License to Ease: Competition Policy and the Conduct of Monetary Policy *

Antoine Camous† Basile Grassi‡ Matteo Samarani§

March 2026

[Last Version]

Abstract

Monetary policy is transmitted through firms' pricing decisions, which are shaped by the competitive environment. Competition policy affects that environment through measures that shape the degree of product-market competition. We study the interaction between competition policy and monetary policy in a New Keynesian model with oligopolistic competition. A competition authority chooses the degree of product-market competition, trading off its benefits against a resource cost, while a central bank conducts monetary policy under discretion. In this environment, competition shapes firms' markups and strategic pricing incentives, thereby affecting both the slope of the Phillips curve and the long-run distortion that gives rise to inflation bias. We characterize the optimal institutional design when a benevolent planner delegates authority to both institutions and chooses their mandates. Our main result is that the optimal degree of central bank conservatism depends systematically on the effectiveness of competition policy. When competition can be promoted more effectively, equilibrium competition is higher, inflation bias is lower, and it is optimal to appoint a less conservative central bank. The analysis implies that competition policy and monetary policy should be jointly designed rather than in isolation.

*The views expressed in this paper are solely those of the authors, and no responsibility for them should be attributed to the Banque de France or the Eurosystem. We would like to thank Andrey Alexandrov, Guido Ascari, Eric Bartelsman, Florin Bilbiie, Job Boerma, Jeffrey Campbell, Maarten de Ridder, Timo Haber, Henri Graul, Luigi Iovino, Virgiliu Midrigan, Tommaso Monacelli, Lukas Nord, Nicola Pavoni, Xavier Ragot, Dmitriy Sergeyev, Guido Tabellini and Antonella Trigari for useful comments. We also thank the organizers and participants of the 4th XAmsterdam Macroeconomic Workshop for useful comments and suggestions.

†Banque de France, antoine.camous@banque-france.fr.

‡Bocconi University, OFCE, IGIER, and CEPR. basile.grassi@unibocconi.it.

§Bocconi University, matteo.samarani2@phd.unibocconi.it.

1 Introduction

The transmission mechanism of monetary policy depends not only on nominal rigidities, but also on the competitive environment in which firms set prices. The intensity of product-market competition shapes the sensitivity of demand to a firm's price and, with it, the extent to which changes in marginal costs are passed through into prices. Competition authorities play a central role in shaping that environment: through merger control, antitrust enforcement, and market regulation, they influence the conditions under which firms set prices.

The relationship between market structure and monetary transmission depends on the institutional design of monetary and competition policy. Monetary and competition policy are delegated to independent institutions, whose mandates are designed to align with social objectives. The effectiveness of this institutional setup depends not only on each mandate in isolation, but also on the interaction between them. By shaping the product-market environment, competition policy affects the incentive structure faced by a central bank. This interdependence raises a broader institutional question: *How should the mandates of a competition authority and a central bank be jointly designed?*

To study this question, we develop a tractable New Keynesian model with oligopolistic competition. The economy features a competition authority that chooses the degree of product-market competition, balancing its benefits against a resource cost, and a central bank that conducts monetary policy under discretion. We use this framework to characterize the optimal institution designed, reflecting both the optimal degree of competition and the mandates of the two authorities. The key insight is that, once competition is treated as a policy instrument, competition policy directly alters the monetary policy problem. Specifically, competition policy affects the slope of the Phillips curve, the central bank's inflation–output trade-off, and the economy's long-run inefficiency, measured by the wedge between the natural and the efficient levels of output.[1] Therefore, the design of monetary delegation cannot be studied in isolation from competition policy. When a benevolent competition authority is *better able* to foster competition, the optimal institutional response is to appoint a *less conservative* central bank: competition effectively provides a *"license to ease"*.

We reach this conclusion in a tractable framework that combines nominal rigidities with strategic price setting. Households consume differentiated varieties grouped into sectors, with nested CES preferences as in Atkeson and Burstein (2008). Within each sector,

[1] The natural level of output is the level that would prevail under flexible prices, given the economy's underlying distortions, whereas the efficient level of output is the level chosen by a planner in the absence of those distortions. The wedge between the two therefore captures the economy's long-run inefficiency.

n firms compete à la Bertrand and face Calvo price rigidities. Because firms internalize the effect of their pricing decisions on both sectoral demand and the sectoral price index, desired markups are endogenous and depend on market structure. Price setting therefore features strategic complementarities, and greater competition strengthens monetary pass-through and the sensitivity of inflation to economic activity. As a result, the slope of the Phillips curve increases with the number of firms and declines with market concentration. Using the state-level Phillips curve approach of Hazell et al. (2022), we provide reduced-form evidence that the slope of the Phillips curve varies systematically with local market structure in the United States: it is steeper in less concentrated states and flatter in more concentrated ones.

In our environment, the competition authority chooses the number of firms operating in each sector. Increasing the number of firms increases product-market competition, but it comes with a cost, which we model as proportional to aggregate output. This specification captures the idea that promoting and sustaining competition—through entry facilitation, antitrust enforcement, or regulatory oversight—absorbs real resources. The degree of competition determines firms' markups and, in turn, the objects most relevant for the conduct of monetary policy: the slope of the Phillips curve and the inflation bias driven by the economy's inefficiency wedge.

The interaction between policy institutions is structured as a sequential Stackelberg game. At an initial stage, an institutional designer chooses the mandates of the competition authority and the central bank. Given these mandates, the competition authority selects a degree of product-market competition. Then, at each time t, the monetary authority sets policy under discretion, given the competition structure and its mandate. This timing reflects the slower-moving nature of institutional design and competition policy relative to monetary policy, which is set each period in response to economic shocks. As a result, conditional on its mandate, the competition authority internalizes how market structure shapes the sequential conduct of monetary policy.

The central bank sets policy to address the inflation-output trade-off emphasized by Clarida et al. (1999), Benigno and Woodford (2003), Benigno and Woodford (2005), and Galí (2015), which is endogenously shaped by market structure. A higher number of firms steepens the Phillips curve, thereby improving the inflation–output stabilization trade-off, while also reducing the inefficient inflation bias associated with market power. As a result, stronger competition both weakens the structural incentive of a monetary authority to generate excessive inflation and increases the central bank's incentive to shift stabilization toward the output gap. The degree of competition therefore shapes monetary policy along two dimensions: cyclical stabilization and structural inflation bias.

These mechanisms have important implications for institutional design. The compe-

tition authority should be assigned a mandate that aligns with social welfare, since its choice affects both long-run efficiency and the monetary policy stabilization problem. By contrast, the optimal degree of central bank conservatism decreases with the effectiveness of competition policy. Economies that can sustain stronger competition at a lower cost require less conservative central banks, while declines in competition call for a more conservative monetary stance. In this sense, we emphasize that the competition policy enforcement not only pursues a microeconomic objective—reducing sectoral markups and product-market distortions—but also a macroeconomic one, because it shapes the economy's inflation–output trade-off and the conduct of monetary stabilization policy.

Related literature. Our paper lies at the intersection of three literatures: oligopolistic competition under nominal rigidities, welfare and optimal monetary policy in New Keynesian models, and monetary discretion and delegation. Our main contribution is to connect these strands in a unified framework in which product-market competition is an endogenous policy variable. In our setting, competition policy affects both the slope of the Phillips curve and the long-run distortion that gives rise to the inflation bias. As a result, the design of monetary delegation cannot be studied independently of competition policy.

We relate first to the literature on oligopolistic competition under nominal rigidities. Etro and Rossi (2015), Mongey (2021), Wang and Werning (2022), Baqaee et al. (2024), Meier and Reinelt (2024), and Gagliardone et al. (2025) show how strategic interactions among firms shape markups, pass-through, and inflation dynamics. We build on this insight, but depart from this literature by treating the degree of competition not as an exogenous feature of the private economy, but as the outcome of policy chosen by a competition authority. This allows us to study how product-market policy feeds back into the conduct of monetary policy.

Our paper also relates to the New Keynesian literature on welfare and optimal monetary policy. Rotemberg and Woodford (1997), Clarida et al. (1999), Benigno and Woodford (2003), and Benigno and Woodford (2005), followed by Gali and Monacelli (2005), Woodford (2011), Coibion et al. (2012), Galí (2015), and Olivi et al. (2023), characterize optimal policy when the relevant trade-offs depend on underlying distortions. We contribute to this literature by showing that competition policy changes both objects that matter for discretionary monetary policy: the slope of the Phillips curve, which governs stabilization, and the long-run inefficiency, which gives rise to the inflation bias. In this sense, market structure is not just a determinant of pricing behavior, but a primitive source of the monetary policy trade-off.

A further related strand studies monetary discretion and delegation. Beginning with

Barro and Gordon (1983) and Rogoff (1985), and continuing through Walsh (1995), Svensson (1997), Backus and Driffill (1985), Canzoneri (1985), Cukierman and Meltzer (1986), Athey et al. (2005), Halac and Yared (2020), Halac and Yared (2022), and Afrouzi et al. (2025), this literature studies how monetary institutions should be designed under discretion. A common feature of this work is that the structure of the economy, and hence the inflation bias and the stabilization trade-off, is taken as given. By contrast, we show that the optimal degree of central bank conservatism depends on the conduct of another delegated authority, because competition policy changes both the central bank's temptation to inflate and the inflation–output trade-off it faces.

Finally, our paper relates to work on strategic interaction across policy institutions. Dixit and Lambertini (2003) study how the interaction between monetary and fiscal authorities depends on the assignment of objectives and on the ability of each authority to commit. We share this emphasis on interacting delegated institutions, but our setting is different in a crucial respect: the second authority is a competition authority rather than a fiscal authority, and its instrument changes the structure of price-setting itself. This allows us to characterize the joint design of monetary and competition policy mandates.

Structure. The rest of the paper proceeds as follows. Section 2 presents the model. Section 3 derives the oligopolistic Phillips curve. Section 4 derives welfare and formulates the policy game. Section 5 presents the main results.

2 Model

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. The economy is populated by a representative household and a unit mass of symmetric sectors. In each sector there is a *finite* number n of firms (varieties). The finiteness of n implies oligopolistic (strategic) interaction within each sector. We focus on symmetric equilibria across sectors and firms. The parameter n will serve as our index of product-market competition.

2.1 Households

As the main focus of the paper lies on the supply side, we adopt a parsimonious specification of the demand side. Accordingly, we assume the economy admits an infinitely-lived representative household who derives utility from consumption and disutility from labor

according to the objective function $\mathbb{E}_0\left[(1-\beta)\sum_{t\in\mathbb{N}}\beta^t u\left(C_t,L_t\right)\right]$, where

$$u\left(C_t,L_t\right)=\begin{cases}\frac{C_t^{1-\sigma}}{1-\sigma}-\frac{L_t^{1+\varphi}}{1+\varphi} & \text{if } \sigma\in(0,1)\cup(1,\infty),\\ \log C_t-\frac{L_t^{1+\varphi}}{1+\varphi} & \text{if } \sigma\to 1,\end{cases} \tag{1}$$

is the period utility function, L_t is hours worked, $\beta\in(0,1)$ the discount factor, and C_t denotes an aggregate consumption bundle composed of sector- and firm-specific varieties. Specifically, the household consumes from a continuum of differentiated sectors $k\in\mathcal{K}\equiv[0,1]$, each of which contains a finite set of varieties $i\in\mathcal{I}\equiv\{1,\ldots,n\}$, with $n>1$. Preferences over these varieties are represented by a nested Constant Elasticity of Substitution (CES) aggregator:

$$C_t\equiv\left(\int_{k\in\mathcal{K}}C_{k,t}^{\frac{\omega-1}{\omega}}dk\right)^{\frac{\omega}{\omega-1}}\text{ and }C_{k,t}\equiv n^{-\frac{1}{\varepsilon-1}}\left(\sum_{i\in\mathcal{I}}C_{ik,t}^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{2}$$

where $C_{ik,t}$ denotes consumption of variety i in sector k, and the elasticities satisfy $\varepsilon>\omega\geq 1$. The scaling factor $n^{-\frac{1}{\varepsilon-1}}$ removes any love-of-variety effect, thereby allowing us to isolate the role of competition. Intuitively, in this formulation, an increase in n corresponds to more firms producing a fixed set of varieties rather than an expansion of the product space. Letting $P_{ik,t}$ denote the price of variety i in sector k, standard optimization implies the following demand system:

$$C_{ki,t}=\frac{1}{n}\left(\frac{P_{ik,t}}{P_{k,t}}\right)^{-\varepsilon}\left(\frac{P_{k,t}}{P_t}\right)^{-\omega}C_t, \tag{3}$$

where the corresponding ideal price indexes satisfy

$$P_t^{1-\omega}\equiv\int_{k\in\mathcal{K}}P_{k,t}^{1-\omega}dk\text{ and }P_{k,t}^{1-\varepsilon}\equiv n^{-1}\sum_{i\in\mathcal{I}}P_{ik,t}^{1-\varepsilon}. \tag{4}$$

The only saving technology available to the household is a one-period nominal risk-free bond. The household also owns all firms in the economy and receives their profits each period as lump-sum income. The resulting budget constraint is

$$P_tC_t+\left(1+i_t\right)^{-1}B_{t+1}\leq B_t+W_tL_t+P_t\Pi_t-P_tT_t, \tag{5}$$

where B_{t+1} denotes bond holdings at the beginning of period $t+1$, i_t the nominal interest rate, W_t the nominal wage, Π_t real profits, and T_t real lump-sum taxes. The solution of the

household's problem is characterized by the following standard optimality conditions:

$$C_t^{-\sigma} = \beta (1+i_t) \mathbb{E}_t \left[C_{t+1}^{-\sigma} \frac{P_t}{P_{t+1}} \right], \tag{6}$$

$$\frac{W_t}{P_t} C_t^{-\sigma} = L_t^{\varphi}, \tag{7}$$

which, together with the household's budget constraint, determine its optimal choices for any given sequence of prices and nominal interest rate.[2]

2.2 Firms: Oligopolistic Calvo pricing

We assume that each variety $i \in \mathcal{I}$ in sector $k \in \mathcal{K}$ is produced by a single firm. The fact that the number of varieties in each sector is a *finite* integer $n > 0$ allows for *non-atomistic strategic interaction*: firms internalize the effect of their decisions on sectoral aggregates. Strategic interaction within each sector is modeled as Bertrand competition.

