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

Jean-Paul L'Huillier, Ernesto Pastén, Gregory Phelan, and William R. Zame*

March 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, Williams College, and UCLA. We thank Olivier Blanchard, Anastasios Karantounias, Yuriy Gorodnichenko, Raphael Schoenle, Slavik Sheremirov, Alp Simsek, Sanjay R. Singh, and Michael Weber for useful conversations. We also thank Zheyuan Cui for helpful research assistance. This paper previously circulated with the title "Flattening of the Phillips Curve and the Mandate of the Central Bank." 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. Bakhshi, Khan, and Rudolf (2007) also derive a general Phillips curve under state-dependent pricing. 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 Quantitative Analysis

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, $\gamma = \frac{2}{\nu+1} \quad \Longleftrightarrow \quad \nu = \frac{2}{\gamma} - 1$, 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.

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.

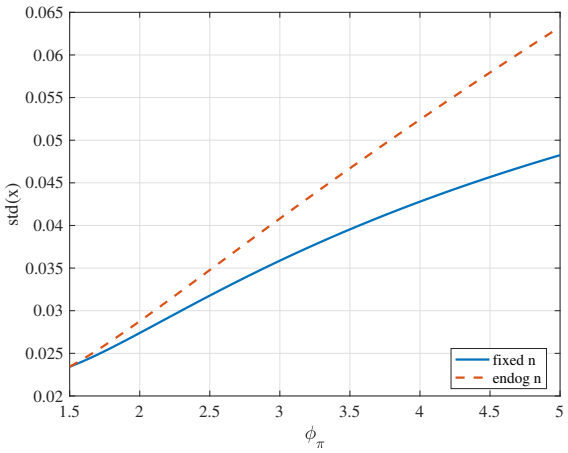
Figure 2 plots equilibrium volatilities and slopes, this time varying ϕ_x. The story changes substantially when monetary policy responds more aggressively to output gaps. In this case, endogenous price adjustment frequency actually increases, and the Phillips Curve steepens. As a result, inflation volatility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

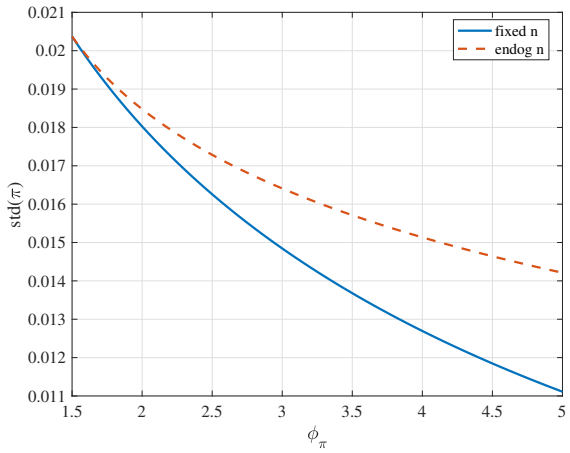
In the model with endogenous price adjustment, the optimal ϕ_π decreases to 1.95. And when $\phi_\pi = 1.5$, the optimal ϕ_x increases to 1.55.

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.

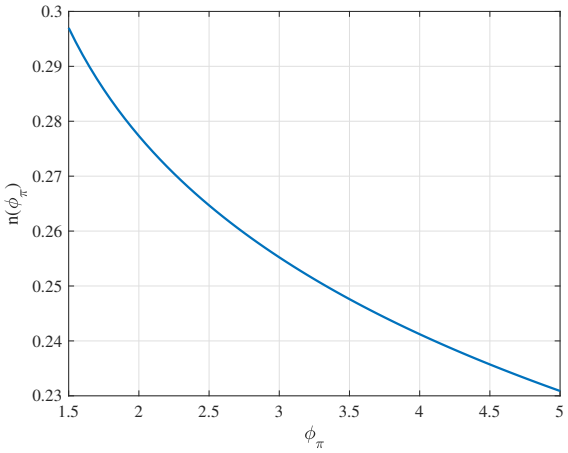
We are currently working on a full-blow quantitative exercise to precisely estimate the quantitative significance of our mechanism in a richer model.



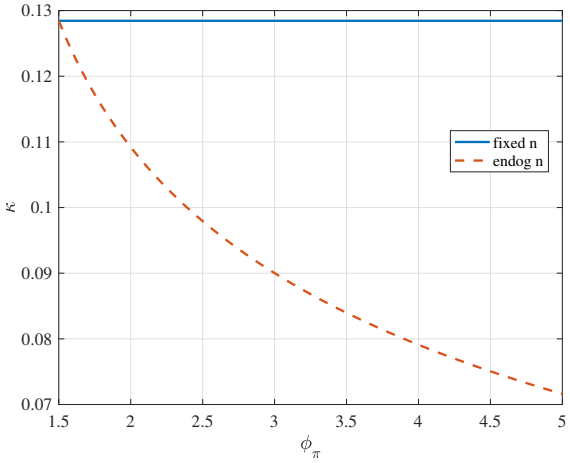
(a) Output gap volatility



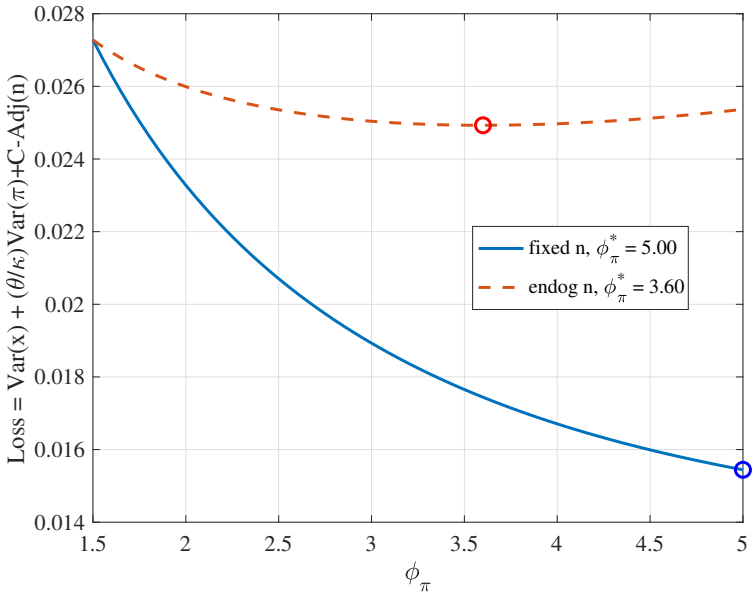
(b) Inflation volatility



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

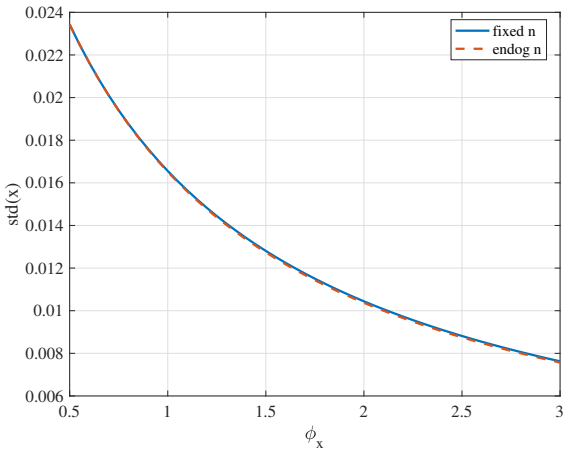


(d) Phillips curve slope κ

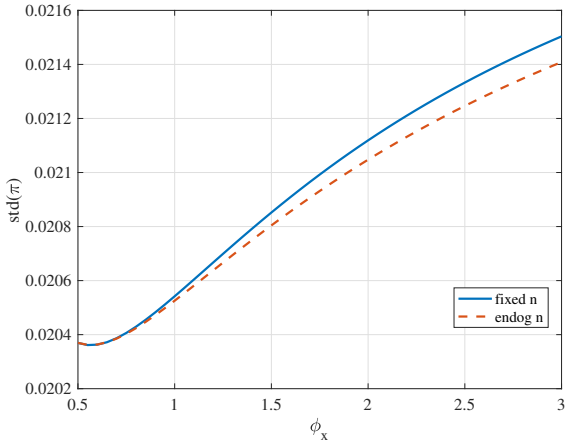


(e) Welfare loss

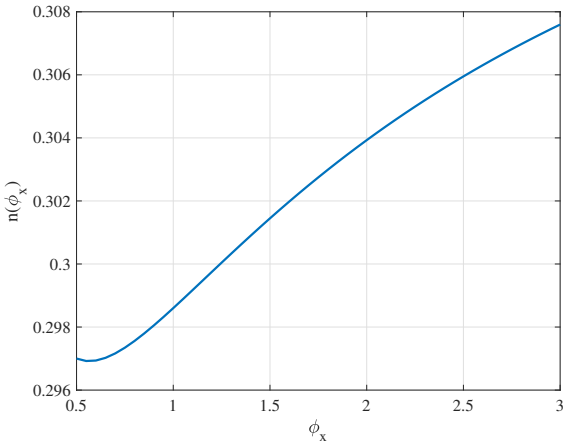
Figure 1: Policy aggressiveness varying ϕ_π, and equilibrium outcomes ($\nu = 2$)



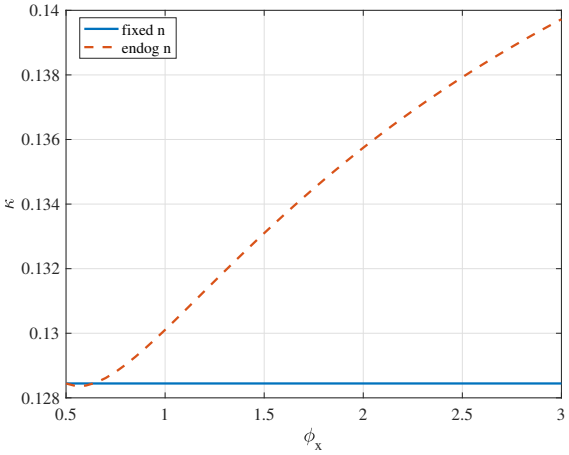
(a) Output gap volatility



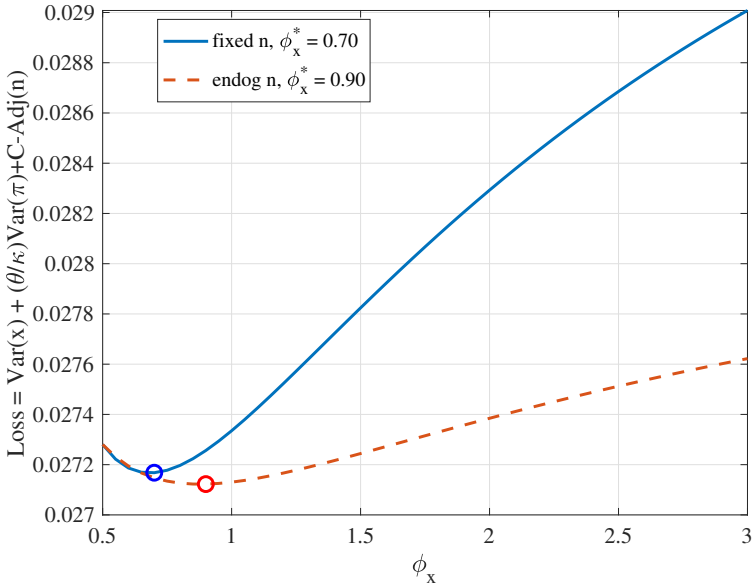
(b) Inflation volatility



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

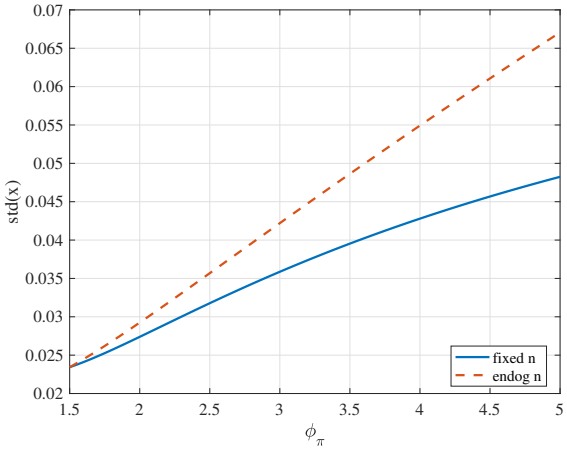


(d) Phillips curve slope κ

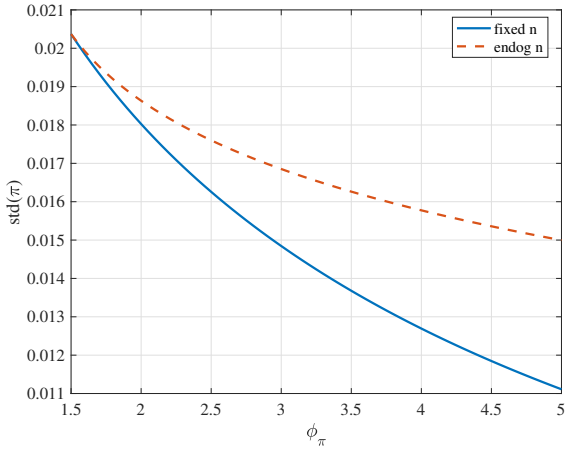


(e) Welfare loss

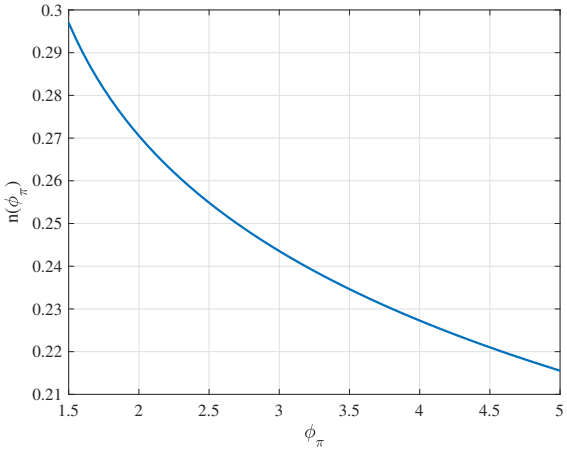
Figure 2: Policy aggressiveness varying ϕ_x, and equilibrium outcomes ($\nu = 2$)



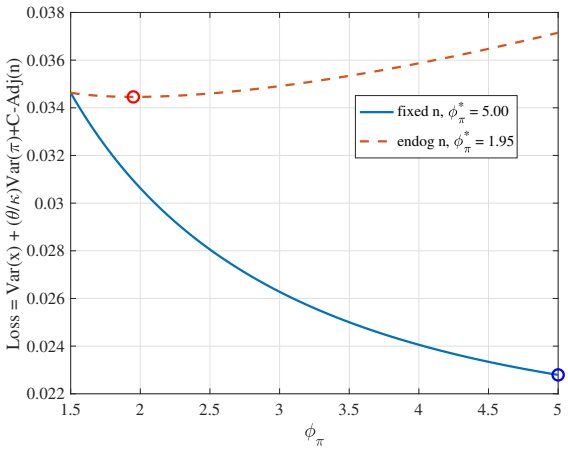
(a) Output gap volatility



(b) Inflation volatility

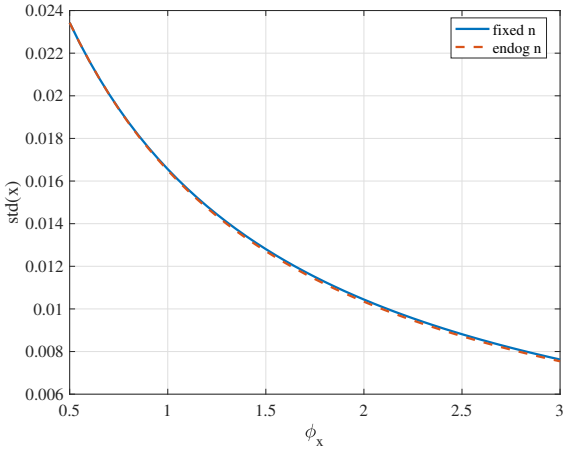


(c) Price adjustment frequency $n(\phi_\pi)$

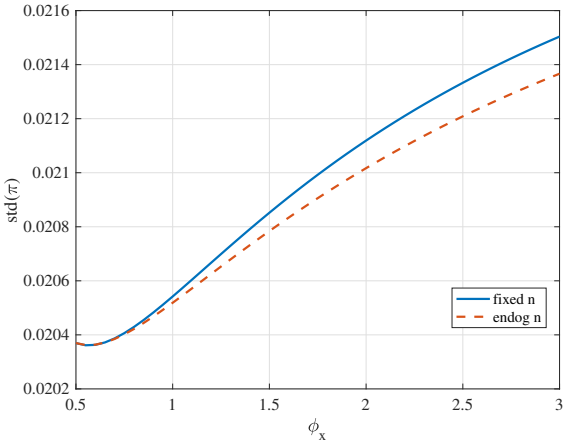


(d) Welfare loss

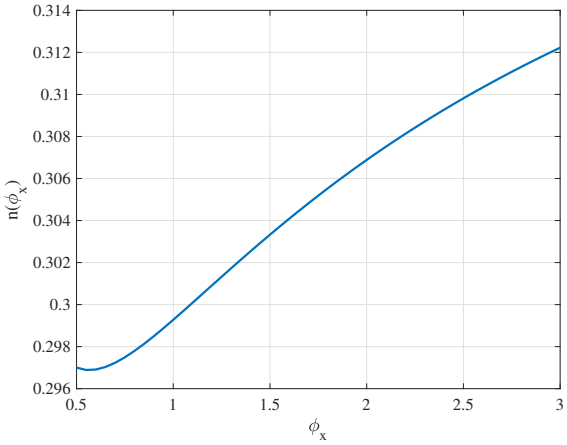
Figure 3: Policy aggressiveness varying ϕ_π, and equilibrium outcomes ($\nu = 1$)



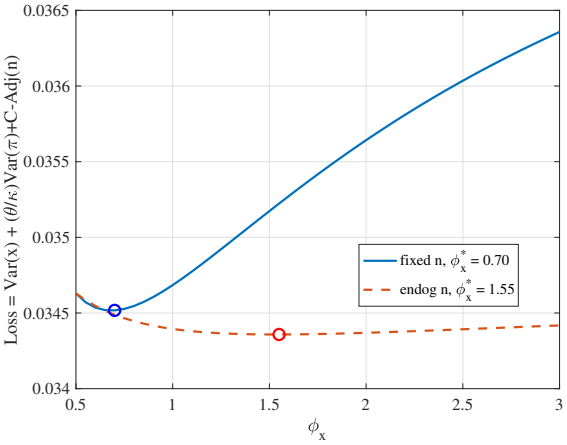
(a) Output gap volatility



(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.

