Hawks & Doves with Risk Premia

Pau Belda*

First version: April 2026.
This version: November 2026.
Preliminary version.
Latest version available here.

Abstract

I study how central bank preferences shape monetary policy when aggregate risk is priced. In response to an inflationary supply shock, policy determines how risk is split between output and inflation, and hence the risk priced by financial markets. The resulting risk premia alter both policy implementation and macroeconomic outcomes. A more inflation-oriented central bank can contract activity more while raising the policy rate less, because greater output risk already dampens demand due to higher premia. Interest-rate smoothing can backfire: delayed adjustment can raise risk premia enough to deepen rather than cushion the output contraction. Finally, when uncertainty is sufficiently high, society optimally appoints a central banker who places more weight on output stabilization than society itself to limit the recession risk generated by aggressive inflation stabilization, reversing Rogoff's conservative-central-banker result.
JEL Codes: E52, E44, G12, E31.
Keywords: Monetary Policy, Central Bank Preferences, Risk Premia, Uncertainty Shocks.

1. Introduction

Central bank mandates are incomplete in a fundamental sense: they specify broad objectives, such as price stability and full employment, but leave policymakers discretion over how to trade them off. The classic result of Rogoff (1985) exploits precisely this distinction: society may prefer a more inflation-conservative central banker whose stabilization preferences differ from its own, rather than one whose objective simply coincides with social welfare, as in standard normative New Keynesian analyses. Recent work by Bocola et al. (2024) shows that these preferences can be highly consequential: they argue that a perceived decline in the Federal Reserve's weight on inflation stabilization contributed substantially to the post-pandemic inflation surge.

The effects of policy preferences on monetary transmission are well understood in models with certainty equivalence, but much less so once aggregate risk pricing becomes relevant for monetary policy. In the canonical log-linear New Keynesian benchmark, a greater weight on inflation stabilization makes policymakers more willing to tolerate output losses after a cost-push shock, while a stronger preference for interest-rate smoothing cushions and delays the output response (Galí, 2008; Woodford, 2003). Certainty equivalence implies that conditional uncertainty does not enter these policy trade-offs at first order. This paper studies how these familiar prescriptions change when risk compensation enters the monetary policy problem.

The first result concerns monetary transmission. Priced risk can reverse the certainty-equivalent mapping from policy preferences to the policy rate. When a more inflation-oriented (*hawkish*) central bank would raise the policy rate more under certainty equivalence, sufficiently high supply uncertainty can instead make it contract activity more while raising the short rate by less. This does not occur because priced risk changes the optimal inflation–output allocation: with an unrestricted policy rate, optimal policy fully offsets the aggregate nominal risk premium through the policy rate, leaving the discretionary allocation unchanged. Priced risk instead changes how that allocation is implemented.

An inflation-oriented central bank absorbs more of a cost-push shock through output, generating greater output risk (variance in real wealth returns); an output-oriented central bank absorbs more through inflation, generating greater inflation risk. Because these exposures enter aggregate wealth returns with opposite signs, policy determines how strongly they hedge one another. The resulting risk premium is U-shaped in the weight on output stabilization and vanishes at an intermediate hedge point. Away from that point, financial markets supply part of the required tightening, reducing the adjustment needed in the policy rate. Very *hawkish* and very *dovish* central bankers therefore rely more on risk repricing to implement the desired stabilization, whereas those whose mandates better

*Research Economist at the Bank of England. This paper was previously circulated under the title "A Note on Energy Shock in a Risk-Centric Model". I have benefited from conversations with Kristina Bluwstein, Ricardo Caballero, Mridula Duggal, Jordi Galí, Michael McLeay, Ricardo Reis, Alp Simsek, Danny Walker. The views expressed in this paper are solely those of the author and do not represent the views of the Bank of England or any of its committees.

balance output and inflation risk rely more directly on the policy-rate path. The policy rate alone is therefore an incomplete measure of monetary stance and policy preferences.

The second result is that interest-rate smoothing can reverse its usual macroeconomic effects. In the certainty-equivalent New Keynesian benchmark, greater inertia dampens the initial rate response to an inflationary supply shock, cushions the output contraction, and allows more inflation. When aggregate risk is priced, smoothing also changes the relative output and inflation exposures of wealth and hence the risk premium. On the inflation-risk side of the hedge point, greater inertia raises priced risk. If supply uncertainty is sufficiently high, the resulting tightening in financial conditions dominates the conventional smoothing channel: greater gradualism then produces a larger output contraction and lower, rather than higher, inflation. Thus a preference for smoother policy rates need not imply a smoother macroeconomic adjustment once delayed rate movements themselves alter the compensation required for bearing aggregate risk.

The third result is that priced risk can overturn the classic Rogoff (1985) conservative-central-banker prescription for optimal delegation. When policy-rate adjustment is costly and supply uncertainty is sufficiently important, society may optimally appoint a central banker who places *more* weight on output stabilization than society itself. The mechanism is a new delegation force toward the preference at which real wealth-return and inflation risks exactly hedge in aggregate wealth. Depending on the location of this hedge point, the risk-hedging force can attenuate, reinforce, or reverse the conventional motive to appoint an inflation-conservative central banker. In particular, when the hedge lies on the output-oriented side of society's preference, sufficiently strong uncertainty reverses the classic result. By contrast, when the policy rate is unrestricted, the central bank can fully offset the risk premium, so risk pricing affects only implementation and the standard Rogoff prescription is unchanged.

I derive these results in a tractable Two-Agent New Keynesian economy with full market segmentation, following the risk-centric approach of Caballero and Simsek (2020, 2022). Hand-to-mouth workers supply labor and hold no financial assets, while capitalists hold financial assets and do not work. The supply side is standard: firms set prices subject to Calvo rigidities and a cost-push disturbance generates the familiar inflation–output trade-off. I introduce a cost-push shock with stochastic volatility, so the model contains both first-moment supply disturbances and time variation in their conditional second moment, which generates time-varying compensation for aggregate risk. The central bank has a standard quadratic stabilization objective over inflation and the output gap, augmented where relevant by a preference for interest-rate smoothing. The stabilization weights and smoothing parameter are treated as primitives in the model.

The effect of aggregate priced risk on demand is summarized by a *Risk-Centric IS curve*,

$$x_t = \mathbb{E}_t x_{t+1} - (i_t - \mathbb{E}_t \pi_{t+1} - \bar{r}_t^n) - \Psi_t, \tag{1}$$

where Ψ_t is the risk premium on aggregate nominal wealth. As in the canonical New Keynesian model, demand depends on the gap between the ex-ante real policy rate and the natural real rate. The risk-centric feature is the additional wedge Ψ_t: a higher risk premium depresses demand one-for-one, just as a higher real policy rate would. Hence, for a given desired output gap, the policy rate required to implement it is weakly lower than under certainty equivalence, by exactly the size of the risk-premium wedge. Financial markets can therefore provide part of the tightening otherwise delivered by the policy rate. Crucially, this wedge is endogenous to stabilization policy itself: by determining how a supply shock is split between output and inflation, policy changes the relative exposures of aggregate wealth to real-return and inflation risk that markets price.

The paper shows that familiar conclusions about monetary-policy preferences and transmission can change once policy stance is understood in terms of broader financial conditions rather than the short-term interest rate alone. This builds on a growing literature arguing that the short-term policy rate is not a sufficient statistic for monetary stance (Caballero and Simsek, 2020, 2022; Caballero, Caravello, and Simsek, 2026; Caballero and Simsek, 2025; Galí, 2026). The model builds closely on the risk-centric framework of Caballero and Simsek (2020, 2022), while imposing full rationality, introducing time-varying supply uncertainty and, crucially, allowing for policy inertia; it also nests the textbook New Keynesian model as a limiting case. Relative to the "Long Taylor Principle" of Galí (2026), the analysis explicitly models endogenous risk premia as a component of financial conditions and traces their implications for stabilization preferences, gradualism, and optimal delegation.

The core mechanism also connects the paper to the literature on endogenous risk premia and monetary policy. Palomino (2012) is a direct antecedent in showing that monetary stabilization changes exposure to markup risk and hence risk compensation in asset returns. Caballero and Simsek (2020, 2022) emphasize the reverse link: risk premia and asset prices affect aggregate demand and therefore the policy rate required to stabilize the economy. I study this two-way interaction when policy-rate adjustment is costly. Because fully offsetting risk premia through the policy rate is then costly, endogenous risk premia feed back into output and inflation, altering the effects of gradualism and optimal delegation. This mechanism is related but distinct from other work that makes risk premia endogenous through liquidity and financial intermediation (Drechsler, Savov, and Schnabl, 2018), redistribution toward agents with different risk-bearing capacities (Kekre and Lenel, 2022; Sun, 2026), endogenous risk aversion (Pflueger and Rinaldi, 2022), or bank risk

taking (Abbate and Thaler, 2023).

The paper connects to the literature on uncertainty shocks (Bloom, 2009; Fernández-Villaverde et al., 2015; Basu and Bundick, 2017) and, more specifically, to work on monetary policy under second-moment supply disturbances. Cho et al. (2021) study how uncertainty affects inflation through precautionary pricing, while Cho and Oh (2023) show that cost-push uncertainty can alter the optimal inflation–output trade-off by generating a time-varying gap between natural and efficient output. I deliberately abstract from these channels by keeping the Phillips curve log-linear. Conditional uncertainty therefore does not directly change the central bank's desired inflation–output allocation. Instead, it determines how strongly financial markets price the exposures generated by that allocation. This isolates a distinct role for supply uncertainty: it governs the strength of the risk-pricing channel and thereby changes the implementation of stabilization, the effects of gradualism, and the optimal choice of central banker.

Finally, the paper also contributes to the literature on central-bank preferences and delegation, from the conservative-central-banker result of Rogoff (1985) to subsequent work on optimal contracts, targets, and policy objectives (Lohmann, 1992; Persson and Tabellini, 1993; Walsh, 1995; Svensson, 1997; Woodford, 2003), and more recent evidence on the importance of central-bank preferences (Favero and Rovelli, 2003; Surico, 2007; Bocola et al., 2024). I revisit these questions when stabilization itself shapes priced aggregate risk, showing that this interaction can alter familiar conclusions about policy preferences and optimal delegation.

2. Model and benchmark monetary policy

2.1. Environment

There are two groups of households and a standard New Keynesian production block. Markets are fully segmented: hand-to-mouth workers supply all labor, hold no financial assets, and consume their disposable labor income; capitalists supply no labor, own firms and financial wealth, and choose consumption and a portfolio of the market claim, a real safe bond, and a nominal safe bond. Capitalists have log utility, full information, and rational expectations.

Let y_t denote log output and y_t^n log natural output, defined as the flexible-price level of output under the same real disturbances. The output gap is $x_t \equiv y_t - y_t^n$.

