Federal Reserve Bank of Cleveland Working Paper Series

Can Supply Shocks Be Inflationary with a Flat Phillips Curve?

Jean-Paul L'Huillier and Gregory Phelan

Working Paper No. 23-36

December 2023

Suggested citation: L'Huillier, Jean-Paul and Gregory Phelan. 2023. "Can Supply Shocks Be Inflationary with a Flat Phillips Curve?" Working Paper No. 23-36. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202336.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Can Supply Shocks Be Inflationary with a Flat Phillips Curve?

Jean-Paul L'Huillier and Gregory Phelan*

This version: October 28, 2023

Abstract

Not in standard models. With conventional pricing frictions, imposing a flat Phillips curve *also* imposes a price level that is rigid with respect to supply shocks. In the New Keynesian model, price markup shocks need to be several orders of magnitude bigger than other shocks in order to fit the data, leading to unreasonable assessments of the magnitude of the increase in costs during inflationary episodes. To account for the facts, we propose a strategic microfoundation of *shock-dependent* price stickiness: prices are sticky with respect to demand shocks but *flexible* with respect to supply shocks. This friction is demand-intrinsic, in line with narrative accounts for the imperfect adjustment of prices. Firms can credibly justify a price increase due to a rise in costs, whereas it is harder to do so when demand increases. Inflation from supply shocks is efficient and does not justify a monetary policy response.

Keywords: Cost-push shocks, shock dependence, price stickiness, output-inflation trade-off.

JEL classification: E31, E52, E58

*L'Huillier: Federal Reserve Bank of Cleveland and Brandeis University, email: jplhuillier2010@gmail.com. Phelan: Department of Economics, Williams College, email: gp4@williams.edu. We are grateful for feedback from Guido Ascari, Susanto Basu, Edouard Challe, William Chen, Pablo Cuba-Borda, Marco Del Negro, Keshav Dogra, Jordi Galí, Laura Gáti, Michael Lee, Michele Lenza, Christian Matthes, Mathieu Pedemonte, Giorgio Primiceri, Ricardo Reis, Alp Simsek, Sanjay R. Singh, Rosen Valchev, and workshop participants at the 2023 NBER Summer Institute, CEBRA and Swiss National Bank, and thank Harvard University for hospitality during the preparation of this manuscript. We thank Martin DeLuca for outstanding research assistance. All errors are our own.

1 Introduction

The current global rise in inflation presents a challenge for macroeconomics. For over three decades, inflation was virtually a non-story in advanced economies. Inflation remained incredibly stable, even amidst several large and global recessions and substantial changes in fiscal and monetary policy. One could say that there was a "great moderation" in inflation even as output remained anything but moderate at times. In other words, the relationship between inflation and demand—the so-called *Phillips curve* (PC)—has been flat, suggesting that demand shocks are not important determinants of variations in inflation.

And yet, inflation globally is now running at levels not seen in decades. In fact, such rates of inflation have not been observed in developed countries since the Great Inflation of the 1970s. Moreover, the conjunction of higher inflation with global supply factors, both in the 1970s and in the current episode, is hard to overlook. Indeed, it suggests that supply is a major determinant of inflation.[1]

Thus, our paper is motivated by two empirical facts regarding the dynamics of inflation.

Fact 1: The Phillips Curve Is Very Flat. There is a significant body of evidence that for advanced economies the PC is incredibly flat, with a slope that is small or close to zero. Hazell, Herreño, Nakamura, and Steinsson (2022) use cross-sectional data to provide evidence of only a modest flattening of the PC since 1990. According to their findings, the PC has *always* been very flat. Del Negro, Lenza, Primiceri, and Tambalotti (2020) find overwhelming evidence in favor of a very flat PC, especially since 1990. Their findings are consistent with other New Keynesian medium-scale DSGE estimations. Both papers estimate slopes on the order of 0.0020.[2]

Fact 2: Supply Shocks Are Inflationary. There is mounting evidence that supply-driven sources explain a significant portion of the current inflation. Känzig (2021) finds that oil news shocks

[1] It is of course plausible that demand factors are *also* a major part of the explanation for the recent rise in inflation, especially in the U.S., which enacted a significant post-COVID fiscal package. However, demand alone does not readily explain inflation as a *global* phenomenon, or its magnitude and cross-country synchronization, especially given flat estimates of the PC. More generally, it is well-known that in quantitative DSGE models, cost-push shocks are the main drivers of inflation dynamics.

[2] To be clear, the arguments in our paper do not necessarily rely on the PC being *very* flat. A *moderately* flat PC leads to the same remarks. We are more explicit about this in Section 2.

alone explain 50% of the forecast error variance decomposition in the U.S. Consumer Price Index (CPI), which provides a lower bound on how much supply disruptions affect inflation. In considering firm-level pricing decisions in the UK following the pandemic, Bunn, Anayi, Bloom, Mizen, Thwaites, and Yotzov (2022) conclude that "it is supply side factors that can explain most of the rise in inflation since 2021." They especially note the role of labor and materials shortages. Similarly, Ball, Leigh, and Mishra (2022) find that energy prices, supply-chain backlogs, and auto-related prices have been most important in explaining U.S. inflation. Papers with more modest findings attribute 40%–45% of the recent inflation to supply shocks (Di Giovanni et al., 2022; Shapiro, 2022).

These two facts represent a challenge for standard models aiming to account for short-run inflation dynamics. Consider first the New-Keynesian (NK) family of models. In these models, because of the flatness of the PC, the degree of price stickiness required to fit the data is very high (Del Negro et al., 2020; Hazell et al., 2022). The "Calvo fairy" does not visit firms very frequently. But then, supply shocks of a reasonable size cannot be inflationary, since the fairy *also* needs to visit firms for them to adjust prices when cost-push shocks hit.

A common misconception is that a flat PC means merely that variations in *demand* do not cause variations in inflation—whereas variations in *supply* could cause inflation. In reality, there are no empirically plausible structural shocks that will create significant variations in inflation when the PC slope is close to zero. NK models cannot deliver periods of high inflation volatility if the PC is flat, unless one is willing to accept the existence of cost shocks several orders of magnitude larger than other structural shocks. It goes without saying that at times, inflation volatility is far from zero. Using a simple calibration based on the estimates by Del Negro, Lenza, Primiceri, and Tambalotti (2020), we find that in order to generate a 1 percentage point (pp) increase in inflation, firms' desired markups need to increase by 500%. A steeper PC reduces this required markup increase somewhat, but actually not to reasonable levels. This observation is a variation on the critiques in Chari, Kehoe, and McGrattan (2009); Bils, Klenow, and Malin (2012), but here we show that the flat PC is the fundamental cause of the implausible price markup shocks that they highlight.

Standard pricing frictions—such as Calvo, Taylor, menu, or Rotemberg costs—all suffer from the same issue. The reason is the common feature of treating rigidity with respect to supply and demand symmetrically. Standard pricing frictions cannot account for significant variations in inflation (such as the recent episode), *and* at the same time predict a flat PC, which is clearly evidenced by the "missing disinflations" during the Great Recession and the onset of the COVID recession, plus the "missing inflations" during the housing boom of the early 2000s, and during the unprecedented expansion of the Federal Reserve's balance sheet of the last decade. Notice that the PC has remained remarkably flat despite *large* movements in aggregate demand.[3]

Our main goal, therefore, is to present a model that can address both Facts 1 and 2. As we observed, the data suggest that firms treat demand and supply disturbances distinctly when setting prices. In other words, the evidence points towards *shock-dependence* in price adjustment. We provide a microfounded model of shock-dependent price stickiness in which inflation can be entirely rigid with respect to demand shocks and yet entirely flexible with respect to supply shocks. In the model, firms face no exogenous or technological constraints to adjust prices, such as menu costs or the blessing of a Calvo fairy. The proposed microfoundation is based on a strategic interaction between firms and consumers. The strategic environment we consider can simultaneously produce a very flat PC while also producing significant inflation in response to supply disturbances.

Our model captures a strategic firm-consumer interaction where firms are better informed than consumers about aggregate conditions. This can be motivated by the view that firms are sophisticated players in many markets, or by the view that firms observe a wider range of transactions in the market they supply. The ability to aggregate information puts firms in a privileged position. For instance, firms may posses superior information about the state of aggregate demand (given their observation of quantities sold), or about costs (given their observation of the cost of inputs). As a result of their superior information, firms optimally evaluate the implications of alternative pricing strategies. The formal macroeconomic setup builds on the models of L'Huillier (2020) and L'Huillier and Zame (2022), with the addition of supply shocks. The central idea of this setup is that prices may convey information possessed uniquely by the firm, and consumers' perception of alternative prices influence their purchase decision, and ultimately firms' profits.[4]

[3]For instance, during the onset of the COVID pandemic in 2020:Q2, U.S. real GDP fell by more than 32% (annualized) and unemployment rose to 15%, but headline CPI inflation was roughly -4%(annualized), and back to positive territory in 2020:Q3.

[4]L'Huillier (2020) shows that this model delivers hump-shaped dynamic responses of both output and inflation, even in the absence of bells and whistles. Thus, these models also deliver realistic predictions for the propagation

More specifically, in the model, there are two aggregate shocks: a shock to aggregate demand and a shock to aggregate supply. A fraction of consumers have less information than firms about the realization of the shocks. Firms are monopolists, and set prices. We use standard preferences and assumptions about agent rationality, and solve for a perfect Bayesian equilibrium of the game played between firms and consumers. We find that prices can be *endogenously* sticky with respect to demand shocks. But, firms always pass on cost increases to consumers (for any degree of information asymmetry). Prices adjust flexibly to supply shocks.

The intuition underlying our results is provided by considering firms' incentives, and how they differ depending on the type of shock.

Consider first the case of demand shocks. A nominal demand shock determines whether it is a "good time to buy" (or not). Suppose that, for exogenous reasons such as strong labor markets or improved income prospects, it is now a good time to buy. If all agents know this, then firms would post a high price in response to high demand. But what if *a subset* of consumers do not know that now is a good time to buy? Call these consumers "uninformed." The demand function of uninformed consumers is, in equilibrium, sensitive to the firm's pricing strategy. The question is whether the firm can, by posting a higher price, convince the uninformed consumers that it is indeed a good time to buy. If this were an equilibrium, the firm would be using the price increase to stimulate total demand: posting a high price means it is a good time to buy, which makes the informed consumers pay higher nominal prices for the same real quantity (hence shifting the demand curve outwards). Crucially, high prices are not necessarily credible: firms are always tempted to stimulate demand (even when it is not a good time to buy). Clearly, such behavior cannot lead to an equilibrium with rational consumers: higher prices would fail to stimulate demand, and lead to lower profits. Hence, a strategic friction emerges. As a result of this strategic friction, for low fractions of informed consumers, the equilibrium is one where firms' prices do not reflect the state of demand, i.e. prices are endogenously sticky. If enough consumers are informed, they provide discipline, and prices can perfectly adjust to reflect the demand shock.[5]

Consider now the case of supply shocks. In our model, a supply shock is real shock that

of shocks. L'Huillier and Zame (2022) show that the price stickiness result is robust to the consideration of optimal mechanisms and contract setting.

[5]Formally, the firm's incentive constraint of posting *low* prices when it is *not* a good time to buy needs to be satisfied in equilibrium. This constraint is satisfied at a cutoff value on the fraction of informed consumers (high prices are credible).

determines the level of firms' marginal costs. The crucial distinction between demand and supply shocks is that supply shocks are *not* payoff-relevant for consumers. By standard arguments, consumers only care about the price they pay, and about their own demand (the realization of the demand shock). Therefore, firms have no incentive to stimulate demand when costs increase. Whether or not consumers know the costs, when costs are high, firms earn higher profits by charging higher prices; and when costs are low, firms earn higher profits by charging lower prices. Profit-maximizing prices do not depend on consumers knowing firms' costs. In this case, there is no strategic friction, and the equilibrium is given by the solution to the standard monopoly problem. When costs rise, prices rise and consumers demand less. Nonetheless, higher prices are necessary because of the higher costs.[6]

To summarize, demand changes lead to a strategic friction. Sticky prices emerge naturally as a credible firm strategy. But supply disruptions do not lead to a strategic friction, and prices are flexible. When both shocks arise simultaneously, the price can fully adjust to the supply shock, but not adjust to the demand shock, an extreme case of shock-dependence in price adjustment.

Thus, our model can simultaneously produce a flat PC and large responses of inflation to supply disturbances. This result helps rationalize the empirical evidence by Bunn et al. (2022), Känzig (2021) and Ball et al. (2022), among others, which favors the interpretation that supply shocks are responsible for the recent inflation. This is precisely what our model predicts.

In general equilibrium, our model delivers a PC with a slope that depends on the degree of nominal rigidity with respect to output gaps, which are determined by the demand-side of the economy, and can be zero. Furthermore, the PC's *intercept* is determined by the supply side of the economy. This is important: the PC can flexibly shift up and down with supply shocks, *and*, at the same time, be entirely flat. This is impossible in standard models, because there the intercept is fixed at steady state when the PC is flat (see Section 2).

The shape of the PC implied by our model has important policy implications. Supply shocks shift the position of the PC, but do not generate output gaps nor price dispersion, resulting in an efficient goods market allocation. An immediate implication is that no policy action is justified

[6]As a matter of fact, our model is one where firms would prefer to communicate that price increases are due to cost increases ("we are raising our prices because our costs have increased"), rather than demand increases ("we are raising our prices because our product has become more popular"). Whereas this is not exactly how firms behave in the model, it is a useful thought to get intuition. It is worth noting that Blinder et al. (1998) provide survey evidence that "when costs increase, [...] customers normally tolerate price increases" (see Table 7.3, p. 157).

by supply disturbances. This offers thought-provoking results regarding the optimal behavior of a central bank facing supply-driven inflation. Following a supply disruption, a rise in the price level is socially optimal and simply reflects firms' lower productivity or higher costs of production. Higher prices result in lower demand, which is efficient. Therefore, there are no welfare losses due to price rigidities. If the central bank raises interest rates to lower inflation, it will generate a negative demand shock. The price level will turn sticky downward, and an inefficient negative output gap will result as a consequence of the extra fall in demand generated by the central bank.

To sum up, in our model, price level fluctuations due to shifts in the supply side of the economy are optimal and should not be actively stabilized. This type of prediction highlights the usefulness of explicitly microfounding the sources of the price friction. Our paper underlines the importance of capturing the shock-dependent response of inflation commanded by the data when thinking about the optimal monetary response to alternative shocks.

Related Literature. There is a classic literature providing evidence that the firm-customer relation is what limits price adjustment, suggesting that nominal price-setting frictions are demand-based (Hall and Hitch, 1939; Okun, 1981; Kahneman et al., 1986; Greenwald and Stiglitz, 1989; Blinder, 1991). Blinder et al. (1998) provide survey evidence that when asked to explain their reluctance to increase prices after an increase in costs, firms' managers usually answer that "price increases cause difficulties with customers." For recent work on the topic, see Rotemberg (2005), Nakamura and Steinsson (2011), Gilchrist, Schoenle, Sim, and Zakrajšek (2017), Gaballo and Paciello (2021), and Dupraz (2017). To be best of our knowledge, none of these works share the prediction that prices are sticky to demand shifts but flexible to supply shifts.

A recent number of papers argues that fairness concerns and other behavioral features constitute bases for price rigidity in the survey evidence cited above (Rotemberg, 2005, 2011; Eyster et al., 2021). Our line of work provides a theoretical foundation for this type of rigidity in a model with standard assumptions on agents' rationality and preferences. In our model, firms behave strategically when setting prices, considering how consumers may perceive a posted price. The key insight is that the degree of information among consumer may limit price adjustment, when firms are strategic and may be tempted to stimulate demand.

There is a robust literature that attempts to reconcile the New Keynesian framework with the data. The first major challenge is explaining (or reinterpreting) the so-called missing inflations during the Great Recession. It constitutes an anomaly within the standard paradigm. Several factors have been considered to explain (or reinterpret) these phenomena, such as inflation expectations (Jorgensen and Lansing, 2019), online retail (Cavallo, 2018), and globalization (Forbes, 2019). See L'Huillier and Schoenle (2023) for related evidence of the link between the frequency of price adjustment and the inflation target.

There are laudable contemporaneous attempts within the literature to improve the ability of the NK model to match inflation dynamics, for example by taking seriously non-linearities or how belief formation affects the PC (e.g., Ascari and Fosso, 2021; Ascari et al., 2022; Harding et al., 2022, 2023). This literature moves the NK model in the right direction because it reconciles the observed flat PC with the rise in inflation observed after 2021:Q2. In this vein, our contribution is targeted at de-linking the tight restriction that symmetric rigidities impose within the NK model. We argue that incorporating shock-dependent price rigidities (with respect to supply and demand disturbances) will provide macroeconomists with the ability to better account for the facts.

