0.0107, \qquad \kappa_{\text{post}} \approx 0.0050,$$

corresponding to a flattening of roughly a factor of two.

Implied policy elasticity of the Phillips curve. Suppose that over the relevant range the Phillips curve slope depends on monetary policy according to the reduced–form relationship

$$\kappa(\phi_\pi) = K\phi_\pi^{-\gamma},$$

where $\gamma > 0$ measures the elasticity of price flexibility with respect to inflation stabilization. Using two regimes, the implied elasticity is

$$\gamma = \frac{\log(\kappa_{\text{pre}}/\kappa_{\text{post}})}{\log(\phi_{\pi,\text{post}}/\phi_{\pi,\text{pre}})}.$$

Substituting the values above yields

$$\gamma \approx \frac{\log(2.14)}{\log(2.59)} \approx 0.80.$$

Thus, under these moments, the Phillips curve flattens as inflation stabilization becomes more aggressive, but less than proportionally in elasticity terms.

In the demand-shock-only, policy–dominant environment analyzed above, the model delivers a tight mapping between the reduced–form elasticity γ and the curvature of the price–adjustment cost function. In particular, $\gamma = \frac{2}{\nu+1} \iff \nu = \frac{2}{\gamma} - 1$, where ν governs the curvature of adjustment costs.

Using $\gamma \approx 0.80$ implies $\nu \approx 1.5$, corresponding to convex adjustment costs. This places the economy in the region of the model in which the Phillips curve flattens with stronger inflation stabilization, but not fast enough to generate a perverse widening of the output gap. As a result, we'll consider $\nu = 1$ and $\nu = 2$ for robustness.

Interpretation and caveat. This exercise highlights both the usefulness and the discipline imposed by the model. Using empirically conservative estimates of Phillips curve flattening (such as those in Hazell et al. (2022) with time fixed effects), the inferred elasticity γ is below unity, and the implied adjustment–cost curvature is convex. In contrast, using estimates that imply dramatically larger flattening—for example, specifications without time fixed effects that attribute most low– frequency inflation movements to slack—would mechanically imply $\gamma > 1$ and hence $\nu < 1$, a region that is difficult to reconcile with economically meaningful adjustment costs.

4.3 Results

Figure 1 plots equilibrium volatilities and slopes, varying ϕ_π, with $\nu = 2$. In the theoretical model with demand shocks alone, a higher ϕ_π will not increase output-gap volatility if $\nu > 1$. However, the presence of cost-push shocks changes this result substantially.

With cost-push shocks, in the standard NK model with exogenous price adjustment frequency, a higher ϕ_π increases the volatility of output gaps because the central bank must more aggressively

respond to inflation shocks, thus creating output gaps.

With endogenous price adjustment, a higher ϕ_π actually increases the volatility of output gaps by even more. A higher ϕ_π does generally decrease the volatility of inflation, and this decreases the endogenous price adjustment frequency. But with a lower endogenous price adjustment frequency, the Phillips Curve flattens, and the consequences is more volatile inflation as well as output gaps.

This has important consequences for welfare. In the standard model, a higher ϕ_π is strictly better for consumer welfare, decreases losses. Hence, cost-push shocks here are not sufficient to overturn the Divine Coincidence over this range of ϕ_π (in this calibration, losses do not flatten out until ϕ_π exceeds 9). With endogenous price frequency, however, the optimal ϕ_π is much lower. With $\nu = 2$, the optimal ϕ_π that minimizes losses is 3.60.

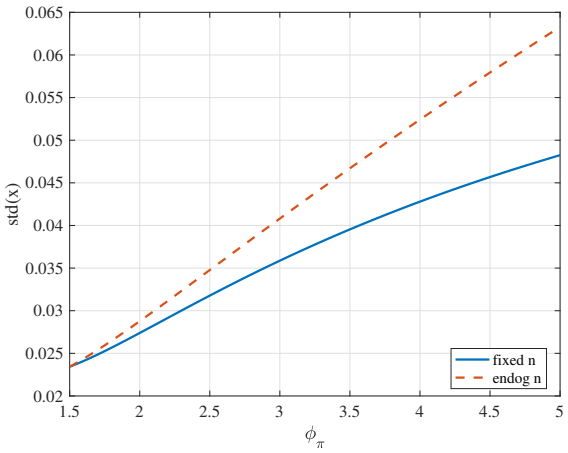
Figure 2 plots equilibrium volatilities and slopes, this time varying ϕ_x. The story changes substantially when monetary policy responds more aggressively to output gaps. In this case, endogenous price adjustment frequency actually increases, and the Phillips Curve steepens. As a result, inflation volatility decreases relative to the standard model. Accordingly, the optimal ϕ_x increases relative to the standard model, increasing from 0.7 to 0.9.