The economy contains a single cost-push disturbance u_t, interpreted as a reduced-form shock that generates an inflation–output trade-off. It follows

$$u_{t+1} = \rho_u u_t + \eta_{t+1}, \qquad \eta_{t+1} \mid \mathcal{I}_t \sim \mathcal{N}(0, s_t), \tag{2}$$

$$s_{t+1} = (1-\rho_s)\bar{s} + \rho_s s_t + \nu_{t+1}^s, \qquad \mathbb{E}_t \nu_{t+1}^s = 0, \tag{3}$$

where $0 < \rho_u, \rho_s < 1$. Volatility innovations have a time-invariant, conditionally symmetric distribution and are conditionally independent of η_{t+1}. To preserve positivity while retaining an affine volatility state, I impose bounded support: for some $d \in (0, \bar{s})$,

$$|\nu_{t+1}^s| \le (1-\rho_s)d \qquad \text{a.s.} \tag{4}$$

Hence the interval $s_t \in [\bar{s}-d, \bar{s}+d]$ is invariant and $s_t > 0$ at all dates. Writing $\hat{s}_t \equiv s_t - \bar{s}$, the stationary variance of the uncertainty state satisfies

$$V_s \equiv \mathrm{Var}(\hat{s}_t) = \frac{\mathrm{Var}(\nu_t^s)}{1-\rho_s^2} \le \bar{V}_s \equiv \frac{1-\rho_s}{1+\rho_s} d^2. \tag{5}$$

The bounded affine specification is a tractability device for generating time variation in priced macroeconomic risk; the economic results below depend on the current conditional variance s_t, not on this particular law of motion. In the baseline, u_t does not move potential output.

Supply side. Monopolistically competitive firms set prices subject to Calvo rigidities. To first order, inflation satisfies the standard New Keynesian Phillips curve

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + u_t, \qquad \kappa > 0. \tag{6}$$

The cost-push disturbance shifts the Phillips curve and creates the usual stabilization trade-off: reducing its inflationary effect requires a negative output gap. The model deliberately combines a first-order New Keynesian supply block with an exact nonlinear demand block. This is not intended as a uniform higher-order approximation to the fully nonlinear economy. Rather, the asymmetry isolates the risk-pricing channel while preserving the familiar inflation–output stabilization problem and analytical tractability. Keeping the Phillips curve log-linear shuts down additional higher-order supply-side mechanisms, such as precautionary price setting (see Cho et al. (2021)), which are conceptually distinct from the financial channel studied here. The nonlinear demand side is discussed in detail in the next section.

2.2. A Risk-Centric IS curve

Equilibrium output. Let P_t be the ex-dividend real price of the market portfolio and let $\alpha \in (0,1)$ denote the share of output accruing to capitalists as aggregate after-tax profits. Workers consume their labor-income share,

$$C_t^w = (1-\alpha) Y_t, \tag{7}$$

while log utility implies that capitalists consume a constant fraction of financial wealth,

$$C_t^c = (1-\beta) R_t^p A_{t-1}. \tag{8}$$

In equilibrium the two safe bonds are in zero net supply and all financial wealth is the market portfolio, so

$$R_t^p A_{t-1} = \alpha Y_t + P_t. \tag{9}$$

Goods-market clearing therefore gives

$$Y_t = \frac{1-\beta}{\alpha\beta} P_t. \tag{10}$$

Thus asset prices affect aggregate demand through capitalist wealth and consumption.

Consider the gross return on wealth $R_{t+1} = (\alpha Y_{t+1} + P_{t+1})/P_t$. The same output–wealth equilibrium implies

$$r_{t+1} = \rho + y_{t+1} - y_t, \qquad \rho \equiv -\log\beta, \tag{11}$$

and hence

$$\mathbb{E}_t r_{t+1} = \rho + \mathbb{E}_t \Delta y_{t+1}. \tag{12}$$

Under the flexible-price allocation, $r^n_{t+1} = \rho + y^n_{t+1} - y^n_t$, so

$$\mathbb{E}_t r^n_{t+1} = \rho + \mathbb{E}_t \Delta y^n_{t+1}. \tag{13}$$

Output and return gaps. Using $y_t = x_t + y^n_t$, equations (12)–(13) imply

$$x_t = \mathbb{E}_t x_{t+1} - \left(\mathbb{E}_t r_{t+1} - \mathbb{E}_t r^n_{t+1}\right). \tag{14}$$

This expression provides a direct link with the standard New Keynesian IS curve. In both cases, the output gap is governed by the gap between the actual and natural values of the intertemporal return relevant for expenditure. In the canonical New Keynesian model, certainty equivalence makes this return the one-period safe real rate. Here, aggregate demand depends on capitalist wealth and wealth is risky, so the relevant return is the expected real return on aggregate wealth. The Risk-Centric IS curve is therefore governed by the gap between actual and natural expected wealth returns.

Output risk and the natural benchmark. The Euler equation for a one-period safe real bond is $1 = \mathbb{E}_t \exp(r^f_t - r_{t+1})$. In the Gaussian benchmark, it is useful to write the associated risk compensation directly in terms of the conditional variance of real wealth returns. Define

$$\mathcal{V}^y_t \equiv \mathrm{Var}_t(r_{t+1}) = \mathrm{Var}_t(\Delta y_{t+1}). \tag{15}$$

Then the real-bond Euler equation implies

$$\mathbb{E}_t r_{t+1} = r^f_t + \frac{1}{2}\mathcal{V}^y_t. \tag{16}$$

Thus $\mathcal{V}^y_t$ measures the output risk borne by aggregate wealth. The factor 1/2 is the usual Gaussian risk-compensation term; the economically relevant object for the mechanism below is the underlying conditional variance itself.

The natural allocation admits the same decomposition. Define

$$\bar{r}^n_t \equiv \mathbb{E}_t r^n_{t+1} = \rho + \mathbb{E}_t \Delta y^n_{t+1}. \tag{17}$$

When natural-return innovations are conditionally Gaussian, let

$$\mathcal{V}^n_t \equiv \mathrm{Var}_t(r^n_{t+1}). \tag{18}$$

Then

$$\bar{r}^n_t = r^{f,n}_t + \frac{1}{2}\mathcal{V}^n_t, \tag{19}$$

and therefore the natural rate (i.e., safe real short rate consistent with demand at potential) is

$$r^{f,n}_t = \bar{r}^n_t - \frac{1}{2}\mathcal{V}^n_t = \rho + \mathbb{E}_t \Delta y^n_{t+1} - \frac{1}{2}\mathcal{V}^n_t. \tag{20}$$

The object $\bar{r}^n_t$ is the natural intertemporal return required to sustain aggregate demand at potential. In the canonical New Keynesian model, it coincides with the natural safe real rate. In the Risk-Centric economy, the same natural demand benchmark applies, but aggregate output risk lowers the natural safe rate by its required compensation.[1]

Output risk, inflation risk, and the Fisher mapping. Because the central bank controls the nominal short safe rate, the next step is to map the policy rate into the expected real return on wealth. Define inflation risk and the output–inflation hedging term as

$$\mathcal{V}^\pi_t \equiv \mathrm{Var}_t(\pi_{t+1}), \qquad \mathcal{H}_t \equiv -\mathrm{Cov}_t(r_{t+1}, \pi_{t+1}). \tag{21}$$

Under conditional joint Gaussianity, the nominal- and real-bond Euler equations imply

$$r^f_t = i_t - \mathbb{E}_t \pi_{t+1} + \frac{1}{2}\mathcal{V}^\pi_t - \mathcal{H}_t. \tag{22}$$

The notation deliberately separates the two primitive risks from their covariance. $\mathcal{V}^\pi_t$ is inflation risk—not itself a risk premium—while $\mathcal{H}_t$ measures the extent to which inflation hedges bad real-return states in nominal wealth. When inflation is high precisely when real wealth returns are low, $\mathcal{H}_t > 0$ and the two risks offset one another. In conventional asset-pricing terminology, $\mathcal{H}_t = -\mathrm{Cov}_t(r_{t+1}, \pi_{t+1})$ is the inflation risk premium on nominal bonds; here its economically useful interpretation is the hedging term.[2]

Combining (22) with (16) gives

$$\mathbb{E}_t r_{t+1} = i_t - \mathbb{E}_t \pi_{t+1} + \frac{1}{2}\mathcal{V}^y_t + \frac{1}{2}\mathcal{V}^\pi_t - \mathcal{H}_t. \tag{23}$$

The Risk-Centric IS curve. The two risks and their hedge can be collected in a single object. Define the nominal return on aggregate wealth as

$$r^N_{t+1} \equiv r_{t+1} + \pi_{t+1}, \tag{24}$$

[1] For example, if natural output follows $y^n_t = y^n_{t-1} + z_t$ with $z_{t+1} \mid \mathcal{I}_t \sim \mathcal{N}(0, \sigma^2_z)$, then $\bar{r}^n_t = \rho$, $\mathcal{V}^n_t = \sigma^2_z$, and $r^{f,n}_t = \rho - \frac{1}{2}\sigma^2_z$.

[2] The term $\frac{1}{2}\mathcal{V}^\pi_t$ is also the Jensen correction that appears when converting the nominal safe rate into a safe real rate under Gaussianity. I label it by its economic source—inflation risk—because the policy mechanism studied below is about how the central bank reallocates conditional variance between output and inflation.

and define the nominal wealth risk premium as

$$\Psi_t \equiv \log \mathbb{E}_t \exp\left[-\left(r^N_{t+1} - \mathbb{E}_t r^N_{t+1}\right)\right]. \tag{25}$$

The nominal-bond Euler equation implies

$$\mathbb{E}_t r^N_{t+1} = i_t + \Psi_t, \tag{26}$$

so Ψ_t is the compensation required for bearing nominal aggregate wealth-return risk. Under conditional joint Gaussianity,

$$\Psi_t = \frac{1}{2}\mathrm{Var}_t(r^N_{t+1}) = \frac{1}{2}\mathcal{V}^y_t + \frac{1}{2}\mathcal{V}^\pi_t - \mathcal{H}_t. \tag{27}$$

Equivalently,

$$\mathbb{E}_t r_{t+1} = i_t - \mathbb{E}_t \pi_{t+1} + \Psi_t. \tag{28}$$

Substituting into (14) yields the compact Risk-Centric IS curve,

$$x_t = \mathbb{E}_t x_{t+1} - (i_t - \mathbb{E}_t \pi_{t+1} - \bar{r}^n_t) - \Psi_t. \tag{29}$$

This representation preserves the canonical New Keynesian structure. The familiar ex-ante real-rate gap continues to govern aggregate demand, but is augmented by a single risk-premium wedge, Ψ_t. Financial markets therefore contribute an additional tightening or loosening of monetary conditions through the compensation required for bearing residual nominal-wealth risk.