Firms are ex-ante homogeneous and hire labor to produce their respective varieties. Specifically, each firm i in sector k produces output with a constant return to scale technology $Y_{ik,t} = A_t L_{ik,t}$, implying a (nominal) marginal cost $\Psi_t \equiv W_t / A_t$, where A_t denotes the common level of total factor productivity (TFP). This, in turn, evolves according to the stochastic process $\log A_t = \rho_a A_{t-1} + \sigma_a \epsilon_{a,t}$, with $\epsilon_{a,t} \sim \mathcal{N}(0,1)$, $\rho_a \in (0,1)$, and $\sigma_a > 0$. In addition, we assume that each firm receives a production subsidy at gross rate $\mathcal{T}_t$, proportional to its marginal cost and financed through lump-sum taxes. The government sets this subsidy according to the stochastic process $\log \mathcal{T}_t = (1-\rho_\tau) \log \mathcal{T} + \rho_\tau \log \mathcal{T}_{t-1} + \sigma_\tau \epsilon_{\tau,t}$, with $\epsilon_{\tau,t} \sim \mathcal{N}(0,1)$, $\rho_\tau \in (0,1)$, $\sigma_\tau > 0$, and where $\mathcal{T} \in (0,1)$ denotes the steady-state subsidy.[3] The distortionary nature of this subsidy gives rise to a stochastic labor wedge, implying that the economy is subject to cost-push shocks that break the divine coincidence. In detail, we define the cost-push as the stochastic process $\tau_t \equiv \log(\mathcal{T}_t / \mathcal{T})$.Following the standard New Keynesian literature (Woodford, 2011; Galí, 2015), this should be interpreted as a parsimonious modeling device for generating a short-run monetary policy trade-off.[4]

Firms adjust prices infrequently. In particular, in each period t, an i.i.d. Calvo (1983) lottery determines who can adjust its price optimally. Each firm is allowed to adjust its price

[2] In addition to the budget constraint we assume the household is subject to a solvency conditioning ruling out Ponzi schemes. That is, $\lim_{s\to\infty} \mathbb{E}_t [\mathcal{Q}_{t,t+s} B_{s+1} / P_{t+s}] \geq 0$, where $\mathcal{Q}_{t,t+s} \equiv \beta^s u_c(C_{t+s}, L_{t+s}) / u_c(C_t, L_t)$ is the representative household's stochastic discount factor.

[3] We will discuss the steady-state properties of this subsidy in more detail later, when we introduce the cost of competition policy.

[4] Another way to introduce a labor wedge is to assume a time-varying elasticity of substitution across varieties. While we maintain this specification as our baseline, we discuss alternative microfoundations in the Appendix.

with probability $(1-\theta) \in (0,1)$, while non-selected firms keep their prices unchanged from the previous period. The combination of nominal rigidities and Bertrand competition generates a *dynamic oligopolistic interaction* within each sector. In principle, in a Markov-perfect equilibrium – such as in Mongey (2021) and Wang and Werning (2022) – a firm internalizes not only the dynamic effects of its pricing decision on its own demand and on the sectoral price index. It also internalizes the fact that, if it is unable to reset its price at some future date $t+s$, the price chosen at time t becomes a state variable for its competitors, who will adjust their behavior in response. To gain tractability, we follow Etro and Rossi (2015) and Gagliardone et al. (2025) in assuming that each firm takes the entire sequence of competitors' prices as given. Under this assumption, firms account for the dynamic implications of their pricing decisions for demand and for the sectoral price index, but not for competitors' strategic responses.

With this, a firm i in sector k selected to optimally reset its price at time t takes the sequence $\left\{P_{t+s}, \left(P_{jk,t+s}\right)_{j\neq i}\right\}_{s\in\mathbb{N}}$ as given and solves

$$\begin{aligned}
&\max_{\left(P^*_{ik,t},\mathbf{x}_{i,t}\right)} \left\{ \mathbb{E}_t \left[\sum_{s=0}^{\infty} \theta^s \frac{\mathcal{Q}_{t,t+\tau}}{P_{t+\tau}} \left(P^*_{ik,t} - \mathcal{T}_{t+s}\Psi_{t+s}\right) Y_{ik,t+s|t} \right] \right\} \\
&\text{s.t. } Y_{ik,t+s|t} = \frac{1}{n} \left(\frac{P^*_{ik,t}}{P_{k,t+s}} \right)^{-\varepsilon} \left(\frac{P_{k,t+s}}{P_{t+s}} \right)^{-\omega} Y_{t+s}, \qquad (8) \\
&P_{k,t+s} = n^{-\frac{1}{1-\varepsilon}} \left(\left(P^*_{ik,t}\right)^{1-\varepsilon} + \sum_{j\neq i} P^{1-\varepsilon}_{jk,t+s} \right)^{\frac{1}{1-\varepsilon}},
\end{aligned}$$

where $\mathbf{x}_{i,t} \equiv \left(\left\{ Y_{ik,t+s|t} \right\}_{s\in\mathbb{N}}, P_{k,t+s} \right)$, Ψ_{t+s} is the nominal marginal cost, $\mathcal{Q}_{t,t+s} \equiv \beta^s u_c\left(C_{t+s}, L_{t+s}\right) / u_c\left(C_t, L_t\right)$ the stochastic discount factor, and $Y_{ik,t+s|t}$ the demand at time $t+s$ for a firm that last reset its price at time t.

Since firms internalize their impact on the sectoral price index, pricing behavior exhibits *strategic complementarities* that, as is well understood, amplify nominal rigidities. To see this, notice that the elasticity of residual demand for firm i in sector k can be written as $\varpi_{ik,t} \equiv -\frac{d\log Y_{ik,t}}{d\log P_{ik,t}} = \varepsilon\left(1-S_{ik,t}\right) + \omega S_{ik,t}$, where $S_{ik,t} \equiv n^{-1}\left(P_{ik,t}/P_{k,t}\right)^{1-\varepsilon}$ is firm i's market share within sector k. This object increases with the sectoral price index and decreases with the intensity of competition n. An increase in the sectoral price index therefore raises $S_{ik,t}$ which, given $\varepsilon > \omega$, lowers the elasticity $\varpi_{ik,t}$. When prices are sticky, a given shock triggers only a limited adjustment in the sectoral price index. This implies that, *ceteris paribus*, the elasticity of residual demand moves less and the effect of the shock is correspondingly dampened, compared to the monopolistic competition benchmark. Importantly, this dampening effect weakens as n increases: in the limit $n \to \infty$, we have

$S_{ik,t} \to 0$, so the elasticity of residual demand becomes constant and equal to ε, as in the standard monopolistic competition model. As we will discuss later, this mechanism lies at the core of our theory of long-run inflation and provides a channel through which the competition authority can, in principle, influence the efficiency properties of the long-run equilibrium.

Since firms are ex-ante homogeneous, we focus on equilibria in which sectors are symmetric, i.e., $P_{k,t} = P_t$ for every $t \in \mathbb{N}$. In addition, conditional on resetting their price optimally at time t, firms share the same marginal cost and therefore choose the same price, i.e., $P^*_{ki,t} = P^*_t$ for each $(i,k) \in \mathcal{I} \times \mathcal{K}$. In the Appendix we show that the optimal price P^*_t is obtained by imposing these restrictions in the system of first-order conditions associated with problem (8). Specifically, this gives

$$\mathbb{E}_t \left[\sum_{s=0}^{\infty} \theta^s \mathcal{Q}_{t,t+\tau} \frac{Y_{t+s|t} \left(\varpi_{t+s|t} - 1 \right)}{P_{t+s}} \left(P^*_t - \mathcal{M}_{t+s|t} \mathcal{T}_{t+s} \Psi_{t+s} \right) \right] = 0, \tag{9}$$

where $\mathcal{M}_{t+s|t} \equiv \varpi_{t+s|t} / \left(\varpi_{t+s|t} - 1 \right)$ denotes the optimal markup at time $t+s$ for a firm that last reset its price at time t. The term $\varpi_{t+s|t}$ is the corresponding elasticity of residual demand, defined as $\varpi_{t+s|t} \equiv -d \log Y_{t+s|t} / d \log P^*_t$. Moreover, under the symmetry assumptions above, the aggregate price index satisfies the standard recursion $P_t^{1-\varepsilon} = (1-\theta) \left(P^*_t \right)^{1-\varepsilon} + \theta P_{t-1}^{1-\varepsilon}$.[5]

Finally, imposing labor market clearing within sectors and varieties yields the aggregate production function $Y_t = \mathcal{D}_t^{-1} A_t L_t$, where $\mathcal{D}_t \equiv \int_{k \in \mathcal{K}} n^{-1} \sum_{i \in \mathcal{I}} \left(P_{ik,t} / P_t \right) dk$ is an index of price dispersion.

2.3 Cost of Competition

We model the cost of competition in a very parsimonious, yet general, way. In our baseline specification, maintaining a competition level n requires the competition authority to devote each period a fraction of aggregate output to this purpose. These resources are financed by lump-sum taxes on the representative household: $T_t = \mathcal{G}(n) Y_t$, where $\mathcal{G}'(n) > 0$, $\mathcal{G}(n) \to 0$ as $n \to 1$, and $\mathcal{G}(n) \to 1$ as $n \to \infty$. Thus, sustaining higher levels of competition requires the authority to allocate an increasingly large share of total output, with this share approaching one as $n \to \infty$. As a result, the monopolistic competition limit is *de facto* non optimal. This assumption implies the following aggregate resource constraint: $C_t = (1 - \mathcal{G}(n)) Y_t$.

[5] In deriving the aggregate price index we are assuming the Law of Large numbers holds at sectoral level, even tough we have a small numbers of firms within each sector. This assumption allows us to abstract from granularity effects. See Burstein et al. (2025) for a model in which this channel is considered.

Notice that in this formulation we can interpret

$$\alpha(n) \equiv \left(\left| \frac{d \log (1 - \mathcal{G}(n))}{d \log n} \right| \right)^{-1} \tag{10}$$

as a measure of the *effectiveness of the competition authority*. A higher value of $\alpha(n)$ implies that a marginal increase in the number of firms induces a proportionally smaller reduction in consumption. In this sense, the competition authority is more effective at limiting the resource cost associated with competition.

Finally, observe that the cost of competition is distortionary: a social planner would recognize that part of labor must be diverted to finance this cost, altering the optimal level of output. To remove this distortion, we assume a production subsidy in the steady state equal to $\mathcal{T} = (1 - \mathcal{G}(n))^{-1}$. This ensures that the cost of competition acts purely as a non-distortionary deadweight loss rather than affecting marginal production decisions.[6]

2.4 Flexible Price Limit

The flexible-price limit is obtained as $\theta \to 0$ when the probability of resetting price is equal to one. In this case, firms set a common price given by $P_t^* = P_t = \mathcal{M}(n) \mathcal{T}_t \Psi_t$, where $\mathcal{M}(n) \equiv \varpi(n) / (\varpi(n) - 1)$ and $\varpi(n) = \varepsilon - (\varepsilon - \omega) n^{-1}$. In other words, firms set their price as a markup over marginal cost. Oligopolistic competition implies that this markup is decreasing in n, i.e., $\mathcal{M}'(n) < 0$, with $\mathcal{M}(n) \to \varepsilon / (\varepsilon - 1)$ as $n \to \infty$. This arises because firms internalize their impact on the sectoral price index and therefore operate along a region of their residual demand curve with lower elasticity. As the number of firms increases, this strategic effect diminishes, and the model converges to the monopolistic-competition limit. Finally, in the flexible-price limit there is no price dispersion, i.e., $\mathcal{D}_t = 1$.