Bakhshi, H., H. Khan, and B. Rudolf (2007). The Phillips curve under state-dependent pricing. *Journal of Monetary Economics 54*(8), 2321–2345.

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*.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018). Really uncertain business cycles. *Econometrica 86*(3), 1031–1065.

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.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020). What's up with the Phillips curve? *BPEA Conference Drafts*.

Diamond, D. W. and P. H. Dybvig (1983). Bank Runs, Deposit Insurance, and Liquidity. *The Journal of Political Economy 91*(3), pp. 401–419.

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*.

Lagos, R. and R. Wright (2005). A unified framework for monetary theory and policy analysis. *Journal of Political Economy 113*(3), 463–484.

L'Huillier, J.-P. (2020). Consumer imperfect information and endogenous price rigidities. *American Economic Journal: Macroeconomics (forthcoming)*.

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.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review 99*(5), 2050–84.

Lucas, R. E. and N. L. Stokey (1987). Money and interest in a cash-in-advance economy. *Econometrica 55*(3), 491–513.

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

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

□

B Strategically Sticky Prices

We present a model in which price stickiness arises from the strategic behavior of firms reacting to an informational asymmetry between firms and consumers. We show how monetary policy influences price stickiness, and hence, the slope of the Phillips curve. For ease of exposition, we offer a two-period, partial-equilibrium model. The same points could be made in a much more complicated infinite-horizon, general-equilibrium model (see Appendix F).

B.1 Environment

There are two dates, the present and the future, which we interpret as the short run and the long run. In the short run, there is production and trading in goods markets will be subject to frictions; in the long run, agents have exogenous endowments and trading will be frictionless. All that follows is

common knowledge.

Setup: Geography, Agents and Markets. The economy is populated by firms, consumers, and a central bank (CB). At each date, firms and consumers trade in a market for a single good. Short-run markets are decentralized; we formalize this by positing a continuum of islands, each served by a single monopolistic firm and populated by a continuum of consumers. Long-run markets are centralized: all consumers trade endowments in a Walrasian, perfectly-competitive market. We denote decentralized-market variables in lowercase and centralized-market variables in uppercase. Thus, the good in the present is c and its price is p; the good in the future is C and its price is P. We normalize the long-run price to $P = 1$. There is a short-run bond market with nominal interest rate i. We posit a cashless economy in which the central bank sets the nominal interest rate i. (Appendix F presents an equivalent model with monetary frictions where the CB sets money supply.) In this partial-equilibrium setup, there is no labor supply.

Aggregate Risk. In the present, there is uncertainty about the aggregate state s, which we model as a shock to consumers' discount rate, denoted by ρ_s, capturing fluctuations in aggregate demand pressure. This modeling device is meant as a proxy for the many possible reasons that "the present" would, all else equal, be a good or bad time for consumers to spend. For simplicity we assume there are only two possible states, low L, and high H, that occur with equal probability. The aggregate state describes the level of demand *in the present*. Hence, demand in the present will be *high* when consumers' discount *rate* is *high* (i.e., $\rho_L < \rho_H$).

The central bank responds to the aggregate state s by setting the nominal interest rate i_s following a Taylor rule that adjusts to deviations in inflation and the output gap:

$$\log(1+i_s^{Taylor}) = \log(1+i_0) + \phi_\pi(\bar{p}_s - \bar{p}_0) + \phi_x(\bar{x}_s - \bar{x}_0) \tag{35}$$

where $\bar{p}_s$ is the aggregate price level (inflation), $\bar{x}_s$ is aggregate output, i_0 is a base interest rate, and ϕ_π and ϕ_x are the Taylor coefficients and $\bar{p}_0$ and $\bar{x}_0$ are base measures of prices and output; hence $\bar{p}_s - \bar{p}_0$ is the deviation of inflation from baseline and $\bar{x}_s - \bar{x}_0$ is the output gap. In equilibrium this

means the CB will endogenously set $i_L \leq i_H$ in response to the endogenous levels of inflation and the output gap.

We could also model the nominal demand shock as a shock to the future price level, or, equivalently, as a shock to money supply that implies a proportional adjustment of prices in the long run, see L'Huillier (2020). We prefer the current formulation in terms of shocks to future marginal utility in order to allow for a standard specification of monetary policy in the quantitative application.

Islands: Consumer Types and Firms. Each island is populated by a continuum of consumers of total mass one and a single monopolistic firm. There are two types of consumers: Insiders (informed consumers) $\iota \in I$ and Outsiders (uninformed consumers) $o \in O$. Insiders are perfectly informed about the state; Outsiders are uninformed about the state but know the probability distribution, and may draw inferences from the price set by the firm with which they trade. The fraction $\alpha \in [0,1)$ of Insiders on a particular island varies across islands. We assume the distribution of α is given by a cdf F whose support is not a singleton and has the property that $\lim_{\alpha \to 1} F(\alpha) = 1$. That is, the fraction of islands on which all consumers are Insiders is 0.

Each island is inhabited by a single monopolistic firm; for simplicity we assume that firms produce the consumption good at zero cost, an assumption we relax in Appendices F and G. All firms know the true state and the fraction of Insiders on their island; firms can condition the price they set on the true state. The assumption that Insiders and firms know the true state is just a convenient abstraction of the idea that they are better informed than Outsiders.

Consumer Problem. We index a typical consumer by j. To simplify notation, we drop state-subscripts when not absolutely necessary. Consumers have a real endowment E in the future and receive firm profits d in the present.

All consumers have the same quasilinear utility function $U(c,C) = (c - c^2/2) + \beta\theta C$, where β is a constant satisfying $0 < \beta < 1$, and θ is a random variable embodying the discount factor shock. Equivalently, one can think about θ as shocks to future marginal utility, an interpretation convenient for preserving imperfect information in the short run. The shock θ is meant as a proxy for the many possible reasons that the present would, all else equal, be a "good" or "bad time"

for consumers to spend, e.g., strong labor markets, rosy expectations of income growth, etc.[6] Outsiders are uncertain about the value of their future nominal marginal utility. This potentially unknown shock to future preferences captures uncertainty about future demand in a similar way that preference shocks in Diamond and Dybvig (1983) capture unforeseen liquidity needs.

Consumer j solves

$$\begin{aligned} \max_{c,C} \quad & \mathbb{E}_j[(c - c^2/2) + \beta\theta C] \\ \text{s.t.} \quad & pc + QC = d + QE, \end{aligned}$$

where $\mathbb{E}_j[\,\cdot\,]$ is consumer j's expectation operator at the present (conditioned on information available to that consumer), p is the nominal price the consumer faces in the short run for goods, and $Q \equiv \frac{1}{1+i}$ is the nominal price for bonds. The natural rate of interest ρ is defined implicitly by the equation $\beta\theta \equiv \frac{1}{1+\rho}$.[7]

Consumer j's optimal short-run demand will be:

$$c^* = 1 - p\,\mathbb{E}_j\left[\frac{\beta\theta_s}{Q_s}\right] \equiv 1 - p\,\mathbb{E}_j\left[\frac{1}{\xi_s}\right]$$

where

$$\frac{1}{\xi_s} \equiv \frac{\beta\theta_s}{Q_s} = \frac{1+i_s}{1+\rho_s}$$

is the ratio of the consumers' and the market discount rates. Note that when ξ_s is high, demand in the present is high, and thus ξ_s acts as a demand shifter. Hence, $\mathbb{E}_j\left[\frac{1}{\xi_s}\right]$ is the net demand shock determining whether demand in the present is strong or weak. In the absence of any additional

[6] We do not take the shock literally, but rather suppose that it captures the various economic mechanisms that would make nominal consumption relatively attractive in the present relative to the future. Quadratic utility is needed for analytical tractability since it generates linear demand functions. We can derive qualitatively similar results with general utility functions.

[7] In order to substantiate the assumption that some consumers may have imperfect information about their future marginal utility (equivalently, the aggregate discount factor), Appendix H.1 offers a more involved model where all consumers perfectly know their future marginal utility (or discount factor), but they are subject to a signal extraction problem with aggregate and idiosyncratic shocks. Both models lead to the same conclusions. In addition, Appendix H.2 offers a microfoundation of our future marginal utility shock based on shocks to future endowments, which then affect the realization of future marginal utility. In that case, one can interpret this shock as mimicking imperfect information about future income, similar to the work by Lorenzoni (2009).

information

$$\mathbb{E}_j\left[\frac{1}{\xi}\right] = \frac{1}{\xi_0}$$

where

$$\xi_0 \equiv \left[\frac{1}{2}\left(\xi_L^{-1} + \xi_H^{-1}\right)\right]^{-1}$$

is the harmonic mean of ξ_L, ξ_H.

B.2 Equilibrium Short-Run Prices

In this section we take a partial-equilibrium approach to solving for equilibrium where we take the nominal interest rates i_L, i_H as given, and thus the demand shifters ξ_L, ξ_H are exogenous. We relax this assumption below when we endogenize monetary policy.

Because each firm is a monopolist on its island, each firm sets the price for the consumption good. If firms faced no strategic concerns, they would simply set the optimal monopoly price given perfect information. It is easy to check that the perfect information monopoly (nominal) price and the resulting demand and (nominal) profit are

$$p_s = \frac{\xi_s}{2}, \quad x_s(p_s) = \frac{1}{2}, \quad d_s(p_s) = \frac{\xi_s}{4}$$

We abuse language and refer to the price schedule $\{p_s\}$ as the *flexible price*.

For each state s, the price p_s maximizes the firm's profit in state s, so the flexible price is optimal among all price schedules when information is perfect. However, if information is not perfect and consumers and firms behave strategically, this is not the only consideration; we must also ask whether the flexible price is consistent with equilibrium in the implicit game between the firm and the consumers. The question is whether adherence to the flexible price is optimal for the firm. For example: would the firm prefer to charge the price p_H even when the true state is L? It is possible that the firm will be tempted to charge p_H instead of p_L to extract more rents, especially if many consumers are Outsiders. In this context, the appropriate notion of equilibrium is Perfect

Bayesian Equilibrium (PBE) so we should ask whether the flexible price is consistent with some PBE. This guarantees that the firm does not deviate from p_L in state L. Proposition 8 provides a sharp answer to this question.

Proposition 8 (PBE with Flexible Prices)**.** *The flexible price* $\{p_s\}$ *is consistent with some PBE if and only if*

$$\alpha \geq \frac{\xi_L}{\xi_H} \equiv \hat{\alpha}$$

When the fraction of Insiders is high enough, the flexible prices are in fact consistent with an equilibrium.

The opposite end of the spectrum from the flexible price is a price schedule that is the same in both states of the world – a *sticky price*. For simplicity, we focus on a particular sticky price, denoted p_0, which is a natural choice for two reasons: it is the price that would be optimal if no consumers were informed ($\alpha = 0$), and it is the price that would be optimal in the absence of a shock. Of course, we require that p_0 be consistent with some PBE as well; Proposition 9 provides a complete characterization.

Proposition 9 (PBE with Sticky Prices)**.**

(a) The sticky price is $p_0 \equiv \xi_0/2$ *and is independent of the monetary policy rule.*

(b) The sticky price p_0 *is consistent with some PBE if and only if*

$$\alpha \leq \frac{\xi_L}{\xi_0}$$

When the fraction of Insiders is low enough, the sticky price is in fact consistent with an equilibrium. Echoing what we said before: if too many consumers know the state, the sticky price is not a sustainable strategy.

Corollary 1.

(a) For every $\alpha \in [0,1)$ at least one of the sticky price and the flexible price schedule is consistent with PBE.

(b) In the region $\{\alpha : \xi_L/\xi_H) \leq \alpha \leq \xi_L/\xi_0\}$ both *the sticky price and the flexible price schedule are consistent with PBE (but the the firm strictly prefers the flexible price schedule).*

B.3 The Phillips Curve

We define a Phillips Curve in this environment as follows. First, it is instructive to suppose that the CB observes the aggregate shock and responds directly, setting the nominal interest rate as a function of the state. Second, once we specify the behavior of the CB, it is easy to define the endogenous aggregate variables.

Monetary Policy. We now suppose that the CB observes the aggregate shock and sets the nominal interest rate as a function of the state. The CB operates by partially responding to the demand shock, and we parameterize the responsiveness of the CB by $\gamma \in [0,1)$ and specify the CB policy according to

$$1 + i_s^{\gamma} = \gamma \left(\frac{1+\rho_s}{\xi_0} \right) + (1-\gamma)(1+i_0) \tag{36}$$

where $1+i_0$ is the base interest rate. We write the interest rate as i_s^{γ} because it is explicitly parameterized by γ, in contrast to i_s^{Taylor} which is defined by the endogenous levels of inflation and the output gap. This implies that for $\gamma \in [0,1)$,

$$\frac{1+\rho_L}{1+i_L} \equiv \xi_L < \xi_H \equiv \frac{1+\rho_H}{1+i_H}$$

so that, when ξ_s is known, early consumption c is high in the H state and low in the L state.

The intuition for this rule is the following. Note that the CB can completely offset the demand shock by setting $1+i_s = \frac{1}{\xi_0}(1+\rho_s)$ so that $\xi_s = \xi_0$. In contrast, if the CB does not respond to the demand shock, then it would choose a constant interest rate i_0 satisfying $1+i_0 = \frac{1}{\xi_0}(1+\rho_0)$, where $1+\rho_0$ is the harmonic mean of $1+\rho_L, 1+\rho_H$, i.e., $\frac{1}{1+\rho_0} = \theta_0$ and θ_0 is the average of θ_L, θ_H. Hence, γ denotes the degree to which the CB adjusts the nominal interest rate from i_0 toward the

consumers' discount rate ρ_s.[8] Finally, note that γ therefore parametrizes the shock $\xi_s(\gamma)$ according to

$$\begin{aligned}\frac{1}{\xi_s(\gamma)} &= \gamma\left(\frac{1}{\xi_0}\right) + (1-\gamma)\left(\frac{1}{\xi_0}\right)\left(\frac{1+\rho_0}{1+\rho_s}\right) \\ &= \gamma\left(\frac{1}{\xi_0}\right) + (1-\gamma)\left(\frac{1}{\xi_0}\right)\left(\frac{\theta_s}{\theta_0}\right)\end{aligned}$$

and the harmonic mean of $\xi_s(\gamma)$ is ξ_0 for all γ. Propositions 8 and 9 and Lemma 2 can all be written equivalently in terms of $\xi_s(\gamma)$ and we can define $\hat{\alpha}(\gamma)$.

As we discuss below, in equilibrium the Taylor coefficients ϕ_π and ϕ_x will endogenously determine how much the CB changes i_s in response to the shock. For now, we take γ as exogenous and show how variations in γ affect equilibrium. When we characterize the equilibrium behavior of monetary policy, we will see how changes in the Taylor coefficients affect equilibrium.

Aggregate Variables. We now characterize the relevant equilibrium aggregate variables for a given CB policy γ. On each island, the firm chooses between the sticky price p_0 and the flexible price schedule $\{p_s\}$ defined above, subject to the requirement that whatever it chooses should be consistent with PBE and delivers the higher profit. On islands where $\alpha \geq \hat{\alpha}$ the firm will choose the flexible price schedule $\{p_s\}$ and demand will be $x_s(p_s) = 1/2$. On islands where $\alpha < \hat{\alpha}$ the firm will choose the sticky price p_0, and demand will depend on the state. Keeping in mind that the Insiders know the state but the Outsiders do not, we see that on these islands demand will be:

$$x_s(p_0) = \alpha\left[1 - p_0\frac{1}{\xi_s}\right] + (1-\alpha)\left[1 - p_0\frac{1}{\xi_0}\right]$$

(The first term is the demand of the Insiders, who know the true state; the second term is the demand of the Outsiders, who do not know the true state; note that $1 - p_0\frac{1}{\xi_0} = 1/2$.) Having defined the local prices and demands in each state s we define the *average* prices $\bar{p}_s$ and demands

[8] Equation (36) can also be written $\frac{1}{Q_s} = \gamma\left(\frac{1}{\xi_0\theta_s}\right) + (1-\gamma)\left(\frac{1}{\xi_0\theta_0}\right)$ so that the bond price is the γ-weighted harmonic mean of the discount factor θ_s and its expectation, divided by ξ_0.