Equivalently, keeping the underlying risks visible,

$$x_t = \mathbb{E}_t x_{t+1} - (i_t - \mathbb{E}_t \pi_{t+1} - \bar{r}^n_t) - \frac{1}{2}\mathcal{V}^y_t - \frac{1}{2}\mathcal{V}^\pi_t + \mathcal{H}_t. \tag{30}$$

This is the decomposition that organizes the mechanism in the rest of the paper. Monetary policy determines whether a supply disturbance is absorbed primarily as output risk, $\mathcal{V}^y_t$, or inflation risk, $\mathcal{V}^\pi_t$. Because bad output states are associated with higher inflation, the covariance term $\mathcal{H}_t$ allows the two risks to hedge one another. The relevant financial wedge is therefore not either variance in isolation, but the residual nominal-wealth risk left after this hedge.

Finally, using

$$\mathbb{E}_t r_{t+1} = r^f_t + \frac{1}{2}\mathcal{V}^y_t, \qquad \mathbb{E}_t r^n_{t+1} = r^{f,n}_t + \frac{1}{2}\mathcal{V}^n_t,$$

the same Risk-Centric IS curve can be written as

$$x_t = \mathbb{E}_t x_{t+1} - \left(r^f_t - r^{f,n}_t\right) - \frac{1}{2}\left(\mathcal{V}^y_t - \mathcal{V}^n_t\right). \tag{31}$$

This third representation separates monetary restrictiveness into a safe real-rate gap and an output-risk gap. The first is the familiar New Keynesian margin, now defined using the genuinely safe real rate. The second is the additional tightening generated when conditional output risk rises relative to its natural level. Financial markets can therefore tighten aggregate demand even when the safe real-rate gap is unchanged.

2.3. Optimal monetary policy under discretion

The central bank is assigned the standard quadratic stabilization objective

$$\mathcal{L}_t = \frac{1}{2}\left(\pi^2_t + \lambda x^2_t\right), \tag{32}$$

where a larger $\lambda > 0$ places more weight on output-gap stabilization relative to inflation stabilization. I interpret (32) as a reduced-form representation of the central bank's mandate and preferences over the inflation–output tradeoff. A closely related quadratic loss function can be derived from a second-order approximation to household welfare in the canonical New Keynesian model, but that derivation typically pins down the relative weight on output stabilization as a function of structural parameters. Such a restriction is useful for normative welfare analysis, but it is not the object of interest here. Actual central-bank mandates specify broad objectives such as price stability and real-activity stabilization without mechanically fixing the weight placed on one objective relative to the other. I therefore leave λ as a free policy-preference parameter, capturing how the central bank chooses to trade off these objectives within its mandate. I first consider discretion and an unrestricted policy rate. The next section adds history dependence to the policy rate.

Optimal allocation. Under discretion, the central bank takes private expectations as given and chooses the current allocation subject to the Phillips curve. The familiar target criterion is

$$\lambda x_t + \kappa \pi_t = 0. \tag{33}$$

Define

$$\chi_u \equiv 1 - \beta\rho_u > 0, \qquad D(\lambda) \equiv \kappa^2 + \lambda\chi_u. \tag{34}$$

The unique bounded discretionary allocation is

$$x^\star_t = -\frac{\kappa}{D(\lambda)}u_t, \qquad \pi^\star_t = \frac{\lambda}{D(\lambda)}u_t. \tag{35}$$

For an inflationary cost-push shock, $u_t > 0$, a more output-oriented central bank accepts a smaller contraction in activity and more inflation. Thus λ has its usual New Keynesian interpretation: it determines how the central bank allocates a supply disturbance between inflation and the output gap.

Risk-adjusted implementation. The presence of priced macroeconomic risk does not alter this allocation when the policy rate is unrestricted. To see this, for any feasible paths (x_t, π_t) define the policy rate that would implement them in the corresponding certainty-equivalent New Keynesian economy,

$$i^{NK}_t(x, \pi) \equiv \bar{r}^n_t + \mathbb{E}_t \pi_{t+1} + \mathbb{E}_t x_{t+1} - x_t. \tag{36}$$

The Risk-Centric IS curve implies the exact identity

$$i_t = i_t^{NK}(x, \pi) - \Psi_t. \tag{37}$$

Hence the risk premium changes the interest rate required to implement a given allocation, but not the set of feasible inflation–output allocations. The central bank chooses the same allocation as in the canonical New Keynesian model. Whenever the allocation generates a positive risk premium, financial markets provide part of the required tightening, and the policy rate adjusts one-for-one in the opposite direction.[3]

For the optimal allocation in (35), the corresponding certainty-equivalent rate is

$$i_t^{NK,\star} = \bar{r}_t^n + \omega_u(\lambda) u_t, \qquad \omega_u(\lambda) \equiv \frac{\kappa(1-\rho_u) + \lambda\rho_u}{D(\lambda)}. \tag{38}$$

Stabilization preferences and priced macroeconomic risk. For a given stabilization preference λ, aggregate risk does not alter the optimal discretionary allocation relative to the New Keynesian benchmark. It does, however, alter the financial conditions required to implement that allocation. The key is that the mandate determines how a cost-push shock is split between output and inflation, and therefore how much aggregate risk is left for financial markets to bear.

Since the cost-push shock does not move potential output, nominal aggregate wealth inherits the exposure

$$b(\lambda) = \underbrace{-\frac{\kappa}{D(\lambda)}}_{\text{output exposure}} + \underbrace{\frac{\lambda}{D(\lambda)}}_{\text{inflation exposure}}. \tag{39}$$

The two components have opposite signs. A positive cost-push shock lowers output, reducing real wealth returns, but raises inflation, increasing the nominal payoff on aggregate wealth. A mandate tilted toward either stabilization objective leaves one of these exposures dominant. An intermediate mandate instead allows output and inflation risk to hedge one another. Stabilization preferences therefore determine not only the macroeconomic allocation, but also the amount of macroeconomic risk that financial markets must absorb and, through it, the mix between risk-premium adjustment and policy-rate adjustment used to implement that allocation.

Proposition 1 (Stabilization mandates, risk premia, and the policy mix)**.** *Under discretionary policy with an unrestricted interest rate, the risk-premium wedge associated with the optimal allocation is*

$$\Psi_t^\star(\lambda) = \frac{1}{2}\left[\frac{\lambda - \kappa}{\kappa^2 + \lambda\chi_u}\right]^2 s_t. \tag{40}$$

It exhibits a U-shaped pattern: it is decreasing in λ for $\lambda < \kappa$, increasing in λ for $\lambda > \kappa$, and reaches its minimum at

$$\lambda^{\text{hedge}} = \kappa, \qquad \Psi_t^\star(\kappa) = 0. \tag{41}$$

Proof. Equation (39) gives $b(\lambda) = (\lambda - \kappa)/D(\lambda)$, and conditional Gaussianity implies $\Psi_t^\star = \frac{1}{2} b(\lambda)^2 s_t$. Moreover, $b'(\lambda) = \kappa(\kappa + \chi_u)/D(\lambda)^2 > 0$, so $b(\lambda)$ crosses zero once, at $\lambda = \kappa$. Hence $\Psi_t^\star$ falls while $b(\lambda) < 0$ and rises once $b(\lambda) > 0$. □

The central result is therefore a non-monotonic mapping from the central bank's stabilization mandate to priced macroeconomic risk. An inflation-oriented central bank absorbs a cost-push shock primarily through activity: inflation moves little, output contracts sharply, and nominal wealth is risky because real wealth returns fall in the adverse state. As the central bank places more weight on output stabilization, the output contraction becomes smaller and inflation absorbs a larger share of the shock. Because output and inflation move nominal wealth in opposite directions, the two exposures increasingly hedge one another.

At $\lambda = \kappa$, this hedge is exact: the fall in real wealth returns is offset by higher inflation, leaving nominal aggregate wealth locally insulated from the shock and eliminating the risk premium. Moving away from this point in either direction weakens the hedge. With a more inflation-oriented mandate, output risk dominates; with a more output-oriented mandate, inflation risk dominates. In both cases nominal wealth becomes risky again and the required premium rises. Hence the U-shape has a simple economic interpretation: an appropriately balanced mandate minimizes priced macroeconomic risk, whereas a mandate concentrated on either inflation or output stabilization generates larger risk premia. Here "balanced" refers specifically to the risk-hedging mandate $\lambda = \kappa$: it need not correspond to equal numerical weights on inflation and output in the loss function.

Figure 1 summarizes this mechanism. The dashed line plots the exposure of nominal wealth to the cost-push shock, $b(\lambda)$, while the solid line plots the associated risk premium, $\Psi_t^\star = \frac{1}{2} b(\lambda)^2 s_t$. As the mandate becomes more output-oriented, wealth exposure rises monotonically from being dominated by output risk to being dominated by inflation risk. At $\lambda = \kappa$, the two exposures exactly offset one another, so $b(\lambda) = 0$ and the risk premium vanishes. Moving away from this hedge point in either direction raises priced macroeconomic risk, although the two tails need not be symmetric.

The two sides of the U-shape need not be symmetric.

[3] This is the same implementation logic emphasized by Caballero and Simsek (2022). It is also related to the spread-adjusted policy logic in Cúrdia and Woodford (2016); the one-for-one offset is exact here because Ψ_t enters aggregate demand as a pure IS wedge with unit coefficient and does not directly affect the Phillips curve or the central bank's objective.

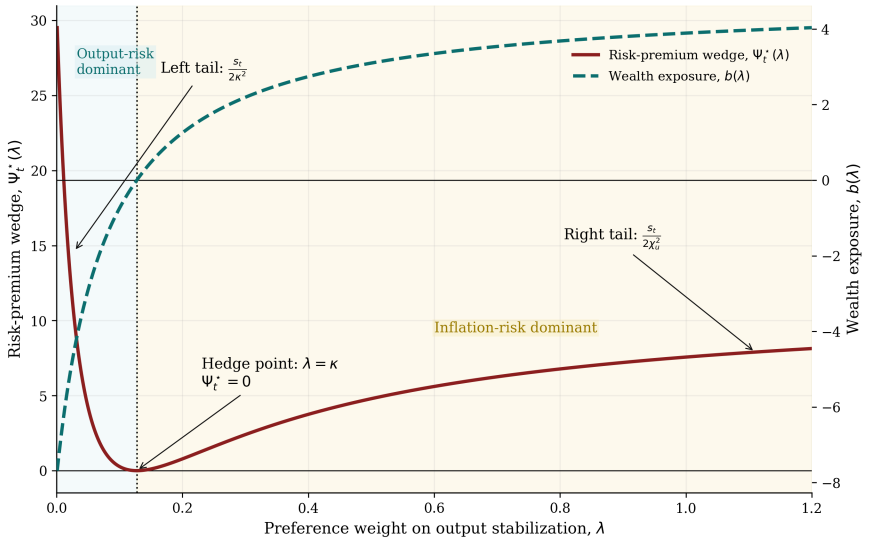


Figure 1: **Central bank preferences and risk premia.** *The solid line shows the risk premium wedge* $\Psi_t^\star(\lambda) = \frac{1}{2}b(\lambda)^2 s_t$*, while the dashed line shows nominal wealth exposure to the cost-push shock,* $b(\lambda) = (\lambda - \kappa)/D(\lambda)$*. For* $\lambda < \kappa$*, output risk dominates nominal-wealth exposure; for* $\lambda > \kappa$*, inflation risk dominates. At the hedge point* $\lambda = \kappa$*, the two exposures exactly offset one another,* $b(\lambda) = 0$*, and the risk premium vanishes. The different limiting values of the two tails illustrate that the U-shape need not be symmetric.*

At the polar, single-objective mandates,

$$\lim_{\lambda \downarrow 0} \Psi_t^\star = \frac{s_t}{2\kappa^2}, \qquad \lim_{\lambda \to \infty} \Psi_t^\star = \frac{s_t}{2\chi_u^2}. \tag{42}$$

The inflation-oriented tail is therefore higher when $\chi_u > \kappa$, whereas the output-oriented tail is higher when $\chi_u < \kappa$. More generally, increasingly output-oriented mandates eventually have diminishing effects on priced risk because nominal-wealth exposure converges to its limiting inflation exposure.