Varying Adjustment Costs Now we set $\nu = 1$. Figure 3 plots equilibrium volatilities and welfare losses varying ϕ_π, and Figure 4 plots equilibrium volatilities and losses varying ϕ_x. While the qualitative results are the same, the quantitative significance increases. In this calibration, firms are much more willing to adjust their price adjustment frequency.

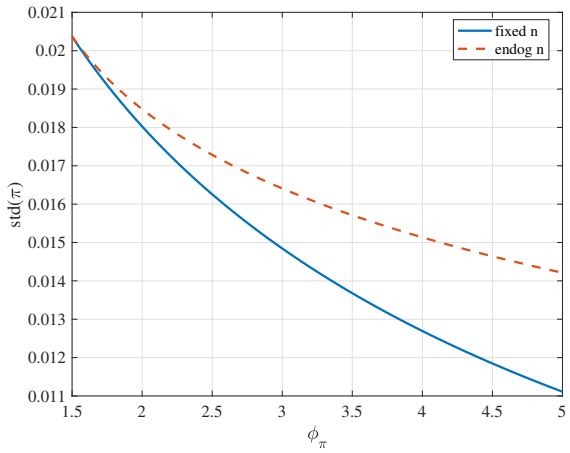
In the model with endogenous price adjustment, the optimal ϕ_π decreases to 1.95. And when $\phi_\pi = 1.5$, the optimal ϕ_x increases to 1.55.

Summary Our textbook calibration has shown that the conditions for our analytical results are actually stronger than required. Even when the price adjustment cost is quadratic, a more aggressive monetary policy response to inflation can sufficiently flatten the PC so that the optimal ϕ_π is much lower than would be true in a model with exogenous frequency of price changes.

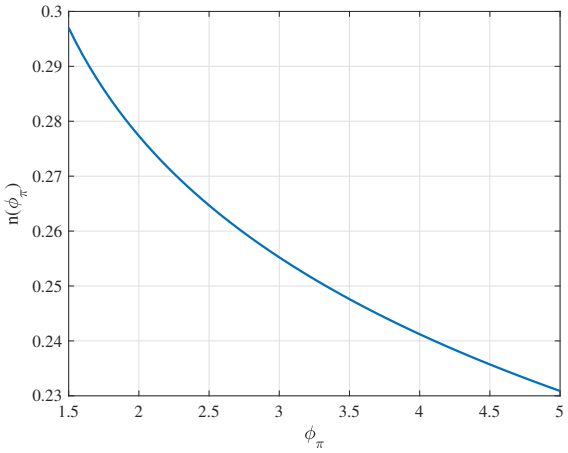
We are currently working on a full-blow quantitative exercise to precisely estimate the quantitative significance of our mechanism in a richer model.