$\bar{x}_s$ in state s to be:

$$\bar{p}_s = \int_0^{\hat{\alpha}} p_0 \, dF(\alpha) + \int_{\hat{\alpha}}^1 p_s \, dF(\alpha) \tag{37}$$

$$\bar{x}_s = \int_0^{\hat{\alpha}} \left(\alpha \left[1 - p_0 \frac{1}{\xi_s} \right] + (1-\alpha) \left[\tfrac{1}{2} \right] \right) dF(\alpha) + \tfrac{1}{2} \left[1 - F(\hat{\alpha}) \right] \tag{38}$$

In this simple model, output is the average demand $\bar{x}_s$.

We define the *Phillips Curve* to be the two points $(\bar{x}_L, \bar{p}_L)$ and $(\bar{x}_H, \bar{p}_H)$. (Note that, as usual, quantities are on the horizontal axis and prices are on the vertical axis.) We define the *slope* $\kappa(\gamma)$ of the Phillips Curve to be the ratio of the (average) price difference to the (average) demand difference:

$$\kappa(\gamma) = \frac{\bar{p}_H - \bar{p}_L}{\bar{x}_H - \bar{x}_L}$$

Alternatively, we might view the Phillips Curve as the line $PC(\gamma)$ through the two points $(\bar{x}_L, \bar{p}_L)$, $(\bar{x}_H, \bar{p}_H)$ and $\kappa(\gamma)$ as the slope of this line.

We now build notation needed for the characterization of the slope of the Phillips Curve. From equations (37) and (38) we can derive more convenient expressions for the average price difference and average demand difference:

$$\bar{p}_H - \bar{p}_L = \frac{1}{2} (\xi_H - \xi_L) (1 - F(\hat{\alpha})) \tag{39}$$

$$\begin{aligned} \bar{x}_H - \bar{x}_L &= ([1 - p_0/\xi_H] - [1 - p_0/\xi_L]) \int_0^{\hat{\alpha}} \alpha \, dF(\alpha) \\ &= \frac{\xi_0}{2} \left(\frac{\xi_H - \xi_L}{\xi_H \xi_L} \right) \int_0^{\hat{\alpha}} \alpha \, dF(\alpha) \end{aligned} \tag{40}$$

The price difference is the difference in the flexible prices for the measure of firms using flexible prices. The output difference is the difference in demand from Insiders on islands with sticky prices. Hence the slope of the Phillips Curve is

$$\kappa(\gamma) = \left(\frac{\xi_H(\gamma) \xi_L(\gamma)}{\xi_0} \right) \frac{1 - F(\hat{\alpha}(\gamma))}{\int_0^{\hat{\alpha}(\gamma)} \alpha \, dF(\alpha)} \tag{41}$$

We show that the Phillips curve flattens (i.e. the slope $\kappa(\gamma)$ decreases) as γ increases, and it becomes perfectly flat (i.e. $\kappa(\gamma) = 0$) in the limit as $\gamma \to 1$. To state this result, we need to introduce some notation. First, define α_0, α_1 to be the lower and upper limits of the support of F

$$\alpha_0 = \sup\{\alpha \in [0,1] : F(\alpha) = 0\}, \quad \alpha_1 = \inf\{\alpha \in [0,1] : F(\alpha) = 1\}$$

Hence, $[\alpha_0, \alpha_1]$ is the smallest closed interval that contains the support of F: α_0 is the fraction of Insiders on the least informed island, and α_1 is the fraction of Insiders on the most informed island. By assumption, the support of F is not a singleton, so $\alpha_0 < \alpha_1$. Note that $\xi_L(0)/\xi_H(0) = (1+\rho_L)/(1+\rho_H)$ and $\xi_L(1)/\xi_H(1) = 1$. Lemma 2 guarantees that the function $\xi_L(\cdot)/\xi_H(\cdot)$ is strictly increasing, so it maps the interval $[0,1)$ onto the interval $[(1+\rho_L)/(1+\rho_H), 1)$.

Lemma 2.

(a) The ratio $\xi_L(\gamma)/\xi_H(\gamma)$ is strictly increasing in γ and

$$\lim_{\gamma \to 1} \frac{\xi_L(\gamma)}{\xi_H(\gamma)} = 1$$

(b) The product $\frac{1}{\xi_H(\gamma)\xi_L(\gamma)}$ is strictly increasing in γ and

$$\lim_{\gamma \to 1} \frac{1}{\xi_H(\gamma)\xi_L(\gamma)} = \left(\frac{1}{\xi_0}\right)^2$$

This means that $\hat{\alpha}(\gamma)$ is increasing in γ. Hence for $k = 0, 1$ we can define γ_k implicitly by the requirement

$$\frac{\xi_L(\gamma_k)}{\xi_H(\gamma_k)} = \max\{\alpha_k, (1+\rho_L)/(1+\rho_H)\}$$

The Phillips Curve will be entirely vertical for very low γ and entirely flat for very high γ. These thresholds are, respectively, γ_0 and γ_1.

We can now state the main result of this Section in Proposition 10: the Phillips curve flattens (i.e. the slope $\kappa(\gamma)$ decreases) as γ increases, and it becomes perfectly flat (i.e. $\kappa(\gamma) = 0$) when

$\gamma \geq \gamma_1$ and (if $\gamma_1 = 1$) in the limit as $\gamma \to 1$.

Proposition 10 (Flattening of the Phillips Curve)**.**

(a) For $\gamma \in [0, \gamma_0)$: $\kappa(\gamma) = +\infty$ (The Phillips Curve is vertical; all islands choose flexible prices.)

(b) For $\gamma \in [\gamma_0, \gamma_1)$: κ is strictly decreasing (The Phillips Curve flattens as the CB response becomes stronger.)

(c) For $\gamma \in [\gamma_1, 1)$: $\kappa(\gamma) = 0$. (The Phillips Curve is flat if the CB response is strong enough; all islands choose the sticky price.)

(d) $\lim_{\gamma \to 1} \kappa(\gamma) = 0$. (In the limit, the Phillips Curve is flat.)

To understand the mechanism underlying Proposition 10, note that, on all islands where sticky prices prevail, the Phillips Curve (for those islands) is horizontal (has slope 0): prices are independent of the state but demands are not. Conversely, on all islands where flexible prices prevail, the Phillips Curve (for those islands) is vertical (has slope $+\infty$): prices depend on the state but demands do not. Propositions 8 and 9 and Lemma 2, taken together, tell us that as the CB increases its commitment to stabilization (i.e. as γ increases), more firms choose sticky prices and fewer firms choose flexible prices. This correctly suggests that as the CB increases its commitment to stabilization the Phillips Curve flattens, and in the limit is completely flat. Moreover, parts (a) and (c) delineate regimes in which the Phillips Curve is vertical or horizontal, respectively.

This intuition is correct, but it is not the whole story. There is also an "intensive margin" which is affected by the stability mandate. The slope of the Phillips Curve as we have defined it is the ratio of the difference $\bar{p}_H - \bar{p}_L$ of (average) prices across states and the difference $\bar{x}_H - \bar{x}_L$ of (average) demands across states. As γ increases, the flexible price prevails on fewer islands and the sticky price prevails on more islands – but the price differences across states and the demand differences across states *both* shrink, and as γ tends to 1 (the CB's commitment to stability becomes perfect), *both* the price differences and the demand differences tend to 0, so the flattening of the Phillips Curve – which is the ratio of these differences – depends not only on the fraction of firms and consumers that respond to the nominal shock but also on the *magnitudes* of those responses.

In sum, the intensity with which the CB adjusts the policy rate to movements of the price level affects the behavior of price setters. This price setter reaction is dictated by the costs of price adjustment, which are endogenously determined via the strategic interaction with consumers. The firms' problem can be thought as a tradeoff between adjusting prices or not. Because uninformed consumers are concerned that the firm will strategically increase prices in order to increase markups, in equilibrium there is an endogenous cost that arises when firms adjust prices. In the presence of many uninformed consumers, it is therefore optimal for the firms to keep prices unchanged. The key insight in the context of CB policy is that, for a given size of nominal demand shocks, more active monetary policy stabilization favors the sticky price equilibrium. The reason is that a given amount of stabilization (quantified by the appropriate parameters in the model) reduces the *effective* size of shocks. For effectively smaller shocks, it is optimal for a larger fraction of firms to keep prices fixed. This flattens the Phillips Curve, generating muted dynamics of the price level following fluctuations in the output gap.

C Welfare and the Divine Coincidence

Having established how monetary policy affects the slow of the Phillips Curve, we now establish the consequences for welfare-relevant measures. We first illustrate the failure of the divine coincidence in our setting. We then explicitly calculate aggregate welfare as a function of the monetary policy rule. We show that welfare can be non-monotonic in the aggressiveness of the central bank. A central bank that is more responsive to shocks can actually worsen welfare outcomes. Section D quantifies these theoretical predictions.

C.1 The Divine Coincidence

The standard NK model features a so-called divine coincidence in which the central bank can close the output gap by setting inflation to zero. In the standard Calvo setting, some firms are unable to adjust prices in response to a shock, and implementing zero inflation is sufficient to ensure that firms with (exogenously) sticky prices do not have price distortions relative to the aggregate. The

same is not necessarily true in our model. Instead, closing the inflation gap can actually increase the output gap.

In our model, a useful measure of the variation in the aggregate price level is difference in average prices across states, $\bar{p}_H - \bar{p}_L$, which we call simply the *price gap*.[9] The price gap is given in equation (39). Similarly, a useful measure of the variation in output is the difference in demand across states, $\bar{x}_H - \bar{x}_L$, which we call simply the *output variation*. Note that the *average output gap* across states is the expected value of $|\bar{x}_s - \bar{x}_0|$. Note that $\mathbb{E}[|\bar{x}_s - \bar{x}_0|] = \frac{\bar{x}_H - \bar{x}_L}{2}$, half our measure of output variation. Output variation is given by equation (40), which is conveniently re-written as

$$\bar{x}_H - \bar{x}_L = \frac{\xi_0}{2}(\zeta_L - \zeta_H)\int_0^{\hat{\alpha}} \alpha \, dF(\alpha), \text{ where } \zeta_s(\gamma) \equiv \frac{1}{\xi_s(\gamma)}.$$

It is trivial to show that increasing γ always decreases the price gap, i.e., more aggressive monetary policy decreases deviations in inflation, because $\hat{\alpha}$ is increasing in γ and $(\xi_H - \xi_L)$ is decreasing in γ. In contrast, increasing γ could *increase* the output variation, and hence the output gap in each state increases. Dropping the $\frac{\xi_0}{2}$ term, we have

$$\frac{d[\bar{x}_H - \bar{x}_L]}{d\gamma} = \underbrace{(\zeta_L' - \zeta_H')\int_0^{\hat{\alpha}} \alpha dF(\alpha)}_{\text{intensive margin}} + \underbrace{(\zeta_L - \zeta_H)\,\hat{\alpha} f(\hat{\alpha})\frac{d\hat{\alpha}}{d\gamma}}_{\text{extensive margin}}.$$

Note that there are two terms. The output gap is caused by sticky price islands. There is an intensive margin: increasing γ decreases the demand gap for informed agents, thus decreasing the output gap on sticky-price islands. There is also an extensive margin: increasing γ increases the set of sticky-price islands. If the extensive margin is sufficiently strong, then the distortion from price stickiness increases and the output gap increases even as the price gap decreases.

Proposition 11 (Divine Coincidence)**.** *The output variation $\bar{x}_H - \bar{x}_L$ is increasing in γ if and only if*

$$(1-\gamma)\,\hat{\alpha} f(\hat{\alpha})\frac{d\hat{\alpha}}{d\gamma} > \int_0^{\hat{\alpha}} \alpha dF(\alpha). \tag{42}$$

[9] Note that the *average price gap* across states is the expected value of $|\bar{p}_s - p_0|$ where p_0 is taken as the base price level. Note that $\mathbb{E}[|\bar{p}_s - p_0|] = \frac{\bar{p}_H - \bar{p}_L}{2}$, half our measure of the price gap.

In this case the output gap is increasing in γ and the divine coincidence fails.

Note that for sufficiently high γ it must be that $(1-\gamma)\,\hat{\alpha} f(\hat{\alpha})\frac{d\hat{\alpha}}{d\gamma} < \int_0^{\hat{\alpha}} \alpha dF(\alpha)$ (since $1-\gamma$ shrinks) and hence for sufficiently high γ the divine coincidence will hold.

We refer to Proposition 11 as a *strong* failure of the divine coincidence, which we contrast with a *weak* failure below (Corollary 2). In this case of a strong failure, the CB moving interest rates by more to offset the shock actually leads to larger output gaps, even though monetary policy offsets the shock by more, thus decreasing the overall residual demand shock ξ_s. The reason for this strong failure is that a greater measure of firms choose sticky prices, thus creating demand distortions. Over this region of γ, more aggressive monetary policy closes the inflation gap but increases the output gap.

C.2 Aggregate Welfare and Monetary Policy

We now calculate aggregate welfare in the model with strategically sticky prices and consider when more aggressive monetary policy (higher γ) is likely to increase welfare. First, note that future consumption is

$$C_s^* = E + \frac{d_s - pc^*}{Q_s}.$$

In equilibrium aggregate future consumption is $C = E$ and we can ignore the distribution of profits due to linearity. We then calculate welfare for agents based on c^*, C^* by islands and aggregate. We summarize our results as follows.

Proposition 12. *Individual welfare in the economy is given as follows*

(a) Agents facing flexible prices have expected utility $E + \frac{3}{8}$,

(b) Uninformed agents facing sticky prices have expected utility $E + \frac{3}{8}$,

(c) Informed agents facing sticky prices have expected utility $E + \frac{1}{2} - \frac{1}{8}\xi_0^2\mathbb{E}[\zeta_s^2] \leq E + \frac{3}{8}$, with strict inequality if $\gamma < 1$.

Furthermore, $\xi_0^2\mathbb{E}[\zeta_s^2]$ is strictly decreasing in γ.

This means that the welfare consequence of sticky prices is to hurt the Insiders. Insiders have lower expected utility on islands with sticky prices compared to islands with flexible prices and would prefer to be on islands with flexible prices.

Aggregating the expected utilities for agents across types and islands, aggregate welfare is

$$\overline{W}(\gamma) = E + \frac{3}{8} + \frac{1}{8}\int_0^{\hat{\alpha}} \alpha \left(1 - \xi_0^2 \mathbb{E}[\zeta_s^2]\right) dF(\alpha) \tag{43}$$

where the integral is strictly negative for $\gamma < 1$. Differentiating (43), the net change welfare from an increase in γ is given by

$$\frac{d\overline{W}(\gamma)}{d\gamma} = \underbrace{\left(1 - \xi_0^2 \mathbb{E}[\zeta_s^2]\right) \hat{\alpha} f(\hat{\alpha}) \frac{d\hat{\alpha}}{d\gamma}}_{\text{extensive margin}<0} + \underbrace{\left|\frac{d\xi_0^2 \mathbb{E}[\zeta_s^2]}{d\gamma}\right| \int_0^{\hat{\alpha}} \alpha dF(\alpha)}_{\text{intensive margin}>0}. \tag{44}$$

There are two consequences on aggregate welfare from increasing γ.

First is an effect from the extensive margin: the sticky price threshold increases, implying more informed agents face sticky prices, which lowers their utility. The extensive margin is given by

$$\hat{\alpha} \left(1 - \xi_0^2 \mathbb{E}[\zeta_s^2]\right) f(\hat{\alpha}) \frac{d\hat{\alpha}}{d\gamma},$$

which is strictly negative. All islands with $\alpha < \hat{\alpha}$ choose sticky prices. As shown, $\hat{\alpha}$ is strictly increasing in γ. Since informed agents have lower utility in islands with sticky prices than with flexible prices, this force decreases aggregate welfare.

Second is an effect from the intensive margin: the welfare of informed agents on islands with sticky prices increases. The intensive margin is given by

$$-\int_0^{\hat{\alpha}} \alpha \frac{d\xi_0^2 \mathbb{E}[\zeta_s^2]}{d\gamma} dF(\alpha),$$

which is strictly positive. This force increases aggregate welfare.

The change in welfare overall is likely to be positive when $\int_0^{\hat{\alpha}} \alpha dF(\alpha)$ is large and $f(\hat{\alpha})$ is small, which is more likely when $\hat{\alpha}$ is small and the distribution of α has a lot of mass at high α.

Thus, the total effect of increasing γ on welfare depends on how many agents are already on islands with sticky prices. If increasing γ is going to decrease welfare, it must be that the extensive margin is very powerful: a lot of islands with informed agents will adopt sticky prices. This implies that $f(\hat{\alpha})$ must be very big. Specifically, we have the following result:

Proposition 13. *Aggregate welfare is globally maximized by setting $\gamma = 1$ but is locally increasing in γ if*

$$(1-\gamma)\,\hat{\alpha} f(\hat{\alpha}) \frac{d\hat{\alpha}}{d\gamma} < 2\int_0^{\hat{\alpha}} \alpha dF(\alpha). \tag{45}$$

Hence, there exists $\bar{\gamma} < 1$ such that welfare is increasing for $\gamma \in (\bar{\gamma}, 1)$.

Local to $\gamma = 1$, it must be that welfare is increasing, since the left side of equation (45) includes $1-\gamma$. Note that comparing equations (42) and (45), if welfare is decreasing in γ, then the divine coincidence also fails.