The allocation of output and inflation risk. The aggregate wedge Ψ_t is residual nominal-wealth risk after output and inflation risk are combined. Under conditional joint Gaussianity,

$$\Psi_t = \underbrace{\frac{1}{2}\mathcal{V}_t^y}_{\text{output risk}} + \underbrace{\frac{1}{2}\mathcal{V}_t^\pi}_{\text{inflation risk}} - \underbrace{\mathcal{H}_t}_{\text{output–inflation hedge}}. \tag{43}$$

The first two objects are the conditional variances of real wealth returns and inflation, respectively. The third is their hedging value. In the present supply-shock environment, output falls when inflation rises, so $\mathrm{Cov}_t(r_{t+1}, \pi_{t+1}) < 0$ and $\mathcal{H}_t > 0$: inflation offsets part of the loss in real wealth returns when wealth is measured in nominal terms. Technically, $\mathcal{H}_t = -\mathrm{Cov}_t(r_{t+1}, \pi_{t+1})$ is the conventional inflation risk premium on nominal bonds. Economically, however, its role in (43) is to determine how much of the two underlying risks is hedged away.

Under the optimal allocation in (35), the three objects are

$$\begin{aligned} \mathcal{V}_t^{y,\star} &= \frac{\kappa^2}{D(\lambda)^2} s_t, \\ \mathcal{V}_t^{\pi,\star} &= \frac{\lambda^2}{D(\lambda)^2} s_t, \\ \mathcal{H}_t^\star &= \frac{\kappa\lambda}{D(\lambda)^2} s_t. \end{aligned} \tag{44}$$

Stabilization preferences therefore determine how the conditional variance generated by the supply shock is allocated between activity and inflation. Output risk declines monotonically with λ, whereas inflation risk rises monotonically:

$$\frac{\partial \mathcal{V}_t^{y,\star}}{\partial \lambda} < 0, \qquad \frac{\partial \mathcal{V}_t^{\pi,\star}}{\partial \lambda} > 0. \tag{45}$$

Their relative size is particularly transparent,

$$\frac{\mathcal{V}_t^{\pi,\star}}{\mathcal{V}_t^{y,\star}} = \left(\frac{\lambda}{\kappa}\right)^2. \tag{46}$$

Hence the two regions in Proposition 1 admit a direct interpretation. For $\lambda < \kappa$, output risk is larger than inflation risk; for $\lambda > \kappa$, inflation risk is larger; and at $\lambda = \kappa$ the two conditional variances are exactly equal.

The covariance term determines how much of those two risks survives in nominal wealth. Under the optimal allocation,

$$\mathcal{H}_t^\star = \sqrt{\mathcal{V}_t^{y,\star}\mathcal{V}_t^{\pi,\star}}. \tag{47}$$

Equivalently,

$$\frac{\mathcal{H}_t^\star}{\frac{1}{2}\mathcal{V}_t^{y,\star}+\frac{1}{2}\mathcal{V}_t^{\pi,\star}}=\frac{2\kappa\lambda}{\kappa^2+\lambda^2}\leq 1, \tag{48}$$

with equality if and only if $\lambda=\kappa$. Thus the hedge point is precisely the mandate at which output risk and inflation risk are equal and the covariance term offsets their entire contribution to nominal-wealth risk. Away from that point, one variance dominates the other, the hedge is incomplete, and the aggregate wedge becomes positive again. This gives the U-shape in $\Psi_t^\star(\lambda)$ its sharpest interpretation: monetary policy reallocates a common supply shock between output risk and inflation risk, and priced risk is minimized where those two exposures exactly hedge.

The stabilization mix. A second implication concerns the *composition* of monetary restraint. Combining (37) and (38), the policy rate implementing the optimal allocation is

$$i_t^\star=\bar{r}_t^n+\omega_u(\lambda)u_t-\Psi_t^\star(\lambda)\leq i_t^{NK,\star}. \tag{49}$$

The first two terms give the certainty-equivalent New Keynesian rate, while the last term is the tightening supplied endogenously by financial markets. At the hedge point, $\lambda=\kappa$, the premium vanishes and the policy rate carries the entire implementation adjustment. Away from the hedge point, a larger risk premium shifts part of the required tightening from the short safe rate to financial markets.

The policy-rate response to the mandate can be decomposed into a certainty-equivalent component and a risk-premium component. Differentiating (49) gives

$$\frac{\partial i_t^\star}{\partial\lambda}=\underbrace{\frac{\kappa}{D(\lambda)^2}\left[\rho_u\kappa-\chi_u(1-\rho_u)\right]u_t}_{\text{certainty-equivalent NK component}}+\underbrace{\frac{\kappa(\kappa+\chi_u)(\kappa-\lambda)}{D(\lambda)^3}s_t}_{\text{risk-premium component}}. \tag{50}$$

The second term is strictly positive on the inflation-oriented side of the hedge point. Hence, priced risk always pushes toward a lower policy rate for a more inflation-oriented central bank: lowering λ raises priced output risk and increases the amount of tightening supplied endogenously by financial markets.

Importantly, however, the sign of the total comparative static need not identify this risk-premium channel. If

$$\rho_u\kappa\geq\chi_u(1-\rho_u), \tag{51}$$

then the certainty-equivalent component is itself nonnegative. In this case, $\partial i_t^\star/\partial\lambda>0$ can obtain even when $s_t=0$: sufficiently persistent cost-push shocks imply that a more inflation-oriented central bank may choose a lower nominal policy rate already in the certainty-equivalent model. Priced risk reinforces this ranking but is not necessary for it.

The distinctive implication of the risk-premium channel arises instead when

$$\rho_u\kappa<\chi_u(1-\rho_u), \tag{52}$$

so that the certainty-equivalent model implies the conventional ordering, $\partial i_t^{NK,\star}/\partial\lambda<0$: a more inflation-oriented central bank raises the policy rate more. The endogenous risk premium reverses this ordering whenever

$$(\kappa+\chi_u)(\kappa-\lambda)s_t>D(\lambda)\left[\chi_u(1-\rho_u)-\rho_u\kappa\right]u_t, \tag{53}$$

or, equivalently, for $u_t>0$,

$$\frac{s_t}{u_t}>\frac{D(\lambda)\left[\chi_u(1-\rho_u)-\rho_u\kappa\right]}{(\kappa+\chi_u)(\kappa-\lambda)}. \tag{54}$$

Above this threshold, a lower λ is associated with both a more contractionary allocation and a lower policy rate, despite the certainty-equivalent model predicting the opposite ranking. The reversal is then entirely due to the endogenous tightening supplied through the risk premium.

For a pure uncertainty shock, $u_t=0$, this mechanism is isolated directly. The certainty-equivalent component in (50) vanishes, while the risk-premium component remains strictly positive for $\lambda<\kappa$. Thus, lowering λ raises the risk premium and lowers the short rate required to implement the optimal allocation.

This result cautions against inferring central-bank preferences or monetary stance directly from movements in the short rate. A more inflation-oriented central bank can implement a more contractionary allocation while moving the policy rate by less, because the induced contraction raises priced aggregate risk and shifts part of the monetary tightening to financial markets. Crucially, this interpretation is strongest in the region characterized by (52)–(54), where the lower policy-rate response is genuinely generated by the risk-premium channel rather than by the certainty-equivalent New Keynesian component.

Level versus uncertainty shocks. The separation between allocation and implementation also clarifies the difference between shocks to the level and to the volatility of the cost-push disturbance. Combining (38) and (40),

$$i_t^\star=\bar{r}_t^n+\omega_u(\lambda)u_t-\frac{1}{2}b(\lambda)^2s_t. \tag{55}$$

A level shock changes the desired inflation–output allocation and therefore the certainty-equivalent policy rate. By contrast, holding the level of the shock fixed, an increase in uncertainty raises the risk premium without changing the desired allocation. The endogenous tightening in financial conditions is then offset one-for-one by a lower policy rate.

This distinction matters when level and uncertainty shocks arrive together. An inflationary cost-push shock ordinarily calls for a higher certainty-equivalent policy rate, but if it is accompanied by a sufficiently large rise in uncertainty, the associated increase in risk premia can provide much of the required tightening. The optimal

policy-rate increase is then smaller than in the New Keynesian benchmark and can, in principle, turn into a rate cut. Thus the relevant policy stance is determined by the combined movement in the policy rate and the risk premium, not by the policy rate alone.[4]

A flat Phillips curve. A flat Phillips curve creates an interesting separation between what stabilization preferences imply for inflation and what they imply for financial conditions. For fixed $\lambda > 0$,

$$\begin{aligned} \frac{\partial \pi_t^\star}{\partial \lambda} &= O(\kappa^2), \qquad \frac{\partial x_t^\star}{\partial \lambda} = O(\kappa), \qquad \kappa \downarrow 0. \\ \frac{\partial \Psi_t^\star}{\partial \lambda} &= O(\kappa), \qquad \frac{\partial i_t^\star}{\partial \lambda} = O(\kappa), \qquad \kappa \downarrow 0. \end{aligned} \tag{56}$$

As the Phillips curve becomes flatter, differences in central-bank preferences therefore become almost invisible in inflation one order faster than they disappear from activity, risk premia, and the policy rate. For small but positive κ, central banks with different stabilization preferences can consequently generate very similar inflation responses while displaying more pronounced differences in financial conditions and in the mix between the policy rate and the risk premium. Hence, a flat Phillips curve makes the financial side of implementation relatively more informative about the central bank's stabilization preference than inflation alone.