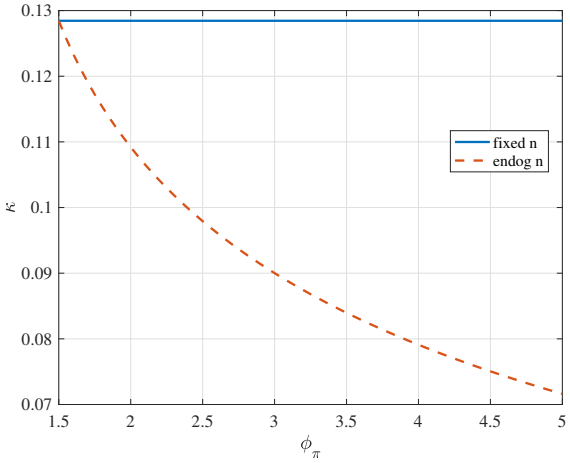
(a) Output gap volatility



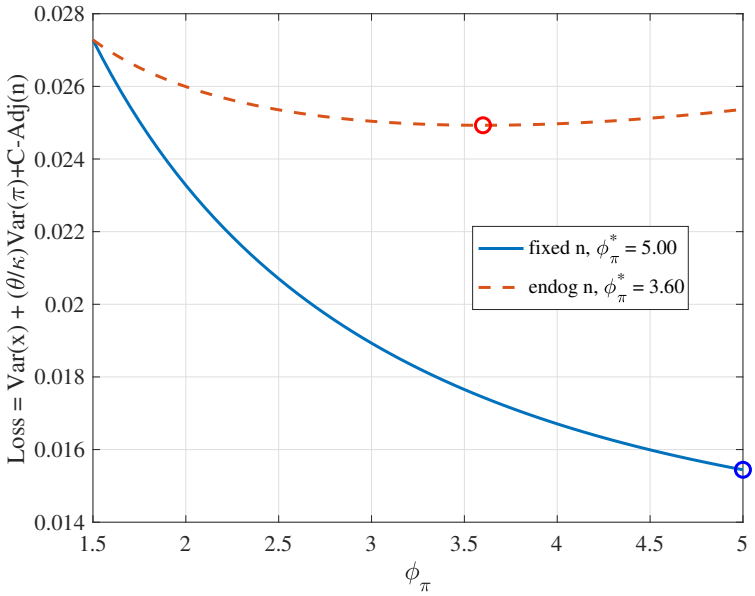
(b) Inflation volatility



(c) Price adjustment frequency $n(\phi_\pi)$

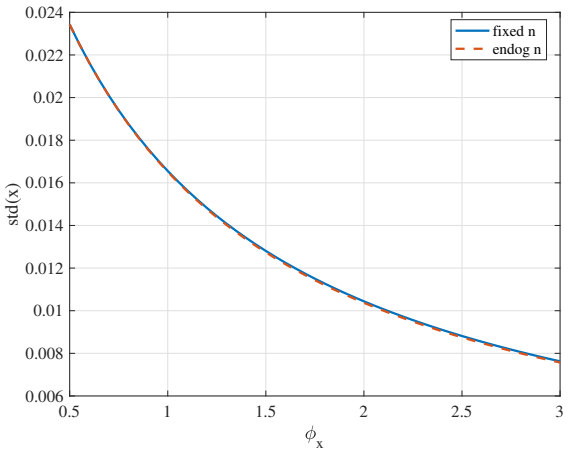


(d) Phillips curve slope κ

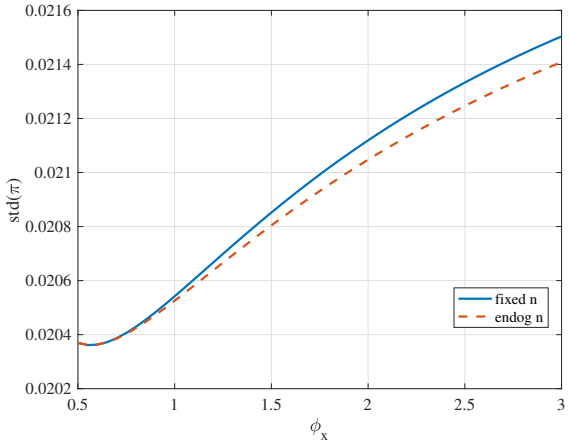


(e) Welfare loss

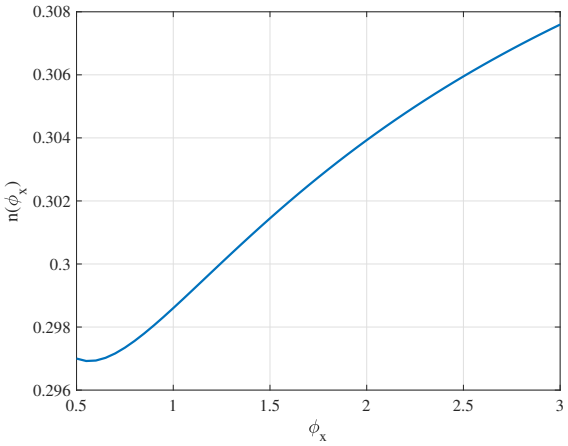
Figure 1: Policy aggressiveness varying ϕ_π, and equilibrium outcomes ($\nu = 2$)