Similarly, motivated by the previous findings, Benigno and Ricci (2011) shows that downward nominal wage rigidities produce nonlinear PCs that flatten at low inflation. Coibion, Gorodnichenko, and Kamdar (2018) consider how belief formations affect the PC. Furlanetto et al. (2023) analyze the role of monetary policy, observing that a firm commitment can generate stable inflation. Blanco et al. (2022) study nonlinear PCs in menu cost economies. Fitzgerald et al. (2020) provide evidence that the PC is stable. McLeay and Tenreyro (2020) show that, away from the zero lower bound, the combination of cost-push and demand shocks leads to difficulties with identification of the slope of the PC.

Many papers recognize that different microfoundations for price stickiness have important implications for inflation and monetary policy (Ascari, 2004; Caballero and Engel, 2007; Karadi et al., 2022; Ascari and Haber, 2022) and that the PC is endogenous, with important implications (Kiley, 2000; Levin and Yun, 2007; Kocherlakota, 2022; Gaballo and Paciello, 2021; Petrosky-Nadeau and Bundick, 2021). Werning (2022) shows that the effect of inflation expectations on aggregate inflation depends on the microfoundation for price stickiness.

There is a robust literature studying information frictions in price setting, pioneered by Mankiw and Reis (2002). Ball, Mankiw, and Reis (2005) show that when price setters are slow to incorporate macroeconomic information into the prices they set, the optimal monetary policy is price level targeting. Acharya (2017) considers a sticky information model in which the endogenous decision of when to acquire new information about different shocks leads prices to change frequently and by large amounts in response to idiosyncratic shocks but sluggishly in response to monetary shocks. Afrouzi and Yang (2021) study an economy with dynamic rational inattention and find that more hawkish monetary policy flattens the PC. Gutiérrez-Daza (2022) considers an economy in which consumers learn from shopping. Bernstein and Kamdar (2022) study optimal monetary policy with central bank inattention. Ilut, Valchev, and Vincent (2020) demonstrate how ambiguity aversion on the side of firms leads to a friction that matches both the micro and macro evidence on price adjustment.

2 Motivation: No Cost-Push Inflation with a Flat Phillips Curve in the New Keynesian Model

The structural New Keynesian (NK) PC is generally written as

$$\widehat{\pi}_t = \beta \mathbb{E}_t[\widehat{\pi}_{t+1}] + \kappa \widehat{x}_t + \lambda \widehat{z}_t, \tag{1}$$

where $\widehat{\pi}_t$ denotes inflation, $\widehat{x}_t$ denotes the output gap, $\widehat{z}_t$ denotes a *structural* cost-push shock, and β, κ and λ are parameters. β is the discount factor. κ, the slope of the PC, measures the sensitivity of inflation to output gap fluctuations. λ measures the sensitivity of inflation to cost-push shocks. (κ and λ are related by a proportionality coefficient.)

Empirical Evidence on κ and λ. Estimates of *both* κ and λ are very small. Hazell, Herreño, Nakamura, and Steinsson (2022) estimate the slope of the PC in the cross section of U.S. states using two distinct instrumental variables (IV) for aggregate demand (which, in their model, causes unemployment to fluctuate). The first IV is given by lagged unemployment and non-tradeable prices, which delivers an estimate of $\kappa = 0.0062$.[7] Del Negro, Lenza, Primiceri, and Tambalotti

[7] The second IV captures tradeable demand spillovers as an instrument for demand (unemployment). Based on the idea that supply shocks in tradeable sectors will differentially affect demand in non-tradeable sectors, they estimate $\kappa = 0.0062$ using their preferred specification (remarkably, both strategies deliver the same estimate). They also

(2020) use an alternative empirical strategy. They employ a medium-scale DSGE model with real rigidities, sticky wages, and financial frictions. Their mean posterior estimates of λ are 0.015 pre-1990 and 0.0015 post-1990. In sum, their empirical estimates imply a value of λ that is no more than 0.015 in general and as small as 0.0015 post-1990.

Structural Cost-Push Shocks, Reduced-Form Cost-Push Shocks, and Inflation. The canonical supply shocks are productivity shocks. In the NK model, negative supply shocks are inflationary because they create a positive output gap. However, productivity shocks cannot create significant inflation because κ is almost zero. This is one reason why the NK literature has adopted other shocks, such as a shocks to markups, as the standard inflationary supply shocks.

A common step in the literature consists of re-scaling markup shocks, especially in practice when it comes to medium-scale DSGE estimation. Defining the *reduced-form shock* $\widehat{v}_t \equiv \lambda \widehat{z}_t$, we can write the NKPC in a form similar to equation (1):

$$\widehat{\pi}_t = \beta \mathbb{E}_t[\widehat{\pi}_{t+1}] + \kappa \widehat{x}_t + \widehat{v}_t. \tag{2}$$

While variations in $\widehat{x}_t$ cannot create inflation because κ and λ are so small, variations in the *reduced-form* shock $\widehat{v}_t$, which affects inflation one-for-one, can.

This last step is problematic when it comes to empirics. Matching the empirical behavior of inflation requires ascribing a great deal of volatility to the $\widehat{v}_t$ term. But, the structural shock is the change in underlying costs $\widehat{z}_t$. If $\lambda = 0.0020$, which is the estimate from Hazell, Herreño, Nakamura, and Steinsson (2022) (or rounding up the post-1990 mean posterior estimate from Del Negro, Lenza, Primiceri, and Tambalotti 2020 at 0.0015), then generating a 1% variation in $\widehat{v}_t$ requires a change in $\widehat{z}_t$ by $1\%/0.002 = 500\%$. If steady-state markups are 12.5%, then a shock of this size requires that desired markups increase to 575%, an implausibly large structural shock.[8] A 1 pp increase in inflation due to a cost-push shock requires implausibly large structural shocks. The standard normalization is far from innocuous. It amounts to an important re-scaling of the

directly estimate $\lambda = 0.0020$ (see Table C.2 in their paper). These estimates point to a very flat PC.

[8]The shock is applied to the gross markup, thus the gross markup with the shock is $1.125(500\%) + 1.125 = 6.75$, which corresponds to a net markup of 575%. Since the gross markup is $\frac{\varepsilon}{\varepsilon-1}$, where ε is the elasticity of substitution, this means that the markup shock corresponds to a drop from the steady-state of 9 to 1.17. If the markup shock is taken literally in terms of log deviations, then the increase in the gross markup is even larger ($1.125e^5$).

markup shock.

The severity of this problem can also be seen by considering an alternative reduced-form way to write the cost-push shock $\widehat{v}_t \equiv \kappa(\widehat{y}_t^e - \widehat{y}_t^n)$, where $\widehat{y}_t^e$ and $\widehat{y}_t^n$ represent the efficient and natural levels of output, respectively (Galí, 2015). With $\kappa = 0.0062$, generating a 1% cost-push shock requires the natural rate of output to differ from the efficient level by $1\%/0.0062 = 161\%$. The welfare implications of these cost-push shocks would, therefore, be enormous.

Our reading of the literature indicates that, while this point is understood, for instance, by Chari, Kehoe, and McGrattan (2009) and Bils, Klenow, and Malin (2012), the novelty here is to point out that this is a direct *economic* implication of the type of pricing friction present in NK models, where the same degree of price rigidity applies to all shocks, symmetrically. In order to generate a very flat PC, the model needs to be parameterized with very high degrees of rigidity. The immediate implication of this is, clearly, that markup shocks cannot easily generate inflation volatility.

How do these observations square with the recent global rise in inflation since the second quarter of 2021? This rise in inflation has been much larger than 1%. For the sake of the argument, consider an increase in costs capable of rising inflation from its steady state value by 5 pp (say from 2% to 7%). To generate such an increase with a value of $\lambda = 0.0020$ requires a change in $\widehat{z}_t$ of 2500%. In this scenario, desired markups increase from 12.5% to 2,812.5%. Equivalently, the natural and efficient levels of output would have to differ by 806%. Both in the 1970s and in the recent episode, the effect of supply shocks on inflation has been large. Therefore, the observed large increases in inflation aggravate the issue raised by the usual normalization of cost-push shocks in the NK model.

The re-scaling $\widehat{v}_t \equiv \lambda \widehat{z}_t$ or $\widehat{v}_t \equiv \kappa(y_t^e - y_t^n)$ is conceptually problematic for the reasons we just gave. However, over time this problem has become acute as there is some evidence that the PC may have flattened over the years (Del Negro, Lenza, Primiceri, and Tambalotti, 2020). Some commentators have suggested that the very recent inflation episode is evidence that the PC has steepened again. We are very cautious about overturning careful empirical work based on decades of data and thoughtful identification strategies on the basis of one episode of high inflation. However, even if the slope has increased, say, ten-fold, the problem remains: based on our first calculation above, structural shocks would need to be at least 50% to generate a 1% reduced-form shock, or a 250% shock to generate a 5% increase in inflation. In terms of output levels, the natural and efficient levels would still have to differ by 16% and 81% to generate those reduced-form shocks.

These conclusions are robust in several dimensions. First, to visualize the robustness of our conclusion to alternative values of the slope of the PC, Figure 1 presents the size of the markup shock needed to generate a 1 pp increase in inflation as a function of λ. The y-axis, which presents the size of the shock, is plotted on a logarithmic scale. The dark-green circle surrounds the most plausible values, based on the Hazell, Herreño, Nakamura, and Steinsson (2022) and Del Negro, Lenza, Primiceri, and Tambalotti (2020) papers. For these values, the conclusion is that the size of the markup shock is orders of magnitude larger than the effect on inflation. Therefore, even important increases in λ, such as Del Negro, Lenza, Primiceri, and Tambalotti (2020)'s pre-1990 estimate or a ten-times larger estimate than their baseline, do not affect our conclusion that the markup shock is very large. Second, real rigidities or sticky wages could, in theory, reconcile a very low κ with a larger λ. As we discuss in Appendix C, the empirical evidence does not favor this possibility.

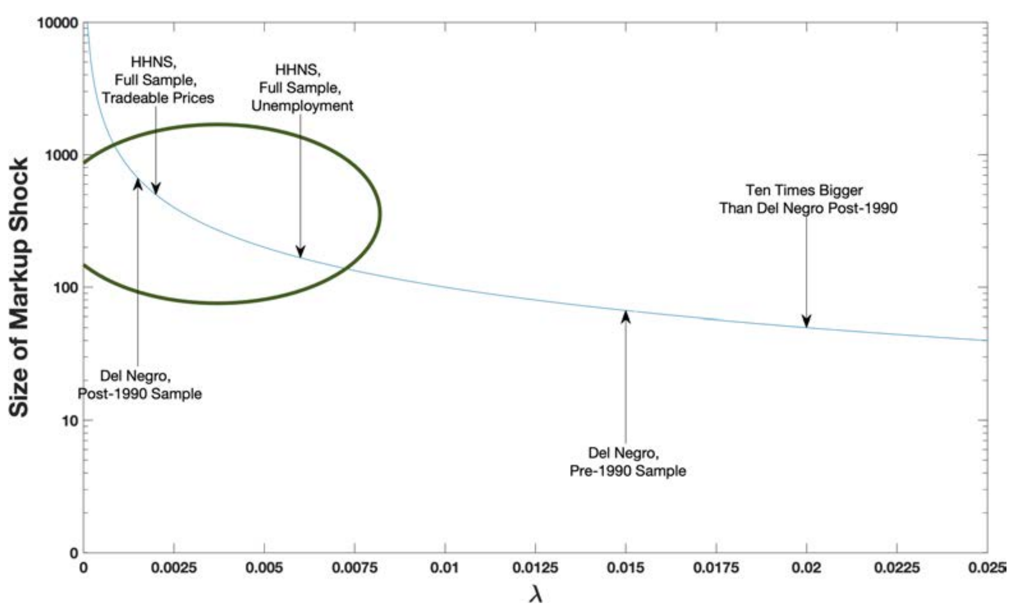

Figure 1: Markup Shock for 1 pp Increase in Inflation (Log Scale)

In the next section, we argue that this observation suggests considering price-setting frictions with shock-dependent price rigidities: firms' prices are flexible with respect to supply shocks but simultaneously sticky with respect to demand disturbances. In essence, we propose a model in which λ and κ are not directly related, and one can have a very low κ without requiring a very low λ, with plausible empirical relationships between marginal costs and output gaps.

3 Strategically Sticky Prices

We now present the microfoundation for shock-dependence in price adjustment. The framework we use is parallel to L'Huillier (2020), with the addition of shocks to supply. For clarity, we first consider the strategic behavior of firms in response to demand shocks and then consider their behavior in response to supply shocks. Finally, we consider both shocks simultaneously. Crucially, in our model, we can use a productivity shock as our canonical supply shock, though our results apply to any type of shock that affects firms' marginal costs.

For ease of exposition, we use a simple two-period, partial-equilibrium model. The same points could be made in a more complicated infinite-horizon, general-equilibrium model (Appendix E).

3.1 Model Setup

There are two dates, the present and the future, which we interpret as the short run and the long run. In the short run, production and trade in goods markets will be subject to frictions; in the long run, agents have exogenous endowments and trade 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; we formalize this by positing that all consumers trade endowments in a Walrasian, perfectly, competitive market. For convenience, we follow Lagos and Wright (2005), and 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 CB sets the nominal interest rate. In this partial-equilibrium setup, there is no labor supply.[9]

Aggregate State. In the present, there is uncertainty about the aggregate state of the economy, denoted by $\{\xi, z\}$. The aggregate state captures two dimensions of the economy: aggregate demand pressure ξ and aggregate supply pressure z. To simplify notation, we drop state-subscripts when not absolutely necessary. We assume that the two shocks are orthogonal.

We model aggregate demand as determined by a shock to consumers' future *nominal* marginal utility, captured by a random variable θ. Under consumption smoothing, shocks to future marginal utility affect optimal nominal spending in the present. 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, e.g., strong labor markets, rosy expectations of income growth, etc. We do not take the preference 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. For simplicity, we assume there are only two possible states, low θ_L and high θ_H, that occur with equal probability.

A virtue of this nominal shock to future marginal utility is that it implicitly defines the consumers' discount rate, denoted by ρ. In the aggregate, it can also be interpreted as the economy's natural rate of interest, since mechanisms affecting nominal demand would affect the economy's natural rate of interest. Demand in the present will be high when consumers' discount rate is high, and therefore ρ_H corresponds to the high state, and ρ_L corresponds to the low state, with $\rho_H > \rho_L$.

The central bank sets the nominal interest rate, denoted i_s. Because monetary policy affects aggregate demand, we define a net demand shock, or *demand pressure*, denoted by ξ. This is given by the two realizations:

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

Thus, the shock ξ captures changes in demand pressure caused by the changes in subjective and market discount factors. Let ξ_0 denote the harmonic mean of ξ_L, ξ_H, that is, $\xi_0 \equiv \left[\frac{1}{2}\left(\xi_L^{-1} + \xi_H^{-1}\right)\right]^{-1}$.

[9] Appendix E presents an equivalent model with labor and monetary frictions where the CB sets money supply.

We normalize $\xi_0 = 1$. The significance of ξ will be clear when we study the consumer problem below.[10]

Aggregate supply pressure determines firms' real marginal costs, denoted by z. We define *supply shocks* as exogenous changes in marginal costs. We can think of variations in z as capturing any variations in marginal costs, such as changes in the prices of intermediate goods or energy (like oil), changes in wages driven either by shocks to labor supply or shocks to wage-bargaining, or changes in productivity. Importantly, and different from the NK model, whether the change in z represents an efficient or inefficient variation is not critical for our analysis. We develop this point more deeply on Section 3.5. We use the terms "supply shocks" and "cost shocks" interchangeably. For simplicity, we suppose that the supply shock can take only two values, high z_H and low z_L, that occur with equal probability, with expectation $\mathbb{E}[z] = z_0$.

We suppose the economy begins at $\{\xi_0, z_0\}$ and then experiences a shock in the present.

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 use this source of heterogeneity to allow for distinct patterns of price adjustment 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. Define α_0, α_1 to be the lower and upper limits of the support of F

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

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

[10]Similar to L'Huillier (2020), it is possible to 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). However, in this paper we prefer the current formulation in terms of shocks to future marginal utility (which could, for instance, also be driven by expectations about future income) in order to allow for a standard specification of monetary policy.

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 $\alpha_0 < \alpha_1$.