3. Monetary Policy Inertia with Risk Premia

Section 2.3 established a one-dimensional mapping from the central bank's stabilization preference, λ, to priced macroeconomic risk. With an unrestricted policy rate, the resulting risk premium changes implementation but not the optimal inflation–output allocation. I now add a second preference parameter, δ, capturing gradualism in the policy rate. This breaks the separation result and turns the preference–risk-premium mapping into a surface over (λ, δ).

3.1. Myopic smoothing and the affine equilibrium

Let $\delta > 0$ measure the central bank's preference for gradual adjustment of the policy rate. Under discretionary reoptimization, the central bank chooses (x_t, π_t, i_t) to minimize

$$\frac{1}{2}\left[\pi_t^2 + \lambda x_t^2 + \delta(i_t - i_{t-1})^2\right] \tag{57}$$

subject to the Risk-Centric IS curve and the New Keynesian Phillips curve.

The formulation in 57 is deliberately myopic: the policymaker internalizes the current cost of changing the policy rate but not the effect of today's rate choice on future policy through the inherited rate. I use this benchmark because it isolates the interaction between costly rate adjustment and endogenous risk compensation while preserving closed-form comparative statics. Appendix E relaxes this assumption and solves the corresponding forward-looking discretionary problem with a delegated interest-rate-smoothing objective in the spirit of Woodford (2003), which induces history dependence under discretion. The continuation value strengthens the effects of current rate movements through the expected policy path, but does not restore allocation–implementation separation. Risk-premium shocks remain only partially offset and therefore affect output and inflation; the continuation value changes their magnitude, not the underlying mechanism.

The first-order conditions imply

$$\lambda x_t + \kappa \pi_t = \delta(i_t - i_{t-1}). \tag{58}$$

Without inertia, the right-hand side is zero and the policy rate is an implementation variable. With inertia, changing the rate is itself costly, so the inherited policy rate becomes a state variable and movements in the risk premium can no longer be neutralized one-for-one without changing the allocation.

Work in deviations from the regime-specific Risk-Centric steady state.[5] The local equilibrium is affine in the inherited policy rate, the cost-push shock, and the risk premium:

$$x_t^{\mathcal{I}} = x_i \hat{i}_{t-1} + x_u u_t + x_\Psi \hat{\Psi}_t^{\mathcal{I}}, \tag{59}$$

$$\pi_t^{\mathcal{I}} = \pi_i \hat{i}_{t-1} + \pi_u u_t + \pi_\Psi \hat{\Psi}_t^{\mathcal{I}}, \tag{60}$$

$$\hat{i}_t^{\mathcal{I}} = q_i \hat{i}_{t-1} + q_u u_t + q_\Psi \hat{\Psi}_t^{\mathcal{I}}. \tag{61}$$

The coefficients are derived in Appendix D. The important point here is that, for $\delta > 0$, $x_\Psi < 0$, $\pi_\Psi < 0$, and $q_\Psi < 0$. A rise in the risk premium tightens private financial conditions and induces the central bank to cut the policy rate, but the cut is only partial: output and inflation also fall. The allocation–implementation separation of Section 2.3 therefore fails.

3.2. Effective stabilization preferences and risk premia

The asset-pricing implications of inertia admit a simple interpretation. Although the primitive stabilization weight remains λ, gradual adjustment changes the inflation–output mix chosen in response to a cost-push shock. This can be summarized by a shock-specific *effective stabilization mandate*.

Because the equilibrium is affine, x_u, π_u, and q_u are the partial responses of (x_t, π_t, i_t) to the current cost-push state u_t, holding the inherited policy rate and the risk-premium state fixed. Differentiating (58) with respect to

[4] This differs from Cho and Oh (2023), where cost-push uncertainty can itself create a stabilization trade-off through precautionary pricing. Here I abstract from that higher-order supply-side channel in order to isolate the demand effect of priced macroeconomic risk.

[5] The main reason for doing this is that, unlike in the no-inertia benchmark, innovations to the uncertainty state affect future allocations and therefore constitute an additional source of nominal-wealth risk. This adds a regime-specific constant to the level of the risk premium. We subtract this constant and focus on the time-varying component.

u_t therefore yields

$$\lambda x_u + \kappa \pi_u = \delta q_u, \tag{62}$$

since i_{t-1} is predetermined. Whenever $x_u \neq 0$, define

$$\lambda_u^{\text{eff}}(\lambda, \delta) \equiv -\kappa \frac{\pi_u}{x_u} = \lambda - \delta \frac{q_u}{x_u}. \tag{63}$$

By construction, the response to the cost-push shock satisfies

$$\lambda_u^{\text{eff}} x_u + \kappa \pi_u = 0, \tag{64}$$

which has exactly the same form as the static stabilization condition in Section 2.3. In the economically relevant region in which an inflationary cost-push shock contracts output and raises the policy rate, $x_u < 0 < q_u$, so

$$\lambda_u^{\text{eff}} = \lambda + \delta \frac{q_u}{-x_u} > \lambda. \tag{65}$$

Interest rate gradualism therefore acts as implicit output stabilization: because a stronger contraction requires a larger rate increase, smoothing induces the central bank to accept relatively more inflation and a smaller output decline.

Let

$$L_u^{\mathcal{I}} \equiv x_u + \pi_u \tag{66}$$

denote nominal-wealth exposure to the cost-push innovation. Using (64),

$$L_u^{\mathcal{I}} = x_u \left(1 - \frac{\lambda_u^{\text{eff}}}{\kappa}\right). \tag{67}$$

Thus the hedging logic of Proposition 1 carries over unchanged once the primitive mandate is replaced by the effective one.

Corollary 1 (Effective preferences and risk premia under inertia)**.** *Under the local stability and regularity conditions stated in Appendix D, the time-varying component of the equilibrium risk premium is*

$$\hat{\Psi}_t^{\mathcal{I}} = G_u^{\mathcal{I}}(\lambda, \delta)\hat{s}_t, \qquad G_u^{\mathcal{I}}(\lambda, \delta) = \frac{1}{2}\left(L_u^{\mathcal{I}}\right)^2 = \frac{1}{2}x_u^2\left(1 - \frac{\lambda_u^{\text{eff}}}{\kappa}\right)^2 \tag{68}$$

The cost-push shock is exactly hedged if and only if

$$\lambda_u^{\text{eff}} = \kappa. \tag{69}$$

In terms of the primitive policy preferences, this condition defines the unique zero-risk locus

$$\lambda_{\mathcal{I}}^H(\delta) = \frac{\kappa^2}{\kappa + \delta}. \tag{70}$$

Moreover,

$$\operatorname{sgn}\left(L_u^{\mathcal{I}}\right) = \operatorname{sgn}\left(\lambda_u^{\text{eff}} - \kappa\right) = \operatorname{sgn}\left[\lambda - \lambda_{\mathcal{I}}^H(\delta)\right], \tag{71}$$

and

$$\operatorname{sgn}\left(\frac{\partial G_u^{\mathcal{I}}}{\partial \delta}\right) = \operatorname{sgn}\left[\lambda - \lambda_{\mathcal{I}}^H(\delta)\right]. \tag{72}$$

The first part of the corollary restores the economics of Proposition 1. If $\lambda_u^{\text{eff}} < \kappa$, output risk dominates nominal-wealth exposure. If $\lambda_u^{\text{eff}} > \kappa$, inflation risk dominates. At $\lambda_u^{\text{eff}} = \kappa$, the contraction in real wealth and the increase in inflation exactly offset one another, so nominal aggregate wealth is locally insulated from the shock and the risk premium vanishes.

What inertia changes is the mapping from the central bank's primitive preferences into this effective mandate. Since gradualism acts as implicit output stabilization, less explicit weight on output is required to attain the same risk-hedging mix. The hedge therefore moves according to

$$\frac{d\lambda_{\mathcal{I}}^H}{d\delta} = -\frac{\kappa^2}{(\kappa + \delta)^2} < 0. \tag{73}$$

The single hedge point $\lambda = \kappa$ in Section 2.3 becomes a downward-sloping zero-risk locus in (λ, δ) space.

Note, though, that the last result in the corollary is stronger than the effective-mandate representation alone. Equation (68) shows that inertia changes not only the stabilization mix but also the scale of the cost-push response, x_u. Nevertheless, these equilibrium changes do not overturn the hedging comparative static in the inertia dimension. Below the zero-risk locus, greater inertia moves nominal-wealth exposure toward zero and lowers priced risk; above it, greater inertia moves exposure away from zero and raises priced risk. Thus the effect of gradualism on risk premia is conditional on the stabilization mandate.

Figure 2 summarizes this interaction. In normalized coordinates, the zero-risk locus is

$$\frac{\lambda}{\kappa} = \frac{1}{1 + \delta/\kappa}. \tag{74}$$

For $\delta = 0$, it intersects the horizontal axis at the benchmark hedge $\lambda = \kappa$, recovering Proposition 1.

3.3. The effects of inertia on the macroeconomy

The previous subsection showed that policy inertia changes the risk borne by financial markets. I now ask whether this endogenous repricing changes the macroeconomic effects of policy itself. In the canonical New Keynesian model, greater interest-rate inertia acts as implicit output stabilization: following an inflationary cost-push shock, the central bank moves the policy rate less, accepts more inflation, and limits the contemporaneous output contraction. With priced aggregate risk, however, this is only part of the equilibrium response. Gradualism also changes the risk premium, and, because that premium is no longer offset one-for-one by the policy rate, the resulting financial tightening feeds back into output and inflation. The effect of gradualism on the allocation therefore combines a standard smoothing channel with an endogenous risk-premium feedback.