This result implies that in the absence of perfect monetary policy that can completely offset aggregate shocks, it is possible that welfare is improved when a central bank is less responsive to shocks hitting the economy. We show in Section D that this may be an empirically relevant result.

D Policy Implementation

We now show that the policy mandate of the central bank can have direct consequences for aggregate volatility and welfare. A dual mandate, targeting output gaps alongside inflation, robustly performs better than a mandate that targets inflation alone. We calibrate our model to quantify the significance of our theoretical mechanisms.

D.1 Policy Implementation and the Phillips Curve

In practice, monetary policy rules typically respond to output gaps and to inflation rather than to a directly observable shock. Hence, because policy makers do not observe shocks directly, the policy rule responds indirectly through observable macro variables. We can write the Taylor rule

as a implicit function of γ:

$$\log(1+i_s^{Taylor}) = \log(1+i_0) + \phi_\pi(\bar{p}_s(\gamma) - p_0) + \phi_x(\bar{x}_s(\gamma) - \bar{x}_0)$$

where the price level $\bar{p}_s(\gamma)$ and output $\bar{x}_s(\gamma)$ are implicitly functions of γ, which is the responsiveness of the CB defined in equation (36).

In equilibrium, the CB responsiveness γ^* is determined by the fixed point of equations (35) and (36). In other words, there exist γ^* and i^* such that at γ^*, $i_s^\gamma(\gamma^*) = i_s^{Taylor}(\gamma^*) = i^*$, where i_s^γ is defined by (36), which is an explicit function of γ, and i_s^{Taylor} is defined by (35), which is an implicit function of γ.

Consider the high state. Equation (36) defines i_s^γ as an increasing function of γ for state H. Furthermore, $\bar{p}_s(\gamma)$ is monotonically decreasing in γ in H. Thus, if ϕ_π is sufficiently large relative to ϕ_x, then i^{Taylor} is decreasing in γ, ensuring that there is a unique equilibrium γ^*. In general, $\bar{x}_s(\gamma)$ need not be monotonically decreasing in γ (when the divine coincidence fails); however, for high γ the output gap converges to zero. Hence, we are guaranteed to have at least one intersection between equations (35) and (36). A sufficient degree of inflation targeting ensures that the intersection is unique. We can state the following result.

Corollary 2 (MP Implementation)**.** *Suppose $\phi_x = 0$ and $\phi_\pi > 0$. Then i_H^{Taylor} is monotonically decreasing in γ and strictly greater than i_0 at $\gamma = 0$ (i_L^{Taylor} is monotonically increasing and strictly less than i_0 at $\gamma = 0$). Hence there is a unique $\gamma^* \leq \gamma_1$ given by the intersection of equations (35) and (36). In addition*

(i) the equilibrium γ^* is increasing in ϕ_π,

(ii) the equilibrium slope of the PC is decreasing in ϕ_π,

(iii) implementing $\gamma^* > \gamma_1$ requires $\phi_x > 0$.

Corollary 2 is important. First, it means that if the CB pursues a stronger commitment to price stability or inflation targeting, then endogenously the PC becomes flatter. Second, in light of the previous results on the divine coincidence (Proposition 11) and welfare (Proposition 13), a stronger

commitment to price stability or inflation targeting can actually increase output gaps and decrease aggregate welfare. Third, achieving the optimal monetary policy requires a dual mandate. Recall that the output gap is decreasing for sufficiently high γ (even when the divine coincidence might fail for lower γ), and thus it is possible to implement high γ^* with a sufficiently high ϕ_x.

The last result describes what we refer to as the *weak* failure on the divine coincidence. Specifically, a greater commitment to price stability (high ϕ_π) can indeed lead to price stability (no variation in aggregate price levels) while significant output gaps remain. Note that at γ_1 there is precisely zero variation in prices—the CB achieves perfect price stability—and yet non-zero output gaps remain. Hence, in this weak failure, the CB achieves price stability without achieving stability in output.

When the divine coincidence holds in the strong sense defined above (see Proposition 11), then the level of output is monotonic in γ as well. In this case, the usual result holds that the CB can respond to shocks by targeting either the price level or the output gap.

Corollary 3. *Suppose the divine coincidence holds. Then i_H^{Taylor} is monotonically decreasing in γ and strictly greater than i_0 at $\gamma = 0$ for any $\phi_\pi \geq 0$ and $\phi_x \geq 0$ with at least one inequality strict (i_L^{Taylor} is monotonically increasing and strictly less than i_0 at $\gamma = 0$). Hence, γ^* is increasing in both Taylor coefficients ϕ_π and ϕ_x, and hence the PC slope is decreasing in both Taylor coefficients ϕ_π and ϕ_x.*

When it comes to implementing optimal monetary policy in this setting, a policy rule must target output gaps directly in order to implement the optimal ($\gamma = 1$). A rule that targets price stability alone would never implement monetary policy with $\gamma > \gamma_1$, because in this range there is zero variation in prices and so there would be zero variation in monetary policy—the central bank would not move interest rates in the right direction at all to mitigate the shock.

Indeed, since the slope of the Phillips Curve is decreasing in γ, it is robustly easier to achieve more aggressive policy by responding to output gaps. As γ increases, the variation in prices falls faster than the variation in output gaps, thus requiring policy rules that respond even more strongly to inflation. If γ_1 is sufficiently small (given the model parameters), then the central bank may robustly achieve price stability without sufficiently stabilizing fluctuations in output; a price stability

mandate alone would not come close to maximizing welfare. Prices would not fluctuate, but output would.

While it is instructive to consider how the central bank can implement the optimal monetary policy rule ($\gamma = 1$) in practice completely mitigating the shock is likely impossible. The result from Proposition 13 that welfare can be decreasing on the margin means that more aggressive monetary policy, achieving greater price stability, could turn out to decrease welfare. We illustrate policy implementation in greater detail below using our calibrated model. This allows us to consider when more aggressive policy could be welfare decreasing in practice.

D.2 A Calibrated Model with Demand Shocks

We now provide a calibrated model to determine the quantitative significance of our theoretical results. To do so, we modify the setup slightly in order to let the data discipline the degree of heterogeneity. Instead of having a distribution of Insiders across islands, we let firms have a distribution of marginal costs. In this way, we can let empirical estimates discipline the distribution of productivity. We then calibrate the fraction of Insiders to match the estimated slope of the Phillips Curve, given particular Taylor coefficients. Appendix G formalizes the results in this setting with a distribution of marginal costs.

We target average markups to be 12.5%, as is standard in the literature. With an average price of $p = \frac{\xi_0(1+k)}{2}$, we set $p/k = 1.125$ and $\xi_0 = 1$ implying an average marginal cost of $k = \frac{0.5}{1.125-0.5} = 0.8$. We calibrate the distribution of productivities (inverse of marginal costs) as log-normal with standard deviation of 5%. Bloom et al. (2018) find that the unconditional standard deviation of micro-productivity shocks is 5.1%. The mean is set so that the average marginal cost equals 0.8.

The household time preference (natural rate) is set to $r_0 = 4\%$ and the discount factor shock (demand shock) is 1%, hence $\theta_L = \frac{1.01}{1+r_0}$ and $\theta_H = \frac{0.99}{1+r_0}$.

We choose α to target the slope of the Phillips Curve in equilibrium for the given Taylor coefficients. We use estimates from Del Negro et al. (2020) who use 1990 as the break in the sample. Their estimates for the relevant Taylor coefficients pre-1990 are $\phi_\pi = 1.5$ and $\phi_x = 0.1$; the estimates for post-1990 are $\phi_\pi = 1.5$ and $\phi_x = 0.22$. The posterior mean, median, and mode

for κ pre-1990 are 0.0157, 0.0145, and 0.0211. The post-1990 posterior mean, median, and mode are 0.00151, 0.00140, and 0.00196. These estimates are similar to what is found in Hazell et al. (2022). We choose α to target these estimates—hence, our model is over-identified, where we are using two moments to discipline a single parameter.

With $\alpha = 0.88$, the model delivers pre- and post-1990 numbers of

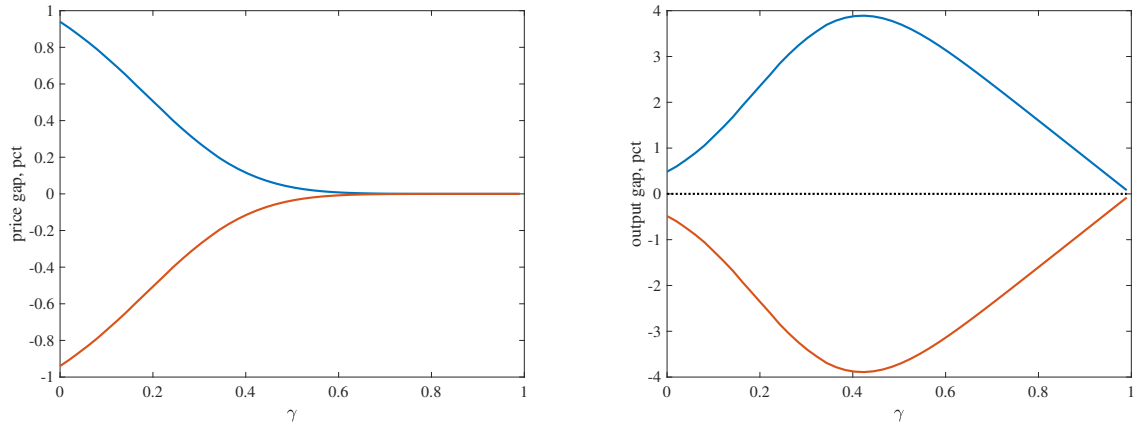
$$\kappa_{pre} = 0.0162, \quad \kappa_{post} = 0.0017$$

Despite its simplicity, the calibrated model can match the decline in the slope of the Phillips Curve based on changes in the Taylor rule using estimated parameters. The equilibrium γ^* is $\gamma^* = 0.456, 0.626$ pre- and post-1990. In this case, welfare improves slightly, from -0.01 to -0.0057 percent, in terms of consumption-equivalent relative to the flexible price equilibrium (hence, really small numbers). Nonetheless, the model exhibits non-monotonic welfare and a strong failure of the divine coincidence for $\gamma < 0.4$. Figure 5 plots equilibrium variables with these parameters.

As is clearly evident with this calibration, increasing γ from 0 initially leads to a large increasing in output fluctuations—the divine coincidence fails for $\gamma < 0.4$. This occurs because the fraction of islands choosing sticky prices increases dramatically. By $\gamma = 0.6$, nearly all islands use sticky prices. Similarly, welfare is non-monotonic in γ, initially declining because higher γ leads to more price stickiness, creating more output distortions. Nonetheless, in this economy, the estimated Taylor coefficients imply that the current Taylor rule is welfare-improving relative to the pre-1990 coefficients.

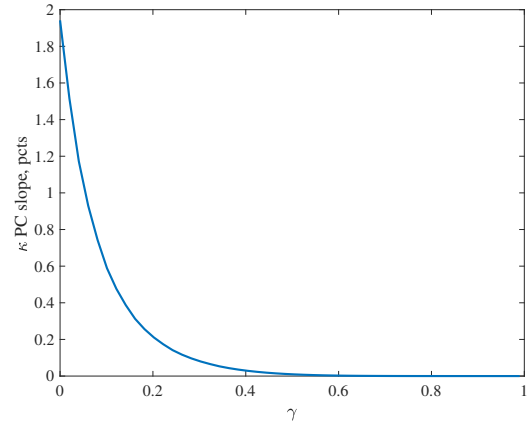
It is instructive to use this calibrated economy to consider the consequences of changes to the policy rule in equilibrium. Our theoretical results imply that a single mandate of targeting inflation alone may not be optimal and may lead to large fluctuations in output even if prices are stable.

We first suppose the CB pursues a single mandate of price stability alone, which we model in the extreme way by setting $\phi_x = 0$. Figure 6 illustrates equilibrium monetary policy in this environment. We consider two sets of Taylor coefficients that both ignore the output gap and respond only to inflation. The blue line has a standard value of $\phi_\pi = 1.5$, and the red-orange line is one hundred

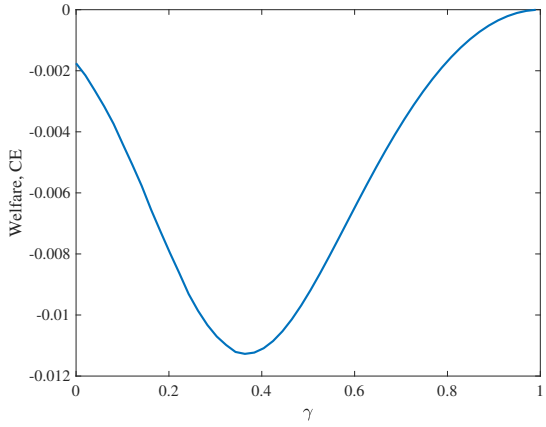


(a) Prices

(b) Output



(c) κ, Phillips Curve slope



(d) Welfare

Figure 5: Aggregate Prices, Output, Phillips Curve, and Welfare.

times more responsive, with $\phi_\pi = 150$. In these cases, the equilibrium responsiveness in each case is respectively $\gamma^* = 0.3255$ and 0.63367. Hence, more aggressive inflation targeting does lead to more aggressive monetary policy. However, it should be clear that a rule that responds to inflation alone will find it very hard to implement $\gamma > 0.8$. Indeed, even if the CB used the incredibly unrealistic rule of $\phi_\pi = 1,500$, the equilibrium γ would still only be 0.739. This illustrates what we call the weak failure of the divine coincidence. The CB pursues a very strong price stability mandate, and achieves price stability, and yet output remains volatile.

Figure 7 illustrates the determination of the CB behavior in this environment when the CB does respond to output gaps. In this case, a dual mandate greatly improves the ability of the CB

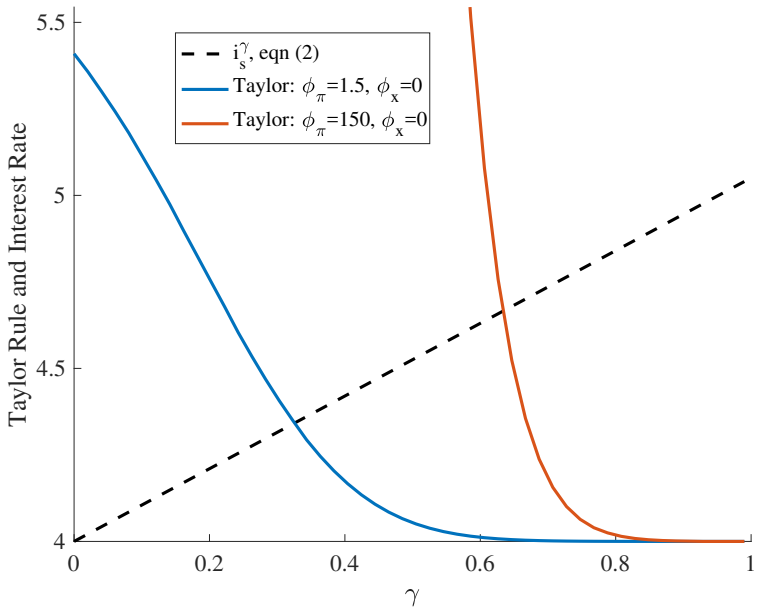


Figure 6: Monetary Policy implementation with a Taylor rule. Greater inflation targeting leads to more aggressive policy, but only to a point. Dashed black line is i_s^γ from equation (36). Solid colored lines are Taylor rules.

to respond to demand shocks. The yellow, purple, and green lines fix the inflation coefficient at $\phi_\pi = 1.5$, the standard value, and vary the responsiveness to the output gap, setting $\phi_x = 0.2$, 0.5, and 1.5. For these Taylor rules, the equilibrium $\gamma^* = 0.626$, 0.7921, and 0.9196, indicating the the CB can achieve much more aggressive responses by targeting the output gap rather than more aggressively targeting inflation.