Each island is inhabited by a single monopolistic firm with real marginal cost z. 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. L'Huillier (2020) motivates the assumption of informed firms by the existence of a 'time 0' where firms observe market transactions.

Consumer Problem. We index a typical consumer by j. 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 the random variable determining the realization of nominal marginal utility in the future. 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.[11]

Consumer j solves

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

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 $\beta\theta \equiv \frac{1}{1+\rho}$. Our assumptions guarantee that, in equilibrium, the consumer's optimal choice in the present will always lead to positive consumption in the future.[12]

[11] L'Huillier (2020) allows for learning about aggregate shocks in the decentralized market for goods. In such a setup, α is a state variable. As a result of consumer learning, the impact of nominal demand shocks slowly grows over time, as the information about future nominal adjustments diffuses among the consumer population. In the long run, α tends to 1. Here, we simplify the setup by having a fixed α in every island, and focus on a static version of the model.

[12] As in Lagos and Wright (2005), quasilinearity in future consumption eliminates wealth effects. Quadratic utility in present consumption is for analytical convenience.

Firm Problem. Each island is populated by a single monopolistic firm. We assume that firms produce the consumption good at a nominal marginal cost $z\xi$, with $z_H < 1$. Let $y(p)$ denote the consumer demand at the island given a price p. The firm profit given price p is $(p - z\xi)\, y(p)$. Note that demand shocks mechanically affect firms' nominal marginal costs, reflecting general equilibrium adjustments under flexible product markets. One microfoundation that generates this relationship would be a working capital assumption: the timing convention that production costs (wages or intermediate goods) are paid at the end of the first period, and therefore the production cost is the discounted value of the price level in the future; see Appendix E for a formalization. In addition, if demand shocks represented nominal shocks to the money supply, then ξ would be the future price level P and the nominal cost zP would directly reflect variations in the long-run price level; see L'Huillier (2020) for a formalization.

Perfect Information Benchmark. The full-information benchmark is straightforward. Consumer j's optimal short-run demand is

$$c^* = 1 - p\frac{\beta\theta}{Q} \equiv 1 - p\frac{1}{\xi}.$$

Note that when ξ is high, demand in the present is high, and thus ξ acts as a demand shifter. The optimal price is $p^* = \frac{\xi(1+z)}{2}$, flexible with respect to both shocks, and the *natural level of output*, denoted y^n, is

$$y^n = \frac{1-z}{2}.$$

Notice, output is fixed regardless of the demand shifter ξ (the demand shock is neutral). However, the supply shock z does affect the natural level of output.

3.2 Demand Shocks and Strategically Sticky Prices

Here, we suppose that firms' real marginal cost is fixed at z and we consider how variations in demand ξ affect the slope of the PC. We denote the state (demand) by s.

Equilibrium Short-Run Prices. With imperfect information, consumer j's optimal short-run demand is

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

Hence, $\mathbb{E}_j\left[\frac{1}{\xi_s}\right]$ is the expectation of the net demand shock, determining whether demand in the present is strong or weak.

When information is not perfect and consumers and firms behave strategically, we must ask whether the full-information price is consistent with equilibrium in the implicit game between the firm and the consumers. Because the full-information price varies with the state, we refer to the full-information prices as the flexible price, with p_s defined by p^* above. 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. Lemma 1 provides a sharp answer to this question.

Lemma 1 (PBE with Flexible Prices)**.** *The flexible price $\{p_s\}$ is consistent with some PBE if and only if $\alpha \geq \overline{\alpha} \equiv \frac{1-z}{(\xi_H/\xi_L)(1+z)-2z}$.*

When the fraction of Insiders is high enough, the flexible prices are in fact consistent with an equilibrium. As a note, $\overline{\alpha}$ is decreasing in z, and $\overline{\alpha} = \frac{\xi_L}{\xi_H}$ when marginal costs are zero.

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 in the absence of a shock (i.e., if $\xi = \xi_0$), and it is the price that would be optimal if no consumers were informed ($\alpha = 0$). Of course, we require that p_0 be consistent with some PBE as well; Lemma 2 provides a complete characterization.

Lemma 2 (PBE with Sticky Prices)**.** *The sticky price is $p_0 \equiv \xi_0(1+z)/2$. For z sufficiently small, p_0 is consistent with some PBE if $\alpha \leq \overline{\alpha}$.*

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. Finally, whenever *both* the sticky price and the flexible price schedule are consistent with PBE, the firm strictly prefers the flexible price schedule.[13]

Aggregate Prices and the Phillips Curve. 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 deliver the higher profit. On islands where $\alpha \geq \overline{\alpha}$, the firm will choose the flexible price schedule $\{p_s\}$ and demand will not depend on the state (it will be at the natural level of output y^n). On islands where $\alpha < \overline{\alpha}$, 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

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

(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.) Having defined the local prices and demands in each state s, aggregate prices $\bar{p}_s$ and demands $\bar{y}_s$ in state s are

$$\bar{p}_s = \int_0^{\overline{\alpha}} p_0 \, \mathrm{d}F(\alpha) + \int_{\overline{\alpha}}^1 p_s \, \mathrm{d}F(\alpha), \tag{3}$$

$$\bar{y}_s = \int_0^{\overline{\alpha}} \left[\alpha\left(1 - p_0 \frac{1}{\xi_s}\right) + (1-\alpha) y^n\right] \mathrm{d}F(\alpha) + y^n \left[1 - F\left(\overline{\alpha}\right)\right]. \tag{4}$$

The baseline aggregate price used to define the PC is therefore given by $\bar{p}_0 = \xi_0(1+z)/2$, which is the flexible price given the baseline demand shock ξ_0, and baseline output is $\bar{y}_0 = y^n$. Define the *output gap* as the difference between actual and natural output: $\bar{x}_s \equiv \bar{y}_s - y^n$. We define the PC to be the line crossing the two points $(\bar{x}_s, \bar{p}_s)$ and $(0, \bar{p}_0)$. We define the *slope* $\bar{\kappa}$ of the PC to be the ratio:

$$\bar{\kappa} = \frac{\bar{p}_s - \bar{p}_0}{\bar{x}_s}.$$

[13] We require that z not be so large because ξ affects both demand and *nominal* costs, and thus, a low z ensures that this latter effect does not dominate.

From equations (3) and (4), the slope of the PC is

$$\bar{\kappa} = \xi_s \left(\frac{1 - F(\overline{\alpha})}{\int_0^{\overline{\alpha}} \alpha \, \mathrm{d}F(\alpha)} \right). \tag{5}$$

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 informed consumers on islands with sticky prices.

The PC flattens (i.e., the slope $\bar{\kappa}$ decreases) as the size of demand shocks decreases, and it becomes perfectly flat (i.e. $\bar{\kappa} = 0$) in the limit as $\xi_H \to \xi_L$. Small net demand shocks could reflect, for example, a hawkish central bank. We can state the following result.

Lemma 3 (Flat Phillips Curve Under Demand Shocks)**.** *The Phillips curve is flat, $\bar{\kappa} = 0$, if demand shocks are sufficiently small i.e., if ξ_H is sufficiently close to ξ_L.*

To understand the mechanism underlying Lemma 3, note that, on all islands where sticky prices prevail, the PC (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 PC (for those islands) is vertical (has slope $+\infty$): prices depend on the state but demands do not. As net demand shocks decrease, more firms choose sticky prices and fewer firms choose flexible prices. Fundamental demand shocks can still be large so long as the central bank policy offsets the shock, (i.e., raising rates in response to positive shocks, see L'Huillier, Phelan, and Zame, 2022).

3.3 Supply Shocks and Strategically Flexible Prices

We now consider the consequences of a supply shock that directly affects firms' costs. Recall that the problem of standard models is having robust responsiveness of inflation to empirically plausible supply shocks when the PC is very flat. To simplify exposition, we suppose there are no demand shocks for now ($\xi = \xi_0 = 1$). We denote the state (supply) by ς.

Let λ denote the marginal costs elasticity of aggregate prices. The response of inflation to supply shocks is given by

$$\lambda \equiv \frac{\partial \bar{p}/p_0}{\partial z_\varsigma / z_0}.$$

As we supposed with the demand shock, the firm and the Insiders know the value of z_ς, but the Outsiders do not. Note that the Outsider demand depends only on the price p, which is observable. Thus, it is a straightforward result that the firm will always choose a flexible price (for any fraction of Insiders).

Proposition 1 (Flexible Prices with Inflationary Shocks)**.** *When the economy experiences shocks to the marginal cost z_ς, all firms choose a flexible price $p_\varsigma = \frac{1+z_\varsigma}{2}$, output is at potential y^n, and $\lambda = 1$, for all α.*

A few remarks are in order. First, in this model, in which prices are endogenously sticky for strategic reasons, demand shocks can lead to price stickiness while supply shocks will never lead to price stickiness. This means that the responsiveness of prices to cost shocks is completely different from the responsiveness to demand shocks.

Second, even though costs are not payoff relevant for consumers, for consistency with the previous section in this PBE consumers can be allowed learn the realization of costs from the price (although one can construct the same equilibrium without this requirement, since z_ς is irrelevant to consumers). Different from the case of demand shock, firms are not subject to an incentive compatibility constraint (there is no strategic friction between firms and consumers).

Third, because the equilibrium is endogenously a flexible-price equilibrium, there is no welfare loss in response to the inflationary shock, and thus the optimal response of the central bank is to do nothing. In this model, there is no need for monetary policy to respond to the inflationary shock. Nonetheless, central banks may have other reasons for wanting to respond to inflation in response to a shock to marginal costs. In the next section we show that if Outsiders do not know how the central bank will respond to an inflationary shock, then firms may adjust their prices in response to the inflationary shock but not in response to the central bank.

3.4 Supply and Demand Shocks

We now introduce both supply and demand shocks, denoting as before the state of supply by ς and the state of demand by s. The following result characterizes equilibrium prices. It characterizes the consequences of shock dependence when both shocks hit the economy.

Proposition 2 (Demand and Supply Shocks)**.** *When the economy experiences both supply and demand shocks and* $\alpha \leq \overline{\alpha}_H \equiv \frac{1-z_H}{(\xi_H/\xi_L)(1+z_H)-2z_H}$, *aggregate prices are flexible with respect to the supply shock but sticky with respect to the demand shock. Firms post* $p_{0,\varsigma} = \frac{1+z_\varsigma}{2}$.

In our shock-dependent model, equilibrium prices are sticky with respect to the demand shock (and therefore firms do *not* condition their prices on it), but equilibrium prices are flexible with respect to the supply shock (and therefore firms condition their prices on it). Stickiness with respect to demand is evident in the lack of dependence of $p_{0,\varsigma}$ on ξ_s (recall the normalization of the state-state demand shocks term, $\xi_0 = 1$). Flexibility with respect to supply is evident in the dependence of $p_{0,\varsigma}$ on the z_ς term.

We have already noted that firms have no strategic incentive to misrepresent the realization of z_ς, and that remains true here. The firm can convey information about the shock z_ς without conveying information about the demand shock ξ.

The Phillips Curve and Supply-Driven Inflation. How sticky aggregate prices are with respect to demand depends on the fraction of firms choosing sticky prices. The aggregate price level given cost z and demand ξ is therefore

$$\bar{p}_{s,\varsigma} = \int_0^{\overline{\alpha}_\varsigma} \frac{\xi_0(1+z_\varsigma)}{2}\,\mathrm{d}F(\alpha) + \int_{\overline{\alpha}_\varsigma}^1 \frac{\xi(1+z_\varsigma)}{2}\,\mathrm{d}F(\alpha) = \left(\frac{1+z_\varsigma}{2}\right)\left(\xi_0 + (\xi - \xi_0)(1-F(\overline{\alpha}_\varsigma))\right), \tag{6}$$

where $\overline{\alpha}_\varsigma$ is parameterized by z_ς. As before, firms choose flexible prices if $\alpha > \overline{\alpha}_\varsigma$. Thus, the aggregate price $\bar{p}_{s,\varsigma}$ is shifted directly by the cost z but also depends on demand through the fraction of firms posting flexible prices, determined by $\overline{\alpha}_\varsigma$.

Because in our model supply shocks do not cause output gaps, the PC will solely capture the relationship between prices and demand pressure. Hence, the appropriate definition is to consider curves parameterized by the supply shock z_ς relating how variations in demand ξ_s affect prices. Demand shocks move along a PC, while supply shocks shift the PC (change the intercept). When the supply shock is z_ς, we have a curve with slope $\bar{\kappa}_\varsigma$ connecting $(\bar{x}_{s,\varsigma}, \bar{p}_{s,\varsigma})$ and $(0, \bar{p}_{0,\varsigma})$, where $\bar{x}_{s,\varsigma}$ denotes the output gap when demand is ξ_s and the supply shock is z_ς. Defining $\overline{\alpha}_\varsigma \equiv \frac{1-z_\varsigma}{\frac{\xi_H}{\xi_L}(1+z_\varsigma)-2z_\varsigma}$, the slopes of the PCs given z_ς are

$$\bar{\kappa}_\varsigma = \xi \left(\frac{1 - F(\overline{\alpha}_\varsigma)}{\int_0^{\overline{\alpha}} \alpha \, \mathrm{d}F(\alpha)} \right).$$

The levels of prices are shifted by the cost z_ς, and the slopes in the states may differ because $\overline{\alpha}$ depends on z_ς. $\overline{\alpha}$ is decreasing in z_ς and therefore the lowest cutoff occurs when the supply shock takes value z_H (i.e., more firms choose sticky prices when the supply shock is low). Define $\underline{\delta}_H \equiv \frac{1-z_H+2z_H\alpha_1}{\alpha_1(1+z_H)}$. If $\delta < \underline{\delta}_H$, then on every island $\alpha < \overline{\alpha}_\varsigma$ for every realization of z_ς. This means that the sticky price is consistent with PBE (there are not enough Insiders to induce firms to choose flexible prices). This means that $\bar{\kappa}_\varsigma = 0$; the PC is perfectly flat with respect to demand shocks, but supply shocks create meaningful inflation. The following is an immediate consequence.

Corollary 1 (Phillips Curve Under Demand and Supply Shocks)**.** *If $\delta < \underline{\delta}_H$, then $\bar{\kappa}_\varsigma = 0$ for all z_ς; the PC is perfectly flat with respect to demand shocks, but supply shocks create inflation ($\lambda = 1$).*

Corollary 1 is crucial. Our model predicts that strategic price stickiness can lead to a completely flat PC with $\kappa = 0$ and also that supply shocks can meaningfully feed through to inflation. Notice that for the sticky price $p_{0,\varsigma}$, the elasticity λ with respect to z_ς is 1. The responsiveness of prices to supply shocks will always exceed the responsiveness to demand shocks. If residual demand shocks are not very volatile, then the PC will be completely flat, with all firms choosing prices that are sticky with respect to demand but flexible with respect to supply. Even if some firms choose flexible prices (i.e., if $\delta > \underline{\delta}_H$), some firms will choose prices that are sticky with respect to demand (and flexible with respect to supply), but all firms will choose prices that are flexible with respect to supply.

Loglinearized PC. To better understand our results, it is useful to derive a loglinearized approximation of the PC to compare it to the standard NK PC. We add time indexes to the equation, consistent with the notation of the infinite horizon model in the appendix (and suppress notation for states s, ς).

Proposition 3 (Loglinearized PC)**.** *If $\delta < \underline{\delta}_H$, we can write inflation in terms of the output gap as*

$$\widehat{\pi}_t = \kappa \widehat{x}_t + \hat{z}_t, \tag{7}$$

where hats denote percentage deviations from steady state, and $\kappa = \frac{1-z_0}{1+z_0}\bar{\kappa}$ is the PC slope defined in terms of percentage deviations around the steady state.

With a flat PC we have $\kappa \approx 0$, and so output gaps do not move inflation while the supply shock $\hat{z}$ does. When the PC is flat, inflation is not at all responsive to demand shocks but inflation remains very responsive to supply shocks in our model. In sum, our model can endogenously produce a perfectly flat PC with inflation that is responsive to structural supply shocks.

Equation (7) is the analog to the standard NK PC, which, dropping inflation expectations (as justified, for instance, in Bilbiie, 2018, 2019), is written $\widehat{\pi}_t = \kappa\widehat{x}_t + \lambda\hat{z}_t$, with $\lambda \approx 0$. Thus, our model produces a PC in which supply shocks are not modified by an additional coefficient that is nearly zero. In fact, in the NK model, a flat PC implies that the intercept is locked-in at the steady state, making it impossible to get inflation out of supply shocks.