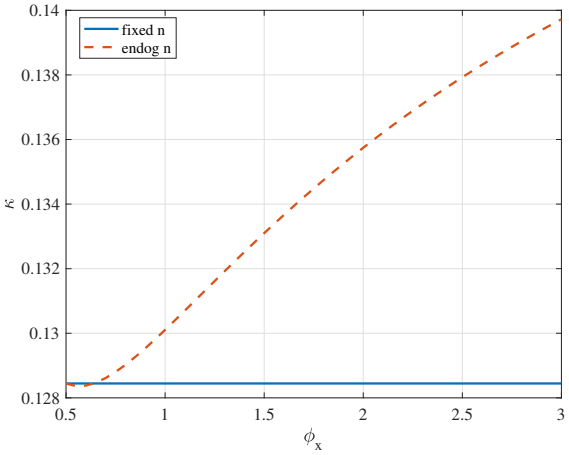
(a) Output gap volatility



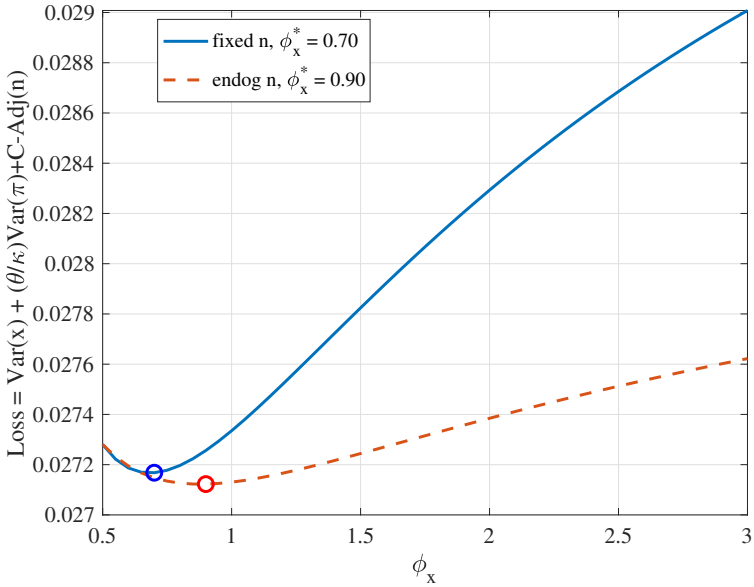
(b) Inflation volatility



(c) Price adjustment frequency $n(\phi_\pi)$

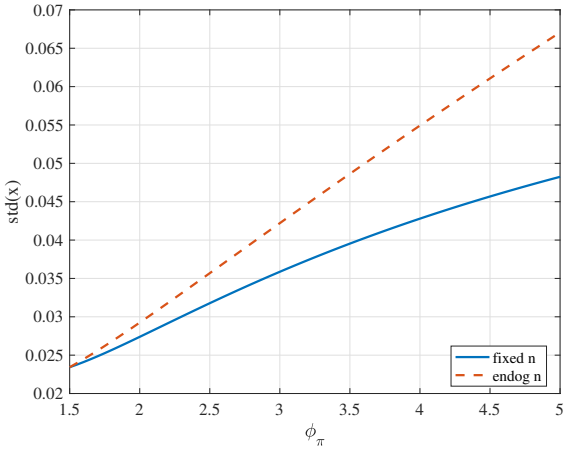


(d) Phillips curve slope κ

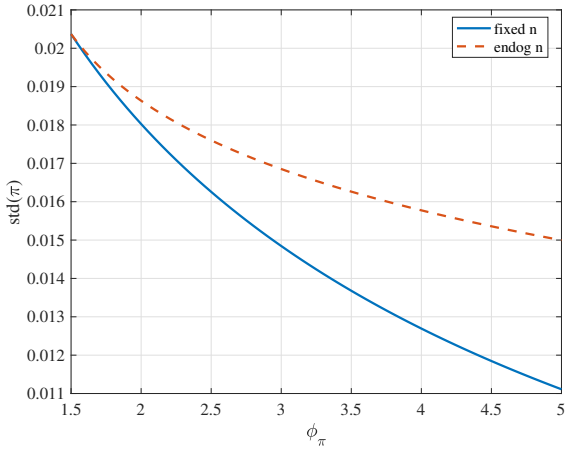


(e) Welfare loss

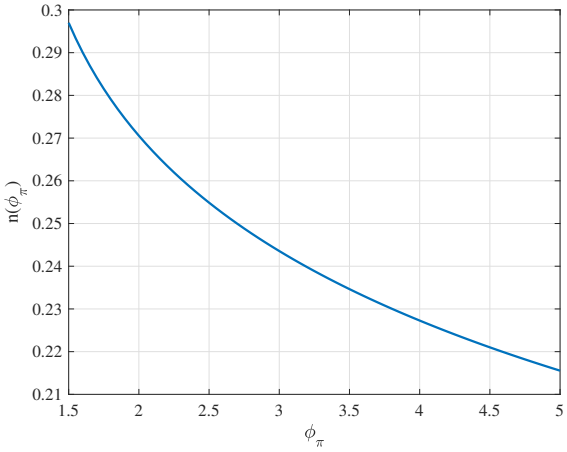
Figure 2: Policy aggressiveness varying ϕ_x, and equilibrium outcomes ($\nu = 2$)