By responding to the output gap (dual mandate), the CB is able to sidestep the weak failure of the divine coincidence and achieve both price and output stability. Importantly, the CB does not need to respond very aggressively to the output gap in order to achieve better implementation. This occurs precisely because the slope of the PC is endogenous, flattening as the CB pursues more aggressive policy. With more aggressive policy, inflation volatility declines, and thus responding to inflation becomes less effective at guiding policy. Instead, the CB can maintain a standard response to inflation and respond somewhat more aggressively to output. Compared to the baseline single-mandate policy with $\phi_\pi = 1.5, \phi_x = 0$, modestly increasing the responsiveness to the output gap by setting $\phi_x = 0.22$ is nearly as effective as maintaining a single mandate and increasing the

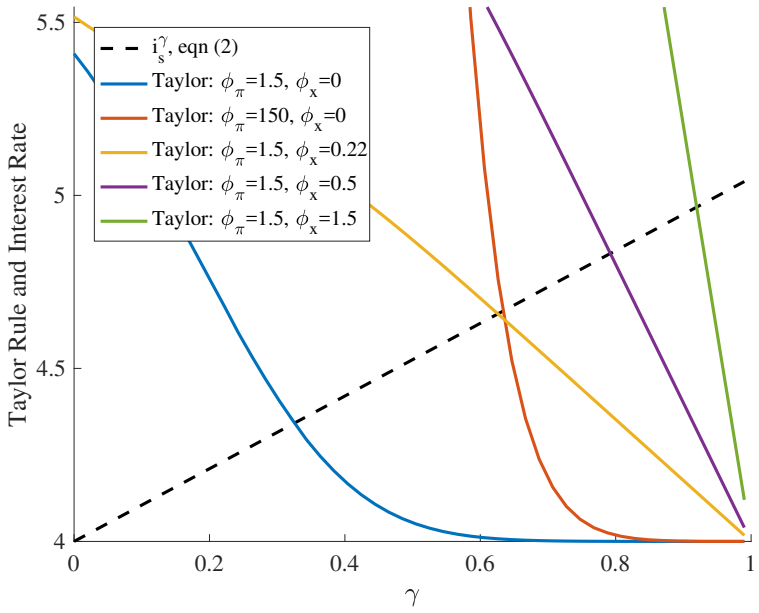
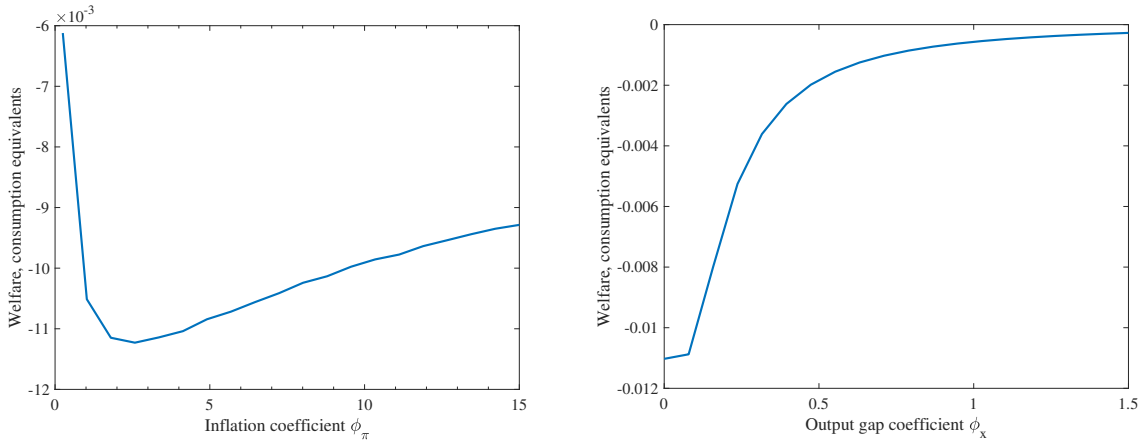


Figure 7: Monetary Policy implementation with a Taylor rule. Targeting the output gap is a robustly better policy.

responsiveness to inflation one-hundred-fold.

Finally, we can illustrate the problem with a single mandate by calculating welfare as a function of ϕ_π. Figure 8(a) plots welfare in equilibrium when the CB varies ϕ_π from 0.25 (low) to 15, which is an order of magnitude higher than empirical estimates. Note that welfare is actually decreasing until ϕ_π is about 2. This means that, with a single mandate alone, increasing welfare requires a very aggressive response to inflation. In contrast, Figure 8(b) plots welfare when the CB fixes $\phi_\pi = 1.5$ and varies ϕ_x from 0 to 1.5, a much smaller range. Welfare is strictly increasing in this case.



(a) Responding to Inflation

(b) Responding to Output Gaps

Figure 8: Monetary Policy and Welfare, the effectiveness of a dual mandate. Varying policy aggressiveness with a single mandate (a) and with a dual mandate (b).

E Proofs in the Strategic Model

Proof. **of Proposition 8** ONLY IF: If the flexible price $\{p_s\}$ is consistent with some PBE, then, in particular, if the true state is L the firm will not prefer to deviate and offer the price p_H rather than the price p_L. Note that if the true state is L and the firm offers p_H, Insiders will know that the true state is L but Outsiders will believe the true state is H. Hence the firm will not want to offer $p_H = \xi_H/2$ rather than $p_L = \xi_L/2$ if and only if:

$$\frac{\xi_L(\gamma)}{2}\left[1-\frac{\xi_L(\gamma)}{2}\frac{1}{\xi_L(\gamma)}\right] \geq \frac{\xi_H(\gamma)}{2}\left\{\alpha\left[1-\frac{\xi_H(\gamma)}{2}\frac{1}{\xi_L(\gamma)}\right] + (1-\alpha)\left[1-\frac{\xi_H(\gamma)}{2}\frac{1}{\xi_H(\gamma)}\right]\right\}$$

After multiplying both sides by $4\frac{1}{\xi_H(\gamma)}$, simplifying and re-arranging, we see that this is equivalent to:

$$\frac{\xi_L(\gamma)}{\xi_H(\gamma)} - 1 \geq \alpha\left[1-\frac{\xi_H(\gamma)}{\xi_L(\gamma)}\right]$$

After dividing both sides by the term in square brackets and noting that this term is negative, we

see that this is equivalent to:

$$\left[1-\frac{\xi_L(\gamma)}{\xi_H(\gamma)}\right] \Big/ \left[\frac{\xi_H(\gamma)}{\xi_L(\gamma)}-1\right] \leq \alpha$$

Finally, simplifying the fraction on the left side, we see that this is equivalent to:

$$\frac{\xi_L(\gamma)}{\xi_H(\gamma)} \leq \alpha$$

which is the desired result.

IF: Given that $\alpha \geq \frac{\xi_L(\gamma)}{\xi_H(\gamma)}$, We must construct a PBE in which prices along the equilibrium path are p_L, p_H. Hence we must show that when the true state is s the firm will not wish to deviate to a price $p \neq p_s$. PBE implies that when the Outsiders see the price p_s, they believe the true state is s, as in (a), (c). However, we are free to assign arbitrary beliefs to Outsiders if they see a price p different from both p_L and p_H, as in (b), (d); in that event we assign to Outsiders the belief that the true state is L. We must rule out four kinds of potentially profitable deviations

(a) The true state is L and the firm offers p_H.

(b) The true state is L and the firm offers $p \neq p_L, p_H$.

(c) The true state is H and the firm offers $p = p_H$.

(b) The true state is H and the firm offers $p \neq p_L, p_H$.

We have posited that when Outsiders see a price $p \neq p_L, p_H$ they believe the state is L and PBE guarantees that when Outsiders see the price p_L they believe the state is L, so we can subsume **(c), (d)** into

(e) The true state is L and the firm offers $p \neq p_H$.

We now verify **(a)** , **(b)** and **(e)** in turn.

(a) This follows immediately by following the steps in the ONLY IF case above, but in reverse order, noting that each inequality is *equivalent* to the one above.

(b) We have posited that when Outsiders see a price $p \neq p_L, p_H$ they believe the state is L. Insiders know the true state so they also believe the state is L. Hence aggregate demand if the firm offers p will be $1 - p\frac{1}{\xi_L(\gamma)}$ and firm profit will be $p[1 - p\frac{1}{\xi_L(\gamma)}]$. By definition, this quantity is maximized when $p = 1/2\frac{1}{\xi_L(\gamma)}$ and the maximum profit will be $1/4\frac{1}{\xi_L(\gamma)}$. However this is the profit when the firm offers $p_L = 1/2\frac{1}{\xi_L(\gamma)}$ so this cannot be a profitable deviation for any such p.

(e) We must show that when the true state is H the firm's profit if it offers p_H is at least as great as when it offers $p \neq p_H$; i.e. we must show

$$\begin{aligned} \frac{1}{4\frac{1}{\xi_H(\gamma)}} &\geq p\Big(\alpha[1 - p\frac{1}{\xi_H(\gamma)}] + (1-\alpha)[1 - p\frac{1}{\xi_L(\gamma)}] \\ &= \alpha p[1 - p\frac{1}{\xi_H(\gamma)}] + (1-\alpha)p[1 - p\frac{1}{\xi_L(\gamma)}] \end{aligned} \tag{46}$$

By definition, $p[1 - p\frac{1}{\xi_H(\gamma)}]$ would be maximized by setting $p = p_H = 1/2\frac{1}{\xi_H(\gamma)}$ and $p[1 - p\frac{1}{\xi_L(\gamma)}]$ would be maximized by setting $p = p_L = 1/2\frac{1}{\xi_L(\gamma)}$ so we must certainly have

$$\alpha p[1 - p\frac{1}{\xi_H(\gamma)}] \leq \alpha\left(\frac{1}{4\frac{1}{\xi_H(\gamma)}}\right) \tag{47}$$

$$(1-\alpha)p[1 - p\frac{1}{\xi_L(\gamma)}] \leq (1-\alpha)\left(\frac{1}{4\frac{1}{\xi_L(\gamma)}}\right) \tag{48}$$

By assumption, $P_L < P_H$ so $1/4\frac{1}{\xi_L(\gamma)} < 1/4\frac{1}{\xi_H(\gamma)}$. Hence if we add the inequalities (47) and (48) and compare with (46) it remains only to verify that

$$\frac{1}{4\frac{1}{\xi_H(\gamma)}} \geq \alpha\left(\frac{1}{4\frac{1}{\xi_H(\gamma)}}\right) + (1-\alpha)\left(\frac{1}{4\frac{1}{\xi_H(\gamma)}}\right)$$

This is a tautology, so we have verified **(e)** .

Having verified **(a)** , **(b)** and **(e)** , the proof is complete. □

Proof. **of Proposition 9** ONLY IF: Assume the sticky price $p_0 = \xi_0/2$ is consistent with some PBE. Suppose that, in that PBE, the true state is L and the firm offers a price $p \neq p_0$. Because the Insiders know the true state, they will demand the quantity $1 - p\frac{1}{\xi_L(\gamma)}$. PBE requires that the Outsiders form some belief about the true state and demand a quantity that is optimal with respect to that belief about the true state; hence the Outsiders will demand $1 - p\mathbb{E}_o\left[\frac{1}{\xi}\right]$ where $\mathbb{E}_o\left[\frac{1}{\xi}\right]$ is their expectation of the future price, given their belief about the true state and the behavior of the CB. The fraction α of consumers are Insiders, so the profit of the firm will be:

$$d_L(p) = p\Big(\alpha\big[1 - p\frac{1}{\xi_L(\gamma)}\big] + (1-\alpha)\big[1 - p\mathbb{E}_o\left[\frac{1}{\xi}\right]\big]\Big)$$

For every $p \neq p_0$, this expression will be minimized if the Outsiders assign probability 1 to the state L, in which case their expectation of the future price will be $E_o(1/P) = \frac{1}{\xi_L(\gamma)}$. Hence if the firm offers $\neq p_0$ we must have

$$d_L(p) \geq p[1 - p\frac{1}{\xi_L(\gamma)}]$$

In particular, this inequality must hold when $p = p_L = 1/2\frac{1}{\xi_L(\gamma)}$; thus we must have

$$d_L\left(\frac{\xi_L(\gamma)}{2}\right) \geq \frac{1}{4\frac{1}{\xi_L(\gamma)}}$$

In PBE the firm has no profitable deviation so it must be that $d_L(p_0) \geq d_L(p)$ for every p; in particular this inequality must hold when $p = p_L = 1/2\frac{1}{\xi_L(\gamma)}$. Recall that $p_0 = \xi_0/2$, so we conclude that

$$\frac{\xi_0}{2}\left(\alpha\left[1 - \frac{\xi_0}{2}\frac{1}{\xi_L(\gamma)}\right] + (1-\alpha)\left[1 - \frac{\xi_0}{2}\mathbb{E}_o\left[\frac{1}{\xi}\right]\right]\right) \geq \frac{1}{4\frac{1}{\xi_L(\gamma)}}$$

Because ξ_0 is the harmonic mean of P_L, P_H, $\mathbb{E}_o\left[\frac{1}{\xi}\right] = 1/\xi_0$; substituting yields

$$\frac{\xi_0}{2}\left(\alpha\left[1-\frac{\xi_0}{2}\frac{1}{\xi_L(\gamma)}\right]+\frac{1-\alpha}{2}\right) \geq \frac{1}{4\frac{1}{\xi_L(\gamma)}}$$

Collecting terms that involve α on the left side, all other terms on the right side, and simplifying leads to

$$\alpha\left[1-\xi_0\frac{1}{\xi_L(\gamma)}\right] \geq \frac{1\xi_0\zeta(\gamma,L)}{\xi_0\frac{1}{\xi_L(\gamma)}}$$

Because $1-\xi_0\frac{1}{\xi_L(\gamma)} \leq 0$, solving for α yields

$$\alpha \leq \frac{1}{\xi_0\frac{1}{\xi_L(\gamma)}}$$

which is the desired result.

IF: To construct a PBE in which p_0 is offered in both states, we have to prescribe the behavior of Outsiders when a price $p \neq p_0$ is offered. As the argument above suggests, we posit that when when a price $p \neq p_0$ is offered, Outsiders believe the true state is L and hence demand $1-p\frac{1}{\xi_L(\gamma)}$. Insiders know the true state s and demand $1-p\zeta(\gamma,s)$ so the profit of the firm is

$$d_s(p) = p\left(\alpha[1-p\zeta(\gamma,s)]+(1-\alpha)[1-p\frac{1}{\xi_L(\gamma)}]\right) \tag{49}$$

If the firm offers the putative equilibrium price $p_0 = \xi_0/2$, the Outsiders' expectation of the future price will be $1/\xi_0$, so the profit of the firm will be

$$d_s(p_0) = \frac{\xi_0}{2}\left(\alpha\left[1-\frac{\xi_0}{2}\zeta(\gamma,s)\right]+\frac{1-\alpha}{2}\right)$$

The equilibrium condition is that

$$d_s(p_0) \geq d_s(p) \tag{50}$$

when $s = L$ and when $s = H$, under the assumption that $\alpha \leq 1/\xi_0 \frac{1}{\xi_L(\gamma)}$.

To check that the inequality (50) is satisfied when the true state $s = L$,we need only follow the argument above in the opposite direction, as we did in the proof of Proposition 8.

To see that (50) is satisfied when the true state $s = H$ is more complicated. First note that simplifying the right side of (49) yields

$$\begin{aligned} d_H(p) &= p\left(1 - p[\alpha \frac{1}{\xi_H(\gamma)} + (1-\alpha)\frac{1}{\xi_L(\gamma)}]\right) \\ &= p\left(1 - p\left[\alpha \frac{1}{\xi_H(\gamma)} + (1-\alpha)\frac{1}{\xi_L(\gamma)}\right]\right) \end{aligned}$$

Write $A = \alpha \frac{1}{\xi_H(\gamma)} + (1-\alpha)\frac{1}{\xi_L(\gamma)}$ and note that $d_H(p) = p[1 - pA]$. Thus, $d_H(p)$ is maximized when $p = 1/2A$ and maximum profit is $d_H(1/2A) = 1/4A$. So it suffices to prove that

$$\frac{\xi_0}{2}\left(\alpha\left[1 - \frac{\xi_0}{2}\frac{1}{\xi_H(\gamma)}\right] + \frac{1-\alpha}{2}\right) \geq \frac{1}{4[\alpha \frac{1}{\xi_H(\gamma)} + (1-\alpha)\frac{1}{\xi_L(\gamma)}]} \tag{51}$$

To this end, multiply both sides by the denominator, simplify, and collect terms. This leaves the following inequality to be verified:

$$\begin{aligned} &\{[1 - \xi_0 \frac{1}{\xi_H(\gamma)}][\frac{1}{\xi_H(\gamma)} - \frac{1}{\xi_L(\gamma)}]\}\xi_0\alpha^2 \\ +\{[\frac{1}{\xi_H(\gamma)} - \frac{1}{\xi_L(\gamma)}] + [p - \xi_0 \frac{1}{\xi_H(\gamma)}\frac{1}{\xi_L(\gamma)}\}\xi_0\alpha \\ &+\left\{\xi_0 \frac{1}{\xi_L(\gamma)} - 1\right\} \geq \ 0 \end{aligned} \tag{52}$$

The left side of (52) is a quadratic polynomial in α; call it $f(\alpha)$. The coefficient of α^2 is negative so the set $\{\alpha : f(\alpha) \geq 0\}$ is a closed interval (or empty). Note that $f(0) = \xi_0 \frac{1}{\xi_L(\gamma)} - 1 > 0$. A straightforward, but long and messy, algebraic manipulation (available from the authors on request) shows that $f(1/\xi_0 \frac{1}{\xi_L(\gamma)}) \geq 0$, so it follows that $f(\alpha) \geq 0$ for $0 \leq \alpha \leq 1/\xi_0 \frac{1}{\xi_L(\gamma)}$. This demonstrates that the inequality (51) is satisfied in the same range of α, and hence that the inequality (50) is satisfied when the true state $s = H$. We conclude that for $0 \leq \alpha \leq 1/\xi_0 \frac{1}{\xi_L(\gamma)}$, there is a PBE in which the sticky price p_0 is offered in both states, as asserted. □

Proof. **of Lemma 2** (a) This requires only simple algebraic manipulation.

$$\begin{aligned}\frac{\frac{1}{\xi_H(\gamma)}}{\frac{1}{\xi_L(\gamma)}} &= \frac{\gamma\left(\frac{1}{\xi_0}\right)+(1-\gamma)\left(\frac{1}{\xi_0}\right)\left(\frac{1+\rho_0}{1+\rho_H}\right)}{\gamma\left(\frac{1}{\xi_0}\right)+(1-\gamma)\left(\frac{1}{\xi_0}\right)\left(\frac{1+\rho_0}{1+\rho_L}\right)} \\ &= \left[\frac{\gamma(1+\rho_H)+(1-\gamma)(1+\rho_0)}{\gamma(1+\rho_L)+(1-\gamma)(1+\rho_0)}\right]\left(\frac{1+\rho_L}{1+\rho_H}\right)\end{aligned}$$

Because $1+\rho_L < 1+\rho_0 < 1+\rho_H$ the numerator of the term in square brackets is strictly increasing in γ and the denominator is strictly decreasing so the term in square brackets is strictly increasing; hence the ratio $\frac{\xi_L(\gamma)}{\xi_H(\gamma)} = \frac{1}{\xi_H(\gamma)}/\frac{1}{\xi_L(\gamma)}$ is also strictly increasing. Moreover as γ tends to 1 the numerator of the left term tends to $1+\rho_H$ and the denominator tends to $1+\rho_L$ so altogether the ratio $\frac{1}{\xi_H(\gamma)}/\frac{1}{\xi_L(\gamma)}$ tends to 1, as asserted.