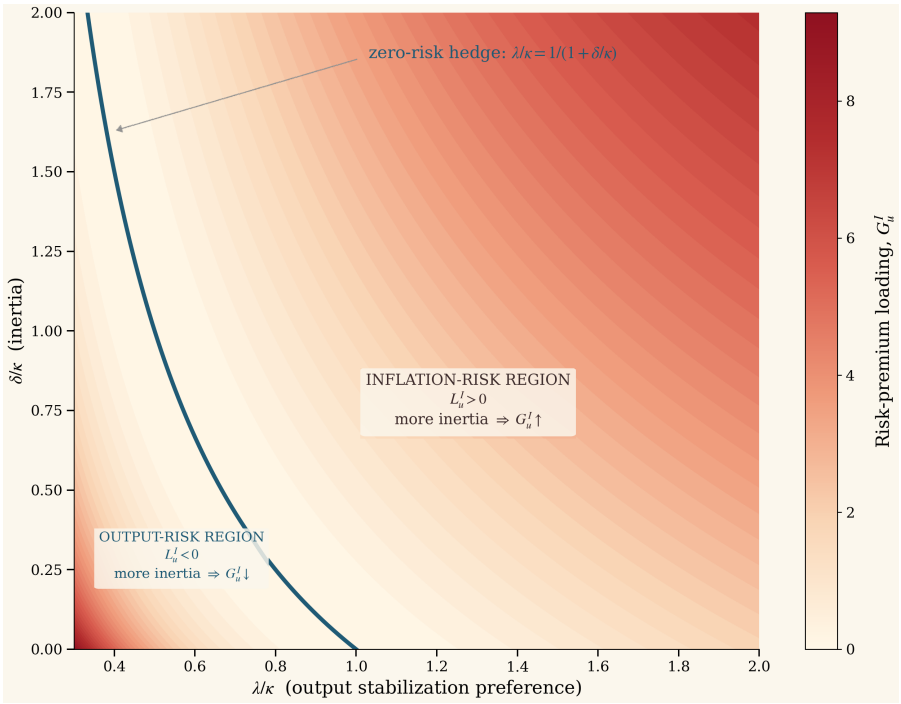


Figure 2: **Effective stabilization mandates and priced macroeconomic risk.** *The heat map reports the equilibrium risk-premium loading $G_u^{\mathcal{I}}(\lambda,\delta) = \frac{1}{2}(L_u^{\mathcal{I}})^2$ as a function of the primitive output-stabilization preference, λ/κ, and interest-rate inertia, δ/κ. The solid line is the zero-risk locus $\lambda/\kappa = 1/(1+\delta/\kappa)$, equivalently $\lambda_u^{\text{eff}} = \kappa$. To the left of the locus, $\lambda_u^{\text{eff}} < \kappa$, output risk dominates and greater inertia lowers priced risk. To the right, $\lambda_u^{\text{eff}} > \kappa$, inflation risk dominates and greater inertia raises priced risk.*

To isolate these forces, consider the impact of an inflationary cost-push event ($u_t > 0, \hat{s}_t \geq 0$) starting from the regime-specific steady state, so that $\hat{i}_{t-1} = 0$. Substituting (68) into (59)–(60) gives

$$x_t^{\mathcal{I}} = x_u u_t + x_\Psi G_u^{\mathcal{I}} \hat{s}_t, \tag{75}$$
$$\pi_t^{\mathcal{I}} = \pi_u u_t + \pi_\Psi G_u^{\mathcal{I}} \hat{s}_t. \tag{76}$$

Their total equilibrium responses to greater gradualism are therefore

$$\frac{dz_t^{\mathcal{I}}}{d\delta} = \underbrace{z_{u,\delta} u_t}_{\text{standard gradualism}} + \underbrace{\left[\underbrace{z_{\Psi,\delta} G_u^{\mathcal{I}}}_{\text{absorption effect}} + \underbrace{z_\Psi G_{u,\delta}^{\mathcal{I}}}_{\text{repricing effect}}\right] \hat{s}_t}_{\text{risk-premium feedback}}, \tag{77}$$

for $z = (x, \pi)$, where $z_{u,\delta} = \frac{\partial z_u}{\partial \delta}$, $z_{\Psi,\delta} = \frac{\partial z_\Psi}{\partial \delta}$, $G_{u,\delta}^{\mathcal{I}} = \frac{\partial G_u^{\mathcal{I}}}{\partial \delta}$. For compactness, define

$$\begin{aligned} \mathcal{R}_x(\lambda,\delta) &\equiv x_{\Psi,\delta} G_u^{\mathcal{I}} + x_\Psi G_{u,\delta}^{\mathcal{I}} \\ \mathcal{R}_\pi(\lambda,\delta) &\equiv \pi_{\Psi,\delta} G_u^{\mathcal{I}} + \pi_\Psi G_{u,\delta}^{\mathcal{I}}. \end{aligned} \tag{78}$$

Proposition 2 (Inertia and allocations with risk premia)**.** *Under the local stability and regularity conditions of Appendix D, and on the economically relevant branch $x_u < 0 < q_u$, the standard gradualism channel satisfies*

$$x_{u,\delta} > 0, \qquad \pi_{u,\delta} > 0. \tag{79}$$

Hence, absent time-varying risk premia, greater gradualism reduces the contemporaneous output contraction and increases inflation following an inflationary cost-push shock.

On the output-risk side of the hedge,

$$\lambda < \lambda_{\mathcal{I}}^H(\delta), \tag{80}$$

greater gradualism lowers priced risk, $G_{u,\delta}^{\mathcal{I}} < 0$. The repricing effect therefore reinforces the conventional gradualism channel, while the absorption effect works in the opposite direction. Hence the sign of the overall risk-premium feedback is in general ambiguous, but the conventional comparative statics,

$$\frac{dx_t^{\mathcal{I}}}{d\delta} > 0, \qquad \frac{d\pi_t^{\mathcal{I}}}{d\delta} > 0,$$

necessarily hold for sufficiently small $\hat{s}_t$.

On the inflation-risk side of the hedge,

$$\lambda > \lambda_{\mathcal{I}}^H(\delta), \tag{81}$$

greater gradualism raises priced risk and both components of the risk-premium feedback oppose the standard channel:

$$\mathcal{R}_x(\lambda,\delta) < 0, \qquad \mathcal{R}_\pi(\lambda,\delta) < 0. \tag{82}$$

Consequently,

$$\frac{dx_t^{\mathcal{I}}}{d\delta} < 0 \quad \Longleftrightarrow \quad \hat{s}_t > \hat{s}_t^{\,x} \equiv \frac{x_{u,\delta} u_t}{-\mathcal{R}_x(\lambda,\delta)}, \tag{83}$$

$$\frac{d\pi_t^{\mathcal{I}}}{d\delta} < 0 \quad \Longleftrightarrow \quad \hat{s}_t > \hat{s}_t^{\,\pi} \equiv \frac{\pi_{u,\delta} u_t}{-\mathcal{R}_\pi(\lambda,\delta)}. \tag{84}$$

Thus, on the inflation-risk side and under sufficiently high supply uncertainty, endogenous risk compensation can overturn the conventional effects of gradualism on both components of the inflation–output allocation.

Proof. See Appendix D.4.

The first terms in (77) are the familiar effects of interest-rate smoothing. Greater δ makes current policy-rate adjustment more costly. In response to an inflationary cost-push shock, the central bank therefore raises the rate less aggressively, accepts more inflation, and limits the contemporaneous contraction in output. This is the sense in which inertia acts as an additional weight on output stabilization.

Priced aggregate risk introduces a second force. The *absorption effect* holds the amount of priced risk fixed and asks how gradualism changes the macroeconomic incidence of a given risk premium. Since

$$x_{\Psi,\delta} < 0, \qquad \pi_{\Psi,\delta} < 0, \tag{85}$$

a more gradual central bank offsets a given premium less through the policy rate. A larger share is consequently absorbed through lower output and lower inflation.

The *repricing effect* operates through the risk premium itself. On the inflation-risk side of the hedge, Corollary 1 implies $G^{\mathcal{I}}_{u,\delta} > 0$. Greater gradualism shifts more of the cost-push disturbance toward inflation, moves nominal wealth exposure farther away from the hedge, and raises the risk adjustment required by financial markets. Since $x_\Psi < 0$ and $\pi_\Psi < 0$, this endogenous repricing reduces both output and inflation. Hence, in the inflation-risk region, the two financial effects reinforce one another:

$$x_{\Psi,\delta} G^{\mathcal{I}}_u < 0, \qquad x_\Psi G^{\mathcal{I}}_{u,\delta} < 0, \tag{86}$$

and

$$\pi_{\Psi,\delta} G^{\mathcal{I}}_u < 0, \qquad \pi_\Psi G^{\mathcal{I}}_{u,\delta} < 0. \tag{87}$$

The resulting tension is the central implication of the proposition. In the canonical model, gradualism delays monetary tightening and therefore protects current activity at the cost of higher inflation. Once aggregate risk is priced, the additional inflation generated by gradualism can itself tighten financial conditions, counteracting the former effect.

The strength of this feedback is proportional to $\hat{s}_t$: for fixed (λ,δ), greater supply uncertainty scales the risk-premium contribution linearly. When uncertainty exceeds $\hat{s}_t^{\,x}$, the risk premium tightening dominates the conventional output-stabilization channel, and gradualism backfires: a stronger preference for smoothing the policy rate produces a larger output contraction. If uncertainty exceeds $\hat{s}_t^{\,\pi}$ as well, even the usual increase in inflation is overturned.

Below the hedge, the repricing force changes sign. There $G^{\mathcal{I}}_{u,\delta} < 0$, so greater gradualism moves nominal-wealth exposure toward the hedge and lowers the risk premium. The resulting financial loosening works against the absorption effect, and the overall risk-premium feedback is therefore generally ambiguous. The possibility of a systematic reversal is specific to the inflation-risk region, where gradualism simultaneously increases priced risk and makes the allocation more sensitive to the financial tightening that follows.

4. Central Bank conservatism with Risk Premia

Sections 2 and 3 treated central-bank preferences as primitives and characterized how they shape macroeconomic allocations and financial conditions. Section 2 showed that the weight on output stabilization determines the risk borne by financial markets and the resulting stabilization mix, while Section 3 showed that interest-rate inertia changes both that risk and the allocation itself. I now reverse the question. Which hawk or dove should society appoint? Following Rogoff (1985), I distinguish society's stabilization preferences from those of the delegated central banker and ask how aggregate risk pricing changes the optimal degree of conservatism. I first recover the familiar inflation bias and Rogoff delegation result in the present model. With an unrestricted policy rate, the answer is immediate from the separation result of Section 2.3: risk premia change the implementation of delegated policy, but not the optimal delegate. The next subsection considers how this conclusion changes once interest-rate inertia breaks that separation.

4.1. Inflation bias

Following Rogoff (1985), suppose a pre-existing distortion makes the socially desired level of activity exceed natural output by a constant $\bar{x} > 0$. A delegated central banker with output-stabilization weight λ^B minimizes

$$\mathcal{L}_t^B = \frac{1}{2}\left[\pi_t^2 + \lambda^B (x_t - \bar{x})^2\right], \qquad \lambda^B > 0. \tag{88}$$

A smaller λ^B therefore corresponds to a more inflation-conservative central banker. Under discretion, the target criterion becomes

$$\lambda^B (x_t - \bar{x}) + \kappa \pi_t = 0. \tag{89}$$

Together with the Phillips curve, this implies a positive steady-state inflation bias

$$\bar{\pi}^B = \frac{\kappa \lambda^B}{\kappa^2 + (1-\beta)\lambda^B}\,\bar{x} > 0, \qquad \frac{\partial \bar{\pi}^B}{\partial \lambda^B} > 0. \tag{90}$$

Thus the familiar time-consistency problem is present in the model: a more output-oriented policymaker has a stronger incentive to exploit nominal rigidities and generates a larger inflation bias, while a more conservative central banker reduces it. The distortion shifts equilibrium means but leaves the response coefficients to the cost-push shock unchanged from (35), with λ replaced by λ^B. All derivations are collected in Appendix F.