Combining the household's intratemporal condition (7) with the above pricing equation under flexible price, and using the aggregate production function together with the resource constraint, we obtain the flexible-price aggregate output

$$\log Y_t^f = \log Y(n) + \left(\frac{1 + \varphi}{\sigma + \varphi} \right) \log A_t - \frac{1}{\sigma + \varphi} \tau_t, \tag{11}$$

[6] In the Appendix, we consider an alternative formulation in which the cost of competition appears as an operating cost − specifically, an additional labor requirement that firms must pay each period in order to produce. This specification generates increasing returns to scale and, in turn, a natural cost of expanding the number of firms. As we show, the main results of the paper are not driven by our baseline assumption, which is adopted solely for tractability.

where $Y(n)$ is the steady-state level of output:

$$\log Y(n) \equiv -\left(\frac{1}{\sigma+\varphi}\right)\{\sigma \log(1-\mathcal{G}(n)) + \log \mathcal{T} + \log \mathcal{M}(n)\}. \tag{12}$$

A social planner would recognize that part of output must be diverted to finance the cost of competition, which alters the optimal level of output. In the Appendix, we show that the efficient level of output is given by $\log Y_t^e = \log Y^e(n) + \left(\frac{1+\varphi}{\sigma+\varphi}\right)\log A_t$, where $\log Y^e(n) \equiv -\left(\frac{\sigma-1}{\sigma+\varphi}\right)\log(1-\mathcal{G}(n))$ is the (log) steady-state efficient level of output.
From the expressions above, it is immediate to verify that

$$\log\left(Y^e(n)/Y(n)\right) = \frac{1}{\sigma+\varphi}\{\log[(1-\mathcal{G}(n))\mathcal{T}] + \log \mathcal{M}(n)\},$$

so the subsidy $\mathcal{T} = (1-\mathcal{G}(n))^{-1}$ ensures that the *only* source of inefficiency in steady-state is the one associated with oligopolistic competition.

3 The Oligopolistic Phillips Curve

In this section, we first derived the Oligopolistic Phillips Curve from the firms' optimal behavior before providing empirical evidence on the relationship between its slope and concentration measures.

3.1 Oligopolistic Pricing and Aggregate Supply

We now study the first-order local dynamics around the inefficient steady state. In what follows, for any variable X_t, we define the log-deviation from the steady state as $x_t \equiv \log X_t/X(n)$.
We begin by analyzing firms' pricing behavior. In particular, we derive a log-linear approximation of the first-order condition governing optimal pricing. The resulting expres-

sion is presented in the following proposition.

Proposition 1 *Up to a first-order approximation around the inefficient steady state, firms' optimal reset price satisfies*

$$p_t^* = (1-\beta\theta)\left(\sum_{s=0}^{\infty}(\beta\theta)^s \mathbb{E}_t\left[\mu_{t+s|t}\right] + \sum_{s=0}^{\infty}(\beta\theta)^s \mathbb{E}_t\left[\psi_{t+s}+\tau_{t+s}\right]\right), \tag{13}$$

where

$$\mu_{t+s|t} = \eta(n)\, s_{t+s|t}, \tag{14}$$

denotes the desired markup of a firm resetting its price at time t, and $s_{t+s|t} = (\varepsilon-1)(p_{t+s}-p_t^*)$ is the corresponding market share. The elasticity of the desired markup with respect to market share is given by

$$\eta(n) \equiv \frac{n(\varepsilon-\omega)}{[n\varepsilon-(\varepsilon-\omega)]\left[n(\varepsilon-1)-(\varepsilon-\omega)\right]}, \tag{15}$$

which satisfies $\eta'(n) < 0$ and $\eta(n) \to 0$ as $n \to \infty$.

Proof *See Appendix A.1.*

The proposition shows that firms set their optimal price equal to the discounted stream of marginal costs (net of the production subsidy), as in the standard monopolistic competition framework, plus the expected discounted path of their desired markup. The latter term is crucial: it captures the strategic complementarities in pricing that dampen the pass-through of nominal shocks to prices.

To see this mechanism more clearly, consider a simple partial equilibrium exercise. Suppose $\tau_t = 0$ for all t, and assume that nominal marginal cost follows a Martingale process, so $\mathbb{E}_t[\psi_{t+s}] = \psi_t$ for every $s > 0$. Under these assumptions, the optimal reset price simplifies to $p_t^* = \psi_t + (1-\beta\theta)\sum_{s=0}^{\infty}(\beta\theta)^s \mathbb{E}_t\left[\mu_{t+s|t}\right]$. The pass-through of a nominal marginal cost shock can be computed as

$$\frac{dp_t^*}{d\psi_t} = \underbrace{1}_{\text{Mon. Competition}} + \underbrace{(1-\beta\theta)\sum_{s=0}^{\infty}(\beta\theta)^s \mathbb{E}_t\left[\frac{d\mu_{t+s|t}}{d\psi_t}\right]}_{\text{Imp. Competition}}, \tag{16}$$

where $d\mu_{t+s|t}/d\psi_t = -(\varepsilon-1)\eta(n)\, dp_t^*/d\psi_t < 0$ for all $s > 0$. In a standard monopolistic competition model, the pass-through would simply be equal to one. Here, however, strategic interactions matter. Because firms understand that raising their price reduces their market share – pushing them onto a more elastic region of the demand curve –

they choose to transmit only a fraction of the initial marginal cost increase into prices. As a result, pass-through is dampened and given by

$$\varrho(n) \equiv \frac{dp_t^*}{d\psi_t} = \frac{1}{1+(\varepsilon-1)\eta(n)} \in \left(\underline{\varrho}, 1\right) \tag{17}$$

with $\underline{\varrho} \equiv \omega(\omega-1) / [\omega(\omega-1)+(\varepsilon-1)(\varepsilon-\omega)]$. Since $\eta'(n) < 0$, it follows that $\varrho'(n) > 0$. Intuitively, as the number of firms increases, each firm's influence on the market becomes smaller, weakening the strategic effect described above. In the limit as $n \to \infty$, we have $\eta(n) \to 0$, so that $\varrho(n) \to 1$, and the model converges to the standard monopolistic competition benchmark.

The general-equilibrium impact of shocks is also attenuated since the firms' pass-through is dampened. This leads to a Phillips curve that is flatter than under the monopolistic competition benchmark. Moreover, since pass-through depends on the degree of competition in the goods market, the slope of the Phillips curve inherits this dependence as well. The following proposition formalizes this result.

Proposition 2 *To a first-order approximation around the inefficient steady state, equilibrium inflation dynamics satisfy*

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa(n) x_t + \kappa(n) u_t, \tag{18}$$

where $x_t \equiv y_t - y_t^e$ *denotes the welfare-relevant output gap, and* $u_t \equiv \tau_t/(\sigma+\varphi)$ *is a normalized cost-push shock. The function*

$$\kappa(n) \equiv \varrho(n)\kappa_m, \qquad \kappa_m \equiv \frac{(1-\theta)(1-\beta\theta)(\sigma+\varphi)}{\theta}, \tag{19}$$

is the slope of the oligopolistic Phillips curve, and κ_m *is the corresponding slope in the monopolistic competition limit. Moreover,* $\kappa'(n) < 0$ *and* $\kappa(n) \to \kappa_m$ *as* $n \to \infty$.

Proof *See Appendix A.2.*

The oligopolistic Phillips curve retains the familiar structure of the New Keynesian Phillips curve but differs in one crucial respect: its slope varies with the degree of product-market competition. In particular, it coincides with the standard monopolistic competition slope, scaled by the pass-through term $\varrho(n)$. Intuitively, κ_m captures the baseline nominal rigidities in the economy, as determined by the frequency of price adjustment. Because oligopolistic competition implies a pass-through below one, the overall responsiveness of inflation to economic slack is proportionally dampened. Consequently, as the number of firms increases and the pass-through approaches one, the Phillips curve becomes

progressively steeper. In the limit as $n \to \infty$, we have $\kappa(n) \to \kappa_m$: strategic interactions vanish and the model converges to the monopolistic competition benchmark.

Another important feature of this Phillips curve is that its slope also modulates the elasticity of inflation to a cost-push shock, given by $d\pi_t/du_t = \kappa(n)$. Intuitively, as the number of firms rises, each firm becomes more responsive in passing shocks through to prices. As a result, exogenous fluctuations in the labor wedge translate more strongly into movements in inflation.

3.2 Phillips Curve Slope and Concentration in the Data

This subsection provides reduced-form evidence that the slope of the Phillips curve is systematically related to local market concentration. The analysis follows the state-level Phillips curve approach in Hazell et al. (2022) and augments it with measures of concentration.

We use the state-level inflation series constructed by Hazell et al. (2022) from the BLS CPI Research Database (micro price quotes across geographic areas, available from 1978 onward). A key advantage of these data is the availability of sectoral price indices, in particular *non-tradable* inflation, which is more tightly linked to local economic conditions than headline CPI inflation in a monetary union. Our baseline outcome is quarterly non-tradable inflation, denoted $\pi^{NT}_{s,t}$ for state s and quarter t.

To proxy for local slack (output gap), we use the seasonally adjusted state unemployment rate $u_{s,t}$ from the BLS Local Area Unemployment Statistics. Compared to detrended state GDP, unemployment is available for all states, is less noisy, and is the standard slack proxy in the regional Phillips-curve literature.

We measure local concentration using Census Business Dynamics Statistics (BDS) employment shares by firm-size bins, merged at the state-year level. We report three standard indicators (constructed within non-tradables): CR1, CR2 and HHI. We discretize states into three groups—Low (L), Medium (M), and High (H) concentration—based on the distribution of the chosen indicator, and define the corresponding dummies $\mathbb{I}^M_s$ and $\mathbb{I}^H_s$ (with Low as the omitted category).

We estimate a heterogeneous-slope Phillips curve with state and time fixed effects:

$$\pi^{NT}_{s,t} = \alpha_s + \gamma_t + \left(\beta_L + \delta_M \mathbb{I}^M_s + \delta_H \mathbb{I}^H_s\right) u_{s,t} + \rho\, \pi^{NT}_{s,t-1} + \varepsilon_{s,t}, \tag{20}$$

where the state fixed-effects α_s captures time-invariant state heterogeneity and the time fixed-effect γ_t absorbs common aggregate shocks. The Phillips-curve slope in Low-concentration states is β_L, while the slopes in Medium and High concentration states are $\beta_L + \delta_M$ and

Table 1: Phillips Curve Slope and Concentration (Non-tradables)

	$CR1_{NT}$	$CR2_{NT}$	HHI_{NT}
δ_M	0.0328^{**}	0.0374^{***}	0.0208
δ_H	0.0528^{***}	0.0342	0.0463^{**}
Slope Low β_L	-0.0600^{***}	-0.0530^{***}	-0.0510^{**}
Slope Medium $(\beta_L + \delta_M)$	-0.0272^{*}	-0.0155	-0.0302^{*}
Slope High $(\beta_L + \delta_H)$	-0.0072	-0.0188	-0.0047
R^2 (within)	0.7819	0.7822	0.7821
Observations	4665	4665	4665

$\beta_L + \delta_H$, respectively.

Table 1 shows a clear pattern: the Phillips curve is *steeper* (more negative) in Low-concentration states and becomes substantially *flatter* as concentration rises. For unemployment-based slack, the estimated slope is around -0.05 to -0.06 in Low-concentration states, but it is close to zero in High-concentration states across measures. The positive estimates of δ_M and δ_H imply that higher concentration attenuates the inflationary response to slack. This reduced-form evidence is consistent with the model mechanism: more market power (lower competition) is associated with a flatter Phillips curve.

4 Welfare and Policy Game

The equilibrium in our model is determined by the interaction between the private economy and the authorities − namely, the monetary authority (m), the competition authority (a), and the institutional designer (d). In this section, we begin by defining the representative household's welfare, then specify the behavior of the authorities, describe the policy game, and finally provide a formal definition of equilibrium.

To simplify the intuition as much as possible, we impose two parametric restrictions.

Assumption 1 *The parameters σ and β satisfy $\sigma \to 1$ and $\beta \to 1$.*

First, we take $\sigma \to 1$, so that utility becomes logarithmic in consumption. Second, we consider the case $\beta \to 1$. Under this assumption, the Phillips curve approaches a "Lucas supply curve", bringing the problem as close as possible to the classic linear-quadratic framework (Barro and Gordon, 1983; Rogoff, 1985). We adopt this assumption in solving the policy game and state explicitly whenever it is invoked.[7]

[7] In the Appendix we show the result for the general model where these restrictions are not imposed.

4.1 Welfare

We begin this section by deriving a second-order approximation of the representative household's welfare function, which serves as both the social planner's objective and our measure of aggregate welfare.

Proposition 3 *When $\sigma \to 1$, to a second-order approximation around the inefficient steady state, and for small steady-state distortions, the welfare of the representative household can be written as*

$$\begin{aligned}\mathcal{W}_0\left(\{x_t,\pi_t\}_{t\in\mathbb{N}},n\right) &= \mathcal{U}(n)+\log\left(1-\mathcal{G}(n)\right)\\ &\quad-\frac{(1+\varphi)(1-\beta)}{\vartheta}\sum_{t=0}^{\infty}\beta^t\mathbb{E}_0\left[\frac{1}{2}\left(\vartheta x_t^2+\pi_t^2\right)-\vartheta\Delta(n)\,x_t\right],\end{aligned} \tag{21}$$

where $\vartheta \equiv \kappa_m/\varepsilon$ measures the weight the representative household places on the welfare-relevant output gap relative to inflation, and

$$\Delta(n) \equiv \log\left(Y^e(n)/Y(n)\right) = (1+\varphi)^{-1}\log\mathcal{M}(n) \tag{22}$$

captures the steady-state distortion arising from imperfect competition. Finally, the term

$$\mathcal{U}(n) \equiv -\frac{1}{1+\varphi}\left(\frac{1+\log\mathcal{M}(n)\,\mathcal{M}(n)}{\mathcal{M}(n)}\right), \qquad \mathcal{U}'(n)>0, \tag{23}$$

represents long-run welfare in a counterfactual economy with $\alpha(n)\to\infty$, i.e., with a infinitely effective competition authority.