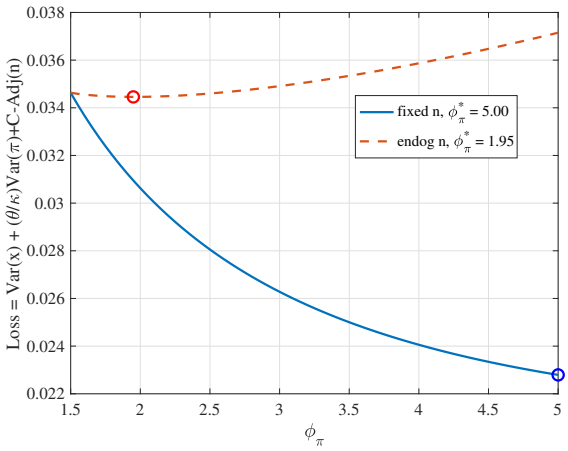
(a) Output gap volatility



(b) Inflation volatility

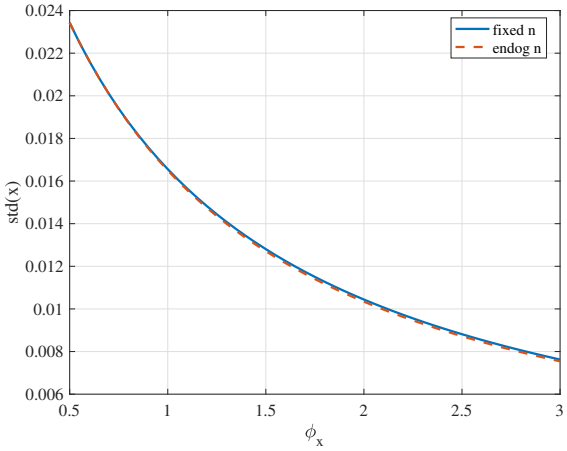


(c) Price adjustment frequency $n(\phi_\pi)$

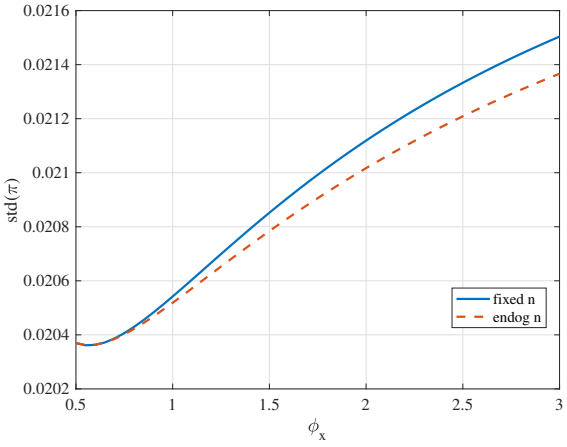


(d) Welfare loss

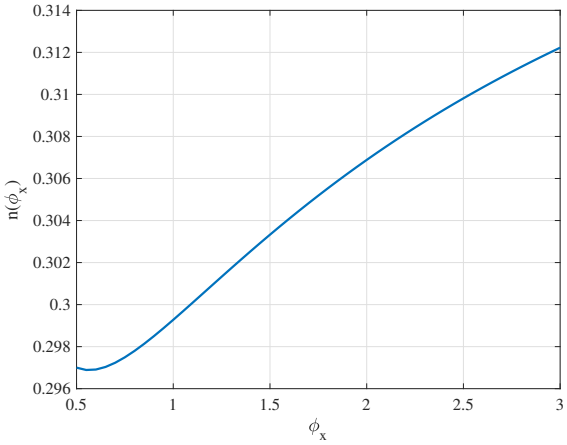
Figure 3: Policy aggressiveness varying ϕ_π, and equilibrium outcomes ($\nu = 1$)



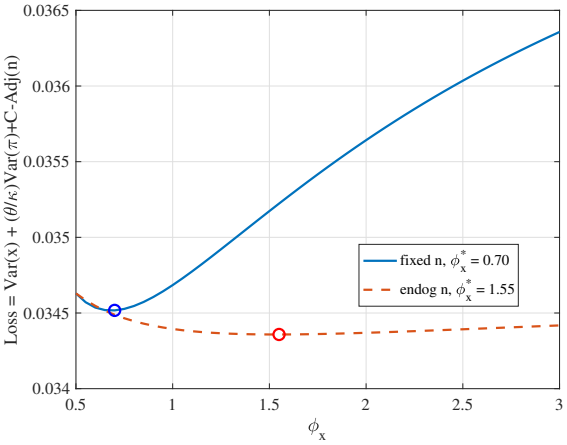
(a) Output gap volatility



(b) Inflation volatility



(c) Price adjustment frequency $n(\phi_\pi)$



(d) Welfare loss

Figure 4: Policy aggressiveness varying ϕ_x, and equilibrium outcomes ($\nu = 1$)

5 Conclusion

What is the optimal mandate for central banks? Our findings challenge the conventional wisdom that price stability alone is sufficient for macroeconomic stability. A mandate focused solely on price stabilization may be problematic because it can inadvertently increase the volatility of output, or lead to persistent output gaps. Moreover, our results suggest that the observed flattening of the Phillips curve may be a consequence of strict price stabilization policies, rather than an exogenous structural change. When inflation becomes very stable, it may lose its effectiveness as a policy target, and price stickiness resulting from strict inflation targeting can amplify output fluctuations. These findings illustrate two distinct ways in which the divine coincidence fails: first, inflation stability does not ensure output stability, and second, policies aimed at price stability may actively worsen output fluctuations. These findings support a dual mandate that explicitly considers output fluctuations alongside inflation.