4.2. Recovering Rogoff's conservative central banker

Let society's own stabilization weight be λ^S, and suppose it evaluates monetary-policy outcomes according to the same class of stabilization objectives

$$\mathcal{L}_t^S = \frac{1}{2}\left[\pi_t^2 + \lambda^S (x_t - \bar{x})^2\right]. \tag{91}$$

Society chooses the central bank's stabilization weight λ^B before shocks are realized. When $\bar{x} > 0$, discretionary policy generates the familiar inflation bias,

$$\bar{\pi}^B = \frac{\kappa \lambda^B}{\kappa^2 + (1-\beta)\lambda^B}\,\bar{x} > 0, \qquad \frac{\partial \bar{\pi}^B}{\partial \lambda^B} > 0, \tag{92}$$

so placing less weight on output-gap stabilization reduces the temptation to generate surprise inflation. At the same time, changing λ^B also changes the central bank's response to shocks. The standard Rogoff trade-off therefore applies. Let λ^R denote an optimal delegated mandate in the corresponding certainty-equivalent economy. As shown in Appendix F.2, the standard Rogoff force is recovered:

$$0 < \lambda^R < \lambda^S. \tag{93}$$

Thus society appoints a central banker who places relatively more weight on inflation stabilization than society itself, but not an infinitely conservative one.

Aggregate risk pricing does not alter this delegation problem when the policy rate is unrestricted. For any given mandate λ^B, the separation identity (37) allows the central bank to offset the endogenous risk premium one-for-one through the safe policy rate. Hence the inflation–output allocation is exactly the same as in the certainty-equivalent economy for every λ^B. Since society evaluates delegated mandates solely through that allocation, its welfare ranking over λ^B is also unchanged. It follows immediately that

$$\lambda_{RP}^{B,*} = \lambda_{CE}^{B,*} = \lambda^R < \lambda^S. \tag{94}$$

Priced aggregate risk therefore affects only the interest rate required to implement Rogoff's delegated allocation:

$$\begin{aligned} i_t^{RP}(\lambda^B) &= i_t^{NK}(\lambda^B) - \Psi_t^\star(\lambda^B), \\ \Psi_t^\star(\lambda^B) &= \frac{1}{2}\left[\frac{\lambda^B - \kappa}{\kappa^2 + \lambda^B \chi_u}\right]^2 s_t. \end{aligned} \tag{95}$$

Thus, when the short rate can adjust freely, pricing the aggregate risk created by stabilization changes the implementation of Rogoff's solution, but not the optimal degree of central-bank conservatism itself. Appendix F provides the corresponding derivations.

4.3. Optimal delegation with priced aggregate risk

When $\delta > 0$, the separation result no longer applies: offsetting a rise in the risk premium requires a movement in the policy rate, and such movements are costly. I interpret δ as a common institutional preference for gradual policy-rate adjustment: society wants monetary policy to stabilize inflation and activity while avoiding abrupt changes in the policy rate, and delegates this same smoothing objective to the central bank. Society therefore evaluates outcomes according to

$$\mathcal{L}_t^S = \frac{1}{2}\left[\pi_t^2 + \lambda^S (x_t - \bar{x})^2 + \delta (i_t - i_{t-1})^2\right], \tag{96}$$

while policy is delegated to a central banker with the same implementation cost δ but a possibly different output-stabilization weight λ^B.

Let $W^{RP}(\lambda^B; V_s, \delta)$ denote society's unconditional expected loss under the equilibrium induced by mandate λ^B when aggregate risk is priced, where

$$V_s \equiv \mathrm{Var}(\hat{s}_t) \tag{97}$$

measures the variance of the time-varying cost-push uncertainty state. Let $W^{CE}(\lambda^B; \delta)$ denote the corresponding loss in the certainty-equivalent economy with the same degree of interest-rate inertia. The priced-risk economy can be written as

$$W^{RP}(\lambda^B; V_s, \delta) = W^{CE}(\lambda^B; \delta) + V_s\, \Omega(\lambda^B, \delta), \tag{98}$$

where

$$\Omega(\lambda^B, \delta) = \frac{1}{8} J_\Psi(\lambda^B, \delta)\left[\mathcal{E}_u^{\mathcal{I}}(\lambda^B, \delta)\right]^4 \left[\lambda^B - \lambda_H(\delta)\right]^4 \geq 0, \tag{99}$$

$J_\Psi > 0$ is the social loss generated by a unit risk-premium disturbance, and $\mathcal{E}_u^{\mathcal{I}} > 0$ maps the delegated mandate into nominal-wealth exposure. The zero-risk hedge is

$$\lambda_H(\delta) = \frac{\kappa^2}{\kappa + \delta}, \tag{100}$$

so that $\Omega(\lambda^B, \delta) = 0$ if and only if $\lambda^B = \lambda_H(\delta)$.

Define the certainty-equivalent Rogoff mandate and the optimal mandate with priced risk as

$$\begin{aligned} \lambda_R^B(\delta) &\in \arg\min_{\lambda^B \geq 0} W^{CE}(\lambda^B; \delta), \\ \lambda_{RP}^{B,*}(V_s, \delta) &\in \arg\min_{\lambda^B \geq 0} W^{RP}(\lambda^B; V_s, \delta). \end{aligned} \tag{101}$$

The standard delegation force implies $\lambda_R^B(\delta) < \lambda^S$: absent priced-risk considerations, society appoints a more inflation-conservative central banker. Proposition 3 revisits this result when aggregate risk is priced.

Proposition 3 (Optimal delegation with priced aggregate risk)**.** *Suppose $\delta > 0$, the maintained stable-branch regularity conditions hold, and the certainty-equivalent delegation problem has the conventional Rogoff ranking $\lambda_R^B(\delta) < \lambda^S$. For any closed set of mandates $\mathcal{A} \subset [0, \infty)$ that is bounded away from the zero-risk hedge $\lambda_H(\delta)$, define*

$$V^\star(\mathcal{A}) \equiv \sup_{\lambda \in \mathcal{A}} \frac{\left[W^{CE}(\lambda_H(\delta); \delta) - W^{CE}(\lambda; \delta)\right]_+}{\Omega(\lambda, \delta)}, \quad (102)$$

where $[z]_+ \equiv \max\{z, 0\}$. Then $V^\star(\mathcal{A}) < \infty$, and whenever $V_s > V^\star(\mathcal{A})$, every optimal delegated mandate lies outside $\mathcal{A}$.[6]

Consequently:

(i) Rogoff reversal. *If*

$$\lambda_R^B(\delta) < \lambda^S < \lambda_H(\delta), \quad (103)$$

let

$$\mathcal{A}_R = [0, \lambda^S], \qquad V_R^\star \equiv V^\star(\mathcal{A}_R). \quad (104)$$

For every admissible $V_s > V_R^\star$,

$$\lambda_{RP}^{B,*}(V_s, \delta) > \lambda^S. \quad (105)$$

Society optimally appoints a central banker who is less inflation-conservative than society itself.

(ii) Attenuated conservatism. *If*

$$\lambda_R^B(\delta) < \lambda_H(\delta) < \lambda^S, \quad (106)$$

let

$$\mathcal{A}_A = [0, \lambda_R^B(\delta)] \cup [\lambda^S, \infty), \qquad V_A^\star \equiv V^\star(\mathcal{A}_A). \quad (107)$$

For every admissible $V_s > V_A^\star$,

$$\lambda_R^B(\delta) < \lambda_{RP}^{B,*}(V_s, \delta) < \lambda^S. \quad (108)$$

The conservative-central-banker prescription survives, but optimal conservatism is weaker than in the certainty-equivalent Rogoff benchmark.

(iii) Amplified conservatism. *If*

$$\lambda_H(\delta) < \lambda_R^B(\delta) < \lambda^S, \quad (109)$$

let

$$\mathcal{A}_C = [\lambda_R^B(\delta), \infty), \qquad V_C^\star \equiv V^\star(\mathcal{A}_C). \quad (110)$$

For every admissible $V_s > V_C^\star$,

$$\lambda_{RP}^{B,*}(V_s, \delta) < \lambda_R^B(\delta). \quad (111)$$

Priced inflation risk calls for a central banker who is even more inflation-conservative than the conventional Rogoff delegate.

Proof. See Appendix F.4.

The proposition gives a simple global characterization of the new delegation problem. Rogoff's inflation-bias motive always pushes the delegated weight on output stabilization below society's own weight. Priced aggregate risk introduces a different force: it makes the zero-risk hedge $\lambda_H(\delta)$ an additional attractor for optimal delegation. The direction of this force is therefore not intrinsically *hawkish* or *dovish*. It depends on the position of the hedge relative to the certainty-equivalent Rogoff mandate $\lambda_R^B(\delta)$ and society's own preference λ^S. Figure 3 illustrates the three cases in Proposition 3 using numerical solutions of the model. In each panel, priced risk creates a force pulling the optimal delegated mandate away from the certainty-equivalent Rogoff benchmark and toward the zero-risk hedge. Once the finite threshold in the corresponding case is crossed, the relative ordering of $\lambda_R^B(\delta)$, λ^S, and $\lambda_H(\delta)$ determines whether this force reverses, attenuates, or amplifies the conventional conservative-central-banker prescription.

The most striking case is the Rogoff reversal in Figure 3, Panel (a). If society itself lies on the output-risk side of the hedge, an even more conservative banker absorbs a larger share of an inflationary supply disturbance through a contraction in activity. This raises the exposure of nominal wealth to bad output states and increases the risk premium. Because $\delta > 0$, the resulting tightening of private financial conditions cannot be offset one-for-one through the policy rate. When priced risk is sufficiently important, society therefore prefers to move the delegated mandate beyond its own preference and toward the hedge: the Rogoff ranking reverses, with $\lambda_{RP}^{B,*} > \lambda^S$. The numerical example confirms that this reversal occurs at a finite, admissible value of V_s.

Panel (b) illustrates the intermediate case in which $\lambda_R^B(\delta) < \lambda_H(\delta) < \lambda^S$. Priced risk again pushes optimal delegation away from the certainty-equivalent Rogoff mandate and toward the hedge, but the hedge remains more inflation-conservative than society. The standard Rogoff ranking therefore survives, although it is attenuated: $\lambda_R^B(\delta) < \lambda_{RP}^{B,*} < \lambda^S$. In Panel (c), the hedge lies to the inflation-conservative side of the Rogoff mandate, $\lambda_H(\delta) < \lambda_R^B(\delta) < \lambda^S$. The same risk-hedging force now points in the opposite direction, so priced inflation risk reinforces the conventional delegation motive and implies $\lambda_{RP}^{B,*} < \lambda_R^B(\delta)$.

Thus the general lesson is not that risk premia favor hawks or doves. They introduce a second target of optimal delegation—the mandate that minimizes the exposure of nominal wealth to aggregate risk. Whether this second force overturns or reinforces Rogoff depends on where $\lambda_H(\delta)$ lies relative to the conventional delegation benchmark and society's own stabilization preference.