Proof *See Appendix A.3.*

The proposition establishes that the welfare of the representative household can be separated into a long-run component and a component capturing business-cycle fluctuations.[8] The long-run component, in turn, can be decomposed into two parts: the term $\mathcal{U}(n)$, which reflects the long-run gains from reducing the steady-state inefficiencies associated with imperfect competition, and the term $\log\left(1-\mathcal{G}(n)\right)$, which captures the deadweight loss arising from the reduction in consumption required to sustain a given level of competition n. Notice that under our assumptions on $\mathcal{G}(n)$, we have $\log\left(1-\mathcal{G}(n)\right)\to-\infty$ as $n\to\infty$, implying that it is never optimal to push the economy toward the monopolistic competition limit.

[8] Notice that the long-run component also appears in the standard New Keynesian model, but it is typically treated as a constant because policy cannot generally influence it within that framework.

The term reflecting the welfare effects of business-cycle fluctuations is standard in both the New Keynesian and linear-quadratic approaches. Nominal rigidities imply that whenever output or inflation deviates from its flexible-price path, relative prices become distorted, leading to inefficient resource allocation across firms and, consequently, to welfare losses. In particular, observe that we can write $x_t = \log Y_t - \log Y_t^e + \Delta(n)$.

Setting output equal to the efficient level implies $x_t = \Delta(n)$, which generates a second-order misallocation loss captured by ϑx_t^2. Thus, to reduce the misallocation costs associated with business cycles, it must be that $\log Y_t < \log Y_t^e$. By contrast, because the steady state is inefficient, running a hot economy by setting $\log Y_t = \log Y_t^e$ yields a first-order welfare gain, represented by the term $\vartheta\Delta(n) x_t$. Importantly, this first-order gain is scaled by the magnitude of the steady-state inefficiency $\Delta(n)$. Hence, the competition authority can, in principle, influence the size of the gains from running the economy hot.

4.2 Mandates and Objectives of Authorities

As shown in Proposition 3, the trade-off between the output gap and inflation for the representative household is governed by the parameter ϑ: higher values of ϑ indicate that society places greater weight on stabilizing the output gap relative to inflation. However, the classic result in (Rogoff, 1985) shows that, depending on the central bank's ability to commit, it may be optimal to appoint a central banker who is more conservative than society – that is, one who assigns a lower weight than ϑ to output stabilization in their objective function. Motivated by this insight, we allow each authority to behave according to a generalized version of the welfare function in (21), in which the degree of conservatism may differ from that of society.

Formally, let $\mathcal{A} \equiv \{d, c, m\}$ denote the set of authorities, where d stands for the institutional designer, c for the competition authority, and m for the monetary authority. For any $a \in \mathcal{A}$, we define its objective function as

$$\begin{aligned}\mathcal{W}_0^a\left(\{x_t, \pi_t\}_{t\in\mathbb{N}}, n, \vartheta_a\right) \equiv\ & \mathcal{U}(n) + \log\left(1 - \mathcal{G}(n)\right) \\ & - \frac{(1+\varphi)(1-\beta)}{\vartheta_a} \sum_{t=0}^{\infty} \beta^t \mathbb{E}_0 \left[\frac{1}{2}\left(\vartheta_a x_t^2 + \pi_t^2\right) - \vartheta_a \Delta(n) x_t\right].\end{aligned} \tag{24}$$

In our policy game, the institutional designer is responsible for choosing the weights that the competition authority and the monetary authority assign to the output-inflation trade-off. A natural assumption, therefore, is that the designer's preferences reflect those

of society.

Assumption 2 *The institutional designer shares preferences with the representative household, that is* $\vartheta_d = \vartheta$.

Given this, we define a mandate as the pair of weights that the competition authority and the monetary authority are instructed to apply in their respective objective functions.

Definition 1 *A mandate is a pair* $(\vartheta_c, \vartheta_m) \in \mathbb{R}^2_+$ *such that the objective functions of the competition authority and the monetary authority are, respectively,* $\mathcal{W}_0^c(\cdot, n, \vartheta_c)$ *and* $\mathcal{W}_0^m(\cdot, n, \vartheta_m)$.

4.3 Game Structure and Equilibrium Definition

We are now ready to specify the structure of the policy game. To do so, we must take a stand on the authorities' ability to commit to their policy choices. Because the institutional designer and the competition authority make decisions at a much lower frequency than the monetary authority, it is natural to assume that they move first and can commit to their choices. As a result, the central bank takes both its mandate and the prevailing degree of competition as given and conducts policy under discretion. Likewise, it is natural to assume that the institutional designer moves before the competition authority, so that the latter treats its mandate as given.

Formally, the game proceeds in three stages involving the institutional designer d, the competition authority c, and the monetary authority m. Define the state at time t as $\mathbf{s}_t \equiv (u_t, r_t^e)$. The vector of the monetary authority's policy functions $\boldsymbol{\phi}_m(\mathbf{s}_t, \vartheta_m, n)$ map this state into (x_t, π_t, i_t). Similarly, let $\phi_c(\vartheta_c, \vartheta_m)$ denote the competition authority's policy function determining the degree of competition. Finally, we denote with $\boldsymbol{\vartheta}^* = (\vartheta_c^*, \vartheta_m^*)$ the institutional designer's optimal choice. With this, we can specify the timing of the game as follows.

Stage 1 (Institutional Designer). The institutional designer chooses the mandate $(\vartheta_c, \vartheta_m)$ under commitment, internalizing the induced behavior of both the competition and monetary authorities. Formally, the optimal mandate solves

$$\begin{aligned} \boldsymbol{\vartheta}^* \equiv \arg \max_{(\vartheta_c, \vartheta_m) \in \mathbb{R}^2_+} & \{\mathcal{W}_0(\mathbf{x}_0, \boldsymbol{\pi}_0, n, \vartheta)\} \\ \text{s.t. } & n = \phi_c(\vartheta_c, \vartheta_m), \\ & (x_t, \pi_t, i_t) = \boldsymbol{\phi}_m(\mathbf{s}_t, \vartheta_m, n) \text{ for all } t. \end{aligned} \tag{25}$$

Stage 2 (Competition Authority). Given the mandate $(\vartheta_c, \vartheta_m)$, the competition authority chooses the degree of competition n under commitment, internalizing the monetary

authority's policy function. Its problem is

$$\begin{aligned}\phi_c\left(\vartheta_c,\vartheta_m\right) \equiv \arg\max_{n>1}\left\{\mathcal{W}_0^c\left(\mathbf{x}_0,\boldsymbol{\pi}_0,n,\vartheta_c\right)\right\} \\ \text{s.t. } \left(x_t,\pi_t,i_t\right)=\boldsymbol{\phi}_m\left(\mathbf{s}_t,\vartheta_m,n\right) \text{ for all } t.\end{aligned} \tag{26}$$

Stage 3 (Monetary Authority). This stage is determined jointly by private agents' expectation formation and the monetary authority's behavior. Specifically, at each $t \in \mathbb{N}$ private agents observe the state $\mathbf{s}_t$ and form expectations according to

$$\mathbb{E}_t\left[x_{t+1}\right] = \int \phi_{m,x}\left(\mathbf{s}_{t+1},\vartheta_m,n\right) d\mathcal{F}\left(\mathbf{s}_{t+1} \mid \mathbf{s}_t\right), \tag{27}$$

$$\mathbb{E}_t\left[\pi_{t+1}\right] = \int \phi_{m,\pi}\left(\mathbf{s}_{t+1},\vartheta_m,n\right) d\mathcal{F}\left(\mathbf{s}_{t+1} \mid \mathbf{s}_t\right), \tag{28}$$

where $\mathcal{F}\left(\mathbf{s}_{t+1} \mid \mathbf{s}_t\right)$ denotes the conditional joint distribution of shocks. Given $\mathbf{s}_t$, and taking expectations as given, the monetary authority chooses $\left(x_t,\pi_t,i_t\right)$ under discretion to solve

$$\begin{aligned}\boldsymbol{\phi}_m\left(\mathbf{s}_t,\vartheta_m,n\right) \equiv \arg\max_{\left(x_t,\pi_t,i_t\right)}\left\{\mathcal{W}_t^m\left(\mathbf{x}_t,\boldsymbol{\pi}_t,n,\vartheta_m\right)\right\} \\ \text{s.t. } x_t=\mathbb{E}_t\left[x_{t+1}\right]-\left(i_t-\mathbb{E}_t\left[\pi_{t+1}\right]-r_t^e\right), \\ \pi_t=\beta\mathbb{E}_t\left[\pi_{t+1}\right]+\kappa\left(n\right)x_t+\kappa\left(n\right)u_t.\end{aligned} \tag{29}$$

An equilibrium consists of an optimal mandate $\boldsymbol{\vartheta}^*$, an induced degree of competition $n^* = \phi_c(\vartheta_c^*,\vartheta_m^*)$, and sequences $\{x_t,\pi_t,i_t,\mathbb{E}_t[x_{t+1}],\mathbb{E}_t[\pi_{t+1}]\}_{t\in\mathbb{N}}$ such that the institutional designer and the competition authority act optimally, the monetary authority behaves optimally under discretion at each date and state, and private-sector expectations are rational.

5 Results

In this section, we characterize the equilibrium of our economy by proceeding via backward induction: we begin with the central bank's problem, then analyze the competition authority's choice, and finally determine the optimal mandate set by the institutional designer. Before turning to this analysis, it is useful to lay out additional assumptions and present the baseline calibration used in the numerical exercises, beginning with the functional form for $\mathcal{G}\left(n\right)$. We assume that the cost of competition function $\mathcal{G}\left(n\right)$ takes the form: $\mathcal{G}\left(n\right)=1-n^{-1/\alpha}$ with $\alpha>0$. We adopt this specification for convenience. It is straightforward to verify that the elasticity $\alpha\left(n\right) \equiv \left(\left|\frac{d\log(1-\mathcal{G}(n))}{d\log n}\right|\right)^{-1} = \alpha$, is constant and therefore independent of the degree of competition. This property greatly simplifies the

comparative statics that follow. In addition, since divine coincidence holds, our baseline analysis focuses on an economy subject only to cost-push shocks. Specifically, we assume $\sigma_a \to 0$, which implies $r_t^e \to \rho$. Finally, Table 2 reports our baseline calibration.

Table 2: Baseline Calibration

Parameter	Value	Parameter	Value
φ	5	α	10
ε	9	ρ_u	0.5
ω	3	σ_u	0.1
θ	0.6		

5.1 Monetary Policy

We start by characterizing the monetary authority's behavior. Since the interest rate i_t is fully determined once x_t and π_t are chosen, it can be eliminated as a separate control variable, and the IS curve does not enter as a constraint.[9] Moreover, because the central bank takes private sector expectations as given and the model contains no predetermined state variables, the policy problem is static. Taking the first-order conditions and combining them, we obtain the central bank's optimal intratemporal trade-off between output gap and inflation:

$$\vartheta_m x_t = \vartheta_m \Delta(n) - \kappa(n)\,\pi_t. \tag{30}$$

The left-hand side captures the marginal cost of output gap, while the right hand side captures its marginal benefit. Importantly, this marginal benefit depends on both the degree of goods market competition, and the central bank's mandate. The degree of goods market competition can thus influence the central bank's trade-off, and it does so through two channels. First, by altering long-run distortions: higher n reduces the steady-state inefficiency $\Delta(n)$, thereby lowering the marginal benefit of a positive output gap. Second, by affecting pass-through and therefore the slope of the Phillips curve. A higher n steepens the Phillips curve, making the output gap more effective in stabilizing inflation and increasing its marginal benefit. This incentive structure, combined with private agents' expectations, yields the solution of the third-stage policy game. The next

[9] The IS curve, obtained by log-linearizing the consumption Euler equation and imposing goods market clearing. Expressed in terms of the welfare-relevant output gap, it takes the form $x_t = \mathbb{E}_t[x_{t+1}] - \sigma^{-1}(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^e)$, where r_t^e denotes the flexible-price real rate of interest associated with the efficient allocation, $r_t^e \equiv \rho + \sigma \mathbb{E}_t[\Delta y_{t+1}^e]$.

proposition characterizes this solution under the assumption $\beta \to 1$.

Proposition 4 *Under assumptions 1, given a mandate ϑ_m and competition n, the monetary authority reaction function (third stage of the policy game) is given by*

$$\begin{aligned} \pi_t &= \overline{\pi}(n, \vartheta_m) + [1 - \delta_x(n, \vartheta_m)] \frac{\kappa(n)}{(1-\rho_u)} u_t, \\ x_t &= -\delta_x(n, \vartheta_m) u_t, \end{aligned}$$

where

$$\delta_x(n, \vartheta_m) \equiv \frac{\kappa(n)^2}{\vartheta_m(1-\rho_u) + \kappa(n)^2} \text{ with } \frac{\partial \delta_x(n, \vartheta_m)}{\partial n} > 0, \text{ and } \frac{\partial \delta_x(n, \vartheta_m)}{\partial \vartheta_m} < 0,$$

is the share of the cost-push shock absorbed by the output gap, and

$$\overline{\pi}(n, \vartheta_m) \equiv \vartheta_m \frac{\Delta(n)}{\kappa(n)} \text{ with } \frac{\partial \overline{\pi}(n, \vartheta_m)}{\partial n} < 0, \text{ and } \frac{\partial \overline{\pi}(n, \vartheta_m)}{\partial \vartheta_m} > 0,$$

is the inflation bias.