3.5 Structural Cost-Push Shocks

Our model, with equation (7), offers a structural theory of cost-push shocks. To see this, it is useful to make the following comparison with the NK PC. Suppose the model features only preference and productivity shocks. In this case, the NK PC is written in terms of output, not an output gap, as

$$\widehat{\pi}_t = \kappa\widehat{y}_t - \kappa\hat{a}_t$$

where, for ease of comparison, we have made abstraction of forward-looking terms, and where $\hat{a}_t$ are shocks to log productivity. A positive productivity shock decreases marginal costs and, therefore, for a given level of output, decreases inflation. In this setting, the output gap is $\widehat{x}_t = \widehat{y}_t - \hat{a}_t$, which implies the familiar equation $\widehat{\pi}_t = \kappa\widehat{x}_t$. This algebra is trivial but provides two important implications. First, since κ is very small, productivity shocks are unable to provide meaningful inflation in the NK model. Second, in the NK model, a negative productivity shock creates inflation by producing a *positive* output gap, even as output falls. But matching the data requires a cost-push shock that creates inflation with a negative output gap. Accordingly, the NK model needs to invoke an ad-hoc shock $\widehat{\nu}_t$ (e.g. a shock to firms' markups):

$$\widehat{\pi}_t = \kappa\widehat{x}_t + \widehat{\nu}_t,$$

to generate inflation together with a negative output gap. The shock $\widehat{v}_t$ creates inflation, and the tightening response of the CB creates a negative output gap. Productivity shocks do not produce that outcome.

This is not the case in our model. Under productivity shocks $\hat{z}_t = -\hat{a}_t$. Equation (7) allows for inflation dependence on both output gaps *and* on cost-push shocks, which originate in movements in productivity. In our model, a negative productivity shock increases prices as marginal costs increase. As we show in the next section, a negative output gap follows if the CB raises rates to offset the rise in inflation. This is precisely the same equilibrium response in the NK model that generate a negative output gap in response to a cost-push shock. Thus, a productivity shock in our model does create inflation with a negative output gap in equilibrium.

This discussion emphasizes that shock-dependence allows for supply shocks, modelled simply as shocks to productivity, to enter (structurally) the PC as cost-push shocks. The reason is that shock-dependence tightly links output gap fluctuations to demand fluctuations. A negative productivity shock shock acts like a cost-push shock in our setup: inflation goes up (prices are flexible), and output goes down. (In the NK model, with very sticky prices, prices do not move, and output does not move if the shock is temporary, either.)

This point echoes the early critique by Chari, Kehoe, and McGrattan (2009). Their critique qualifies shocks to markups as "non-structural", since they do not have an evident counterpart in reality. It also explains our focus, in Section 2, on "inefficient" (markup) shocks within the NK model, and reconciles it with our usage of "efficient" (productivity) shocks in the current section.

3.6 Discussion: Relation to Other Frictions

In closing the section, the following remarks are in order. First, shock dependence is a unique and novel feature of our model. Conventional models have the common feature of *not* distinguishing between demand and supply shocks when predicting price adjustment, which, as emphasized earlier, provides a challenge for quantitatively matching aggregate facts. We develop this theme further in Appendix D, where we show qualitatively that, within this model setup, quadratic costs, Calvo-Taylor, and fixed costs are all unable to obtain for cost-push inflation when the PC is entirely flat. Our discussion in the appendix is somewhat simplistic and stylized; a full quantitative analysis in the context of other models is out of the scope of our contribution. However, the intuition we propose is general, and applies to all that models that treat demand and supply symmetrically.

Second, besides the analysis in the appendix, a specific discussion of the connection to state-dependent models is warranted. In state-dependent models, larger shocks imply more flexible prices, a point recognized at least since Ball, Mankiw, and Romer (1988). State-dependence, and the implied non-linearities that it offers, constitutes a complementary and promising channel to understand the unprecedented global rise in inflation, where economies around the world have been buffeted by shocks of immense magnitude. However, there, whether the shock is a demand or a supply shock is irrelevant. By way of implication, non-linearities in price adjustment that depend on the size of the shock also imply sizeable adjustment of prices to large demand shocks. A priori, logic suggests that state-dependent models cannot simultaneously predict a flat Phillips curve for large demand shocks (such as the Great Recession) and lots of inflation from supply shocks. In our model, the *nature* of the shock, regardless of whether the shock is large or small, is what determines the degree of adjustment. Therefore, prices are completely flexible with respect supply shocks, and at the same time, completely rigid with respect to demand shocks.

4 Quantitative Implications and Monetary Policy

We now provide a calibrated model to determine the quantitative significance of our theoretical results. We have shown that it is theoretically possible to have a flat PC in this setting and for supply shocks to cause inflation. We now show that our proposed mechanism produces an empirically realistic PC given a reasonable calibration. We then use our calibrated model to consider the aggregate consequences of supply shocks when monetary policy responds to supply shocks.

4.1 Setup and Calibration

We modify the setup slightly in order to let the data discipline the degree of heterogeneity. First, 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. Second, we explicitly model the behavior of the central bank using a Taylor rule that determines interest rates in equilibrium. Finally, we calibrate the fraction of Insiders to match the estimated slope of the PC, given particular Taylor coefficients disciplined by the data.

Firm Heterogeneity. 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 $z \leq 1$. Let the idiosyncratic component of marginal costs be distributed according to G, where G is a distribution with the usual properties. Because α is constant across islands, one can show that there is a cutoff for sticky or flexible prices now depends on the marginal cost (see Lemma 4 in Appendix B).

Monetary Policy. We suppose the CB responds to the aggregate state s by setting the nominal interest rate i_s according to a standard Taylor rule, which we write

$$i_t^{Taylor} = i_0 + \phi_\pi \widehat{\pi}_t + \phi_x \widehat{x}_t, \tag{8}$$

where i_0 is a base interest rate, ϕ_π and ϕ_x are the Taylor coefficients. 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. In the model, the effects of monetary policy are akin to the effect of demand shocks, leading endogenous to imperfect price adjustment to policy.[14]

In practice, monetary policy rules typically respond to output gaps and to inflation, rather than to a directly observable shock. Therefore, because policy makers do not observe shocks directly, the policy rule responds indirectly through observable macro variables. The Taylor coefficients ϕ_π and ϕ_x will endogenously determine how much the CB changes i_t in response to the shock.

Solving for equilibrium requires solving a fixed-point problem. Note that the interest rate i_t determines ξ, which affects average prices and output—which are themselves the inputs in the Taylor Rule. All else being equal, when the Taylor coefficients are large, the CB will respond aggressively to variations in output and prices (offsetting the demand shock), which will endogenously lead to smaller fluctuations in output and demand. In this way, the slope of the PC is determined by the Taylor rule (see L'Huillier, Phelan, and Zame, 2022).

[14]We suppose that Outsiders do not know the stance of monetary policy but they do know the equilibrium distribution of possible interest rates. In this way, consistent with the model in Section 3, Outsiders do not know ξ even as the CB responds to the supply shock.

Calibration. We let the data discipline the degree of heterogeneity in the model, thus calibrating the slope of the PC in equilibrium. We target average markups to be 12.5%, as is standard in the literature. With an average price of $p = \frac{\xi_0(1+z)}{2}$, we set $p/z = 1.125$ and $\xi_0 = 1$, implying an average marginal cost of 0.8. We calibrate the idiosyncratic 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 of the distribution of productivities is set so that the average marginal cost equals 0.8. The entire distribution of marginal costs shifts with the aggregate supply shock ς by 1% (up or down). The household time preference (natural rate) is set to $r_0 = 4\%$, and the discount factor shock (demand shock) is 1%. Therefore, $\theta_L = \frac{1.01}{1+r_0}$ and $\theta_H = \frac{0.99}{1+r_0}$.

We choose α to target the slope of the PC in equilibrium for the given Taylor coefficients. We use estimates from Del Negro, Lenza, Primiceri, and Tambalotti (2020) who use 1990 as the break in the sample. Their estimates for the relevant Taylor coefficients post-1990 are $\phi_\pi = 1.5$ and $\phi_x = 0.22$. The posterior mean, median, and mode for κ post-1990 are 0.00151, 0.00140, and 0.00196. These estimates are similar to what is found in Hazell, Herreño, Nakamura, and Steinsson (2022). With $\alpha = 0.88$, equilibrium features an equilibrium PC slope of $\kappa = 0.0017$. In other words, the calibrated model can match the slope of the PC given the estimated the Taylor rule parameters.

4.2 Results

A calibrated version of our model can produce a flat PC consistent with empirical estimates. Equilibrium features endogenous monetary policy responses and firms with heterogeneous productivities disciplined by empirical estimates. We consider the aggregate dynamics of a 1% structural demand shock and a 1% structural supply shock (an increase in marginal costs of 1%), each decaying at a rate of 0.9. We shut down learning dynamics so that $\alpha = 0.88$ is constant across time. (The simulation is obtained by simply repeating the short run equilibrium presented above.)

In this model, demand shocks produce virtually no change in average prices. In contrast, a 1% aggregate productivity shock, which would change marginal costs by 1%, changes aggregate prices by 0.45% on impact. Thus, empirically plausible supply shocks can produce meaningful inflation even while demand shocks produce a very flat PC.

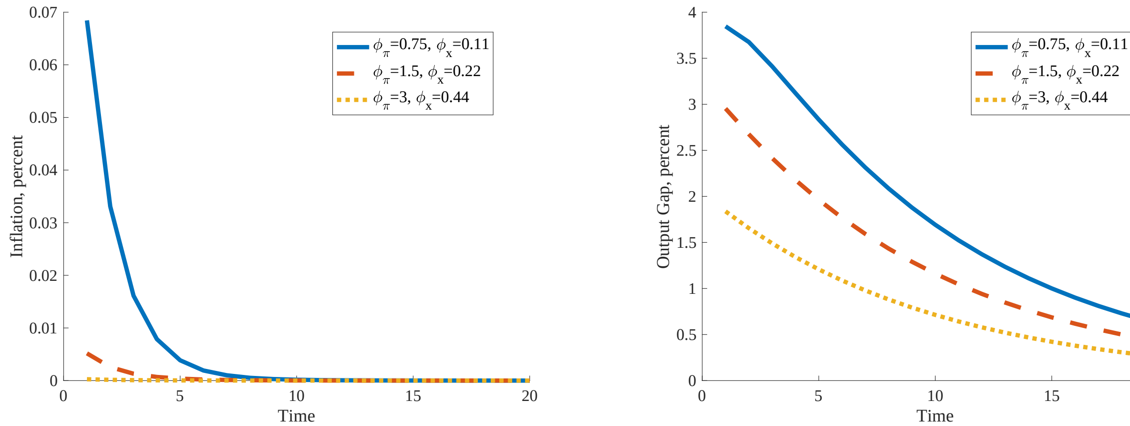

(a) Inflation

(b) Output Gap

Figure 2: Aggregate consequences of a demand shock for inflation and output, varying the monetary policy rule. Red: Baseline Taylor rule; Blue: Dovish MP with Taylor coefficients times 1/2; Yellow: Hawkish MP with Taylor coefficients times 2.

We first consider a decaying demand shock. Figure 2 plots inflation and the output gap given a 1% demand shock. We also consider variation in monetary policy, by considering more and less hawkish Taylor rules. A dovish monetary policy (MP) response follows a Taylor rule with coefficients multiplied by 1/2, and a hawkish MP response follows a Taylor rule with coefficients multiplied by 2. The demand shock generates a positive output gap and negligible inflation (note the scale on the y-axis). In response to the demand shock, the CB raises rates by almost 70 basis points (bps) on impact with the baseline Taylor rule, decaying to about 10 bps (not shown). In this model, the output gap is a “sufficient statistic” for welfare losses (whether the gap is positive or negative). Because the demand shock leads to an output gap (the CB does not completely offset the shock with MP), the shock creates welfare losses.

Figure 3 instead considers the aggregate consequences of a structural supply shock. In this case, we vary the aggressiveness of the central bank in responding to inflation by varying ϕ_π alone while fixing ϕ_x at the baseline. In the hawkish case, the central bank has twice the response to inflation without changing its weight on the output gap. In the dovish case, we set $\phi_\pi = 0$ and in equilibrium there is no monetary policy response to a supply shock (interest rates are constant).

The theoretical result of our paper is that a supply shock leads to a flexible change in prices with respect to the cost shock, but prices can remain sticky with respect to changes in demand, which

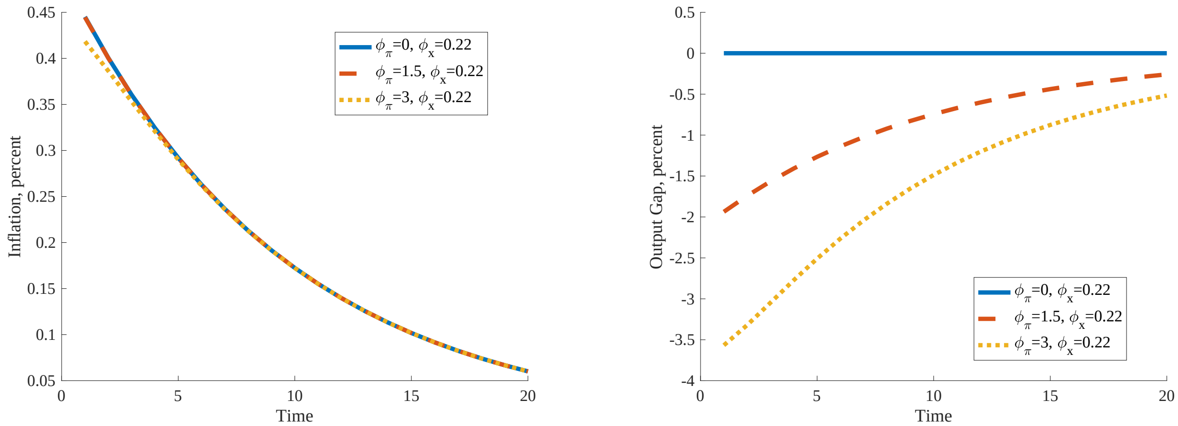

(a) Inflation

(b) Output Gap

Figure 3: Aggregate consequences of a supply shock for inflation and output, varying the hawkishness of the monetary policy rule. Red: Baseline Taylor rule; Blue: Dovish MP, $\phi_\pi = 0$; Yellow: Hawkish MP, $\phi_\pi = 3$.

in this case would be variations in monetary policy. Accordingly, a 1% structural supply shock leads to 0.45% inflation on impact (orders of magnitude larger than the effects from a demand shock) and meaningful negative output gaps in response to the increase in interest rates. The negative output gap leads to welfare losses. In this case, more aggressive monetary policy is not welfare-improving. Crucially, the inflation outcomes in all three scenarios are virtually identical: the central bank response has virtually no effect on inflation. Importantly, an aggressive monetary policy is not enough to bring down inflation.

When the central bank does not respond at all to the supply shock, there is no output gap (flexible-price equilibrium), and therefore there are no welfare losses. Notice the absence of price dispersion in this model, which is key for this result. Moreover, even a very hawkish response that puts twice the baseline weight on inflation leads to almost no change in inflation. An extremely hawkish response, tripling of the Taylor coefficient to $\phi_\pi = 4.5$, would raise rates in equilibrium by 60 bps in response to a 1% supply shock and would still yield 0.32% inflation on impact while creating a negative output gap of almost 4%. Our results suggest that a central bank attempting to bring down inflation in response to a supply shock faces a daunting task. Because the PC is very flat, a very aggressive response in interest rates is likely to have a large negative effect on output without a significant effect on inflation.

In our simple benchmark model, our theoretical results suggest that the increase in prices in response to a supply shock would be efficient; the central bank should not respond by raising rates. In reality, there are likely to be other frictions and rigidities in the economy so that inflation may be costly. We have left out, for example, the possibility of embedded inflation expectations responding in adverse ways. Our analysis nonetheless highlights the policy challenges in responding to inflation in the event that supply shocks lead to efficient inflation. To the extent that prices are less rigid in response to supply shocks, the welfare losses due to nominal rigidities would necessarily be lower. The next section provides empirical evidence that, indeed, aggregate prices are more flexible with response to cost shocks, suggesting that the welfare considerations of supply shocks are lower than those of demand shocks.

5 Empirical Evidence

In this section, we provide U.S. time-series evidence supporting the view that cost-push shocks lead to stronger short-run inflation responses than demand shocks. Specifically, we identify both of these shocks using a state-of-the-art procedure. For a similar effect on U.S. industrial production (IP), we show that whereas demand shocks lead to a relatively small inflation response in the short run, cost-push shocks deliver a response that is about 2.5 times larger over two years and more than 5 times larger over one year.