(b) This too requires only simple algebraic manipulation. For convenience, write $h(\gamma) = \frac{1}{\xi_H(\gamma)}\frac{1}{\xi_L(\gamma)}$. Multiplying out, collecting terms and recalling that ξ_0 is the harmonic mean of ξ_L, ξ_H yields:

$$\begin{aligned}h(\gamma) &= \left(\frac{\gamma}{\xi_0}\right)^2 + \frac{\gamma(1-\gamma)}{\xi_0\xi_L} + \frac{\gamma(1-\gamma)}{\xi_0\xi_H} + \frac{(1-\gamma)^2}{\xi_L\xi_H} \\ &= \left(\frac{\gamma}{\xi_0}\right)^2 + \frac{2(\gamma(1-\gamma)}{\xi_0^2} + \frac{(1-\gamma)^2}{\xi_L\xi_H}\end{aligned}$$

It is therefore immediate that $h(1) = \frac{1}{\xi_0}^2$. Differentiating and collecting terms yields

$$\begin{aligned}h'(\gamma) &= \frac{2\gamma}{\xi_0^2} + \frac{(2-4\gamma)}{\xi_0^2} - \frac{2(1-\gamma)}{\xi_L\xi_H} \\ &= (2-2\gamma)\left[\frac{1}{\xi_0^2} - \frac{1}{\xi_L\xi_H}\right]\end{aligned}$$

To see that $h'(\gamma) > 0$, and hence that $h(\gamma)$ is strictly increasing, it suffices to show that $1/\xi_0^2 -$

$1/\xi_L\xi_H > 0$. To see this plug in to the definition of $1/\xi_0$ and expand

$$\begin{aligned}
\frac{1}{\xi_0^2} - \frac{1}{\xi_L\xi_H} &= \tfrac{1}{2}^2\left[\frac{1}{\xi_L} + \frac{1}{\xi_H}\right]^2 - \frac{1}{\xi_L\xi_H} \\
&= \tfrac{1}{4}\left[\left(\frac{1}{\xi_L}\right)^2 + \frac{2}{\xi_L\xi_H} + \left(\frac{1}{\xi_H}\right)^2\right] - \frac{1}{\xi_L\xi_H} \\
&= \tfrac{1}{4}\left[\left(\frac{1}{\xi_L}\right)^2 - \frac{2}{\xi_L\xi_H} + \left(\frac{1}{\xi_H}\right)^2\right] \\
&= \left(\tfrac{1}{2}\left[\frac{1}{\xi_L} - \frac{1}{\xi_H}\right]\right)^2
\end{aligned}$$

This is certainly strictly positive so the proof is complete. □

Proof. **of Proposition 10**

All of the desired assertions follow quickly. First, from Lemma 2 the term $\left(\frac{\xi_H(\gamma)\xi_L(\gamma)}{\xi_0}\right)$ is strictly decreasing in γ.

(a) For $\gamma \in [0, \gamma_0)$: the numerator $1 - F(\hat{\alpha}(\gamma))$ is strictly positive and the denominator $\int_0^{\hat{\alpha}(\gamma)} \alpha \, dF(\alpha)$ is 0 so $\kappa(\gamma) = +\infty$.

(b) For $\gamma \in [\gamma_0, \gamma_1)$: the numerator is strictly decreasing in γ and the denominator is strictly increasing, so $\kappa(\gamma)$ is strictly increasing.

(c) For $\gamma \in [\gamma_1, 1)$: the numerator is 0 and the denominator is strictly positive so $\kappa(\gamma) = 0$.

(d) As γ converges to 1, the numerator converges to 0 and the denominator increases to something strictly positive, so $\lim_{\gamma \to 1} \kappa(\gamma) = 1$.

This completes the proof. □

Proof of Proposition 11. Since $\zeta_s(\gamma) = \gamma\left(\frac{1}{\xi_0}\right) + (1-\gamma)\left(\frac{1}{\xi_0}\right)\left(\frac{\theta_s}{\theta_0}\right)$, we have

$$\zeta_L - \zeta_H = (1-\gamma)\left(\frac{\theta_L}{\theta_0\xi_0} - \frac{\theta_H}{\theta_0\xi_0}\right).$$

and so

$$\left(\zeta_L' - \zeta_H'\right) = -\left(\frac{\theta_L - \theta_H}{\theta_0 \xi_0}\right) < 0,$$

since $\theta_L > \theta_H$ ($\rho_L < \rho_H$). This confirms that the intensive margin is negative. Since $\frac{d\hat{\alpha}}{d\gamma} > 0$, the extensive margin is positive. We want to show when the extensive margin exceeds the intensive margin. From the previous derivations, we can write $\frac{d\mathbb{E}[|\bar{x}_s - \bar{x}_0|]}{d\gamma}$ as

$$\frac{d\mathbb{E}\left[|\bar{x}_s - \bar{x}_0|\right]}{d\gamma} = \left(\frac{\theta_L - \theta_H}{\theta_0 \xi_0}\right)\left(\underbrace{(1-\gamma)\,\hat{\alpha} f(\hat{\alpha})\frac{d\hat{\alpha}}{d\gamma}}_{\text{extensive margin}} - \underbrace{\int_0^{\hat{\alpha}} \alpha dF(\alpha)}_{\text{intensive margin}}\right).$$

This completes the proof. □

Proof of Proposition 12. On islands with flexible prices, all agents are informed and face $p_s = \frac{\xi_s}{2}$, implying demand $c^* = \frac{1}{2}$. Because all agents consume $1/2$, firms with flexible prices have profits $d_s = \xi_s/4$. It remains to solve for C_s^* and expected utility. Future consumption is $C^* = E + (d - pc^*)/Q$. But note that $d = pc^*$, and therefore we have $C_s^* = E$. Hence, since $c^* = 1/2, C^* = E$, utility on islands with flexible prices equals

$$\mathbb{E}[U] = E + \frac{3}{8}.$$

On islands with sticky prices, all agents face $p = \xi_0/2$. Denoting consumption by Insiders (informed) and Outsiders (uninformed) agents by ι and o. Uninformed agents consume $c_o^* = \frac{1}{2}$, which, notably, is the same present consumption as agents on islands with flexible prices. Informed agents demand $c_\iota^* = 1 - \frac{\xi_0 \zeta_s}{2}$. Dropping the state subscript, future consumption for agent j is $C_j^* = E + (d - pc_j^*)/P$, but it is not the case that for both types of agents $d = pc_j^*$. However, since firm profits are just p_0 times aggregate consumption $(\alpha c_\iota^* + (1-\alpha)c_o^*)$, it is the case that profits equal

the sum of pc^* across agents. Thus,

$$\begin{aligned}\alpha C_i^* + (1-\alpha)C_o^* &= \alpha\left(E + (d - pc_i^*)/Q\right) + (1-\alpha)\left(E + (d - pc_o^*)/Q\right), \\ &= E + \frac{d - p\left(\alpha c_i^* + (1-\alpha)c_o^*\right)}{Q} = E.\end{aligned}$$

Hence, by linearity in C^*, the distribution of firm profits does not matter in the aggregate. To simplify our calculations and for expositional clarity, we set $C = E$ for both sets of agents.

As before, expected utility for uninformed agents (Outsiders) is given by $\mathbb{E}[U] = E + \frac{3}{8}$, which is the same as expected utility on islands with flexible prices. In state s, informed agents have utility

$$\begin{aligned}U(c_s, C_s) &= c - \frac{c^2}{2} + E, \\ &= 1 - \frac{\xi_0\zeta_s}{2} - \frac{1}{2}\left(1 - \frac{\xi_0\zeta_s}{2}\right)^2 + E, \\ &= 1 - \frac{\xi_0\zeta_s}{2} - \frac{1}{2}\left(1 - \xi_0\zeta_s + \frac{\xi_0^2\zeta_s^2}{4}\right) + E, \\ &= E + \frac{1}{2} - \frac{\xi_0^2}{8}\zeta_s^2.\end{aligned}$$

Thus to calculate the expected utility we need $\mathbb{E}[\zeta_s^2]$. Note that $\mathbb{E}[\zeta_s] = \mathbb{E}[1/\xi_s] = 1/\xi_0$. By convexity we have $\mathbb{E}[\zeta_s^2] > 1/\xi_0^2$, which implies that if $\gamma < 1$ then for insiders $\mathbb{E}[U] < E + \frac{3}{8}$. Recall that

$$\zeta_s = \frac{1}{\xi_0}\left(\gamma + (1-\gamma)\frac{\theta_s}{\theta_0}\right),$$

and therefore

$$\zeta_s^2 = \frac{1}{\xi_0^2}\left(\gamma^2 + 2\gamma(1-\gamma)\frac{\theta_s}{\theta_0} + (1-\gamma)^2\frac{\theta_s^2}{\theta_0^2}\right).$$

Taking expectations we have

$$\mathbb{E}[\zeta_s^2] = \frac{1}{\xi_0^2}\left(\gamma^2 + 2\gamma(1-\gamma) + (1-\gamma)^2\frac{\mathbb{E}\left[\theta_s^2\right]}{\theta_0^2}\right),$$
$$\implies \xi_0^2\mathbb{E}[\zeta_s^2] = \gamma^2 + 2\gamma(1-\gamma) + (1-\gamma)^2\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right].$$

It is easy to show that $\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right] > 1$ by convexity. This also means that $\xi_0^2\mathbb{E}[\zeta_s^2] > 1$ if $\gamma \in (0,1)$ and so utility is less than $E + \frac{3}{8}$ for $\gamma \in (0,1)$.

We can show that expected utility for Insiders is strictly increasing in γ. Note that

$$\frac{d\mathbb{E}[U]}{d\gamma} = -\frac{1}{8}\frac{\xi_0^2 d\mathbb{E}[\zeta_s^2]}{d\gamma}.$$

We have

$$\frac{\xi_0^2 d\mathbb{E}[\zeta_s^2]}{d\gamma} = 2\left(1-\gamma-(1-\gamma)\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]\right) < 0,$$

which is strictly negative because the expectation exceeds 1. Thus, the utility of informed agents facing sticky prices is strictly increasing in γ. □

Proof of Proposition 13. We can write aggregate welfare as

$$\begin{aligned}\overline{W}(\gamma) = &\int_0^{\hat{\alpha}}\left[\alpha\left(E + \frac{1}{2} - \frac{1}{8}\xi_0^2\mathbb{E}[\zeta_s^2]\right) + (1-\alpha)\left(E+\frac{3}{8}\right)\right]dF(\alpha)\\ &+\left(E+\frac{3}{8}\right)(1-F(\hat{\alpha})),\end{aligned} \tag{53}$$

which we can simplify to

$$\begin{aligned}\overline{W}(\gamma) &= E + \frac{3}{8} + \int_0^{\hat{\alpha}}\left(\alpha\left(\frac{1}{8} - \frac{1}{8}\xi_0^2\mathbb{E}[\zeta_s^2]\right)\right)dF(\alpha),\\ &= E + \frac{3}{8} + \frac{1}{8}\int_0^{\hat{\alpha}}\alpha\left(1-\xi_0^2\mathbb{E}[\zeta_s^2]\right)dF(\alpha).\end{aligned} \tag{54}$$

Global maximization at $\gamma = 1$ follows from equation (43) and the fact that $\frac{d\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]}{d\gamma} < 0$.

Since welfare can only be strictly less than $e+\frac{3}{8}$ for $\gamma<1$ and equals $e+\frac{3}{8}$ at $\gamma=1$, welfare is maximized at $\gamma=1$.

We can write $\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]-1$ as

$$\begin{aligned}\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]-1 =&\gamma^2+2\gamma(1-\gamma)+(1-\gamma)^2\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]-(\gamma+(1-\gamma))^2\\ =&(1-\gamma)^2\left(\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]-1\right)\end{aligned}$$

which implies that

$$\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]=(1-\gamma)^2\left(\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]-1\right)+1$$

Thefore we can write $\frac{d\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]}{d\gamma}$ as

$$\frac{d\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]}{d\gamma}=-2(1-\gamma)\left(\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]-1\right)=2(1-\gamma)\left(1-\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]\right)$$

This allows us to rewrite the marginal change in welfare as

$$\begin{aligned}\frac{d\overline{W}(\gamma)}{d\gamma}=&\underbrace{(1-\gamma)^2\left(1-\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]\right)\hat{\alpha}f(\hat{\alpha})\frac{d\hat{\alpha}}{d\gamma}}_{\text{extensive margin}<0}\\ &\underbrace{+2(1-\gamma)\left(\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]-1\right)\int_0^{\hat{\alpha}}\alpha dF(\alpha)}_{\text{intensive margin}>0}\end{aligned}$$

which we can simplify to

$$\frac{d\overline{W}(\gamma)}{d\gamma}=(1-\gamma)\left(\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]-1\right)\left(-(1-\gamma)\hat{\alpha}f(\hat{\alpha})\frac{d\hat{\alpha}}{d\gamma}+2\int_0^{\hat{\alpha}}\alpha dF(\alpha)\right)\tag{55}$$

Note that $(1-\gamma)\left(\mathbb{E}\left[\left(\frac{\theta_s}{\theta_0}\right)^2\right]-1\right)>0$. Hence, welfare is increasing in γ if

$$(1-\gamma)\,\hat{\alpha}f(\hat{\alpha})\frac{d\hat{\alpha}}{d\gamma}<2\int_0^{\hat{\alpha}}\alpha dF(\alpha)$$

which is the result. □

F General Equilibrium Framework

This section lays down a general equilibrium framework in which our model of price stickiness can be embedded. We have two goals. The first is to clarify that the earlier results can be obtained in a model with labor supply (and no endowments). The second is to clarify that the earlier results can be obtained in a model with where money plays an essential role.

The setup's pieces are standard. However, putting the pieces together is quite involved. Therefore we start its description with a preview. We subsequently fully describe every piece of the model.

F.1 Preview

The setup is based on the foundational papers by Lagos and Wright (2005) and Lucas and Stokey (1987). As Lagos and Wright (2005), we exploit quasilinearity and periods that are divided in a day and a night to be able to handle agent (informational) heterogeneity. As Lucas and Stokey (1987), we use a cash-in-advance model with credit and cash goods. The presence of credit goods is key for specifying trading in goods markets with partially informed consumers.

The population of the economy is composed by a unit mass of households. These households own a unit mass of firms, which operate in different and segmented geographic locations called islands. There is a unit mass of islands, and on each island there is a single firm.

Households are divided into workers and consumers. Workers supply labor; consumers shop for consumption goods.

As in Lagos and Wright (2005), each period is divided into two subperiods: a day and a night.

All the action of interest takes place during the day; the night is simply introduced as a technical device to close the model. Trading of credit goods takes place during the day; trading of cash goods takes place during the night.

The exogenous aggregate state of the economy is given by a preference shock θ, which is the discount factor between the day and the night. As in most of the literature, this preference shock is a modeling device to generate fluctuations in nominal aggregate demand. As in the simple model presented in previous sections' of the paper, a key assumption of our setup is that there is household heterogeneity in the information about this aggregate shock. Some households may be imperfectly informed about the value of discount factor at night.[10] We model this by making the sharp assumption that a fraction of households is perfectly informed about the realization of the shock and the complement is uninformed about the realization of the shock.

Firms, by assumption, are informed about the preference shock. We motivate this simplifying assumption by a story in which firms are able to aggregate consumer demand via goods market trading. So long as a non-zero mass of each firm's consumers are informed, their demand then reveals the aggregate preference shock to firms. To simplify the exposition, here we simply assume that firms are informed right from the start. On the other hand, imperfectly informed consumers learn by looking at firms' prices.[11] In fact, firms and consumers play a sequential game. Consumers and firms meet in decentralized locations. Each firm posts a price, consumers observe the price, and then post their demand.

The central bank controls money supply, which determines relative price between the night and day. The central bank uses a rule to determines its policy. This rule depends on deviations of inflation from a target and on the output gap.

[10] One can also think about this shock representing a shift in marginal utility at night. Under this interpretation, the assumption is that, during the day, imperfectly informed households do not receive full information about marginal utility at night.

[11] Notice that, our informational assumptions force us to move away from monopolistic competition (or other forms of centralized goods markets).

F.2 Full Model

Population and Geography. The economy is populated by households, firms, and a Central Bank (CB). The geography is given by a unit mass of islands, and a mainland. Each island is populated by a continuum of households of mass one and is served by a single monopolistic firm. The mainland is visited by all consumers in the economy at given dates, and is served by a competitive representative firm. Households are divided in workers and consumers.

Time Structure. Time is discreet. Similar to Lagos and Wright (2005), periods are divided into two subperiods, called "day" and "night". Following their notation, we will denote day variables in lower case, and night variables in upper case. Subperiods are indexed by t: $t = 0$ signifies the day, and $t = 1$ signifies the night. (However, to simplify the notation, we skip t notation when possible.) Periods are indexed by τ and run from $\tau = 0$ to infinity.