Most analyses of optimal monetary policy rely on the New Keynesian framework, which assumes an exogenous Phillips curve slope. Our analysis endogenizes price stickiness, demonstrating that monetary policy postures shape the trade-off between inflation and output stabilization. We concentrate on a model in which long-run policy expectations affect the slope of the Phillips curve and long-run inflation expectations *jointly* (in the NK model, there is no *interaction* between the slope and long-run inflation expectations).

References

Alves, S. A. L. (2014). Lack of divine coincidence in new keynesian models. *Journal of Monetary Economics 67*, 33–46.

Auclert, A. (2019, June). Monetary policy and the redistribution channel. *American Economic Review 109*(6), 2333–67.

Benigno, P. and G. Eggertsson (2023). It's baaack: The inflation surge of 2020s and the return of the non-linear phillips curve. *NBER Working Paper No. 31197*.

Benigno, P. and L. A. Ricci (2011, June). The inflation-output trade-off with downward wage rigidities. *American Economic Review 101*(4), 1436–66.

Blanchard, O. and J. Gali (2007). Real wage rigidities and the New Keynesian model. *Journal of Money, Credit and Banking 39*, 35–65.

Blanco, A., C. Boar, C. Jones, and V. Midrigan (2024a). The inflation accelerator. *NBER Working Paper No. 32531*.

Blanco, A., C. Boar, C. Jones, and V. Midrigan (2024b). Non-linear inflation dynamics in menu cost economies. *NBER Working Paper No. 32094*.

Boivin, J. and M. P. Giannoni (2006, 08). Has monetary policy become more effective? *The Review of Economics and Statistics 88*(3), 445–462.

Clarida, R., J. Gali, and M. Gertler (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *The Quarterly Journal of Economics 115*(1), 147–180.

Cúrdia, V. and M. Woodford (2016). Credit frictions and optimal monetary policy. *Journal of Monetary Economics 84*, 30–65.

Galí, J. (2015). *Monetary policy, inflation and the business cycle: An introduction to the New Keynesian framework*. Princeton University Press.

Giannoni, M. and M. Woodford (2004). Optimal inflation-targeting rules. In *The Inflation-Targeting Debate*, pp. 93–172. University of Chicago Press.

Giannoni, M. P. (2014). Optimal interest-rate rules and inflation stabilization versus price-level stabilization. *Journal of Economic Dynamics and Control 41*, 110–129.

Gitti, G. (2024). Nonlinearities in the regional phillips curve with labor market tightness. *Mimeo*.

Gitti, G. and A. Cerrato (2022). Inflation since covid: Demand or supply? *Mimeo*.

Gornemann, N., K. Kuester, and M. Nakajima (2016). Doves for the rich, hawks for the poor? distributional consequences of monetary policy.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2022, 02). The Slope of the Phillips Curve: Evidence from U.S. States. *The Quarterly Journal of Economics 137*(3), 1299–1344.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

Karadi, P., A. Nakov, G. Nuño, E. Pastén, and D. Thaler (2024). Strike while the iron is hot: Optimal monetary policy with a nonlinear phillips curve. *Mimeo*.

L'Huillier, J.-P., G. Phelan, and W. R. Zame (2022). Flattening of the phillips curve and the mandate of the central bank. Working paper, Brandeis University.

Mackowiak, B., F. Matejka, and M. Wiederholt (2021). Rational inattention: A review. *Journal of Economic Literature 59*(3), 706–773.

Mackowiak, B. and M. Wiederholt (2009). Optimal sticky prices under rational inattention. *American Economic Review 99*(3), 769–803.

Reis, R. (2013, November). Central bank design. *Journal of Economic Perspectives 27*(4), 17–44.

Svensson, L. E. (1999). Inflation targeting as a monetary policy rule. *Journal of Monetary Economics 43*(3), 607 – 654.

Svensson, L. E. O. (2003, June). What is wrong with taylor rules? using judgment in monetary policy through targeting rules. *Journal of Economic Literature 41*(2), 426–477.