Proof *See Appendix A.4.*

The central bank's policy functions can be decomposed into a long-run component, capturing the inflation bias, and a short-run component, reflecting stabilization motives. Figure 1 provides a graphical illustration of these policy functions for different levels of competition: the solid lines correspond to a baseline level n, while the dashed lines correspond to a higher level $n' > n$.

Competition and Central Bank's Incentives. As discussed above, greater product-market competition has two complementary effects on the discretionary inflation bias. First, it reduces the steady-state inefficiency gap, lowering the central bank's incentive to create inflation in order to push output above its natural level. Second, it steepens the Phillips curve, which raises the inflation cost of any attempt to stimulate activity. When inflation becomes both less desirable and more costly to use as a stabilization tool, the policymaker's temptation to inflate diminishes. Consequently, the inflation bias, $\overline{\pi}(n, \vartheta_m)$, is a decreasing function of the degree of competition. This is illustrated in Figure 1 where the policy function for n' features a lower intercept for π_t, reflecting the reduction in the discretionary inflation bias as competition rises.

Turning to the stabilization component of the policy functions, note that – using the

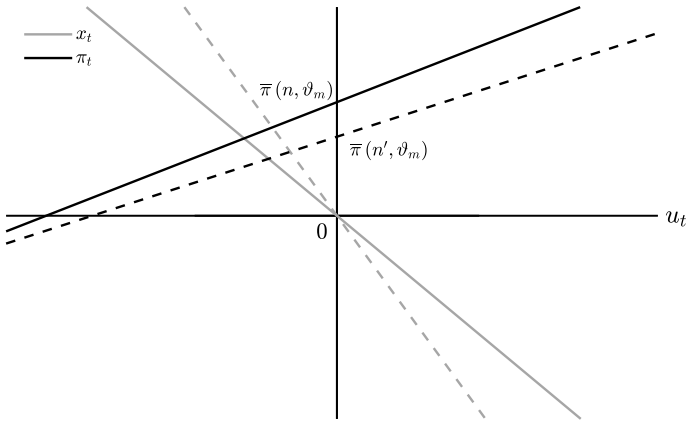


Figure 1: Central Bank's Policy Functions for Different Competition Levels ($n < n'$).

Phillips curve – we can decompose the cost-push shock as follows:

$$u_t = \underbrace{\left[1 - \delta_x(n, \vartheta_m)\right] u_t}_{=\pi_t - \mathbb{E}_t[\pi_{t+1}]} + \underbrace{\delta_x(n, \vartheta_m)\, u_t}_{=-\kappa(n) x_t}.$$

Thus, the central bank optimally absorbs a fraction $\delta_x(n, \vartheta_m)$ of the shock through adjustments in the output gap, and the remaining fraction $1 - \delta_x(n, \vartheta_m)$ through (net) inflation. As shown in Proposition 4, the share $\delta_x(n, \vartheta_m)$ is decreasing in n. Intuitively, as n increases and the Phillips curve becomes steeper, the marginal benefit of using the output gap as a stabilization instrument rises. The central bank therefore finds it optimal to rely more heavily on output-gap movements and less on inflation to absorb shocks. This mechanism shapes the policy function for x_t, implying higher n is associated with a higher output gap volatility, as shown in Figure 1 where the policy function becomes steeper for $n' > n$.

For the inflation policy function, two opposing channels must be considered. On the one hand, as n increases and the central bank relies more heavily on the output gap for stabilization, the share $1 - \delta_x(n, \vartheta_m)$ declines, reducing the fraction of the shock that is absorbed by inflation. On the other hand, a higher n steepens the Phillips curve, making inflation more sensitive to a given cost-push disturbance.

It is straightforward to show that for low values of n, the second effect dominates: when the output gap is an ineffective stabilization instrument, a marginal increase in n barely changes the central bank's incentives, and the direct steepening of the Phillips curve makes inflation more reactive. By contrast, once n is sufficiently high and the output

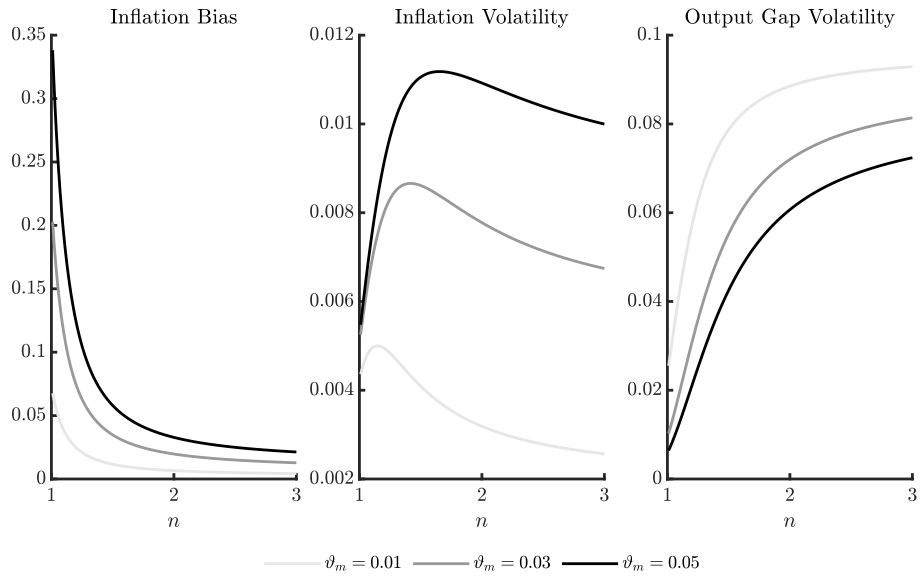


Figure 2: Inflation Bias and Volatility as a Function of n.

gap becomes an effective tool, a marginal increase in n induces the central bank to shift strongly toward output-gap adjustment, so the first effect prevails and inflation responds less to shocks. In Figure 1 we illustrate the case in which this latter effect dominates: moving from n to n' flattens the policy function for inflation.

The Role of the Central Bank's Mandate. The central bank's long- and short-run trade-offs are likewise influenced by the mandate it is required to pursue. Proposition 4, as well as the standard result in (Rogoff, 1985), shows that the inflation bias increases in ϑ_m: as the central banker becomes more dovish $-$ placing greater weight on the output gap $-$ its incentive to generate inflation to reduce the inefficiency gap rises, leading to higher long-run inflation. This relationship is illustrated in the left panel of Figure 2, which shows that higher values of ϑ_m are associated with systematically larger inflation biases for every level of n.

Regarding the effect of the mandate on the short-run stabilization trade-off, the middle and right panels of Figure 2 plot the volatilities of inflation and the output gap implied by the third stage of the policy game.[10] These volatilities are shown as functions of n for different mandates. Notice that an increase in ϑ_m increases the volatility of inflation and reduces the volatility of output gap. Moreover, the mandate also shifts the point at which inflation volatility becomes a decreasing function of n. Intuitively, a more dovish

[10]Notice that the volatilities of inflation and the output gap are simply given by the absolute value of the slope of their respective policy functions. Specifically,$\sigma_\pi \equiv \sqrt{\mathbb{V}[\pi_t]} = [1-\delta_x(n,\vartheta_m)]\frac{\kappa(n)}{(1-\rho_u)}\sigma_u$ and $\sigma_x \equiv \sqrt{\mathbb{V}[x_t]} = \delta_x(n,\vartheta_m)\sigma_u$.

central banker requires a highly effective output gap channel (i.e., a relatively larger initial n) before a marginal increase in n induces a reallocation of volatility from inflation to the output gap strong enough to offset the higher inflation volatility generated by a stronger cost-push shock. In the middle panel, this is illustrated by the rightward shift of the curve's maximum as ϑ_m increases. In the left panel, it is further reflected in the fact that, for any given n, the slope of the curve is systematically lower when ϑ_m is higher.

5.2 Competition Policy

We now examine the second stage of the game, in which we characterize the behavior of the competition authority. Substituting the central bank's policy functions into the authority's objective function allows us to rewrite the problem as

$$\begin{aligned}\mathcal{W}_0^c(n,\vartheta_c) = \mathcal{U}(n) + \log(1-\mathcal{G}(n)) \\ - \frac{(1+\varphi)}{\vartheta_c}\left(\vartheta_c\sigma_x^2(n,\vartheta_m) + \sigma_\pi^2(n,\vartheta_m) + \overline{\pi}(n,\vartheta_m)^2\right),\end{aligned}$$

where $\sigma_x^2(n,\vartheta_m)$ and $\sigma_\pi^2(n,\vartheta_m)$ denote the output-gap and inflation volatilities implied by the monetary authority's behavior in the third stage. The competition authority chooses the degree of competition n to maximize $\mathcal{W}_0^c(n,\vartheta_c)$. The optimal choice, denoted $n^* = \phi_c(\vartheta_c,\vartheta_m)$, satisfies the first-order condition

$$\mathcal{U}'(n) - \frac{\mathcal{G}'(n)}{1-\mathcal{G}(n)} - \frac{(1+\varphi)}{\vartheta_c}\left(\vartheta_c\frac{\partial\sigma_x^2(n,\vartheta_m)}{\partial n} + \frac{\partial\sigma_\pi^2(n,\vartheta_m)}{\partial n} + \frac{\partial\overline{\pi}(n,\vartheta_m)^2}{\partial n}\right) = 0. \quad (31)$$

Figure 3 displays the competition authority's objective function under the baseline calibration (solid black line), along with the optimal degree of competition implied by the first-order condition (black dot).

As discussed in the previous section, the degree of competition influences the central bank's policy trade-off. The competition authority internalizes this interaction and optimally adjusts the level of competition in response to the central bank's mandate. The left panel of Figure 3 illustrates this relationship through a comparative static exercise in which we vary the central bank's mandate. As shown, there is a positive relationship between ϑ_m and n^*: when the central bank is more dovish − that is, when it places a relatively higher weight on output stabilization − its behavior generates a larger inflation bias. Anticipating this, the competition authority responds by pursuing a stricter competition policy, thereby increasing n^* to mitigate the inflationary consequences of the central bank's mandate.

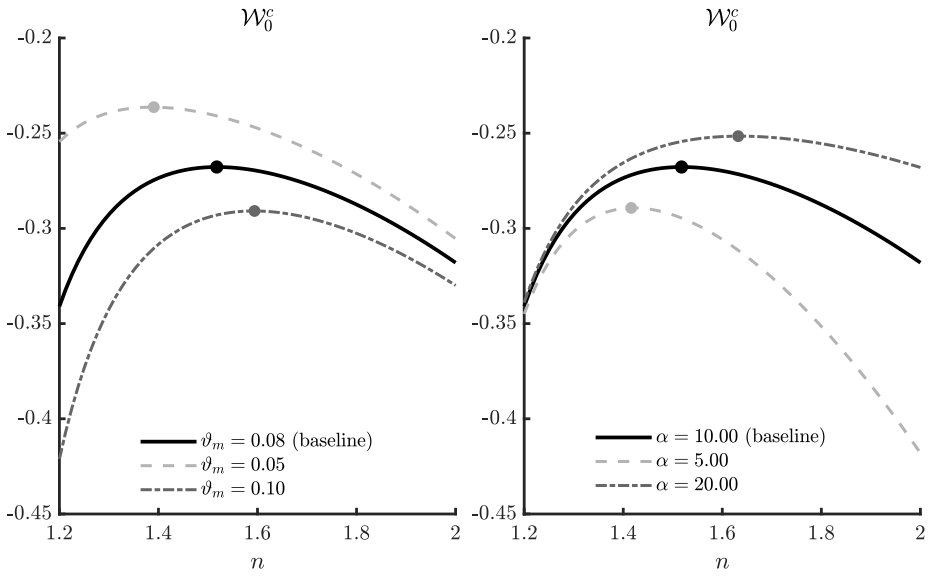


Figure 3: Optimal Competition Policy − Comparative Statics (baseline: $\vartheta_m = 0.08$, $\alpha = 10$).

In addition, a higher α implies that the competition authority is more effective at mitigating the resource costs associated with competition. Consequently, economies with a more effective competition authority have stronger incentives to pursue a stricter competition policy. This is illustrated in the right panel of Figure 3. Starting from the baseline calibration, an increase in α raises the optimal degree of competition, whereas a decrease in α leads the authority to choose a lower level of competition.

5.3 Institutional Design: License to Ease

We now turn to the question of optimal institutional design after having analyzed the interaction between the competition and monetary authorities in the third and second stages of the game above. In particular, we ask: what mandate should a benevolent social planner assign to each authority, taking into account their equilibrium interaction?