Goods Markets. We start by describing day-time trading in the decentralized market. Each mass of consumers are served by a price-setting monopolist (on a given island), which sells good c at a nominal price p. These decentralized goods are bought on credit.

We now describe the functioning of the night-time, centralized, market. At night, all consumers are sent to the mainland. There, they consume an aggregate good C, produced by a competitive firm, and sold at an aggregate nominal price P. We also refer to this aggregate price P as the night price level. This good is sold in exchange for cash.

Labor Markets. During both the day and the night labor markets are open. During the day, workers supply labor in a centralized labor market. Local firms hire workers from this centralized market. At night, labor is supplied in the mainland. Both (day and night) markets are competitive. Daytime labor is denoted l; nighttime labor is denoted L. We denote wages as w and W, respectively.

Credit, Financial, and Money Markets. During the day, all transactions take place on credit. Consumers buy consumption goods on credit, workers bring back wages, and firms pay profits (the

firm is owned by local households).

At night, goods are bought in cash. (Labor is supplied on credit.)

The money market opens only at (the end of the) night. Similar to Lucas and Stokey (1987), all credit transactions are settled at this moment. A (long-term) bond is available across periods. These are trades in exchange of money holdings for the next period $\tau+1$. Long term bonds and cash holdings are denoted B and M respectively.

Exogenous Aggregate State. The exogenous aggregate state is given by the realization of a preference shock θ_τ. We specify the process for the preference shock below.

Central Bank. The Central Bank sets the money supply following the commitment a rule, which we specify below. The money supply determines the price level during the night P.

Information Structure. Consumers are heterogeneous in terms of the information they possess. There are two types of consumers: Insiders (informed consumers) $\imath \in I$ and Outsiders (uninformed consumers) $o \in O$. Insiders are perfectly informed about the state θ_τ; Outsiders are uninformed about the state but know the probability distribution, and may draw inferences from the price set by the firm with which they trade. The fraction $\alpha \in [0,1)$ of Insiders on a particular island varies across islands. We assume the distribution of α is given by a cdf F whose support is not a singleton and has the property that

$$\lim_{\alpha \to 1} F(\alpha) = 1$$

That is, the fraction of islands on which all consumers are Insiders is 0. All other agents in the economy have perfect information.

All of the above is common knowledge.

Household Optimization. We start by presenting an inner problem of the household. In this problem, household solve for all variables that trade in credit. This is the "day-to-night" problem where the action happens. (The outer problem is presented below.)

We index a typical household by j. The inner problem at date τ consists in solving

$$\max_{c_\tau, l_\tau, L_\tau} \mathbb{E}_{j\tau}\left[\left(u(c_\tau) - l_\tau\right) + \theta_\tau\left(U(\bar{C}) - L_\tau\right)\right]$$

where choices variables have been defined above. The variable $\bar{C}$ denotes a fixed allocation of the nighttime consumption good. Since this good is traded in cash, its consumption is fixed in the inner problem (the outer problem will determine this quantity). The random variable θ_τ is the exogenous aggregate state, which determines the discount factor between the day and the night. Following the previous sections, we specify the process for θ_τ to be very simple, by assuming it follows an i.i.d. binary Markov chain with two values, θ_L and θ_H, with $Pr(\theta_L) = Pr(\theta_H) = 1/2$ and $\mathbb{E}[\theta_\tau] = \theta < 1$. The realization θ_H corresponds to the high state, and the realization θ_L corresponds to the low state. The household values daytime consumption relatively more in the high state (and hence demand is higher than in the low state). Hence, the realizations are such that $\theta_L > \theta_H$.[12] The utility functions $u(\cdot)$ and $U(\cdot)$ are assumed to be twice continuously differentiable on $\mathbb{R}^{++}$, strictly increasing, and strictly concave. Below, we make an assumption on $u(\cdot)$ such that the monopolist's problem has a solution.[13] We assume that there is a value of C such that $U'(C) = 1/\beta$. The expectation operator is indexed by j to signify the household member's information set at the time they make a choice.

This problem is subject to a constraint given by

$$p_\tau c_\tau + P_\tau \bar{C} = \pi_\tau + w_\tau l_\tau + W_\tau L_\tau \tag{56}$$

where prices have been defined above and π_τ are profits.

Denoting by λ_τ the Lagrange multiplier of the constraint (56), the first-order condition for daytime consumption c is

$$u'(c_\tau) = \mathbb{E}_{j\tau}\left[\lambda_\tau p_\tau\right]$$

It is important to emphasize that, depending on the index j, this condition may be taken under

[12] The model allows for richer specifications of the exogenous process for the state, such as persistent Markov chains and AR(1). (To simplify the notation of this GE framework, we omit the subscript s to denote the state as in the previous sections, and simply use the notation θ_τ.)

[13] The earlier sections assume quadratic utility, which is a convenient assumption for welfare calculations. Here in the GE framework we aim to show that this particular restriction is not needed to find a general equilibrium solution.

imperfect information. Indeed, it defines the choice, in a decentralized market, by a consumer that can be an Outsider. (Insiders have full information about the product $\lambda_\tau p_\tau$.) Notice however that Outsiders observe the price of the firm they meet, p_τ, and hence the expectation conditional on this price. Therefore, it can be taken out of the expectation operator.

We obtain the following set of first-order conditions:

$$u'(c_\tau) = p_\tau \mathbb{E}_{j\tau}[\lambda_\tau]$$
$$1 = \lambda_\tau w_\tau$$
$$\theta_\tau = \lambda_\tau W_\tau$$

(The remaining optimality conditions are taken under perfect information, since they involve choices in centralized markets.)

From the second equation we observe that $\lambda_\tau = 1/w_\tau$. Further manipulating the equations above we can summarize the set of first-order conditions to

$$u'(c_\tau) = p_\tau \mathbb{E}_{j\tau}\left[\frac{\theta_\tau}{W_\tau}\right] \tag{57}$$
$$\frac{1}{w_\tau} = \theta_\tau \frac{1}{W_\tau} \tag{58}$$

We continue by presenting the outer problem of the household. This is the problem solved from one day to the other. (Below we will formally establish the relation between both problems.) This problem will give rise to an explicit role for money and hence allows us to define the monetary policy instrument.

Define

$$\mathscr{U}(c_\tau, l_\tau, C_\tau, L_\tau) = \Big(u(c_\tau) - l_\tau\Big) + \theta_\tau\Big(U(C_\tau) - L_\tau\Big)$$

In the outer problem, the household needs to solve

$$\max_{c_\tau, l_\tau, C_\tau, L_\tau, M_\tau, B_\tau} \mathbb{E}_{j\tau}\left[\sum_{\tau=0}^{\infty} \beta^\tau \, \mathscr{U}(c_\tau, l_\tau, C_\tau, L_\tau)\right]$$

which involves choosing infinite sequences of consumption, labor supply, money and bond holdings subject to

$$p_\tau c_\tau + P_\tau C_\tau + B_\tau + M_\tau = w_\tau l_\tau + W_\tau L_\tau + M_{\tau-1} + T_\tau + \left(1 + i_\tau^{LT}\right) B_{j\tau-1} + \pi_\tau \tag{59}$$

where i_τ^{LT} is a long-term nominal interest rate, and T_τ is a lump-sum cash transfer set by the Central Bank. Purchases of the cash good are also subject to a cash-in-advance (CIA) constraint

$$P_\tau C_\tau \leq M_{\tau-1} + T_\tau \tag{60}$$

Denoting by χ_τ the multiplier on the budget constraint (59) and by ψ_τ the multiplier on the CIA constraint (60), we get the set of first-order conditions

$$\beta^\tau u'(c_\tau) = \mathbb{E}_{j\tau}\left[\chi_\tau p_\tau\right] \tag{61}$$

$$\beta^\tau = \chi_\tau w_\tau \tag{62}$$

$$\beta^\tau \theta_\tau U'(C_\tau) = (\chi_\tau + \psi_\tau) P_\tau \tag{63}$$

$$\beta^\tau \theta_\tau = \chi_\tau W_\tau \tag{64}$$

$$\chi_\tau = \mathbb{E}_\tau\left[\chi_{\tau+1} + \psi_{\tau+1}\right] \tag{65}$$

$$\chi_\tau = \left(1 + i_{\tau+1}^{LT}\right) \mathbb{E}_\tau\left[\chi_{\tau+1}\right] \tag{66}$$

where it is important to notice the presence of two different expectation operators, the daytime expectation operator $\mathbb{E}_{j\tau}[\,\cdot\,]$ (conditional on consumer j's information), and the nighttime expectation operator $\mathbb{E}_\tau[\,\cdot\,]$ (conditional on full information, which is available in the centralized market).

From (62) and (64), we observe that $\chi_\tau = \beta^\tau / w_\tau$ and $\chi_\tau = \beta^\tau \theta_\tau / W_\tau$. Thus,

$$\frac{1}{w_\tau} = \theta_\tau \frac{1}{W_\tau}$$

which is the same as (58). Also, plugging in the expression for χ_τ obtained from (64) into (61), we

get

$$u'(c_\tau) = p_\tau \mathbb{E}_{j\tau}\left[\frac{\theta_\tau}{W_\tau}\right] \tag{67}$$

which is the same as (57).

The remaining conditions (determining the demand for money and bonds) can be simplified as follows. Equation (63), one period forward, is $\beta^{\tau+1}\theta_{\tau+1}U'(C_{\tau+1}) = (\chi_{\tau+1} + \psi_{\tau+1})P_{\tau+1}$. Solving for $\chi_{\tau+1} + \psi_{\tau+1}$, and using the expression for χ_τ, equation (65) becomes

$$\frac{\theta_\tau}{W_\tau} = \beta\ \mathbb{E}_\tau\left[\frac{\theta_{\tau+1}}{P_{\tau+1}}U'(C_{\tau+1})\right] \tag{68}$$

Finally, equation (66) is equivalent to

$$\frac{\theta_\tau}{W_\tau} = \beta\left(1 + i^{LT}_{\tau+1}\right)\mathbb{E}_\tau\left[\frac{\theta_{\tau+1}}{W_{\tau+1}}\right] \tag{69}$$

Production. All firms in the economy have a linear technology and produce using only labor. Within every period, monopolist of the decentralized market produces c according to the production function $c_\tau = Al_\tau$. For simplicity, we assume that $k \equiv 1/A$ is commonly known.

The competitive firm produces C according to the production function $C_\tau = L_\tau$, where productivity has been normalized to 1.

Game in the Decentralized Market. The equilibrium notion for the game played between consumers and firms is the one described in full detail in the body. Below we shall prove that, this setup is tractable in the following sense: Any equilibrium of this game is part of a general equilibrium for the whole economy.

Central Bank. The central bank sets the money supply M^S_τ. An increase of the money supply (away from its steady state value) is expansionary since it is increases aggregate demand, and vice versa. The central banks behaves by adjusting money supply as a function of inflation and the

output gap, according to the following rule:

$$M_\tau^S = M_0 \left(\hat{p}_\tau\right)^{-\phi_\pi} \left(\hat{x}_\tau\right)^{-\phi_x} \tag{70}$$

where M_0 is the natural level of the money supply, $\hat{p}_\tau$ is inflation, defined as the percentage deviation of the price level P_τ away from steady state p_0: $\hat{p}_\tau = \int p(\alpha)dF(\alpha)/p_0$, and $\hat{x}_\tau$ is the output gap, defined as the percentage deviation of aggregate output from steady state $y_0 = 1/2$: $\hat{x}_\tau = \int y(\alpha)dF(\alpha)/(1/2)$. Below we show that this policy rule can be expressed as a rule for a short-term interest rate rate. This rule is equivalent to the Taylor rule (35) presented in the simple version of the model.

We are finally in a position where we can define a general equilibrium for the economy.

Definition of Equilibrium. A (general) equilibrium of this economy is given by consumption allocations, labor supply, bond holdings and money demand (for each household) $\{c_\tau, C_\tau, l_\tau, L_\tau, B_\tau, M_\tau\}$, labor demand (for each firm) $\{l_{j\tau}^D, L_\tau^D\}$, profits $\{\pi_\tau\}$, money supply $\{M_\tau^S\}$, nominal transfers $\{T_\tau\}$, nominal prices $\{p_\tau, P_\tau\}$, nominal wages $\{w_\tau, W_\tau\}$, long-term nominal interest rates $\{1 + i_\tau^{LT}\}$, for all τ, such that:

1. Households' conditions for optimality and corresponding constraints are satisfied;
2. The price-setting game is solved as specified above;
3. The representative firm maximizes profits taking the price as given;
4. The CB sets money supply as specified by the rule above;
5. Goods, labor, bonds, and money markets clear.

General Equilibrium Characterization. First, we conjecture that C_τ is constant in equilibrium. If so, then the price of this good is pinned down by the cash in advance constraint. We denote this constant $C_\tau = \bar{C}$. Second, we conjecture that $M_\tau = M_\tau^S$, for all $_\tau$. Then, $P_\tau = M_\tau/\bar{C}$.

By the optimality condition for the production of the representative firm, the nominal wage $W_\tau = P_\tau$ (since productivity is normalized to 1). Thus,

$$P_\tau = W_\tau = \frac{M_\tau}{\bar{C}} \tag{71}$$

Now, taking equation (68) and writing it as

$$\frac{\theta_\tau}{M_\tau} = \beta U'(\bar{C})\, \mathbb{E}_\tau \left[\frac{\theta_{\tau+1}}{M_{\tau+1}} \right]$$

reveals that as long as $U'(\bar{C}) = 1/\beta$ and θ_τ / M_τ is a martingale, equation (68) is satisfied. Corollaries 2 and 3 show that a monotonic rule can be mapped into a degree of monetary policy adjustment γ. Hence, we write the rule

$$\frac{1}{M_\tau} = \gamma \frac{1}{\theta_\tau} + (1-\gamma)\frac{1}{M_0}$$

According to this rule, when $\gamma = 1$, there is full adjustment ($M_\tau = \theta_\tau$), and when $\gamma = 0$, there is no adjustment ($M_\tau = M_0$). This rule can be written

$$\frac{\theta_\tau}{M_\tau} = \gamma + (1-\gamma)\frac{\theta_\tau}{M_0}$$

Taking the expectation of θ_τ / M_t shows, trivially, that this ratio is a martingale, and thus equation (68) is satisfied.

A similar argument for $i_\tau^{LT} = 1/\beta - 1$ shows that equation (69) is satisfied.

Since this is a closed economy with a zero net supply of bonds, we can simply set $B_\tau = 0$ for all households. It remains to check that the labor markets clear. The centralized market clears when each household supplies $L_\tau = C_\tau$. In the decentralized market, each household's labor supply is set to satisfy their respective budget constraints. Aggregating the budget constraint gives the economy's resource constraint, and from this one can establish that the labor market clears in each island. (Notice, this implies that any equilibrium solution to the game played between firms and consumers is a GE. This ensure a tractable and isolated treatment of the game.)

Finally, set $T_\tau = M_\tau - M_{\tau-1}$. At this point, we are able to verify our money demand and cen-

tralized good consumption conjectures. This completes the characterization of the GE framework.

Equivalence Results. In order to understand the sense in which the program of the household admits an inner and an outer problem, notice first that the first order condition for c_τ in both problems are the same (equations (57) and (67)). Also, since the equilibrium in the outer problem requires $L_\tau = C_\tau$, and since $W_\tau = P_\tau$, $M_\tau = M_{\tau-1} + T_\tau$ and $B_\tau = 0$, then, setting $E = \bar{C}$ the budget constraints in both problems reduce to

$$p_\tau c_\tau = \pi_\tau + w_\tau l_\tau$$

leading to the same choice of l_τ in both problems. The following result has then just been established.

Lemma 3. *The equilibrium allocations of c_τ, l_τ, L_τ in the inner problem are the same as in the outer problem. Moreover, the equilibrium allocation C_τ is an admissible endowment E of the inner problem.*

To obtain the simple, partial equilibrium, model in section B, interpret the cash good as a numeraire good. Since the credit good and the cash good are purchased in subsequent periods (call them period 0 and period 1), the price P_τ can be interpreted as the price of an asset traded at period 0, that pays 1 unit of the numeraire good in period 1. Denote this price Q_τ. Then, $Q_\tau = P_\tau$. Moreover, since in this simple model marginal costs are zero, take the limit $A \to \infty$, which implies zero labor demand.

In the simple model, the choice of C is determined by the budget constraint. Since households are heterogeneous in terms of their information and implied choice of c, nothing guarantees that C is equal to E. However, by the linearity, one can verify that the aggregate quantity of C is indeed equal to E.

In order to obtain the rule (35), notice that $M_\tau = E \cdot Q_\tau$. So the rule (70) is

$$E \cdot Q_\tau = E \cdot Q_0 \left(\hat{p}_\tau\right)^{-\phi_\pi} \left(\hat{x}_\tau\right)^{-\phi_x}$$

which is

$$1+i_\tau = (1+i_0)\left(\hat{p}_\tau\right)^{\phi_\pi}\left(\hat{x}_\tau\right)^{\phi_x}$$

In logs

$$\log(1+i_\tau) = \log(1+i_0) + \phi_\pi \hat{p}_\tau + \phi_x \hat{x}_\tau$$

Finally, the simple model can be written using two periods only, which allows to drop the τ index and keep only the lower case and upper case notation for $t=0$ and $t=1$.

Thereby, the following lemma establishing the alleged equivalence has been proven.