Our empirical approach is fairly off-the-shelf. The simple time-series exercise we present essentially collects and combines findings from recent studies. Both types of shocks are identified using external instruments. In order to identify the effect of demand shocks on inflation and output, we consider well-identified monetary shocks based on the instrument proposed by Gertler and Karadi (2015). We follow their procedure closely by running a VAR on log IP, the log consumer price index (CPI), the one-year government bond rate (as the policy indicator), and a credit spread, and by using the three-month-ahead funds rate future surprise to identify monetary policy shocks.

In order to identify the effect of cost-push shocks on inflation and output, we consider well-identified oil news shocks based on Känzig (2021). Here, we also follow his procedure closely by running a VAR on the real price of oil, world oil production, world oil inventories, world IP, U.S. industrial production, and the log CPI, and by using his series of high-frequency surprises around OPEC announcements to identify oil shocks.

For both exercises, the data are monthly and the sample spans 1983:4 through 2017:12. Both VARs have 12 lags. Having identified the shocks, we compute the impulse response functions of IP and inflation. As expected, following a monetary shock, IP and inflation co-move, but they do not following a cost-push shock. For both types of shocks, we consider a shock that raises inflation (i.e., an expansionary monetary shock and a contractionary cost-push shock). We set the size of an oil news shock to one standard deviation and compute the responses of IP and CPI. We then scale the size of a monetary shock as follows: we compute the IP response after the cost-push shock at a horizon of 24 months, and then we re-size the monetary shock to deliver the same IP response at the same horizon (in absolute value).

Figure 4 presents the results. It plots the impulse response functions (IRFs) at the point estimate, and the corresponding 68% error bands. Looking at the monetary shock (the dashed, red line), we see that IP rises gradually and reaches a 0.60% increase in 24 months. The CPI raises by roughly 0.10% on impact of the shock, and then stays roughly constant over the horizon considered. For the cost-push shock (the solid, blue line), we estimate a gradual decline in IP, with a fairly rapid rise in the CPI that peaks at 0.30% at 12 months. The inflation response drops slightly thereafter, settling at 0.25% after 24 months.

Overall, for a similar effect on IP, we note that the response of inflation in the case of the cost-push is roughly 2.5 times larger than in the case of the monetary shock over 24 months and more than 5 times larger over 12 months. Moreover, the response in the case of the cost-push shock is statistically different from zero at all horizons, whereas the response to the monetary shock is actually not different from zero for the majority of the response.

6 Conclusion

Inflation data point towards a very flat Phillips curve, and also to meaningful inflation caused by supply disturbances. These two features suggest that a model with shock-dependent friction is a promising candidate for understanding inflation dynamics.

Motivated by this idea, we provide a microfoundation of price stickiness based on strategic behavior of informed firms that can simultaneously produce both a very flat PC and also inflation in response to supply shocks. When demand shocks hit, firms strategically choose sticky prices

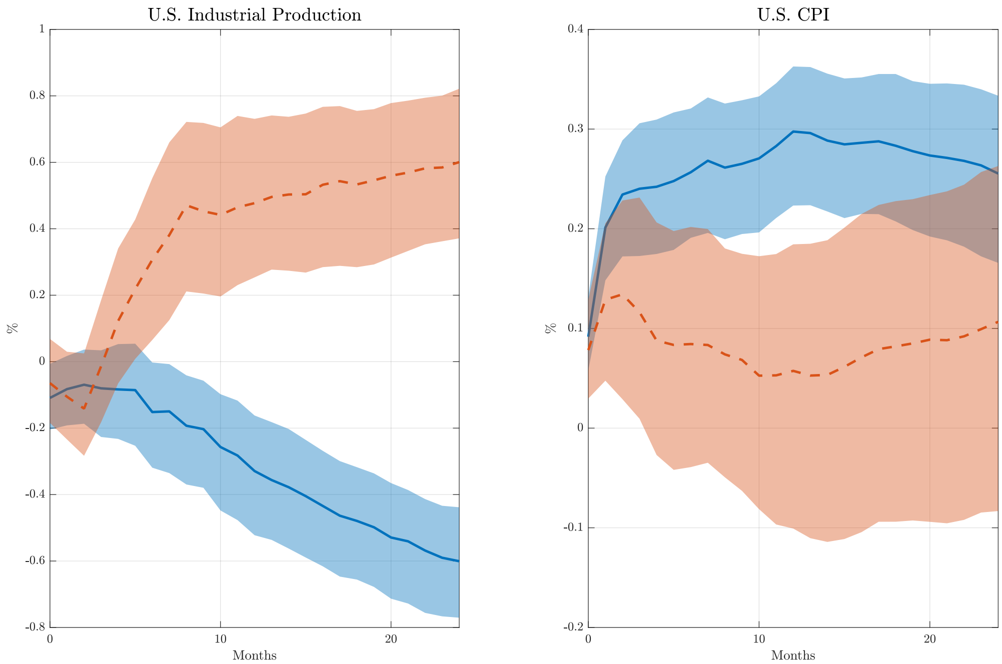

Figure 4: Estimated responses of U.S. Industrial Production and the Consumer Price Index to a monetary shock (the dashed, red line), and to a cost-push shock (the solid, blue line). The cost-push shock is re-sized to deliver the an IP response of the same size at 24 months.

that do not fluctuate. In contrast, firms have no strategic incentive to choose prices that do not fluctuate in response to marginal costs, and therefore prices are flexible with respect to changes in supply.

Our model is able to reconcile the facts that the PC. It does not need to resort to implausibly large structural shocks to firms' pricing decisions. Prices can be completely flexible with respect to supply shocks and yet remain rigid with respect to demand shocks, including changes in monetary policy. A calibrated version of our model shows that a 1% structural cost shock could increase inflation on impact by 0.45% (orders of magnitude larger than a standard New Keynesian model) and suggests that central banks might not want to respond to supply shocks.

More research is required on the broader consequences of shock-dependent rigidities. Our baseline model does not incorporate any other frictions aside from nominal rigidities at the level of firms' pricing decisions. Future work ought to consider the role of wage rigidities, which are unlikely to share the same microfoundation as we have proposed at the level of product prices.

References

ACHARYA, S. (2017): "Costly Information, Planning Complementarities, and the Phillips Curve," *Journal of Money, Credit and Banking*, 49, 823–850.

AFROUZI, H. AND C. YANG (2021): "Dynamic rational inattention and the Phillips curve," Tech. rep., Columbia University.

ASCARI, G. (2004): "Staggered prices and trend inflation: some nuisances," *Review of Economic Dynamics*, 7, 642–667.

ASCARI, G., P. BONOMOLO, AND Q. HAQUE (2022): "The Long-Run Phillips Curve is... a Curve," *Mimeo*.

ASCARI, G. AND L. FOSSO (2021): "The Inflation Rate Disconnect Puzzle: On the International Component of Trend Inflation and the Flattening of the Phillips Curve," *Mimeo*.

ASCARI, G. AND T. HABER (2022): "Non-linearities, state-dependent prices and the transmission mechanism of monetary policy," *The Economic Journal*, 132, 37–57.

BALL, L., D. LEIGH, AND P. MISHRA (2022): "Understanding US inflation during the COVID era," *NBER Working Paper No. 30613*.

BALL, L., N. G. MANKIW, AND R. REIS (2005): "Monetary policy for inattentive economies," *Journal of Monetary Economics*, 52, 703–725.

BALL, L., N. G. MANKIW, AND D. ROMER (1988): "The New Keynesian Economics and the Output-Inflation Trade-Off," *Brookings Papers on Economic Activity*, 1988, 1–82.

BENIGNO, P. AND L. A. RICCI (2011): "The Inflation-Output Trade-Off with Downward Wage Rigidities," *American Economic Review*, 101, 1436–66.

BERNSTEIN, J. AND R. KAMDAR (2022): "Rationally Inattentive Monetary Policy," *Review of Economic Dynamics*, 48, 265–296.

BILBIIE, F. O. (2018): "Monetary policy and heterogeneity: An analytical framework," *CEPR Discussion Paper No. DP12601*.

——— (2019): "Optimal forward guidance," *American Economic Journal: Macroeconomics*, 11, 310–345.

BILS, M., P. J. KLENOW, AND B. A. MALIN (2012): "Reset price inflation and the impact of monetary policy shocks," *American Economic Review*, 102, 2798–2825.

BLANCO, A., C. BOAR, C. JONES, AND V. MIDRIGAN (2022): "Nonlinear Inflation Dynamics in Menu Cost Economies," *Mimeo*.

BLINDER, A. (1991): "Why Are Prices Sticky? Preliminary Results from an Interview Study," *American Economic Review*, 81, 89–96.

BLINDER, A., E. CANETTI, D. LEBOW, AND J. RUDD (1998): *Asking About Prices: A New Approach to Understanding Price Stickiness*, New York, Russell Sage Foundation.

BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2018): "Really uncertain business cycles," *Econometrica*, 86, 1031–1065.

BUNN, P., L. S. ANAYI, N. BLOOM, P. MIZEN, G. THWAITES, AND I. YOTZOV (2022): "Firming up price inflation," *Mimeo*.

CABALLERO, R. J. AND E. M. ENGEL (2007): "Price stickiness in Ss models: New interpretations of old results," *Journal of monetary economics*, 54, 100–121.

CAVALLO, A. (2018): "More Amazon effects: online competition and pricing behaviors," *NBER Working Paper No. 25138*.

CHARI, V. V., P. J. KEHOE, AND E. R. MCGRATTAN (2009): "New Keynesian Models: Not Yet Useful for Policy Analysis," *American Economic Journal: Macroeconomics*, 1, 242–66.

COIBION, O., Y. GORODNICHENKO, AND R. KAMDAR (2018): "The formation of expectations, inflation, and the phillips curve," *Journal of Economic Literature*, 56, 1447–91.

DEL NEGRO, M., M. LENZA, G. E. PRIMICERI, AND A. TAMBALOTTI (2020): "What's up with the Phillips curve?" *Brookings Papers on Economic Activity*, 301–357.

DI GIOVANNI, J., Ş. KALEMLI-ÖZCAN, A. SILVA, AND M. A. YILDIRIM (2022): "Global Supply Chain Pressures, International Trade, and Inflation," *NBER Working Paper No. 30240*.

DIAMOND, D. W. AND P. H. DYBVIG (1983): "Bank Runs, Deposit Insurance, and Liquidity," *The Journal of Political Economy*, 91, pp. 401–419.

DUPRAZ, S. (2017): "A Kinked-Demand Theory of Price Rigidity," *Mimeo*.

EYSTER, E., K. MADARASZ, AND P. MICHAILLAT (2021): "Pricing Under Fairness Concerns," *Econometrica*, 19, 1853–1898.

FITZGERALD, T., C. JONES, M. KULISH, AND J. P. NICOLINI (2020): "Is There a Stable Relationship between Unemployment and Future Inflation?" *Federal Reserve Bank of Minneapolis, Staff Report 614*.

FORBES, K. J. (2019): "Inflation dynamics: dead, dormant, or determined abroad?" *Brookings Papers on Economic Activity*, 257–338.

FURLANETTO, F., D. BERGHOLT, AND E. VACCARO-GRANGE (2023): "Did Monetary Policy Kill the Phillips Curve? Some Simple Arithmetics," *Mimeo*.

GABALLO, G. AND L. PACIELLO (2021): "Spending Allocation under Nominal Uncertainty: A Model of Effective Price Rigidity," *CEPR Discussion Paper No. DP16101*.

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

GERTLER, M. AND P. KARADI (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

GILCHRIST, S., R. SCHOENLE, J. SIM, AND E. ZAKRAJŠEK (2017): "Inflation dynamics during the financial crisis," *American Economic Review*, 107, 785–823.

GREENWALD, B. AND J. E. STIGLITZ (1989): "Toward a theory of rigidities," *The American Economic Review*, 79, 364–369.

GUTIÉRREZ-DAZA, Á. (2022): "Business Cycles when Consumers Learn by Shopping," *Mimeo*.

HALL, R. L. AND C. J. HITCH (1939): “Price theory and business behaviour,” *Oxford economic papers*, 12–45.

HARDING, M., J. LINDÉ, AND M. TRABANDT (2022): “Resolving the missing deflation puzzle,” *Journal of Monetary Economics*, 126, 15–34.

——— (2023): “Understanding Post-Covid Inflation Dynamics,” *Journal of Monetary Economics, forthcoming*.

HAZELL, J., J. HERREÑO, E. NAKAMURA, AND J. STEINSSON (2022): “The Slope of the Phillips Curve: Evidence from U.S. States,” *The Quarterly Journal of Economics*, 137, 1299–1344.

ILUT, C., R. VALCHEV, AND N. VINCENT (2020): “Paralyzed by Fear: Rigid and Discrete Pricing under Demand Uncertainty,” *Econometrica*, 88, pp. 1899–1938.

JORGENSEN, P. L. AND K. J. LANSING (2019): “Anchored inflation expectations and the flatter Phillips curve,” *Federal Reserve Bank of San Francisco Working Paper No. 2019-27*.

KAHNEMAN, D., J. L. KNETSCH, AND R. H. THALER (1986): “Fairness as a constraint on profit seeking: Entitlements in the market,” *American Economic Review*, 76, 728–741.

KÄNZIG, D. R. (2021): “The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements,” *American Economic Review*, 111, 1092–1125.

KARADI, P., R. SCHOENLE, AND J. WURSTEN (2022): “Price Selection in the Microdata,” *Mimeo*.

KILEY, M. T. (2000): “Endogenous price stickiness and business cycle persistence,” *Journal of Money, Credit and Banking*, 28–53.

KOCHERLAKOTA, N. R. (2022): “The Unavoidability of Low Inflation–Low Output Traps,” *The Review of Economic Studies*, 89, 3410–3435.

LAGOS, R. AND R. WRIGHT (2005): “A Unified Framework for Monetary Theory and Policy Analysis,” *Journal of Political Economy*, 113, 463–484.

LEVIN, A. AND T. YUN (2007): "Reconsidering the natural rate hypothesis in a New Keynesian framework," *Journal of Monetary Economics*, 54, 1344–1365.

L'HUILLIER, J.-P. (2020): "Consumer Imperfect Information and Endogenous Price Rigidities," *American Economic Journal: Macroeconomics*, 12, 94–123.

L'HUILLIER, J.-P., G. PHELAN, AND W. R. ZAME (2022): "Flattening of the Phillips Curve and the Mandate of the Central Bank," *Mimeo*.

L'HUILLIER, J.-P. AND R. SCHOENLE (2023): "Raising the Inflation Target: How Much Extra Room Does It Really Give?" *Mimeo*.

L'HUILLIER, J.-P. AND W. R. ZAME (2022): "Optimally sticky prices: Foundations," *Journal of Economic Dynamics and Control*, 104397.

LUCAS, R. E. AND N. L. STOKEY (1987): "Money and interest in a cash-in-advance economy," *Econometrica*, 55, 491–513.

MANKIW, N. G. AND R. REIS (2002): "Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

MCLEAY, M. AND S. TENREYRO (2020): "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual*, 34, 199–255.

NAKAMURA, E. AND J. STEINSSON (2011): "Price Setting in a Forward-Looking Customer Markets," *Journal of Monetary Economics*, 58, 220–233.

OKUN, A. M. (1981): *Prices and Quantities: A Macroeconomic Analysis*, Washington, DC: Brookings Institution.

PETROSKY-NADEAU, N. AND B. BUNDICK (2021): "From Deviations to Shortfalls: The Effects of the FOMC's New Employment Objective," *Federal Reserve Bank of Kansas City Working Paper No. 21-04*.

ROTEMBERG, J. J. (1982): "Monopolistic price adjustment and aggregate output," *Review of Economic Studies*, 49, 517–531.

——— (2005): "Customer anger at price increases, changes in the frequency of price changes and monetary policy," *Journal of Monetary Economics*, 52, 829–852.

——— (2011): "Fair pricing," *Journal of the European Economic Association*, 9, 952–981.

SHAPIRO, A. H. (2022): "How Much do Supply and Demand Drive Inflation?" *FRBSF Economic Letter*, 15.

WERNING, I. (2022): "Expectations and the Rate of Inflation," *NBER Working Paper No. 30260*.