A first important observation is that, because the competition authority operates under commitment, its optimal mandate should naturally align with social preferences. To see this, note that, using the competition authority's policy function, social welfare can be written as

$$\begin{aligned}\mathcal{W}_0\left(n^*, \vartheta_c\right) = \mathcal{U}\left(n^*\right) + \log\left(1 - \mathcal{G}\left(n^*\right)\right) \\ - \frac{(1+\varphi)}{\vartheta}\left(\vartheta\sigma_x^2\left(n^*, \vartheta_m\right) + \sigma_\pi^2\left(n^*, \vartheta_m\right) + \overline{\pi}\left(n^*, \vartheta_m\right)^2\right),\end{aligned} \tag{32}$$

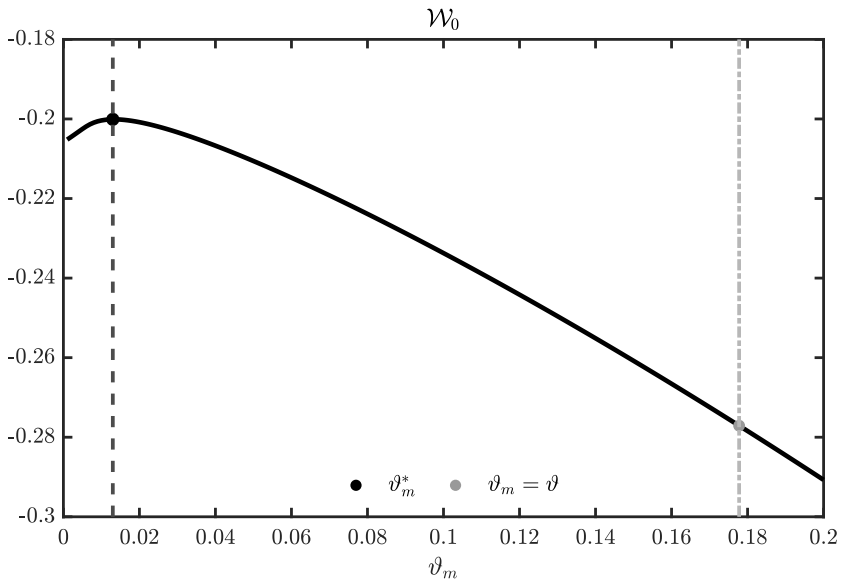


Figure 4: Optimal Monetary Authority Mandate – Baseline Calibration

where $n^* = \phi_c(\vartheta_c, \vartheta_m)$. The institutional designer chooses $(\vartheta_c, \vartheta_m)$ to maximize $\mathcal{W}_0(n^*, \vartheta_c)$, as defined in (32). The first-order conditions with respect to ϑ_m and ϑ_c are, respectively,

$$\left[\mathcal{U}' - \frac{\mathcal{G}'}{1-\mathcal{G}} - \frac{(1+\varphi)}{\vartheta}\left(\vartheta\frac{\partial\sigma_x^2}{\partial n} + \frac{\partial\sigma_\pi^2}{\partial n} + \frac{\partial\overline{\pi}^2}{\partial n}\right)\right]\frac{\partial n^*}{\partial\vartheta_m} - \frac{(1+\varphi)}{\vartheta}\left(\vartheta\frac{\partial\sigma_x^2}{\partial\vartheta_m} + \frac{\partial\sigma_\pi^2}{\partial\vartheta_m} + \frac{\partial\overline{\pi}^2}{\partial\vartheta_m}\right) = 0,$$

$$\left[\mathcal{U}' - \frac{\mathcal{G}'}{1-\mathcal{G}} - \frac{(1+\varphi)}{\vartheta}\left(\vartheta\frac{\partial\sigma_x^2}{\partial n} + \frac{\partial\sigma_\pi^2}{\partial n} + \frac{\partial\overline{\pi}^2}{\partial n}\right)\right]\frac{\partial n^*}{\partial\vartheta_c} = 0,$$

where, for notational convenience, all functions of (n^*, ϑ_m) are written without explicit arguments.

Using the competition authority's first-order condition (31), the first-orer condition of the institutional designer with respect to ϑ_c can be rewritten as $\frac{1}{\vartheta_c}\left(\frac{\partial\sigma_\pi^2}{\partial n} + \frac{\partial\overline{\pi}^2}{\partial n}\right) = \frac{1}{\vartheta}\left(\frac{\partial\sigma_\pi^2}{\partial n} + \frac{\partial\overline{\pi}^2}{\partial n}\right)$, which immediately implies $\vartheta_c^* = \vartheta$. This yields the following proposition.

Proposition 5 *The competition authority's optimal mandate requires its objective function to coincide with social welfare. In particular, the optimal mandate satisfies* $\vartheta_c^* = \vartheta$.

With this result in hand, we now turn to the determination of the optimal mandate for the monetary authority. Figure 4 plots social welfare – the institutional designer's objective – as a function of ϑ_m. The dashed light-gray line and its corresponding dot represent the case in which the monetary authority's mandate is perfectly aligned with social preferences, i.e., $\vartheta_m = \vartheta$. Because the monetary authority operates under discretion, however, the optimal mandate calls for appointing a central banker who is more conservative than

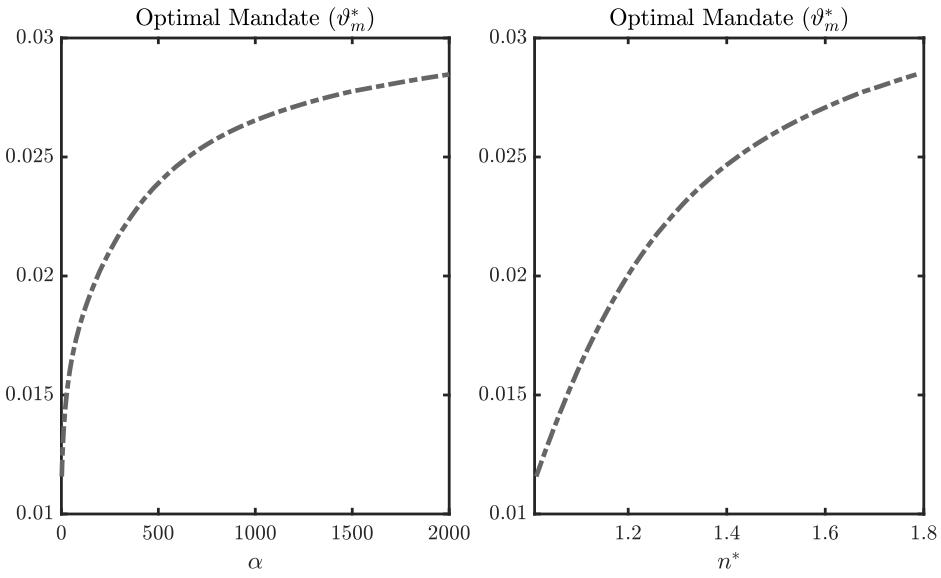


Figure 5: Optimal Mandate and Implied Degree of Competition for Different Competition Authority Effectiveness

society – namely, one who places a higher relative weight on inflation stabilization than implied by social preferences, in line with Rogoff's (1985) classic result. The key question is whether the effectiveness of the competition authority influences the optimal degree of conservatism of the central banker. Put differently, can economies with systematically more effective competition authorities afford to appoint a less conservative central banker?

The left panel of Figure 5 plots the optimal monetary mandate as a function of the competition authority's effectiveness, α. As shown, as the competition authority becomes more effective, the optimal mandate indeed calls for a less conservative central banker. The intuition is straightforward: a more effective competition authority implements a stricter competition policy, resulting in a higher degree of competition in the economy. This reduces the central bank's temptation to use monetary policy to systematically offset the efficiency gap, thereby limiting the inflation bias. As a result, the institutional designer can assign a higher weight to output-gap stabilization. In other words, an effective competition authority serves as a "license to ease", allowing the monetary authority to place relatively more emphasis on stabilizing output and employment. As illustrated in the right panel of Figure 5, this mechanism implies that, in equilibrium, more competitive economies should be associated with less conservative monetary mandates. Conversely, our theory indicates that economies with higher market concentration and larger markups should optimally be assigned more conservative central bankers.

6 Conclusion

This paper studies how competition policy and monetary policy should be designed when the competitive environment in which firms set prices is itself shaped by policy. In a New Keynesian model with oligopolistic competition, a competition authority chooses the degree of product-market competition, while a central bank conducts monetary policy under discretion. In this environment, competition is not merely a background characteristic of the economy. By shaping firms' markups and strategic pricing incentives, it changes the slope of the Phillips curve, the inflation–output trade-off faced by the central bank, and the long-run distortion that gives rise to inflation bias.

This observation has a clear institutional implication. Once competition is treated as a policy instrument, the design of monetary delegation cannot be studied in isolation from product-market policy. Competition policy affects the environment in which discretionary monetary policy operates, and therefore changes the mandate that should be assigned to the central bank. Our analysis shows that the optimal mandate for the competition authority is to align its objective with social welfare. By contrast, the optimal degree of central bank conservatism depends systematically on the effectiveness of competition policy. When competition can be fostered more effectively, the equilibrium degree of competition is higher, the Phillips curve is steeper, the inflation bias is lower, and it becomes optimal to appoint a less conservative central bank. Conversely, when competition is weaker or more costly to promote, the case for a more conservative central bank becomes stronger.

The broader message of the paper is therefore that competition policy and monetary policy are linked not only through outcomes, but through institutional design. Product-market policy has macroeconomic consequences because it shapes the transmission of monetary policy and the incentives of a discretionary central bank. In this sense, preserving and promoting competition is not only a microeconomic objective aimed at reducing markups and distortions; it is also part of the institutional architecture that supports macroeconomic stabilization.

More generally, our framework suggests that changes in market structure should have implications for the design of monetary institutions. Economies characterized by rising concentration and weaker competition may require more conservative monetary delegation, while economies able to sustain stronger competition can afford a more flexible and employment-oriented monetary policy. The central implication is that the mandates of the competition authority and the central bank should be chosen jointly rather than in isolation.

References

Afrouzi, Hassan, Marina Halac, Kenneth S Rogoff, and Pierre Yared, "Monetary policy without commitment," Technical Report, National Bureau of Economic Research 2025.

Athey, Susan, Andrew Atkeson, and Patrick J Kehoe, "The optimal degree of discretion in monetary policy," *Econometrica*, 2005, *73* (5), 1431–1475.

Atkeson, Andrew and Ariel Burstein, "Pricing-to-market, trade costs, and international relative prices," *American Economic Review*, 2008, *98* (5), 1998–2031.

Backus, David and John Driffill, "Inflation and reputation," *The American Economic Review*, 1985, *75* (3), 530–538.

Baqaee, David R, Emmanuel Farhi, and Kunal Sangani, "The supply-side effects of monetary policy," *Journal of Political Economy*, 2024, *132* (4), 1065–1112.

Barro, Robert J and David B Gordon, "A positive theory of monetary policy in a natural rate model," *Journal of political economy*, 1983, *91* (4), 589–610.

Benigno, Pierpaolo and Michael Woodford, "Optimal monetary and fiscal policy: A linear-quadratic approach," *NBER macroeconomics annual*, 2003, *18*, 271–333.

_ **and** _ , "Inflation Stabilization and Welfare: The Case of a Distorted Steady State," *Journal of the European Economic Association*, 2005, *3* (6), 1185–1236.

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

Canzoneri, Matthew B, "Monetary policy games and the role of private information," *The American Economic Review*, 1985, *75* (5), 1056–1070.

Clarida, Richard, Jordi Gali, and Mark Gertler, "The science of monetary policy: a new Keynesian perspective," *Journal of economic literature*, 1999, *37* (4), 1661–1707.

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland, "The optimal inflation rate in New Keynesian models: should central banks raise their inflation targets in light of the zero lower bound?," *Review of Economic Studies*, 2012, *79* (4), 1371–1406.

Cukierman, Alex and Allan H Meltzer, "A theory of ambiguity, credibility, and inflation under discretion and asymmetric information," *Econometrica: journal of the econometric society*, 1986, pp. 1099–1128.

Dixit, Avinash and Luisa Lambertini, "Interactions of Commitment and Discretion in Monetary and Fiscal Policies," *American Economic Review*, 2003, *93* (5), 1522–1542.

Etro, Federico and Lorenza Rossi, "New-Keynesian Phillips curve with Bertrand competition and endogenous entry," *Journal of Economic Dynamics and Control*, 2015, *51*, 318–340.

Gagliardone, Luca, Mark Gertler, Simone Lenzu, and Joris Tielens, "Anatomy of the Phillips curve: micro evidence and macro implications," *American Economic Review*, 2025, *115* (11), 3941–3974.

Galí, Jordi, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press, 2015.

Gali, Jordi and Tommaso Monacelli, "Monetary policy and exchange rate volatility in a small open economy," *The Review of Economic Studies*, 2005, *72* (3), 707–734.

Halac, Marina and Pierre Yared, "Commitment versus flexibility with costly verification," *Journal of Political Economy*, 2020, *128* (12), 4523–4573.

_ and _, "Fiscal rules and discretion under limited enforcement," *Econometrica*, 2022, *90* (5), 2093–2127.

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

Meier, Matthias and Timo Reinelt, "Monetary policy, markup dispersion, and aggregate tfp," *Review of Economics and Statistics*, 2024, *106* (4), 1012–1027.

Mongey, Simon, "Market structure and monetary non-neutrality," Technical Report, National Bureau of Economic Research 2021.

Olivi, Alan, Vincent Sterk, and Dajana Xhani, "Optimal monetary policy during a cost of living crisis," Technical Report, mimeo, UCL. 1 2023.

Rogoff, Kenneth, "The optimal degree of commitment to an intermediate monetary target," *The quarterly journal of economics*, 1985, *100* (4), 1169–1189.

Rotemberg, Julio J and Michael Woodford, "An optimization-based econometric framework for the evaluation of monetary policy," *NBER macroeconomics annual*, 1997, *12*, 297–346.

Svensson, Lars E. O., "Optimal Inflation Targets, "Conservative" Central Banks, and Linear Inflation Contracts," *American Economic Review*, 1997, *87* (1), 98–114.

Walsh, Carl E., "Optimal Contracts for Central Bankers," *American Economic Review*, 1995, *85* (1), 150–167.

Wang, Olivier and Ivan Werning, "Dynamic Oligopoly and Price Stickiness," *American Economic Review*, 2022, *112* (8).