Lemma 4. *The model presented in section B has the same equilibrium allocation of c_τ as the full GE model. Also, the aggregate consumption of C in the simple model is equal to $E=\bar{C}$. Moreover, rules (35) and (70) lead to the same policies.*

Finally, we note it is straightforward to extend this result, based on the full GE model, to the case of positive marginal costs with finite A and $k=1/A$.

G Model with a Distribution of Marginal Costs

We formalize the analysis of the version of the baseline model with a distribution of marginal costs k rather than a distribution of the fraction of informed agents. The model behaves essentially the same.

Let the setup be modified as follows. The fraction of insiders α is assumed to be constant across islands, but now firms differ in their marginal costs k. Let $k \sim G$, where G is the CDF with the usual properties.

We restate the main results in this setting. The proofs are in the following subsection.

G.1 Short-run Equilibrium, Phillips Curve, and Welfare

Let the firm have marginal cost k. With flexible prices, the firm maximizes

$$(1-p/\xi)(p-k\xi) \tag{72}$$

In the absence of shocks, the optimal price is

$$p = \frac{\xi(1+k)}{2}$$

and total demand is

$$x = 1 - \frac{\xi(1+k)}{2\xi} = \frac{1-k}{2}.$$

Similarly, the sticky-price equilibrium is given by $p_0 = \frac{\xi_0(1+k)}{2}$, with informed agents demanding $x_\iota = 1 - \frac{\xi_0(1+k)}{2\xi_s}$ and uninformed agents demanding $\frac{1-k}{2}$.

We can now restate the original results from the paper.

Proposition 14 (PBE with Flexible Prices and Marginal Costs)**.** *The flexible price* $\{p_s\}$ *is consistent with some PBE if and only if*

$$k \geq \frac{1 - \alpha\xi_H/\xi_L}{\alpha(\xi_H/\xi_L - 2) + 1} \equiv \hat{K}$$

When the fraction of Insiders is high enough, the flexible prices are in fact consistent with an equilibrium. Recall that $\xi_H/\xi_L \to 1$ as $\gamma \to 1$. Note that as $\xi_H/\xi_L \to 1$ we have $\hat{k} \to \frac{1-\alpha}{1-\alpha} = 1$, and so we have all firms choosing flexible prices since the model requires $k \leq 1$.

The Proposition with sticky prices follows with $p_0 = \xi_0(1+k)/2$, as does the Corollary 1.

Having defined the local prices and demands in each state s we define the *average* prices $\bar{p}_s$ and demands $\bar{x}_s$ in state s. In contrast to the model with zero marginal costs, the sticky and flexible prices are now functions of k on each island.

$$\bar{p}_s = \int_0^{\hat{K}} p_0(k)\, dG(k) + \int_{\hat{K}}^1 p_s(k)\, dG(k) \tag{73}$$

$$\bar{x}_s = \int_0^{\hat{K}} \left(\alpha \left[1 - p_0(k)\frac{1}{\xi_s} \right] + (1-\alpha) \left[\frac{1-k}{2} \right] \right) dG(k) + \int_{\hat{K}}^1 \frac{1-k}{2} dG(k) \tag{74}$$

Plugging in we can write these equations as

$$\bar{p}_s = \int_0^{\hat{K}} \xi_0 \frac{1+k}{2}\, dG(k) + \int_{\hat{K}}^1 \xi_s \frac{1+k}{2}\, dG(k) \tag{75}$$

$$\bar{x}_s = \int_0^{\hat{K}} \left(\alpha \left[1 - \frac{1+k}{2}\frac{\xi_0}{\xi_s} \right] + (1-\alpha)\left[\frac{1-k}{2} \right] \right) dG(k) + \int_{\hat{K}}^1 \frac{1-k}{2} dG(k) \tag{76}$$

We can derive more convenient expressions for the average price difference and average demand difference:

$$\bar{p}_H - \bar{p}_L = \frac{1}{2}(\xi_H - \xi_L) \int_{\hat{K}}^1 (1+k)\, dG(k) \tag{77}$$

$$\begin{aligned} \bar{x}_H - \bar{x}_L &= \int_0^{\hat{K}} \alpha \left[\left(1 - \frac{1+k}{2}\frac{\xi_0}{\xi_H} \right) - \left(1 - \frac{1+k}{2}\frac{\xi_0}{\xi_L} \right) \right] dG(k) \\ &= \int_0^{\hat{K}} \alpha \left[\frac{1+k}{2}\frac{\xi_0}{\xi_L} - \frac{1+k}{2}\frac{\xi_0}{\xi_H} \right] dG(k) \\ &= \frac{\alpha\xi_0}{2} \left(\frac{\xi_H - \xi_L}{\xi_H \xi_L} \right) \int_0^{\hat{K}} (1+k)\, dG(k) \end{aligned} \tag{78}$$

Hence the slope of the Phillips Curve is

$$\kappa(\gamma) = \left(\frac{\xi_H \xi_L}{\alpha \xi_0} \right) \frac{\int_{\hat{K}}^1 (1+k)\, dG(k)}{\int_0^{\hat{K}} (1+k)\, dG(k)} \tag{79}$$

Note that $\hat{K}$ is strictly increasing in γ and $\xi_H \xi_L$ is strictly decreasing and converges to ξ_0^2. The result that increasing γ flattens the Phillips Curve immediately follows. As before, the coefficient in front decreases with $\gamma \to 1$, converging to ξ_0/α. Second, as γ increases, $\hat{K}$ increases, implying the numerator decreases and the denominator increases. If $G(\hat{K}) = 1$, then all firms choose sticky prices, the numerator is zero and $\kappa = 0$ (the Phillips Curve is entirely flat); this occurs for sufficiently high γ. If $G(\hat{K}) = 0$, then all firms choose flexible prices, the denominator is zero and $\kappa = \infty$ (the Phillips Curve is vertical); this occurs for sufficiently low γ. Finally, since the limit of $\hat{K} = 1$, this means that asymptotically all firms use sticky prices and so $\kappa \to 0$. Hence, the main result goes through exactly.

Since we can write

$$\bar{x}_H - \bar{x}_L = \frac{\alpha\xi_0}{2}(\zeta_L - \zeta_H)\int_0^{\hat{K}}(1+k)\,dG(k)$$

Dropping the $\frac{\alpha\xi_0}{2}$ term, we have

$$\frac{d[\bar{x}_H - \bar{x}_L]}{d\gamma} = \underbrace{(\zeta_L' - \zeta_H')\int_0^{\hat{K}}(1+k)dG(k)}_{\text{intensive margin}} + \underbrace{(\zeta_L - \zeta_H)(1+\hat{K})g(\hat{K})\frac{d\hat{K}}{d\gamma}}_{\text{extensive margin}}.$$

Proposition 15 (Divine Coincidence with MC)**.** *The output variation $\bar{x}_H - \bar{x}_L$ is increasing in γ if and only if*

$$(1-\gamma)(1+\hat{K})g(\hat{K})\frac{d\hat{K}}{d\gamma} > \int_0^{\hat{K}}(1+k)dG(k). \tag{80}$$

In this case the output gap is increasing in γ and the divine coincidence fails.

The proof follows immediately as for the analogous result without marginal costs. Note that for sufficiently high γ it must be that $(1-\gamma)(1+\hat{K})g(\hat{K})\frac{d\hat{K}}{d\gamma} < \int_0^{\hat{K}}(1+k)dG(k)$ (since $1-\gamma$ shrinks) and hence for sufficiently high γ the divine coincidence will hold, just as is the case in the model with $F(\alpha)$.

Finally, consider welfare. With costs k the expected utility with flexible price, or for uninformed agents is given by

$$\frac{1-k}{2} - \frac{1}{2}\left(\frac{1-k}{2}\right)^2 + E$$

which is

$$E + \frac{3}{8} - \frac{k}{4} - \frac{k^2}{8} = E + \frac{3}{8} - \frac{k}{4}\left(1 + \frac{k}{2}\right).$$

For informed agents facing sticky prices, expected utility is

$$E + \frac{1}{2} - \frac{1}{8}\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2](1+k)^2$$

which we can write

$$E + \frac{3}{8} - \frac{1}{8}\left(\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2](1+k)^2 - 1\right).$$

Hence, informed agents utility is lower by

$$\frac{1}{8}\left(\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2](1+k)^2-1\right)-\frac{k}{4}-\frac{k^2}{8},$$

or equivalently, multiplying by 8,

$$\begin{aligned}&\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2](1+k)^2-1-2k-k^2,\\&\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2](1+k)^2-(1+k)^2,\\&(1+k)^2\left(\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]-1\right),\end{aligned}$$

which is the welfare loss in the baseline model, times $(1+k)^2$.

We can write aggregate welfare as

$$\begin{aligned}\overline{W}(\gamma)=&\int_0^{\hat{K}}\left[\alpha\left(E+\frac{1}{2}-\frac{1}{8}\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2](1+k)^2\right)+(1-\alpha)\left(E+\frac{3}{8}-\frac{k}{4}\left(1+\frac{k}{2}\right)\right)\right]dG(k)\\&+\int_{\hat{K}}^1\left(E+\frac{3}{8}-\frac{k}{4}\left(1+\frac{k}{2}\right)\right)dG(k),\end{aligned}\tag{81}$$

which we can simplify to

$$\begin{aligned}\overline{W}(\gamma)=&E+\frac{3}{8}-\int_0^1\frac{k}{4}\left(1+\frac{k}{2}\right)dG(k)\\&-\int_0^{\hat{K}}(1+k)^2\left(\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]-1\right)dG(k).\end{aligned}\tag{82}$$

In particular, the welfare loss due to sticky prices is given by

$$\int_0^{\hat{K}}(1+k)^2\left(\xi_0^2\mathbb{E}[\zeta(\gamma,s)^2]-1\right)dG(k)$$

which is just $\int_0^{\hat{K}}(1+k)^2dG(k)$ times the identical welfare loss in the baseline model.

It immediately follows that we derive an analogous condition for when welfare is increasing or decreasing in γ.

G.2 Proofs for this setting

Proof of Proposition 14. Incentive-compatibility for the flexible-price equilibrium requires

$$\left(\frac{\xi_L(1-k)}{2}\right)\left(\frac{1-k}{2}\right) \geq \left(\frac{\xi_H(1+k)-2\xi_L k}{2}\right)\left(\alpha\left(1-\left(\frac{1+k}{2}\right)\frac{\xi_H}{\xi_L}\right)+(1-\alpha)\left(1-\left(\frac{1+k}{2}\right)\right)\right)$$

Simplifying:

$$\xi_L\frac{(1-k)}{2}\left(\frac{1-k}{2}\right) \geq \left(\frac{\xi_H(1+k)-2\xi_L k}{2}\right)\left(\alpha\left(\frac{2-(1+k)\xi_H/\xi_L}{2}\right)+(1-\alpha)\left(\frac{1-k}{2}\right)\right)$$

$$\xi_L(1-k)(1-k) \geq (\xi_H(1+k)-2\xi_L k)(\alpha(2-(1+k)\xi_H/\xi_L)+(1-\alpha)(1-k))$$

Letting $\delta = \xi_H/\xi_L$ and dividing both sides by ξ_L

$$(1-k)(1-k) \geq (\delta(1+k)-2k)(\alpha(2-(1+k)\delta)+(1-\alpha)(1-k))$$
$$(1-k)(1-k) \geq (\delta(1+k)-2k)((1-k)+\alpha(2-(1+k)\delta-(1-k)))$$
$$(1-k)(1-k) \geq (\delta(1+k)-2k)((1-k)+\alpha(1-(1+k)\delta+k))$$
$$(1-k)(1-k) \geq (\delta(1+k)-2k)((1-k)+\alpha(1+k)(1-\delta))$$

Note that $\delta(1+k)-2k > 1-k$ and $\delta > 1$ so that $1-\delta < 0$. Then rearranging we have

$$\alpha(1+k)(\delta-1)(\delta(1+k)-2k) \geq (\delta(1+k)-2k)(1-k)-(1-k)(1-k)$$
$$\alpha(1+k)(\delta-1)(\delta(1+k)-2k) \geq (\delta(1+k)-2k-(1-k))(1-k)$$
$$\alpha(1+k)(\delta-1)(\delta+\delta k-2k) \geq (\delta-1+\delta k-k)(1-k)$$
$$\alpha(1+k)(\delta-1)(\delta+\delta k-2k) \geq (\delta-1)(1+k)(1-k)$$
$$\alpha \geq \frac{1-k}{\delta+\delta k-2k} = \hat{\alpha}$$

Note that if $k=0$ (baseline model), then we get $\hat{\alpha} = 1/\delta = \xi_L/\xi_H$, which is the baseline result. In addition, we have that $\hat{\alpha} \to \frac{1-k}{1-k} = 1$ as $\gamma \to 1$ because $\delta \to 1$. So the model behaves as we

would like, with all firms choosing sticky prices when $\gamma = 1$.

Inverting the requirement (if k varies), then the cutoff for marginal cost solves

$$\alpha(\delta + \delta k - 2k) \geq 1 - k$$
$$\alpha\delta + \alpha(\delta - 2)k \geq 1 - k$$
$$(\alpha(\delta - 2) + 1)k \geq 1 - \alpha\delta$$

Note that $1 + \alpha\delta > 1 + \alpha > 2\alpha$. Hence, $\alpha(\delta - 2) + 1 > 0$, and so we have

$$k \geq \frac{1 - \alpha\delta}{\alpha(\delta - 2) + 1} = \hat{k}$$

□

H Foundations for the Demand Shock

H.1 Idiosyncratic Preference Shocks

To allow for imperfect information about the aggregate discount rate, we introduce idiosyncratic preference shocks. All consumers are perfectly informed about their own preference shock, but do not observe the aggregate preference shock. The aggregate preference shock determines the price of the bond Q (pinned down, in equilibrium, by monetary policy).

Specifically, let, as in the body, θ_H denote the *aggregate* high state, and θ_L denote the *aggregate* low state. Each consumer's individual θ_j is subject to an idiosyncratic component $u_j \sim N(\mu_u, \sigma_u^2)$:

$$\log \theta_j = \log \theta_s + u_j$$

where μ_u is the expectation of the idiosyncratic shock required for $\xi_0 = 1$. As in the body, a fraction α of consumers are Insiders and know the realization aggregate state, and a fraction $1 - \alpha$ are Outsiders and do not know the realization of the aggregate state (but update their beliefs conditional on their idiosyncratic θ_j).

The results of the body go through in this setup. Given the linearity of demand, it is straightforward to aggregate individual demands. The flexible price equilibrium replicates the results of perfect information benchmark, and Outsiders learn the aggregate shock from the price of the firm. In the sticky price equilibrium, we need to check the firms' IC constraints, which are now affected by consumer learning from their own signals. This shifts Outsiders' demand towards the realized shock. However, similar to the model in the body, one can show that there is a cutoff in α, bounded away from zero, for which sticky prices are an equilibrium.

H.2 Shocks to Future Income

A shock to future endowments can be used to microfound the main model. Here, we provide a general-equilibrium setting that shows how to write this microfoundation. Uninformed consumers know their preferences but they do not know their future income; this setting provides an isomorphism to the main model.

H.2.1 Consumers

Let consumers have utility

$$c - \frac{1}{2}c^2 - l + \beta V(C),$$

where c and C represent consumption and l is labor supply and V is a strictly increasing and concave function. The budget constraint is

$$pc + QC = d + wl + QE,$$

where w is the wage, Q is the bond price, d is the firm profit (dividend) and E is the future endowment. Let μ denote the Lagrange multiplier. Household optimization yields

$$1 - c = \mu p, \tag{83}$$

$$1 = \mu w, \tag{84}$$

$$\beta V'(C) = Q\mu. \tag{85}$$

Notice that we have $\mu = \frac{1}{w}$ and $\mu = \frac{\beta V'(C)}{Q}$. Define

$$\theta \equiv V'(C),$$

which is endogenous (we will pin it down soon). Notice that we therefore have

$$\mu = \frac{\beta\theta}{Q} \implies \mu = \frac{1}{\xi},$$

as we defined the demand shock in our paper. As before we have

$$c^* = 1 - \frac{p}{\xi}.$$

H.2.2 Firms

Firms have productivity A so $y = Al$. Thus, the marginal cost of producing y is $\frac{w}{A}$ because producing y units requires $1/A$ units of labor. Let $z \equiv 1/A$. From above we have $w = \frac{1}{\mu} = \xi$, and therefore we have marginal costs

$$\frac{w}{A} = z\xi,$$

just as in our model. The firm profit is

$$d = y(p - z\xi),$$

and the labor bill is

$$wl = \frac{wy}{A} = ywz = yz\xi.$$

H.2.3 Household Collective and Equilibrium

We suppose that consumers belong to a household that distributes profits at the end of the period. The household collective receives d and then distributes the total dividend to each consumer in the household. This is a trivial problem in a flex-price equilibrium. In that case, $c = y$ for each consumer, and so $pc = d + wl$ and $C = E$.

Now consider a sticky-price island. Let c_0 and c_1 denote consumption for the uninformed and informed. Let the household collective distribute dividends so that

$$pc_0 = d_0 + wl, \quad pc_1 = d_1 + wl,$$

i.e., each consumer type receives a portion of firm profits so that present earnings (wages plus profits) are sufficient to cover consumption. Since the firm sells $\alpha c_1 + (1-\alpha)c_0$ units, the dividend distribution scheme is feasible. Thus, $C = E$ for each agent. And that means that, indeed,

$$\theta = V'(E)$$

in equilibrium for each agent. Thus, we can let variations in E drive variations in θ.