Appendices

A Main Proofs

Proof of Lemma 1 . ONLY IF: To find the cutoff $\overline{\alpha}$ we need to confirm that in the low state the firm would rather charge the low flexible price than the high flexible price (which would fool the uninformed agents). If the flexible price $\{p_s\}$ is consistent with some PBE, then 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 $p_H = \xi_H(1+z)/2$ rather than $p_L = \xi_L(1+z)/2$ if and only if:

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

Simplifying:

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

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

$$(1-z)(1-z) \geq (\delta(1+z)-2z)(\alpha(2-(1+z)\delta)+(1-\alpha)(1-z))$$
$$(1-z)(1-z) \geq (\delta(1+z)-2z)((1-z)+\alpha(1+z)(1-\delta)).$$

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

$$\alpha(1+z)(\delta-1)(\delta(1+z)-2z) \geq (\delta(1+z)-2z)(1-z)-(1-z)(1-z)$$
$$\alpha \geq \frac{1-z}{\delta+\delta k-2z} \equiv \overline{\alpha},$$

which is the desired result. Note that if $z = 0$, then we get $\overline{\alpha} = 1/\delta$.

IF: Given that $\alpha \geq \overline{\alpha}$, 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 $p \neq p_L, p_H$.

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 $1 - p\frac{1}{\xi_L}$ and firm profit will be $(p - z\xi_L)[1 - p\frac{1}{\xi_L}]$. By definition, this quantity is maximized when $p = 1/2\frac{(1+z)}{\xi_L}$ and the maximum profit will be $1/4\frac{(1-z)^2}{\xi_L}$. However this is the profit when the firm offers p_L 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

$$\begin{aligned}\frac{\xi_H(1-z)^2}{4} &\geq (p - z\xi_H)\Big(\alpha[1 - p\frac{1}{\xi_H}] + (1-\alpha)[1 - p\frac{1}{\xi_L}] \\ &= \alpha(p - z\xi_H)[1 - p\frac{1}{\xi_H}] + (1-\alpha)(p - z\xi_H)[1 - p\frac{1}{\xi_L}] \qquad (9)\end{aligned}$$

By definition, $(p - z\xi_H)[1 - p\frac{1}{\xi_H}]$ would be maximized by setting $p = p_H$ and $(p - z\xi_H)[1 -$

$p\frac{1}{\xi_L}]$ would be maximized by setting $p = p_L$ so we must certainly have

$$\alpha(p - z\xi_H)[1 - p\frac{1}{\xi_H}] \leq \alpha\left(\frac{\xi_H(1-z)^2}{4}\right) \quad (10)$$

$$(1-\alpha)(p - z\xi_H)[1 - p\frac{1}{\xi_L}] \leq (1-\alpha)\left(\frac{\xi_L(1-z)^2}{4}\right) \quad (11)$$

The result follows by adding the inequalities (10) and (11) together with $\xi_H > \xi_L$ so we have verified **(e)** . Having verified **(a)** , **(b)** , and **(e)** , the proof is complete. □

Proof of Lemma 2 . Assume the sticky price $p_0 = \xi_0(1+z)/2$ 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 $1 - p\frac{1}{\xi_L}$. 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 shock. The profit of the firm will be:

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

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 shock will be $E_o(\frac{1}{\xi_s}) = \frac{1}{\xi_L}$. Hence if the firm offers $p \neq p_0$ we must have

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

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 $p = p_L$. We conclude that

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

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

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

The LHS is decreasing in α; we must set α sufficiently low. Note that if $\alpha = 0$ then the LHS equals $(\xi_0(1+z) - 2z\xi_L)(1-z)$ and $\xi_0(1+z) - 2z\xi_L = \xi_0 + (\xi_0 - \xi_L)z - z\xi_L > \xi_L(1-z)$ since $\xi_0 > \xi_L$. Thus, there exists a threshold $\overline{\alpha}_0 > 0$ such that this inequality holds if $\alpha \leq \overline{\alpha}_0$. Let $\delta_{0L} \equiv \frac{\xi_0}{\xi_L}$. Then

$$\begin{aligned}
(\delta_{0L}(1+z) - 2z)\left(\alpha\left[2 - \delta_{0L}(1+z)\right] + (1-\alpha)(1-z)\right) &\geq (1-z)^2 \\
(\delta_{0L}(1+z) - 2z)\left(1 - z + \alpha\left[2 - \delta_{0L}(1+z) - (1-z)\right]\right) &\geq (1-z)^2 \\
(1-z) - \alpha(\delta_{0L} - 1)(1+z) &\geq \frac{(1-z)^2}{\delta_{0L}(1+z) - 2z}
\end{aligned}$$

Recall that $\delta_{0L} > 1$ since $\xi_L < \xi_0$, and therefore we have

$$\begin{aligned}
(1-z) - \frac{(1-z)^2}{\delta_{0L}(1+z) - 2z} &\geq \alpha(\delta_{0L} - 1)(1+z) \\
\frac{(1-z)(\delta_{0L} - 1)(1+z)}{\delta_{0L}(1+z) - 2z} &\geq \alpha(\delta_{0L} - 1)(1+z)
\end{aligned}$$

And therefore we can simplify to $\alpha \leq \frac{1-z}{\delta_{0L}(1+z) - 2z} \equiv \overline{\alpha}_0$. Recall that

$$\overline{\alpha} \equiv \frac{1-z}{\delta(1+z) - 2z} < \frac{1-z}{\delta_{0L}(1+z) - 2z}$$

since $\delta \equiv \xi_H/\xi_L > \delta_{0L} \equiv \xi_0/\xi_L$. Thus, if $\alpha < \overline{\alpha}$, it follows that $\alpha < \overline{\alpha}_0$ and therefore p_0 is a PBE.

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 $1 - p\frac{1}{\xi_L}$. Insiders know the true state s and demand $1 - p\frac{1}{\xi_s}$ so the profit of the firm is

$$d_s(p) = (p - z\xi_s)\left(\alpha[1 - p\frac{1}{\xi_s}] + (1-\alpha)[1 - p\frac{1}{\xi_L}]\right) \tag{12}$$

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

$$d_s(p_0) = \left(\frac{\xi_0(1+z)}{2} - z\xi_s\right)\left(\alpha\left[1 - \frac{\xi_0(1+z)}{2}\frac{1}{\xi_s}\right] + \frac{(1-\alpha)(1-z)}{2}\right)$$

The equilibrium condition is that

$$d_s(p_0) \geq d_s(p) \tag{13}$$

when $s = L$ and when $s = H$, under the assumption that $\alpha \leq \overline{\alpha}$. That the inequality (13) is satisfied when the true state $s = L$ follows from the exercise we just did. To see that (13) is satisfied when the true state $s = H$ is more complicated. First note that simplifying the right side of (12) yields

$$d_H(p) = (p - z\xi_H)\left(1 - p\left[\alpha\frac{1}{\xi_H} + (1-\alpha)\frac{1}{\xi_L}\right]\right)$$

Define $\xi_\alpha \equiv \left(\alpha\frac{1}{\xi_H} + (1-\alpha)\frac{1}{\xi_L}\right)^{-1}$ as the α-weighted harmonic mean of ξ_s. Since $\xi_\alpha < \xi_H$,

$$d_H(p) < (p - z\xi_\alpha)\left(1 - p\left[\alpha\frac{1}{\xi_H} + (1-\alpha)\frac{1}{\xi_L}\right]\right)$$

and the RHS is maximized by setting $p = \frac{\xi_\alpha(1+z)}{2}$ and equals $\frac{\xi_\alpha(1-z)^2}{4}$. Thus, it suffices to show that for $\alpha \leq \overline{\alpha}_0$

$$(p_0 - z\xi_H)\left(1 - p_0\left[\alpha\frac{1}{\xi_H} + (1-\alpha)\frac{1}{\xi_0}\right]\right) \geq \frac{\xi_\alpha(1-z)^2}{4}.$$

Note first that this is satisfied for $\alpha = 0$ but not for $\alpha = 1$; in the first case there are no Informed agents, so setting $p = p_0$ is strictly dominant; in the second case there are only Informed agents so the flexible price is optimal. Rearranging we have

$$(p_0 - z\xi_H)\left(1 - p_0\frac{1}{\xi_0} + p_0\alpha\left(\frac{1}{\xi_0} - \frac{1}{\xi_H}\right)\right) \geq \frac{(1-z)^2}{4\left(\frac{1}{\xi_L} - \alpha\left(\frac{1}{\xi_L} - \frac{1}{\xi_H}\right)\right)}.$$

Note that the LHS is increasing linearly in α since $\frac{1}{\xi_0} - \frac{1}{\xi_H} > 0$. The RHS is increasing with α.

Multiplying we have

$$\left(\frac{1}{\xi_L}-\alpha\left(\frac{1}{\xi_L}-\frac{1}{\xi_H}\right)\right)(p_0-z\xi_H)\left(1-p_0\frac{1}{\xi_0}+p_0\alpha\left(\frac{1}{\xi_0}-\frac{1}{\xi_H}\right)\right)\geq\frac{(1-z)^2}{4},$$

which is a quadratic equation in α with a negative coefficient on α^2. Thus, if this holds at $\overline{\alpha}_0$ it holds for all $\alpha \leq \overline{\alpha}_0$. L'Huillier, Phelan, and Zame (2022) verify this condition holds for $z=0$. By continuity it holds for z sufficiently small. □

Proof of Lemma 3 . First, to derive $\bar{\kappa}$, from equations (3) and (4), we can derive more convenient expressions for the average price and demand difference:

$$\bar{p}-\bar{p}_0=\frac{1}{2}(\xi-\xi_0)(1+z)(1-F(\overline{\alpha})), \tag{14}$$

$$\bar{x}_s=\frac{(1+z)}{2}\left(\frac{\xi-\xi_0}{\xi}\right)\int_0^{\overline{\alpha}}\alpha\,\mathrm{d}F(\alpha). \tag{15}$$

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.

Define $\delta\equiv\xi_H/\xi_L$. Note that $\delta\geq 1$ and decreases as $\xi_H\to\xi_L$. Define $\underline{\delta}\equiv\frac{1-z+2z\alpha_1}{\alpha_1(1+z)}$. We show that $\bar{\kappa}=0$ whenever $\delta<\underline{\delta}$ First, $\overline{\alpha}$ is decreasing in δ and equals 1 when $\delta=1$; the fraction of islands posting sticky prices increases when δ is lower, converging to 1. Thus, the numerator decreases, the denominator increases, and δ decreases to 1. Therefore, $\bar{\kappa}$ is increasing in δ; as $\delta\to 1$ $\overline{\alpha}\to 1$ and all islands post sticky prices. If $\delta<\underline{\delta}$, then $\overline{\alpha}>\alpha_1$, which means that $F(\overline{\alpha})=1$, all islands post sticky prices, and $\bar{\kappa}=0$. Second, $\xi\to\xi_0$ as $\delta\to 1$. □

Proof of Proposition 1. The proof is immediate. Note that conditional on the price p, Insiders and Outsiders have the same demand, $x=1-p\xi_0$. Thus, profit maximization means choosing p to maximize $(1-p\xi_0)\left(p-z_\varsigma\right)$, which yields the desired result. □

Online Appendices

B Additional Proofs

B.1 Proofs for Section 3.4, Strategic Model with Supply and Demand Shocks

Proof of Proposition 2. The proof follows quickly from earlier results. First, the firm can choose prices that communicates information about ξ_s, thus achieving full information. Let $p_{s,\varsigma}$ denote the price when demand is ξ_s and supply is z_ς. The full-information prices are given by

$$p_{H,H} = \frac{\xi_H(1+z_H)}{2}, p_{L,H} = \frac{\xi_L(1+z_H)}{2}, p_{H,L} = \frac{\xi_H(1+z_L)}{2}, p_{L,L} = \frac{\xi_L(1+z_L)}{2}.$$

Second, the firm can choose prices that do not communicate information about ξ_s. Since the demand shock is orthogonal to the shock to z_ς, it is easy to show that the following prices maximize profits (as in the previous proof) without communicating information regarding ξ_s:

$$p_{0,H} = \frac{\xi_0(1+z_H)}{2}, \quad p_{0,L} = \frac{\xi_0(1+z_L)}{2}.$$

Except by coincidence, these six prices are all distinct. Hence, in the high cost state (z_H), the firm can choose a price that is sticky with respect to ξ_s by offering $p_{0,H} = \frac{\xi_0(1+z_H)}{2}$, the firm can choose a price that is flexible with respect to ξ_s by offering $p_{s,H} = \frac{\xi_s(1+z_H)}{2}$. The cutoff for choosing the sticky or flexible price is given by the threshold in Lemma 1 for a given z_ς. Recall that $\overline{\alpha}$ is decreasing in z_ς, so $\overline{\alpha}_H < \overline{\alpha}_L$ and therefore $\alpha < \overline{\alpha}$ for both realizations of the supply shock.

In this way, the firm can choose a price that is flexible with respect to both shocks by offering $p_{s,\varsigma} = \frac{\xi_s(1+z_\varsigma)}{2}$ or a price that is flexible with respect to the cost shock only by offering $p_{0,\varsigma} = \frac{\xi_0(1+z_\varsigma)}{2}$, which is sticky with respect to demand. The remaining results follow immediately from Lemma 3.

Conditional on the price p, agents have the demand, $x = 1 - p\mathbb{E}\left[\frac{1}{\xi_s}\right]$. Hence, any incentive for firms to convey or hide information via p can only operate through the expectation on ξ_s.

In the proposed equilibrium, Outsiders see a price $p_{0,\varsigma}$ and conclude it conveys information about the marginal cost alone ($p_{0,\varsigma} = \frac{1+z_\varsigma}{2}\xi_0$); they do not learn the demand shock or update their demands. Given these beliefs, on islands with a small number of Insiders, the firms' problem is equivalent to determining the profit-maximizing price as if the demand shock is ξ_0, which provides the price above. By our assumption that $\delta < \underline{\delta}_H$, no island has a sufficiently high number of Insiders for the firm to want to set a price that responds to the demand shock. On each island, the fraction of Outsider is large enough so that firms choose the price $p_{0,\varsigma} = \frac{1+z_\varsigma}{2}\xi_0$ taking as given that Outsiders will set their demand with $x_{0,\varsigma} = 1 - p_{0,\varsigma}\mathbb{E}\left[\frac{1}{\xi_0}\right]$. Thus, no firm on any island adjusts prices in response to demand but every firm on every island adjusts prices in response to costs. $\square$

Proof of Proposition 3. Since prices are generally proportional to $1+z$, it is useful to define a supply shock as the value $1+z$. Let $\bar{p}_0 = \frac{1+z_0}{2}\xi_0$ denote the steady-state price level. We suppose the economy starts at steady state, and therefore this is the prior price level. Inflation is $\widehat{\pi} \equiv \frac{\bar{p}-\bar{p}_0}{\bar{p}_0}$. It is useful to define the shocks $\widehat{\xi} \equiv \frac{\xi-\xi_0}{\xi_0}$, and $\hat{z} \equiv \frac{\Delta z}{1+z_0}$, and the percentage output gap $\widehat{x} = \frac{\bar{x}-\bar{x}_0}{\bar{x}_0}$.

Rearranging $\bar{\kappa} \equiv \frac{\bar{p}-p_0}{\bar{x}-\bar{x}_0}$, we have $\bar{p} - p_0 = \bar{\kappa}(\bar{x}-\bar{x}_0)$. Note that $z = z_0 + \Delta z$ and therefore

$$p_0 \equiv \frac{1+z}{2}\xi_0 = \frac{1+z_0+\Delta z}{2}\xi_0 = \bar{p}_0 + \frac{\Delta z}{2}\xi_0.$$

Alternatively, $\frac{p_0}{\bar{p}_0} = 1 + \hat{z}$. Therefore $\frac{\bar{p}-\bar{p}_0}{\bar{p}_0} = \frac{\bar{p}-p_0}{\frac{1+z_0}{2}\xi_0} + \hat{z}$. Using $\bar{p} - p_0 = \bar{\kappa}(\bar{x}-\bar{x}_0)$, we have

$$\widehat{\pi} = \frac{\bar{\kappa}}{\frac{1+z_0}{2}\xi_0}(\bar{x}-\bar{x}_0) + \hat{z} = \frac{2\bar{\kappa}}{(1+z_0)\xi_0}(\bar{x}-\bar{x}_0) + \hat{z}.$$

To write this in terms of a percentage output gap $\widehat{x}$,

$$\begin{aligned}
\widehat{\pi} &= \frac{2\bar{\kappa}}{(1+z_0)\xi_0}\bar{x}_0\left(\frac{\bar{x}-\bar{x}_0}{\bar{x}_0}\right) + \hat{z}, \\
&= \frac{2\bar{\kappa}}{(1+z_0)\xi_0}\left(\frac{1-z}{2}\right)\widehat{x} + \hat{z}, \\
&= \bar{\kappa}\frac{1-z_0}{(1+z_0)\xi_0}\widehat{x} - \bar{\kappa}\frac{\hat{z}}{\xi_0}\widehat{x} + \hat{z} = \bar{\kappa}\frac{1-z_0}{(1+z_0)\xi_0}\widehat{x} + \hat{z},
\end{aligned}$$

where the last line follows because the product of two small changes, $\hat{z}\widehat{x} \approx 0$. κ is the PC defined in terms of percentage deviations from the steady state. Let $z = z_0$ so that $p_0 = \bar{p}_0$. Then

$$\kappa \equiv \frac{(\bar{p}-\bar{p}_0)/\bar{p}_0}{(\bar{x}-\bar{x}_0)/\bar{x}_0} = \frac{\bar{x}_0}{\bar{p}_0}\xi\left(\frac{1-F(\overline{\alpha})}{\int_0^{\overline{\alpha}} \alpha \, \mathrm{d}F(\alpha)}\right),$$
$$= \frac{(1-z_0)/2}{(1+z_0)\xi_0/2}\bar{\kappa} = \frac{1-z_0}{1+z_0}\bar{\kappa},$$

where $\xi_0 = 1$. As expected, the PC slope is indeed the PC slope, appropriately defined. Note that with a flat PC we have $\kappa \approx 0$, and so output gaps do not move inflation while the supply shock $\hat{z}$ does.