Wicksell, K. (1898). *Interest and prices*. English translation by R. F. Khan. London: Macmillan.

Woodford, M. (2003). *Interest and prices*. Princeton University Press.

Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics 56*(S1), S100–S124.

Online Appendices

A Proofs

Proof of Proposition 2. Differentiating the loss function with respect to ϕ_π,

$$\frac{d\mathscr{L}}{d\phi_\pi} = \frac{(\sigma+\kappa\phi_\pi+\phi_x)^2\,\varepsilon\kappa' - (1+\varepsilon\kappa)2\,(\sigma+\kappa\phi_\pi+\phi_x)\,(\kappa+\phi_\pi\kappa')}{(\sigma+\kappa\phi_\pi+\phi_x)^4}. \tag{30}$$

We can focus on the numerator. Factoring out $(\sigma+\kappa\phi_\pi+\phi_x)$, which is positive, we have

$$(\sigma+\kappa\phi_\pi+\phi_x)\,\varepsilon\kappa' > 2(1+\varepsilon\kappa)(\kappa+\phi_\pi\kappa').$$

Now we suppose that $\kappa' < 0$. Dividing by κ,

$$(\sigma+\kappa\phi_\pi+\phi_x)\,\varepsilon\frac{\kappa'}{\kappa} > 2(1+\varepsilon\kappa)\left(1+\frac{\phi_\pi\kappa'}{\kappa}\right).$$

Now the term

$$\frac{\phi_\pi\kappa'}{\kappa}$$

is the elasticity. Since $\kappa = K\phi_\pi^{-\gamma}$, then $\kappa' = -\gamma K\phi_\pi^{-\gamma-1}$ and so

$$\frac{\phi_\pi\kappa'}{\kappa} = \frac{-\phi_\pi\gamma K\phi_\pi^{-\gamma-1}}{K\phi_\pi^{-\gamma}} = -\gamma,$$

and thus we have a super elastic PC when $\gamma > 1$. Plugging this in above, our condition requires

$$\begin{aligned}(\sigma+\kappa\phi_\pi+\phi_x)\,\varepsilon\frac{-\gamma}{\phi_\pi} &> 2(1+\varepsilon\kappa)\,(1-\gamma)\,,\\ -\gamma(\sigma+\kappa\phi_\pi+\phi_x)\,\varepsilon &> 2\phi_\pi(1+\varepsilon\kappa)\,(1-\gamma)\,.\end{aligned}$$

Solving for γ:

$$\gamma > \frac{2\phi_\pi(1+\varepsilon\kappa)}{2\phi_\pi(1+\varepsilon\kappa)-\varepsilon(\sigma+\kappa\phi_\pi+\phi_x)} \equiv \bar{\gamma}.$$

The inequality requires γ sufficiently greater than 1, since the denominator is smaller than the numerator.

□

Proof of Lemma 1. Let J denote the number of periods since the firm's last adjustment. Under a constant adjustment hazard, the distribution of J is geometric:

$$\Pr(J = j) = n(1-n)^j, \qquad j = 0, 1, 2, \ldots,$$

with mean

$$E[J] = \frac{1-n}{n}.$$

Conditional on $J = j$,

$$r_{i,t} = -(u_t + u_{t-1} + \cdots + u_{t-j+1}),$$

so

$$\mathrm{Var}(r_{i,t} \mid J = j) = j\sigma^2.$$

Taking expectation over J yields the stationary cross-sectional price dispersion:

$$\mathrm{Var}(r) = E[J]\sigma^2 = \sigma^2 \frac{1-n}{n}. \tag{31}$$

Thus higher volatility σ^2 and lower adjustment frequency n both increase dispersion in relative prices. □

Proof of Proposition 4. The first-order condition is

$$\frac{\partial L}{\partial n} = \xi \phi n^{\phi-1} + \Lambda \sigma^2 \frac{d}{dn}\left(\frac{1-n}{n}\right) = 0.$$

Since

$$\frac{1-n}{n} = \frac{1}{n} - 1 \quad \Longrightarrow \quad \frac{d}{dn}\left(\frac{1-n}{n}\right) = -\frac{1}{n^2},$$

the FOC becomes

$$\xi \phi n^{\phi-1} - \frac{\Lambda \sigma^2}{n^2} = 0,$$

or equivalently

$$\xi \phi \, n^{\phi+1} = \Lambda \sigma^2.$$

□

Proof of Proposition 5. Substituting the optimal $n(\sigma^2)$ from (8) into (9) yields an endogenous Phillips curve slope:

$$\kappa(\sigma^2) = \Xi \frac{n(\sigma^2)}{1-n(\sigma^2)} \left(1 - \beta\left(1 - n(\sigma^2)\right)\right).$$