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2011.

A Theory Appendix

A.1 Proof of Proposition 1

Start from equation (9) and write it as

$$\mathbb{E}_t\left[\sum_{s=0}^{\infty}(\theta\beta)^s \underbrace{\exp\left\{\log X_{t+s|t}+\log P_t^*\right\}}_{\text{(a)}}\right]$$

$$=$$

$$\mathbb{E}_t\left[\sum_{s=0}^{\infty}(\theta\beta)^s \underbrace{\exp\left\{\log X_{t+s|t}+\log \mathcal{M}_{t+s|t}+\log \mathcal{T}_{t+s}+\log \Psi_{t+s}\right\}}_{\text{(b)}}\right]$$

where $X_{t+s|t} \equiv \frac{u_c(C_{t+s},L_{t+s})}{u_c(C_t,L_t)}\frac{Y_{t+s|t}\left(\varpi_{t+s|t}-1\right)}{P_{t+s}}$. Up to a first order, we can write

$$\begin{aligned}
\text{(a)} &\simeq XP\left(1+x_{t+s|t}+p_t^*\right),\\
\text{(b)} &\simeq X\mathcal{M}\mathcal{T}\Psi\left(1+x_{t+s|t}+\mu_{t+s|t}+\tau_{t+s}+\psi_{t+s}\right).
\end{aligned}$$

Using this in the main equation and noting that $P=\mathcal{M}\mathcal{T}\Psi$ we end-up with

$$\mathbb{E}_t\left[\sum_{s=0}^{\infty}(\theta\beta)^s x_{t+s|t}+p_t^*\right]=\mathbb{E}_t\left[\sum_{s=0}^{\infty}(\theta\beta)^s \mathcal{M}\mathcal{T}\Psi\left(x_{t+s|t}+\mu_{t+s|t}+\tau_{t+s}+\psi_{t+s}\right)\right],$$

or, rearranging

$$p_t^*=(1-\beta\theta)\left[\sum_{s=0}^{\infty}(\beta\theta)^s\mathbb{E}_t\left[\mu_{t+s|t}\right]+\sum_{s=0}^{\infty}(\beta\theta)^s\mathbb{E}_t\left[\tau_{t+s}+\psi_{t+s}\right]\right].$$

Using the definitions of markup and market share, we have

$$\mathcal{M}_{t+s|t}\equiv \varpi_{t+s|t}/\left(\varpi_{t+s|t}-1\right),\ \varpi_{t+s|t}=\varepsilon-(\varepsilon-\omega)S_{t+s|s},\text{ and } S_{t+s|s}=\left(\frac{P_t^*}{P_{t+s}}\right)^{-\varepsilon}.$$

Taking a log-linear approximation and combining we obtain

$$\mu_{t+s|t} = \eta(n)\, s_{t+s|t},$$

where $s_{t+s|t} = (\varepsilon - 1)(p_{t+s} - p_t^*)$ and

$$\eta(n) := \frac{n(\varepsilon - \omega)}{[n\varepsilon - (\varepsilon - \omega)][n(\varepsilon - 1) - (\varepsilon - \omega)]}.$$

A.2 Proof of Proposition 2

A first-order approximation of the price index around the steady state is

$$p_t = (1 - \theta)\, p_t^* + \theta p_{t-1}.$$

In addition, combining the conditional markup with the conditional market share derived in Appendix A.1, we obtain

$$\mu_{t+s|t} = \eta(n)(\varepsilon - 1)(p_{t+s} - p_t^*).$$

Combining this expression with equation (13) yields

$$p_t^* = (1 - \beta\theta)\, \varrho(n) \left(\sum_{s=0}^{\infty} (\beta\theta)^s \frac{1}{\varrho(n)} \mathbb{E}_t [p_{t+s}] + \sum_{s=0}^{\infty} (\beta\theta)^s \mathbb{E}_t [mc_{t+s} + \tau_{t+s}] \right),$$

where $\varrho(n) \equiv (1 + \eta(n)(\varepsilon - 1))^{-1}$ denotes pass-through, and $mc_{t+s} := \psi_{t+s} - p_{t+s}$ denotes the real marginal cost. Substituting this expression into the recursion for the price index, we obtain

$$p_t = (1 - \theta)(1 - \beta\theta) \sum_{s=0}^{\infty} (\beta\theta)^s \mathbb{E}_t [p_{t+s}] + (1 - \theta)(1 - \beta\theta)\, \varrho(n) \sum_{s=0}^{\infty} (\beta\theta)^s \mathbb{E}_t [mc_{t+s} + \tau_{t+s}] + \theta p_{t-1},$$

which can be written recursively as

$$p_t = (1 - \theta)(1 - \beta\theta)\, p_t + \beta\theta \mathbb{E}_t [p_{t+1}] + \theta p_{t-1} - \beta\theta^2 p_t + (1 - \theta)(1 - \beta\theta)\, \varrho(n)(mc_t + \tau_t).$$

Defining the inflation rate as $\pi_t \equiv p_t - p_{t-1}$ and rearranging, we obtain

$$\pi_t = \beta \mathbb{E}_t [\pi_{t+1}] + \varrho(n) \frac{(1 - \theta)(1 - \beta\theta)}{\theta} (mc_t + \tau_t).$$

Standard steps (see Galí (2015) Chapter 3) imply that the real marginal cost can be written as

$$mc_t = w_t - p_t - a_t = (\sigma + \varphi)\, y_t - (1 + \varphi)\, a_t,$$

and therefore

$$\pi_t = \beta \mathbb{E}_t \left[\pi_{t+1}\right] + \varrho\left(n\right) \frac{(1-\theta)\,(1-\beta\theta)\,(\sigma+\varphi)}{\theta} \left(y_t - \underbrace{\frac{1+\varphi}{\sigma+\varphi} a_t}_{y_t^e} \right) + \varrho\left(n\right) \frac{(1-\theta)\,(1-\beta\theta)}{\theta} \tau_t.$$

Recognizing the definition of the efficient level of output and defining $u_t \equiv (\sigma+\varphi)^{-1} \tau_t$, we obtain

$$\pi_t = \beta \mathbb{E}_t \left[\pi_{t+1}\right] + \kappa\left(n\right) x_t + \kappa\left(n\right) u_t,$$

where

$$\kappa\left(n\right) \equiv \varrho\left(n\right) \kappa_m, \text{ and } \kappa_m \equiv \frac{(1-\theta)\,(1-\beta\theta)\,(\sigma+\varphi)}{\theta}$$

is the slope of the Phillis Curve under monopolistic competition.

A.3 Proof of Proposition 3

We begin by collecting a set of results that will be useful for the main derivation.

From the main text, the steady-state efficient level of output is given by

$$Y^e = (1-\mathcal{G})^{\frac{1-\sigma}{\sigma+\varphi}},$$

whereas the natural level of output is

$$Y^n = (1-\mathcal{G})^{-\frac{\sigma}{\sigma+\varphi}} (\mathcal{T}\mathcal{M})^{-\frac{1}{\sigma+\varphi}},$$

where $\mathcal{T}$ denotes the gross steady-state production subsidy. Taking the ratio of the two expressions yields

$$\frac{Y^e}{Y^n} = \left((1-\mathcal{G})\,\mathcal{T}\mathcal{M}\right)^{\frac{1}{\sigma+\varphi}}.$$

Using the assumption $\mathcal{T} = (1-\mathcal{G})^{-1}$, we obtain

$$\Delta\left(n\right) \equiv \log\left(\frac{Y^e}{Y^n}\right) = \frac{1}{\sigma+\varphi} \log \mathcal{M}\left(n\right),$$

which is the steady-state wedge between the efficient and natural levels of output.

In steady state, labor market clearing implies

$$\frac{1}{\mathcal{MT}} = C^{\sigma} L^{\varphi} \Longleftrightarrow \frac{Y}{\mathcal{MT}} = C^{\sigma} L^{1+\varphi}.$$

Using $Y = (1-\mathcal{G})\, C$ and $\mathcal{T} = (1-\mathcal{G})^{-1}$, we obtain

$$1 - \Phi = \frac{L^{1+\varphi}}{C^{1-\sigma}},$$

where $\Phi = 1 - \mathcal{M}^{-1}$ is an index of the steady-state distortion.

Now consider the period utility function,

$$U_t = \frac{1}{1-\sigma} C_t^{1-\sigma} - \frac{1}{1+\varphi} L_t^{1+\varphi}.$$

A second-order Taylor approximation of $\frac{1}{1-\sigma} C_t^{1-\sigma}$ and $\frac{1}{1+\varphi} L_t^{1+\varphi}$ around the zero-inflation steady state gives

$$\frac{1}{1-\sigma} C_t^{1-\sigma} \simeq \frac{1}{1-\sigma} C^{1-\sigma} + C^{1-\sigma} \left(\frac{C_t - C}{C} \right) - \frac{\sigma}{2} C^{1-\sigma} \left(\frac{C_t - C}{C} \right)^2,$$

and

$$\frac{1}{1+\varphi} L_t^{1+\varphi} \simeq \frac{1}{1+\varphi} L^{1+\varphi} + L^{1+\varphi} \left(\frac{L_t - L}{L} \right) + \frac{\varphi}{2} L^{1+\varphi} \left(\frac{L_t - L}{L} \right)^2.$$

For any variable Z_t, a second-order log approximation implies

$$\log Z_t \simeq \log Z + \frac{Z_t - Z}{Z} - \frac{1}{2} \left(\frac{Z_t - Z}{Z} \right)^2 \Longleftrightarrow z_t + \frac{1}{2} \left(\frac{Z_t - Z}{Z} \right)^2 = \frac{Z_t - Z}{Z},$$

which in turn implies

$$\frac{Z_t - Z}{Z} = z_t + \frac{1}{2} z_t^2.$$

Hence, we can rewrite the previous approximations as

$$\frac{1}{1-\sigma} C_t^{1-\sigma} \simeq \frac{1}{1-\sigma} C^{1-\sigma} + C^{1-\sigma} c_t + \frac{1-\sigma}{2} C^{1-\sigma} c_t^2,$$

and

$$\frac{1}{1+\varphi} L_t^{1+\varphi} \simeq \frac{1}{1+\varphi} L^{1+\varphi} + L^{1+\varphi} l_t + \frac{1+\varphi}{2} L^{1+\varphi} l_t^2.$$

It follows that the utility function can be approximated as

$$U_t - U = C^{1-\sigma}\left(c_t + \frac{1-\sigma}{2}c_t^2\right) - L^{1+\varphi}\left(l_t + \frac{1+\varphi}{2}l_t^2\right).$$

Dividing both sides by $C^{1-\sigma}$ and using the result above, we obtain

$$\frac{U_t - U}{C^{1-\sigma}} = \left(c_t + \frac{1-\sigma}{2}c_t^2\right) - (1-\Phi)\left(l_t + \frac{1+\varphi}{2}l_t^2\right).$$

Assuming that the Law of Large Numbers holds across sectors, the main text implies that

$$Y_t = \mathcal{D}_t^{-1} A_t L_t,$$

where

$$\mathcal{D}_t \equiv \int_0^1 \frac{1}{n}\sum_{j=1}\left(\frac{P_{ik,t}}{P_t}\right)^{-\omega} dk$$

is an index of price dispersion. Therefore,

$$y_t = a_t + l_t - d_t.$$

Using this relation, we can write

$$\begin{aligned}
\frac{U_t - U}{C^{1-\sigma}} &\simeq y_t + \frac{(1-\sigma)}{2}y_t^2 - (1-\Phi)\left[l_t + \frac{(1+\varphi)}{2}l_t^2\right] \\
&= y_t + \frac{(1-\sigma)}{2}y_t^2 - (1-\Phi)\left[y_t - a_t + d_t + \frac{(1+\varphi)}{2}\left(y_t - a_t + d_t\right)^2\right] \\
&= y_t + \frac{(1-\sigma)}{2}y_t^2 - (1-\Phi)\left[-a_t + d_t + \frac{(1+\varphi)}{2}\left(y_t - a_t\right)^2\right] \\
&= \Phi y_t + \frac{(1-\sigma)}{2}y_t^2 - (1-\Phi)\left[d_t + \frac{(1+\varphi)}{2}\left(y_t^2 + a_t^2 - 2y_t a_t\right) - a_t\right] \\
&= \Phi y_t + \frac{(1-\sigma)}{2}y_t^2 - (1-\Phi)\left[d_t + \frac{(1+\varphi)}{2}y_t^2 + \frac{(1+\varphi)}{2}a_t^2 - \frac{(1+\varphi)}{2}2y_t a_t - a_t\right] \\
&= \Phi y_t + \frac{(1-\sigma)}{2}y_t^2 - \left[d_t + \frac{(1+\varphi)}{2}y_t^2 + \frac{(1+\varphi)}{2}a_t^2 - \frac{(1+\varphi)}{2}2y_t a_t - a_t\right] \\
&\quad + \Phi\left[d_t + \frac{(1+\varphi)}{2}y_t^2 + \frac{(1+\varphi)}{2}a_t^2 - \frac{(1+\varphi)}{2}2y_t a_t - a_t\right].
\end{aligned}$$