Furthermore, we can write the PC in terms of structural shocks directly as

$$\widehat{\pi} = \widehat{\xi}(1-F(\overline{\alpha}))+\hat{z}. \tag{16}$$

The aggregate price level given z and ξ is

$$\bar{p} = \int_0^{\overline{\alpha}} \frac{\xi_0(1+z)}{2}\,\mathrm{d}F(\alpha) + \int_{\overline{\alpha}}^1 \frac{\xi(1+z)}{2}\,\mathrm{d}F(\alpha),$$
$$= \left(\frac{1+z}{2}\right)(\xi_0+(\xi-\xi_0)(1-F(\overline{\alpha}))),$$
$$= \left(\frac{1+z}{2}\right)\xi_0\left(1+\widehat{\xi}(1-F(\overline{\alpha}))\right).$$

Dividing by the initial price level $\bar{p}_0$ and substituting $\hat{z}$

$$\frac{\bar{p}}{\bar{p}_0} = \left(\frac{1+z}{1+z_0}\right)\left(1+\widehat{\xi}(1-F(\overline{\alpha}))\right),$$
$$= (1+\hat{z})\left(1+\widehat{\xi}(1-F(\overline{\alpha}))\right),$$
$$= 1+\widehat{\xi}(1-F(\overline{\alpha}))+\hat{z}+\hat{z}\widehat{\xi}(1-F(\overline{\alpha})).$$

Using that the product of hat variables is zero, the last term is annihilated. We therefore have

$$\frac{\bar{p}}{\bar{p}_0}-1 = \widehat{\xi}(1-F(\overline{\alpha}))+\hat{z},$$

and therefore we have $\widehat{\pi} = \widehat{\xi}(1-F(\overline{\alpha}))+\hat{z}$. $\square$

B.2 Proof for Section 4

Lemma 4 (PBE with Flexible Prices and Marginal Costs)**.** *The flexible price* $\{p_s\}$ *is consistent with some PBE if and only if*

$$z \geq \bar{z} \equiv \frac{1-\alpha\xi_H/\xi_L}{\alpha(\xi_H/\xi_L-2)+1}.$$

The *average* prices $\bar{p}_s$ and demands $\bar{y}_s$ in state s can be written

$$\bar{p}_s = \int_0^{\bar{z}} \xi_0 \frac{1+z}{2}\,\mathrm{d}G(z) + \int_{\bar{z}}^1 \xi_s \frac{1+z}{2}\,\mathrm{d}G(z), \tag{17}$$

$$\bar{y}_s = \int_0^{\bar{z}} \left(\alpha\left[1-\frac{1+z}{2}\frac{\xi_0}{\xi_s}\right] + (1-\alpha)\left[\frac{1-z}{2}\right]\right)\mathrm{d}G(z) + \int_{\bar{z}}^1 \frac{1-z}{2}\mathrm{d}G(z). \tag{18}$$

For symmetry, we define the PC comparing the outcomes in the high and low states, and therefore, the slope of the PC is

$$\bar{\kappa} = \frac{\bar{p}_H-\bar{p}_L}{\bar{y}_H-\bar{y}_L} = \left(\frac{\xi_H\xi_L}{\alpha\xi_0}\right)\frac{\int_{\bar{z}}^1(1+z)\,\mathrm{d}G(z)}{\int_0^{\bar{z}}(1+z)\,\mathrm{d}G(z)}. \tag{19}$$

Compared to the baseline model, we aggregate using G instead of F; the model behaves similarly.[15]

Proof of Lemma 4. Incentive-compatibility for the flexible-price equilibrium requires

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

Simplifying:

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

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

[15] As an aside, note that in this setting a supply shock that increases the distribution of z would increase the fraction of flexible-price firms, which would all else being equal steepen the PC. Because this is not the main point of our paper, we do not emphasize this result, but leave it to later work to further investigate this prediction.

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

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

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

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

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

$$\alpha(\delta+\delta k-2z) \geq 1-z, \implies (\alpha(\delta-2)+1)k \geq 1-\alpha\delta.$$

Note that $1+\alpha\delta > 1+\alpha > 2\alpha$. Hence, $\alpha(\delta-2)+1 > 0$, and so we have $z \geq \frac{1-\alpha\delta}{\alpha(\delta-2)+1} = \bar{z}$. $\square$

C Real Rigidities in the NK Model

A possible way to generate a flat short-term PC is using various real rigidities. κ and λ are linked to each other by a proportionality coefficient:

$$\kappa \equiv e \cdot \lambda, \tag{20}$$

where e measures how marginal costs co-move with the output gap: $\widehat{mc}_t = e \cdot \widehat{x}_t$. $\widehat{mc}_t$ denotes deviations in marginal costs (or equivalently, deviations in firms' desired markup). The parameter e is determined by real rigidities.

Many real rigidities show up in λ, while some (like sticky wages) show up in e. Specifically, high degrees of real rigidity in the mapping from the output gap to marginal costs can deliver a low value of the parameter e (a low elasticity of marginal costs to the output gap). This raises the question of whether λ could be large while κ small, resolving, , in the context of the NK model, the issue that we bring up. As previewed, because empirical estimates of both κ and λ (separately) are small, real rigidities do not present a plausible resolution to the puzzle we highlight. In other words, the empirical evidence suggests that the culprit for the observed flat PC is λ, and not e.

In a nutshell, the reason for this is strong evidence that marginal costs move whereas prices remain largely fixed. To expand on this point, recall that the standard NKPC can equivalently be written in linearized form as

$$\widehat{\pi}_t = \beta \mathbb{E}_t[\widehat{\pi}_{t+1}] + \lambda \widehat{mc}_t = \lambda \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t[\widehat{mc}_{t+s}]. \tag{21}$$

The second equality follows by iterating forward. Hence, inflation is caused by changes in firms' marginal costs or desired markups. Importantly, the parameter λ measures how changes in marginal costs translate into changes in prices. In general, λ is a function of the degree of price stickiness, the elasticity of substitution across goods, the capital share in production, and real rigidities caused by, for example, fixed production costs or kinked consumer demand (i.e., Kimball).

Marginal costs can increase due to demand or supply shocks. In response to demand shocks, output gaps can increase marginal costs whenever there is curvature in production or labor supply, leading to an increase in wages or capital costs. In response to supply shocks, marginal costs can increase for a given level of production because of higher input costs or productivity shocks, or firms may have a higher desired price for a given marginal cost. When marginal costs increase due to an output gap, the PC is most often written as in (1).

In the NK model, therefore, to get a flat PC, either e or λ must be very small. The behavior of inflation is entirely determined by the behavior of marginal costs and how much inflation responds to marginal costs. A low co-movement between inflation and output gaps can be the result of low co-movement between marginal costs and output (i.e., a low e) or a low co-movement between prices and marginal costs (i.e., a low λ). To the extent that marginal costs move sufficiently without output gaps, generating a flat PC requires a low responsiveness of prices to marginal costs.

In the textbook NK model, the equilibrium relationship between output gaps and marginal costs is given by $e = \sigma^{-1} + \frac{\varphi^{-1}+\zeta}{1-\zeta}$, where σ is the elasticity of intertemporal substitution (EIS), φ is the Frisch elasticity, and ζ is the curvature of the production function (Galí, 2015). For the plausible values of $\sigma = 1$, $\varphi = 1/5$, and $\zeta = 1/3$, then $e = 9$, implying that λ is an order of magnitude *smaller* than κ in the standard NK model. An alternative textbook benchmark is a model with

GHH preferences and constant returns to scale, in which case $e = \varphi^{-1}$. Micro-level estimates typically find φ to be less than 0.5, whereas macro-level estimates using aggregate variables find φ in the range of 2-4. In either case, a very low κ necessitates a very low λ.

Most importantly, both Hazell et al. (2022) and Del Negro et al. (2020) provide clear empirical evidence that, in the data, marginal costs fluctuate, with a very low pass-through to prices. Specifically, Hazell et al. (2022)'s second IV strategy relies on marginal costs providing enough exogenous variation to identify κ and λ. In particular, the estimation of λ is based on direct measures of marginal costs variation embodied in tradeable prices, which are used as inputs of production for non-tradeables. Similarly, Del Negro et al. (2020) also find considerable levels of aggregate price rigidity, which they attribute to a low λ. Their model explicitly allows for the possibility that e could be estimated to be very small, for example due to sticky wages. However, the estimation does not support the view that a flat PC is due to a weak co-movement of marginal costs with output gaps. Indeed, they find that the co-movement of all real variables and indicators of firms' cost pressures has been remarkably stable since 1964, and therefore the flat PC is most likely caused by a low sensitivity of price changes to marginal costs (see the discussion on pp. 316-7).

Altogether, the evidence points to a flat PC that comes not from a low e but from a low λ.

D Standard Pricing Frictions

This section solves the baseline model from Section 3 with perfect information but with standard pricing frictions. We consider quadratic adjustment costs (Rotemberg, 1982), contract frictions (Calvo or Taylor), and fixed adjustment costs. None of these standard settings produce a flat PC *as well as* significant responsiveness to supply shocks. This is due to the symmetry embedded in these models regarding the reasons for the stickiness. In these settings, κ and λ move together, as is the case in the NK model (for a given e). With standard frictions, a very flat PC requires very large exogenous frictions to change prices. With standard price-setting frictions, variations in inflation require implausibly large structural supply shocks, just as in the standard NK model. The intuition is straightforward: if it is costly to change prices, then it is costly to change prices in response to supply shocks as well as in response to demand shocks.

Quadratic Costs. Let the firm face quadratic costs to adjust prices from p_0 to p, given by $\frac{\varphi}{2}(p-p_0)^2$. Let $p_0 = \frac{1-z}{2}$ be the base price. The firm maximizes

$$(p - z\xi)(1 - p/\xi) - \frac{\varphi}{2}(p-p_0)^2. \tag{22}$$

We can summarize equilibrium as follows:

Lemma 5 (Quadratic Costs)**.** *With quadratic adjustment costs for prices, all firms set a price $p^* = \xi\frac{1+z+\varphi p_0}{2+\varphi\xi}$, and demand at each island is $x = \frac{1-z+\varphi(\xi-p_0)}{2+\varphi\xi}$. The PC slope in this economy is given by $\kappa = \frac{2}{\varphi}$.*

Proof of Lemma 5. Profits are $p - z\xi - p^2/\xi + pz - \frac{\varphi}{2}(p-p_0)^2$. The first-order condition is

$$1 + z - \frac{2}{\xi}p - \varphi(p - p_0) = 0, \tag{23}$$

and solving for p, the optimal price is

$$p^* = \frac{1+z+\varphi p_0}{\frac{2}{\xi}+\varphi} = \xi\frac{1+z+\varphi p_0}{2+\varphi\xi}.$$

and total demand is

$$\bar{y} = 1 - \frac{1+z+\varphi p_0}{2+\varphi\xi} = \frac{1-z+\varphi(\xi-p_0)}{2+\varphi\xi}.$$

The PC is therefore given by

$$\kappa = \frac{\xi_H\frac{1+z+\varphi p_0}{2+\varphi\xi_H} - \xi_L\frac{1+z+\varphi p_0}{2+\varphi\xi_L}}{\left(1 - \frac{1+z+\varphi p_0}{2+\varphi\xi_H}\right) - \left(1 - \frac{1+z+\varphi p_0}{2+\varphi\xi_L}\right)} = \frac{2(\xi_H - \xi_L)}{\varphi\xi_H - \varphi\xi_L},$$

which simplifies finally to $\kappa = \frac{2}{\varphi}$. □

A flat PC corresponds to a very high φ (i.e., a high cost of adjusting prices). Asymptotically, if we want $\kappa \to 0$, we have $\varphi \to \infty$. In this setting, we cannot get inflation in response to cost shocks with a flat PC. We calculate

$$\frac{\partial p^*}{\partial z} = \frac{\xi}{2+\varphi\xi}.$$

Thus, letting $\varphi \to \infty$, we have $\lambda \to 0$. If the PC is flat, a cost shock will not lead to a large change in prices. The intuition is simple: a flat PC corresponds to a high cost of changing prices. It's still costly to change prices, whether responding to demand or supply shocks, and so firms don't change prices much.

Calvo-Taylor. Now suppose a fraction $1-\varepsilon$ of firms cannot change their price. Thus, a fraction ε set their price to p^*, and the rest keep p_0 (i.e., there are no strategical complementarities arising from monopolistic *competition* à la Dixit-Stiglitz). This setting could correspond to probabilistic price changes à la Calvo or staggered price setting à la Taylor. In either case, the fraction of firms that adjust their price (whether probabilistic or pre-determined) is ε. The aggregate price and output are thus

$$\bar{p} = \varepsilon\xi(1+z)/2 + (1-\varepsilon)p_0, \quad \bar{y} = \varepsilon\frac{1-z}{2} + (1-\varepsilon)\left(1 - \frac{p_0}{\xi}\right).$$

Lemma 6 (Calvo-Taylor)**.** *With Calvo-Taylor frictions, the PC slope is* $\kappa = \frac{\varepsilon}{1-\varepsilon}\xi$.

Proof of Lemma 6. The PC slope is given by

$$\kappa = \frac{\varepsilon(\xi - \xi_0)(1+z)/2}{(1-\varepsilon)\left(\frac{p_0}{\xi_0} - \frac{p_0}{\xi}\right)} = \frac{\varepsilon}{1-\varepsilon}\frac{\xi\xi_0(1+z)}{2p_0}.$$

Since $p_0 = \xi_0(1+z)/2$ this simplifies to $\kappa = \frac{\varepsilon}{1-\varepsilon}\xi$. □

A flat PC corresponds to a very low ε (i.e., a low probability of price adjustment). Asymptotically, if we want $\kappa \to 0$, we have $\varepsilon \to 0$. Again, we cannot get inflation in response to cost shocks with a flat PC. We calculate

$$\frac{\partial \bar{p}}{\partial z} = \frac{\varepsilon\xi}{2}.$$

Letting $\varepsilon \to 0$, we have $\lambda \to 0$. If the PC is flat due to Calvo- or Taylor-style frictions, a cost shock will not lead to a large change in prices. The intuition is simple: a flat PC means few firms have the opportunity to change prices at all. If almost no firms can change prices, whether responding to demand shocks or supply shocks, aggregate prices won't change much in response to a cost shock.

Fixed Costs. Suppose firms pay a fixed cost $k_i \in [\underline{k}, \overline{k}]$ to adjust prices, with costs distributed according to CDF F. Adjusting firms choose $p^* = \frac{\xi(1+z)}{2}$, whereas not adjusting yields demand $x = 1 - p_0/\xi$, for a profit $d_0 = (p_0 - z\xi)(1 - p_0/\xi) = p_0(1+z) - z\xi - p_0^2/\xi$. A firm will adjust prices whenever

$$d^* \geq d_0 \implies k_i \leq k^*(\xi) \equiv \xi(1-z)^2/4 - (p_0(1+z) - z\xi - p_0^2/\xi).$$

Thus, the fraction of firms adjusting will be $F(k^*)$.