For small volatility, $n(\sigma^2)$ is small, so $1 - n(\sigma^2) \approx 1$ and

$$\kappa(\sigma^2) \approx \Xi(1-\beta) n(\sigma^2) = \Xi(1-\beta) \left(\frac{\Lambda}{\xi \phi}\right)^{\frac{1}{\phi+1}} (\sigma^2)^{\frac{1}{\phi+1}}.$$

Consider the small-volatility region in which $n(\sigma^2)$ is small. For small n,

$$1 - n \approx 1, \qquad \frac{n}{1-n} \approx n, \qquad 1 - \beta(1-n) = 1 - \beta + \beta n \approx 1 - \beta.$$

Hence, for low σ^2,

$$\kappa(\sigma^2) \approx \Xi(1-\beta) n(\sigma^2) = \Xi(1-\beta) \left(\frac{\Lambda}{\xi \phi}\right)^{\frac{1}{\phi+1}} (\sigma^2)^{\frac{1}{\phi+1}}.$$

In other words, $\kappa(\sigma^2) \propto (\sigma^2)^{\frac{1}{\phi+1}}$ for small volatility. □

Proof of Proposition 6. Because ϕ_π is large

$$\sigma + \kappa \phi_\pi + \phi_x \approx \kappa \phi_\pi.$$

In the pricing block, the optimal adjustment frequency satisfies

$$n(\sigma^2) = B\left(\sigma^2\right)^{\frac{1}{\phi_c+1}}, \qquad \kappa(n) \approx Kn,$$

so the fixed point is

$$\kappa = A\left(\frac{\kappa}{\sigma+\kappa\phi_\pi}\right)^{\frac{2}{\phi_c+1}}, \qquad A \equiv KB\,\sigma_r^{\frac{2}{\phi_c+1}}, \tag{32}$$

with a unique positive solution for κ. Then (32) yields

$$\kappa(\phi_\pi) \propto \phi_\pi^{-\frac{2}{\phi_c+1}}, \qquad \frac{\partial \ln \kappa}{\partial \ln \phi_\pi} = -\frac{2}{\phi_c+1}.$$

The denominator of x_0 becomes

$$D(\phi_\pi) = \sigma + C\,\phi_\pi^{\frac{\phi_c-1}{\phi_c+1}},$$

so the response to a demand shock behaves as follows:

With convex costs ($\phi_c > 1$), $D(\phi_\pi) \to \infty$ as $\phi_\pi \to \infty$, hence $x_0(\phi_\pi) \to 0$. Hence, a stronger rule improves stabilization because the PC does not flatten very quickly. With linear costs ($\phi_c = 1$), the exponent is zero and so $D(\phi_\pi)$ tends to a constant: $x_0(\phi_\pi) \approx$ constant in this case. The output gap is locally invariant to ϕ_π (knife-edge case). The DC fails, but welfare increases as ϕ_π increases because a lower κ decreases misallocation losses. Finally, with concave costs ($0 < \phi_c < 1$), the exponent is negative; $D(\phi_\pi) \to \sigma$ and so $x_0(\phi_\pi) \to \frac{1}{\sigma} r_0^n$. In this case, a stronger rule *worsens* the output gap because the PC flattens sufficiently fast. This yields the "perverse" comparative static. □

Proof of Proposition 7. When ϕ_x is large, the denominator of ψ_r is dominated by ϕ_x:

$$\sigma + \kappa\phi_\pi + \phi_x \approx \phi_x.$$

Thus the fixed point simplifies to

$$\kappa \approx A\left(\frac{\kappa}{\phi_x}\right)^{\frac{2}{\phi_c+1}},$$

which yields

$$\kappa(\phi_x) \propto \phi_x^{-\frac{2}{\phi_c+3}}. \tag{33}$$

The elasticity is therefore

$$\frac{\partial \ln \kappa}{\partial \ln \phi_x} = -\frac{2}{\phi_c+3}. \tag{34}$$

For any $\phi_c > 0$, the elasticity lies strictly between -1 and 0. With demand shocks only, the output gap satisfies

$$x_0 = \frac{1}{\sigma + \kappa(\phi_x)\phi_\pi + \phi_x} r_0^n.$$

Using (33), the denominator behaves asymptotically like

$$D(\phi_x) = \phi_x + C\phi_x^{1-\frac{2}{\phi_c+3}} = \phi_x \left(1 + C\phi_x^{-\frac{2}{\phi_c+3}}\right),$$

so that

$$x_0(\phi_x) \sim \frac{r_0^n}{\phi_x} \quad \text{as} \quad \phi_x \to \infty.$$

□