Notice that

$$\Phi\left[d_t + \frac{(1+\varphi)}{2}y_t^2 + \frac{(1+\varphi)}{2}a_t^2 - \frac{(1+\varphi)}{2}2y_t a_t\right]$$

is of third order. Hence,

$$\begin{aligned}
\frac{U_t - U}{C^{1-\sigma}} &\simeq \Phi y_t + \frac{(1-\sigma)}{2} y_t^2 - d_t - \frac{(1+\varphi)}{2} y_t^2 - \frac{(1+\varphi)}{2} a_t^2 + \frac{(1+\varphi)}{2} 2 y_t a_t + (1-\Phi)\, a_t \\
&= \Phi y_t + \left[\frac{(1-\sigma)}{2} - \frac{(1+\varphi)}{2} \right] y_t^2 - d_t - \frac{(1+\varphi)}{2} a_t^2 + \frac{(1+\varphi)}{2} 2 y_t a_t + (1-\Phi)\, a_t \\
&= \Phi y_t - \frac{1}{2} (\sigma + \varphi)\, y_t^2 - d_t - \frac{(1+\varphi)}{2} a_t^2 + \frac{(1+\varphi)}{2} 2 y_t a_t + (1-\Phi)\, a_t \\
&= \Phi y_t - \frac{1}{2} (\sigma + \varphi)\, y_t^2 - d_t + \frac{(1+\varphi)}{2} 2 y_t a_t - \frac{(1+\varphi)}{2} a_t^2 + (1-\Phi)\, a_t \\
&= \Phi y_t - \frac{1}{2} \left[(\sigma + \varphi)\, y_t^2 - 2 (1+\varphi)\, y_t a_t + 2 d_t \right] - \frac{(1+\varphi)}{2} a_t^2 + (1-\Phi)\, a_t \\
&= \Phi y_t - \frac{1}{2} \left[(\sigma + \varphi) \left\{ y_t^2 - 2 \left(\frac{1+\varphi}{\sigma + \varphi} \right) y_t a_t \right\} + 2 d_t \right] - \frac{(1+\varphi)}{2} a_t^2 + (1-\Phi)\, a_t \\
&= \Phi y_t - \frac{1}{2} \left[(\sigma + \varphi) \left\{ y_t^2 - 2 y_t^e y_t \right\} + 2 d_t \right] - \frac{(1+\varphi)}{2} a_t^2 + (1-\Phi)\, a_t \\
&= \Phi y_t - \frac{1}{2} \left[(\sigma + \varphi) \left\{ y_t^2 - 2 y_t^e y_t + (y_t^e)^2 - (y_t^e)^2 \right\} + 2 d_t \right] - \frac{(1+\varphi)}{2} a_t^2 + (1-\Phi)\, a_t \\
&= \Phi y_t - \frac{1}{2} \left[(\sigma + \varphi) \left\{ (y_t - y_t^e)^2 - (y_t^e)^2 \right\} + 2 d_t \right] - \frac{(1+\varphi)}{2} a_t^2 + (1-\Phi)\, a_t \\
&= \Phi y_t - \frac{1}{2} \left[(\sigma + \varphi)\, (y_t - y_t^e)^2 + 2 d_t \right] + \frac{1}{2} (\sigma + \varphi)\, (y_t^e)^2 - \frac{(1+\varphi)}{2} a_t^2 + (1-\Phi)\, a_t .
\end{aligned}$$

Defining

$$x_t \equiv y_t - y_t^e,$$

we can therefore write

$$\frac{U_t - U}{C^{1-\sigma}} \simeq \Phi x_t - \frac{1}{2} \left[(\sigma + \varphi)\, x_t^2 + 2 d_t \right] + \mathcal{H}_t,$$

where

$$\mathcal{H}_t \equiv \Phi y_t^e + \frac{1}{2} (\sigma + \varphi)\, (y_t^e)^2 - \frac{(1+\varphi)}{2} a_t^2 + (1-\Phi)\, a_t .$$

Let us focus on $\mathcal{H}_t$ and note that, using the definition of y_t^e, we can write

$$\begin{aligned}
\mathcal{H}_t &= \Phi \left(\frac{1+\varphi}{\sigma + \varphi} \right) a_t + \frac{1}{2} (\sigma + \varphi) \left(\left(\frac{1+\varphi}{\sigma + \varphi} \right) a_t \right)^2 - \frac{(1+\varphi)}{2} a_t^2 + \left(1 - \widetilde{\Phi} \right) a_t \\
&= \left[\Phi \left(\frac{1+\varphi}{\sigma + \varphi} \right) + \left(1 - \widetilde{\Phi} \right) \right] a_t + \frac{1}{2} \frac{(1+\varphi)^2}{\sigma + \varphi} a_t^2 - \frac{(1+\varphi)}{2} a_t^2 \\
&= \left[\Phi \left(\frac{1+\varphi}{\sigma + \varphi} \right) + \left(1 - \widetilde{\Phi} \right) \right] a_t + \left[\frac{1}{2} \frac{(1+\varphi)^2}{\sigma + \varphi} - \frac{(1+\varphi)}{2} \right] a_t^2 \\
&= \frac{\Phi (1-\sigma) + (\sigma + \varphi)}{\sigma + \varphi} a_t + \frac{1}{2} \frac{(1+\varphi)(1-\sigma)}{\sigma + \varphi} a_t^2 .
\end{aligned}$$

Equivalently,

$$\mathcal{H}_t = \frac{\Phi(1-\sigma)+(\sigma+\varphi)}{\sigma+\varphi} a_t + \frac{1}{2}\frac{(1+\varphi)(1-\sigma)}{\sigma+\varphi} a_t^2.$$

Assuming that a_t follows an AR(1) process, we can write representative-household welfare as

$$\begin{aligned}\mathcal{W}_0 = U(C,L) &+ \frac{1}{2} C^{1-\sigma} \frac{(1+\varphi)(1-\sigma)\sigma_a^2}{(\sigma+\varphi)(1-\rho_a^2)} \\ &- \frac{1}{2} C^{1-\sigma} \mathbb{E}_0 \left[(1-\beta) \sum_{t=0}^{\infty} \beta^t \left\{ (\sigma+\varphi) x_t^2 + 2d_t - 2\Phi x_t \right\} \right].\end{aligned}$$

We now state a lemma that allows to rewrite the welfare function in terms of inflation.

Lemma 1 *Up to a second order, the following relation holds:*

$$\sum_{t=0}^{\infty} \beta^t d_t = \frac{1}{2}\frac{\varepsilon}{\lambda} \sum_{t=0}^{\infty} \beta^t \pi_t^2$$

where $\lambda \equiv (1-\beta)(1-\beta\theta)/\theta = \kappa_m/(\sigma+\varphi)$.

Using this lemma, we can write

$$\begin{aligned}\mathcal{W}_0 = U(C,L) &+ \frac{1}{2} C^{1-\sigma} \frac{(1+\varphi)(1-\sigma)\sigma_a^2}{(\sigma+\varphi)(1-\rho_a^2)} \\ &- \frac{1-\beta}{2} C^{1-\sigma} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left\{ (\sigma+\varphi) x_t^2 + \frac{\varepsilon}{\lambda} \pi_t^2 - 2\Phi x_t \right\} \right].\end{aligned}$$

For small steady-state distortions, it is easy to show that $\Phi = \log \mathcal{M} = (\sigma+\varphi)\Delta(n)$, hence we can write

$$\begin{aligned}\mathcal{W}_0 = U(C,L) &+ \frac{1}{2} C^{1-\sigma} \frac{(1+\varphi)(1-\sigma)\sigma_a^2}{(\sigma+\varphi)(1-\rho_a^2)} \\ &- \frac{(1-\beta)(\sigma+\varphi)}{2} C^{1-\sigma} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left\{ x_t^2 + \vartheta \pi_t^2 - 2\Delta(n) x_t \right\} \right],\end{aligned}$$

where $\vartheta \equiv \varepsilon/\kappa_m$. Letting $\sigma \to 1$, we obtain

$$\mathcal{W}_0 = U(C,L) - \frac{(1-\beta)(1+\varphi)}{2} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left\{ x_t^2 + \vartheta \pi_t^2 - 2\Delta(n) x_t \right\} \right].$$

Finally, using the Nautral level of output

$$
\begin{aligned}
U(C,L) &= \log C - \frac{1}{1+\varphi} L^{1+\varphi} \\
&= \log\left((1-\mathcal{G}(n))\, Y(n)\right) - \frac{1}{1+\varphi} Y(n)^{1+\varphi} \\
&= -\frac{1}{1+\varphi}\left[\log\left(\mathcal{M}(n)\right) - \frac{1}{\mathcal{M}(n)}\right] + \log\left(1-\mathcal{G}(n)\right) \\
&= \mathcal{U}(n) + \log\left(1-\mathcal{G}(n)\right),
\end{aligned}
$$

where

$$
\mathcal{U}(n) \equiv -\frac{1}{1+\varphi}\left[\log\left(\mathcal{M}(n)\right) - \frac{1}{\mathcal{M}(n)}\right].
$$

A.4 Proof of Proposition 4

Under Assumption 1, the Phillips curve can be written as

$$
\pi_t = \mathbb{E}_t[\pi_{t+1}] + \kappa(n) x_t + \kappa(n) u_t.
$$

We conjecture that inflation takes the form

$$
\pi_t = \alpha_\pi(n) + \zeta_\pi(n) u_t,
$$

where $\alpha_\pi(n)$ and $\zeta_\pi(n)$ are coefficients to be determined. Since u_t follows an AR(1) process with persistence ρ_u, expected inflation is

$$
\mathbb{E}_t[\pi_{t+1}] = \alpha_\pi(n) + \rho_u \zeta_\pi(n) u_t.
$$

Substituting this conjecture into the Phillips curve yields

$$
\pi_t = \alpha_\pi(n) + \rho_u \zeta_\pi(n) u_t + \kappa(n) x_t + \kappa(n) u_t.
$$

From the main text, the first-order condition associated with the optimal monetary policy problem is

$$
x_t = \Delta(n) - \frac{\kappa(n)}{\vartheta_m} \pi_t.
$$

Substituting this expression for x_t into the Phillips curve gives

$$
\pi_t = \alpha_\pi(n) + \rho_u \zeta_\pi(n) u_t + \kappa(n)\left[\Delta(n) - \frac{\kappa(n)}{\vartheta_m}\pi_t\right] + \kappa(n) u_t.
$$

Rearranging,

$$\pi_t = \frac{\vartheta_m}{\vartheta_m + \kappa(n)^2}\alpha_\pi(n) + \frac{\vartheta_m \kappa(n)}{\vartheta_m + \kappa(n)^2}\Delta(n) + \frac{\vartheta_m \left[\rho_u \zeta_\pi(n) + \kappa(n)\right]}{\vartheta_m + \kappa(n)^2} u_t.$$

Matching coefficients with the conjectured law of motion for inflation implies that

$$\alpha_\pi(n) = \frac{\vartheta_m}{\vartheta_m + \kappa(n)^2}\alpha_\pi(n) + \frac{\vartheta_m \kappa(n)}{\vartheta_m + \kappa(n)^2}\Delta(n),$$

so that

$$\alpha_\pi(n) = \vartheta_m \frac{\Delta(n)}{\kappa(n)}.$$

Similarly, $\zeta_\pi(n)$ must satisfy

$$\zeta_\pi(n) = \frac{\vartheta_m \left[\rho_u \zeta_\pi(n) + \kappa(n)\right]}{\vartheta_m + \kappa(n)^2},$$

which implies

$$\zeta_\pi(n) = \frac{\vartheta_m \kappa(n)}{\vartheta_m(1-\rho_u) + \kappa(n)^2}.$$

Therefore, inflation is given by

$$\pi_t = \vartheta_m \frac{\Delta(n)}{\kappa(n)} + \frac{\vartheta_m \kappa(n)}{\vartheta_m(1-\rho_u) + \kappa(n)^2} u_t.$$

Substituting this expression back into the first-order condition yields

$$x_t = -\frac{\kappa(n)^2}{\vartheta_m(1-\rho_u) + \kappa(n)^2} u_t.$$

Finally, defining

$$\delta_x(n, \vartheta_m) \equiv \frac{\kappa(n)^2}{\vartheta_m(1-\rho_u) + \kappa(n)^2},$$

we obtain the result stated in Proposition 4.

B Empirical Appendix

This appendix replicates the baseline state-level Phillips-curve exercise of Section 3.2, but using *headline CPI inflation* (all industries) instead of the non-tradables component. The specification and identification are unchanged: we regress state inflation on state unemployment, including state and time fixed effects, and allow the Phillips-curve slope to vary across Low/Medium/High concentration states (constructed from Census BDS concentration measures). The only difference relative to the baseline is the inflation outcome (CPI instead of non-tradables).

Table 3: Oligopolistic Phillips curve in the data (All industries / CPI inflation)

	CR1	CR2	HHI
δ_M	0.0271^{*}	0.0216^{*}	0.0204
δ_H	0.0364^{**}	0.0161	0.0324^{*}
Slope Low β_L	-0.0540^{***}	-0.0403^{***}	-0.0489^{**}
Slope Medium $(\beta_L + \delta_M)$	-0.0269^{**}	-0.0187	-0.0285^{**}
Slope High $(\beta_L + \delta_H)$	-0.0176	-0.0242	-0.0165
R^2 (within)	0.7878	0.7880	0.7879
Observations	4665	4665	4665

Overall, the qualitative pattern is the same as in the baseline: higher concentration is associated with a *flatter* Phillips curve (i.e. the unemployment coefficient is less negative in Medium/High concentration states), consistent with the model mechanism linking market power to a lower inflation response to slack.