Lemma 7 (Fixed Costs)**.** *With fixed costs of adjusting prices, the PC slope in this economy is*

$$\kappa = \frac{F(k^*(\xi))}{1 - F(k^*(\xi))} \left(\frac{\xi(\xi - \xi_0)(1+z)}{1 + z - \xi_0} \right). \tag{24}$$

Proof of Lemma 7. Recall that $p_0 = \frac{\xi_0(1+z)}{2}$. Hence

$$\begin{aligned} k^*(\xi) &= \xi(1-z)^2/4 + p_0^2/\xi + z\xi - p_0(1+z), \\ &= \frac{\xi}{4}(1-z)^2 + \frac{\xi_0^2(1+z)^2}{4\xi} + z\xi - \frac{\xi_0(1+z)^2}{2}, \\ &= \frac{1}{4}\left(\xi(1-z)^2 + \frac{\xi_0^2(1+z)^2}{\xi} \right) - \frac{\xi_0(1+z)^2}{2} + z\xi. \end{aligned}$$

and we have

$$\frac{dk^*(\xi)}{d\xi} = \frac{(1-z)^2}{4} - \frac{1}{4}\frac{\xi_0^2(1+z)^2}{\xi^2} + z = \frac{1}{4}\left((1-z)^2 - \frac{\xi_0^2(1+z)^2}{\xi^2} \right) + z.$$

Thus, $k^*(\xi_0) = 0$ and is the global minimum. The average price is

$$\bar{p} = F(k^*(\xi))\frac{\xi(1+z)}{2} + (1 - F(k^*(\xi)))p_0,$$

and the level of output is

$$\bar{y} = F(k^*(\xi))\frac{1-z}{2} + (1 - F(k^*(\xi)))\left(1 - \frac{p_0}{\xi}\right).$$

Hence, we can write κ as

$$\kappa = \frac{F(k^*(\xi))\frac{\xi(1+z)}{2} + (1-F(k^*(\xi)))p_0 - p_0}{F(k^*(\xi))\frac{1-z}{2} + (1-F(k^*(\xi)))\left(1-\frac{\xi_0}{2\xi}\right) - \frac{1-z}{2}}$$
$$= \frac{F(k^*(\xi))\left(\frac{\xi-\xi_0}{2}\right)(1+z)}{(1-F(k^*(\xi)))\left(\frac{1+z-\xi_0}{2\xi}\right)}$$
$$= \frac{F(k^*(\xi))}{1-F(k^*(\xi))}\left(\frac{\xi(\xi-\xi_0)(1+z)}{1+z-\xi_0}\right).$$

□

In this setting, we can get a flat PC from very small demand shocks or very large fixed adjustment costs. It is easy to show that κ is increasing in ξ. The numerator is increasing because $\xi(\xi-\xi_0)$ is, and since k^* is increasing, so is $F(k^*)$. If $k^* < \underline{k}$, then no firms will adjust prices and the PC will be perfectly flat. The denominator is decreasing by the same argument. Hence, smaller shocks lead to a flatter PC as fewer firms adjust.

Suppose firms start with marginal cost z_0, which then changes to z. This implies a cutoff $k^*(z)$, which is a different cost threshold for price-adjusting firms.

Lemma 8 (Inflation with Fixed Costs)**.** *With fixed costs of adjusting prices and a shock to marginal costs, $k^*(z) = \frac{\xi_0}{4}(z_0 - z)^2$, which implies $\frac{\partial \bar{p}}{\partial z} = F(k^*(z))\frac{\xi}{2}$.*

Proof of Lemma 8. k^* is defined as above. Now we modify $p_0 = \frac{\xi_0(1+z_0)}{2}$. Hence we have

$$k^*(z) = \xi(1-z)^2/4 + p_0^2/\xi + z\xi - p_0(1+z),$$
$$= \frac{\xi}{4}(1-z)^2 + \frac{\xi_0^2(1+z_0)^2}{4\xi} + z\xi - \frac{\xi_0(1+z_0)(1+z)}{2},$$
$$= \frac{1}{4}\left(\xi(1-z)^2 + \frac{\xi_0^2(1+z_0)^2}{\xi}\right) - \frac{\xi_0(1+z_0)(1+z)}{2} + z\xi.$$

To consider a cost shock, we set $\xi = \xi_0$ to get

$$k^*(z) = \frac{1}{4}\left(\xi_0(1-z)^2 + \frac{\xi_0^2(1+z_0)^2}{\xi_0}\right) - \frac{\xi_0(1+z_0)(1+z)}{2} + z\xi_0 = \frac{\xi_0}{4}(z_0-z)^2.$$

□

Note that $k^*(\xi)$ is linear in ξ, while $k^*(z)$ is quadratic in z. Thus, if shocks are small (percentages), then $(z-z_0)^2$ is likely to be smaller than ξ, and thus for small shocks, $k^*(z) < k^*(\xi)$. This means that we are likely to get even less response to a comparably sized supply shock than to a demand shock. If few firms were adjusting in response to a demand shock, then few firms would be adjusting in response to a (comparable) cost shock. If z is near z_0, then the cutoff k^* is small and so few firms will adjust. Thus, it is completely plausible to have no response in this case as well. If $k^* < \underline{k}$, then no firms will adjust prices in response to a change in z.

E 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 in which money plays an essential role. The setup's ingredients are standard. However, putting the pieces together is quite involved. Therefore we start the model description with a preview. We subsequently fully describe every piece of the model.

E.1 Preview

The setup is based on the foundational papers by Lagos and Wright (2005) and Lucas and Stokey (1987). As in Lagos and Wright (2005), we exploit quasi-linearity and periods that are divided in a day and a night to be able to handle agent (informational) heterogeneity. As in 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, while 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, and by firms' level of productivity A. As in most of the literature, the preference shock is a modeling device to generate fluctuations in nominal aggregate demand. As in the simple model presented in the paper, a key assumption of our setup is that there is household heterogeneity in the information about the aggregate shocks. Some households may be imperfectly informed about the value of discount factor at night.[16] 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. As discussed in the benchmark model, consumers' knowledge of firms' productivity is irrelevant in equilibrium.

Firms, by assumption, are informed about the preference shock as well as their productivity. 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 can learn by looking at firms' prices. Notice that, our informational assumptions force us to move away from monopolistic competition (or other forms of centralized goods markets). 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 determine its policy. This rule depends on deviations of inflation from a target and on the output gap.

[16] 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.

E.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 discrete. 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 τ: $\tau = 0$ signifies the day, and $\tau = 1$ signifies the night. (However, to simplify the notation, we skip τ notation whenever possible.) Periods are indexed by t and run from $t = 0$ 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. Labor markets are open during both the day and the night. 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 $t+1$. Long term bonds and cash holdings are denoted B and M respectively.

Exogenous Aggregate State. The exogenous aggregate state is given by the realizations of a preference shock θ_t and aggregate productivity A_t. We specify the processes for these shocks below.

Central Bank. The money supply set by the central bank 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) and Outsiders (uninformed consumers). Insiders are perfectly informed about the state $\{\theta_t, A_t\}$; 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, the household solves for all variables that trade in credit. This is the "day-to-night" problem where the action happens. (The outer problem is presented below.) In contrast to the setting in the main body, and without loss of generality, here we suppose that time discounting with β occurs only across periods (i.e., the outer problem) rather than within the period (across sub-periods).

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

$$\max_{c_t, l\_t, L_t} \mathbb{E}_{j\tau}\left[\left(u(c_t)-l_t\right)+\theta_t\left(U(\bar{C})-L_t\right)\right]$$

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_t determines the discount factor between the day and the night. Following the previous sections, we specify the process for θ_t to follow 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 $\mathbb{E}[\theta_t] = \theta < 1$. 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 $\theta_L > \theta_H$.[17] 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.[18] 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

$$p_t c_t + P_t \bar{C} = d_t + w_t l_t + W_t L_t \tag{25}$$

where prices have been defined above and d_t are profits.

Denoting by λ_t the Lagrange multiplier of the constraint (25), the first-order condition for daytime consumption c is

$$u'(c_t) = \mathbb{E}_{j\tau}[\lambda_t p_t]$$

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

[17] 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_t.)

[18] 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 $\lambda_t p_t$.) Notice however that Outsiders observe the price of the firm they meet, p_t, 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:

$$\begin{aligned} u'(c_t) &= p_t \mathbb{E}_{j\tau}[\lambda_t] \\ 1 &= \lambda_t w_t \\ \theta_t &= \lambda_t W_t \end{aligned}$$

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

From the second equation we observe that $\lambda_t = 1/w_t$. Further manipulating the equations above we can summarize the set of first-order conditions as

$$u'(c_t) = p_t \mathbb{E}_{j\tau}\left[\frac{\theta_t}{W_t}\right] \tag{26}$$

$$\frac{1}{w_t} = \theta_t \frac{1}{W_t} \tag{27}$$

We continue by presenting the outer problem of the household. This is the problem solved from one day to the next. This problem will give rise to an explicit role for money and hence allows us to define the monetary policy instrument.

Define

$$\mathscr{U}(c_t, l_t, C_t, L_t) = \Big(u(c_t) - l_t\Big) + \theta_t \Big(U(C_t) - L_t\Big)$$

In the outer problem, the household needs to solve

$$\max_{c_t, l_t, C_t, L_t, M_t, B_t} \mathbb{E}_{j\tau}\left[\sum_{t=0}^{\infty} \beta^t \, \mathscr{U}(c_t, l_t, C_t, L_t)\right]$$

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

$$p_t c_t + P_t C_t + B_t + M_t = w_t l_t + W_t L_t + M_{t-1} + T_t + \left(1 + i_t^{LT}\right) B_{jt-1} + d_t \tag{28}$$

where i_t^{LT} is a long-term nominal interest rate, and T_t 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

$$P_t C_t \leq M_{t-1} + T_t \tag{29}$$

Denoting by χ_t the multiplier on the budget constraint (28) and by ψ_t the multiplier on the CIA constraint (29), we get the set of first-order conditions

$$\beta^t u'(c_t) = \mathbb{E}_{j\tau}\left[\chi_t p_t\right] \tag{30}$$

$$\beta^t = \chi_t w_t \tag{31}$$

$$\beta^t \theta_t U'(C_t) = (\chi_t + \psi_t) P_t \tag{32}$$

$$\beta^t \theta_t = \chi_t W_t \tag{33}$$

$$\chi_t = \mathbb{E}_t\left[\chi_{t+1} + \psi_{t+1}\right] \tag{34}$$

$$\chi_t = \left(1 + i_{t+1}^{LT}\right) \mathbb{E}_t\left[\chi_{t+1}\right] \tag{35}$$

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 $\mathbb{E}_t[\,\cdot\,]$ (conditional on full information, which is available in the centralized market).

From (31) and (33), we observe that $\chi_t = \beta^t / w_t$ and $\chi_t = \beta^t \theta_t / W_t$. Thus,

$$\frac{1}{w_t} = \theta_t \frac{1}{W_t}$$

which is the same as (27). Also, plugging in the expression for χ_t obtained from (33) into (30), we get

$$u'(c_t) = p_t \mathbb{E}_{j\tau}\left[\frac{\theta_t}{W_t}\right] \tag{36}$$

which is the same as (26).

The remaining conditions, determining the demand for money and bonds, can be simplified as follows. Equation (32), one period forward, is $\beta^{t+1}\theta_{t+1}U'(C_{t+1}) = (\chi_{t+1}+\psi_{t+1})P_{t+1}$. Solving for $\chi_{t+1}+\psi_{t+1}$, and using the expression for χ_t, equation (34) becomes

$$\frac{\theta_t}{W_t} = \beta\ \mathbb{E}_t\left[\frac{\theta_{t+1}}{P_{t+1}}U'(C_{t+1})\right] \tag{37}$$

Finally, equation (35) is equivalent to

$$\frac{\theta_t}{W_t} = \beta\left(1+i^{LT}_{t+1}\right)\mathbb{E}_t\left[\frac{\theta_{t+1}}{W_{t+1}}\right] \tag{38}$$

Production. All firms in the economy have a linear technology and produce using only labor. Within every period, a monopolist of the decentralized market produces c according to the production function $c_t = A_t l_t$. The real marginal cost is $z_t \equiv 1/A_t$. Following the previous sections, we specify the process for z_t to follow an i.i.d. binary Markov chain with two values, z_L and z_H, with $Pr(z_L) = Pr(z_H) = 1/2$.

The competitive firm produces C according to the production function $C_t = L_t$, 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_t^S. 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:

$$M_t^S = M_0\left(\hat{\pi}_t\right)^{-\phi_\pi}\left(\hat{x}_t\right)^{-\phi_x} \tag{39}$$

where M_0 is the natural level of the money supply, $\hat{\pi}_t$ is inflation, defined as the percentage deviation of the price level away from steady state $\bar{p}_0$: $\hat{\pi}_t = \frac{\int \bar{p}_t dF(\alpha) - \bar{p}_0}{\bar{p}_0}$, and $\hat{x}_t$ is the output gap, defined as the percentage deviation of aggregate output from the natural level y^n: $\hat{x}_t = \frac{\int y_t dF(\alpha) - y^n}{y^n}$. Below we show that this policy rule can be expressed as a rule for a short-term interest rate rate.

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) $\{c_t, C_t, l_t, L_t, B_t, M_t\}$, labor demand (for each firm) $\{l^D_{jt}, L^D_t\}$, profits $\{d_t\}$, money supply $\{M^S_t\}$, nominal transfers $\{T_t\}$, nominal prices $\{p_t, P_t\}$, nominal wages $\{w_t, W_t\}$, long-term nominal interest rates $\{1 + i^{LT}_t\}$, for all t, 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_t 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 $C_t = \bar{C}$. Second, we conjecture that $M_t = M^S_t$, for all t. Then, $P_t = M_t / \bar{C}$.

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

$$P_t = W_t = \frac{M_t}{\bar{C}} \tag{40}$$

Now, taking equation (37) and writing it as

$$\frac{\theta_t}{M_t} = \beta U'(\bar{C}) \, \mathbb{E}_t \left[\frac{\theta_{t+1}}{M_{t+1}} \right]$$

reveals that as long as $U'(\bar{C}) = 1/\beta$ and θ_t/M_t is a martingale, equation (37) is satisfied. A monotonic rule can be mapped into a degree of monetary policy adjustment γ. Hence, we write the rule

$$\frac{1}{M_t} = \gamma\frac{1}{\theta_t} + (1-\gamma)\frac{1}{M_0}$$

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

$$\frac{\theta_t}{M_t} = \gamma + (1-\gamma)\frac{\theta_t}{M_0}$$

Taking the expectation of θ_t/M_t shows, trivially, that this ratio is a martingale, and thus equation (37) is satisfied.

A similar argument for $i_t^{LT} = 1/\beta - 1$ shows that equation (38) is satisfied.

Since this is a closed economy with a zero net supply of bonds, we can simply set $B_t = 0$ for all households. It remains to check that the labor markets clear. The centralized market clears when each household supplies $L_t = C_t$. 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 ensures a tractable and isolated treatment of the game.)

Finally, set $T_t = M_t - M_{t-1}$. At this point, we are able to verify our money demand and centralized 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_t in both problems are the same (equations (26) and (36)). Also, since the equilibrium in the outer problem requires $L_t = C_t$, and since $W_t = P_t$, $M_t = M_{t-1} + T_t$ and $B_t = 0$, then, setting $E = \bar{C}$ the budget constraints in both problems reduce to

$$p_t c_t = d_t + w_t l_t$$

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

Lemma 9. *The equilibrium allocations of c_t, l_t, L_t in the inner problem are the same as in the outer problem. Moreover, the equilibrium allocation C_t is an admissible endowment E of the inner problem.*

To obtain the simple, partial-equilibrium model in the body, 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_t 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_t. Then, $Q_t = P_t$.

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 for each consumer, nothing guarantees that C is equal to E *for each consumer*. However, by linearity, one can verify that the aggregate quantity of C is indeed equal to E (i.e., markets clear).

Notice that $M_t = E \cdot Q_t$. So the rule (39) is

$$E \cdot Q_t = E \cdot Q_0 \left(\hat{\pi}_t\right)^{-\phi_\pi} \left(\hat{x}_t\right)^{-\phi_x}$$

which is

$$1 + i_t = (1 + i_0) \left(\hat{\pi}_t\right)^{\phi_\pi} \left(\hat{x}_t\right)^{\phi_x}$$

In logs

$$\log(1 + i_t) = \log(1 + i_0) + \phi_\pi \hat{\pi}_t + \phi_x \hat{x}_t$$

Finally, the simple model can be written using two periods only, which allows to drop the t 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 10. *The model presented in the body has the same equilibrium allocation of c_t as the full GE model. Also, the aggregate consumption of C in the simple model is equal to $E = \bar{C}$.*