Proposition 3 is a finite-threshold result. The numerical

[6] Under the bounded-support affine volatility process in (3)–(4), the stationary variance satisfies $V_s \leq \bar{V}_s$, with $\bar{V}_s$ given in (5). Thus each finite-threshold result below requires $V_j^\star < \bar{V}_s$ and an admissible $V_s \in (V_j^\star, \bar{V}_s]$. This is a technical restriction of the bounded affine specification used to preserve tractability, rather than part of the economic mechanism. A positive volatility specification with unbounded support, such as log variance, does not impose the same upper bound.

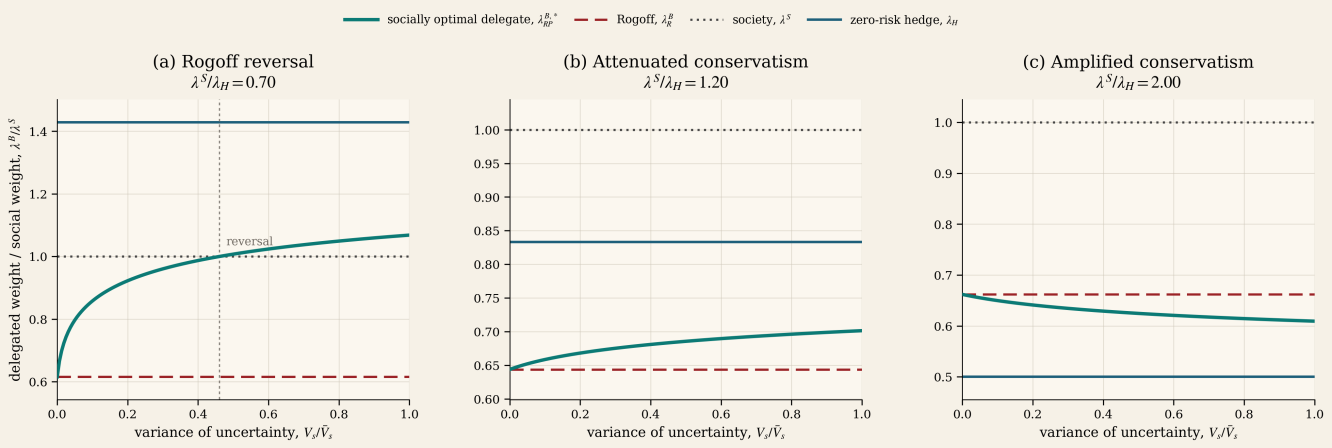


Figure 3: **Optimal delegation with priced aggregate risk.** *The figure provides a numerical illustration of Proposition 3, with optimal delegation computed by minimizing the model-implied social loss. The solid green line denotes the socially optimal delegated mandate $\lambda_{RP}^{B,*}(V_s, \delta)$, the dashed red line the certainty-equivalent Rogoff mandate $\lambda_R^B(\delta)$, the dotted line society's preference λ^S, and the solid blue line the zero-risk hedge $\lambda_H(\delta)$. The horizontal axis measures the variance of uncertainty, V_s, normalized by its maximum admissible value $\overline{V}_s$, while the vertical axis measures the delegated mandate relative to society's preference, λ^B/λ^S. As priced aggregate risk becomes more important, the risk-hedging force pulls optimal delegation toward $\lambda_H(\delta)$. Panel (a) shows the Rogoff reversal, $\lambda_R^B(\delta) < \lambda^S < \lambda_H(\delta)$. Panel (b) shows attenuated conservatism, $\lambda_R^B(\delta) < \lambda_H(\delta) < \lambda^S$. Panel (c) shows amplified conservatism, $\lambda_H(\delta) < \lambda_R^B(\delta) < \lambda^S$. All panels share the same structural parameters and differ only in society's stabilization preference relative to the hedge. The parameter values are illustrative and are not intended as an empirical calibration.*

solutions illustrate optimal delegation moving toward the hedge as V_s increases, but the proposition does not require this movement to be monotonic for every intermediate value of V_s; it only requires priced risk to be sufficiently important to dominate mandates on the relevant side of the hedge.

5. Conclusion

This paper revisits familiar questions about monetary-policy preferences once the macroeconomic risk generated by stabilization itself is priced by financial markets. In the benchmark with an unrestricted policy rate, risk pricing does not change the discretionary inflation–output allocation: the central bank can offset the financial wedge one-for-one through the short rate. It does, however, change the way that allocation is implemented. By determining how a cost-push disturbance is split between output and inflation, the central bank's stabilization preference determines the exposure of nominal wealth to aggregate risk and hence the equilibrium risk premium. This generates a U-shaped relation between the weight on output stabilization and priced risk, with a unique mandate at which output and inflation exposures hedge one another. Away from that hedge, part of the desired monetary tightening is supplied endogenously by financial markets. As a result, a more inflation-oriented central bank can produce a larger contraction in activity while moving the policy rate by less, so the short rate alone need not reveal either the stance or the preferences of monetary policy.

Once policy-rate adjustment is costly, risk pricing also changes the allocation itself. Interest-rate gradualism continues to act as implicit output stabilization through the conventional channel, but it simultaneously changes the risk borne by financial markets. On the inflation-risk side of the hedge, greater gradualism raises the risk premium; under sufficiently high supply uncertainty, this endogenous tightening can dominate the standard smoothing effect and make gradualism contractionary rather than stabilizing. The same failure of separation changes optimal delegation. In the certainty-equivalent economy, the usual Rogoff force favors a central banker who is more inflation-conservative than society. With priced aggregate risk, a second force pulls the delegated mandate toward the zero-risk hedge. Depending on where that hedge lies relative to society's preference and the conventional Rogoff mandate, priced risk can attenuate, reinforce, or overturn the conservative-central-banker prescription. In particular, when the hedge lies on the output-oriented side of society's own preference, sufficiently important aggregate risk leads society to appoint a central banker who places more weight on output stabilization than society itself.

The broader message is that certainty equivalence suppresses a set of interactions between monetary stabilization, asset pricing, and policy design that become economically rich once aggregate risk is priced. The present note isolates these mechanisms in a deliberately stylized environment: a single supply disturbance, segmented households, Gaussian risk, a log-linear supply block, and a reduced-form cost of policy-rate adjustment. A natural next step is therefore to study how the results survive in richer environments with endogenous risk-bearing capacity, heterogeneous households and intermediaries, several

sources of aggregate risk, and fully microfounded policy-adjustment costs. More broadly, the analysis suggests that revisiting standard New Keynesian prescriptions outside the certainty-equivalence benchmark may reveal policy trade-offs that are invisible in first-order linear-quadratic models but central when monetary policy itself shapes the risk that financial markets are asked to bear.

References

Abbate, Angela, and Dominik Thaler. 2023. "Optimal Monetary Policy with the Risk-Taking Channel." *European Economic Review* 152: 104333.

Basu, Susanto, and Brent Bundick. 2017. "Uncertainty Shocks in a Model of Effective Demand." *Econometrica* 85 (3): 937–958.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–685.

Bocola, Luigi, Alessandro Dovis, Kasper Jørgensen, and Rishabh Kirpalani. 2024. "Bond Market Views of the Fed." *NBER Working Paper* 32620.

Caballero, Ricardo J., Tomás E. Caravello, and Alp Simsek. 2026. "FCI-Star." *NBER Working Paper* 33952. First circulated June 2025; revised January 2026.

Caballero, Ricardo J., and Alp Simsek. 2020. "A Risk-Centric Model of Demand Recessions and Speculation." *Quarterly Journal of Economics* 135 (3): 1493–1566.

Caballero, Ricardo J., and Alp Simsek. 2022. "A Monetary Policy Asset Pricing Model." *NBER Working Paper* 30132.

Caballero, Ricardo J., and Alp Simsek. 2025. "Financial Conditions Targeting in a Multi-Asset Open Economy." Working paper, November 2025.

Cho, Daeha, Yoonshin Han, Joonseok Oh, and Anna Rogantini Picco. 2021. "Uncertainty Shocks, Precautionary Pricing, and Optimal Monetary Policy." *Journal of Macroeconomics* 69: 103343.

Cho, Daeha, and Joonseok Oh. 2023. "The Source of Uncertainty and Optimal Monetary Policy." *Economics Letters* 227: 111131.

Cúrdia, Vasco, and Michael Woodford. 2016. "Credit Frictions and Optimal Monetary Policy." *Journal of Monetary Economics* 84: 30–65.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl. 2018. "Liquidity, Risk Premia, and the Financial Transmission of Monetary Policy." *Annual Review of Financial Economics* 10: 309–328.

Favero, Carlo A., and Riccardo Rovelli. 2003. "Macroeconomic Stability and the Preferences of the Fed: A Formal Analysis, 1961–98." *Journal of Money, Credit and Banking* 35 (4): 545–556.

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Keith Kuester, and Juan Rubio-Ramírez. 2015. "Fiscal Volatility Shocks and Economic Activity." *American Economic Review* 105 (11): 3352–3384.

Galí, Jordi. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton, NJ: Princeton University Press.

Galí, Jordi. 2026. "Rethinking Monetary Policy Rules." Working paper, CREI, March 2026.

Kekre, Rohan, and Moritz Lenel. 2022. "Monetary Policy, Redistribution, and Risk Premia." *Econometrica* 90 (5): 2249–2282.

Lohmann, Susanne. 1992. "Optimal Commitment in Monetary Policy: Credibility versus Flexibility." *American Economic Review* 82 (1): 273–286.

Palomino, Francisco. 2012. "Bond Risk Premiums and Optimal Monetary Policy." *Review of Economic Dynamics* 15 (1): 19–40.

Persson, Torsten, and Guido Tabellini. 1993. "Designing Institutions for Monetary Stability." *Carnegie-Rochester Conference Series on Public Policy* 39: 53–84.

Pflueger, Carolin, and Gianluca Rinaldi. 2022. "Why Does the Fed Move Markets so Much? A Model of Monetary Policy and Time-Varying Risk Aversion." *Journal of Financial Economics* 146 (1): 71–89.

Rogoff, Kenneth. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (4): 1169–1189.

Sun, Rui. 2026. "Risk Capacity and Optimal Monetary Policy." Working paper.

Surico, Paolo. 2007. "The Fed's Monetary Policy Rule and U.S. Inflation: The Case of Asymmetric Preferences." *Journal of Economic Dynamics and Control* 31 (1): 305–324.

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

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

Woodford, Michael. 2003. "Optimal Interest-Rate Smoothing." *Review of Economic Studies* 70 (4): 861–886.

A. Model Setup

This environment follows Caballero and Simsek (2022). There are two household groups and a standard New Keynesian production block. Workers consume the final good, supply all labor, and do not participate in asset markets. Capitalists consume the final good, supply no labor, own firms and financial wealth, and choose their portfolio directly among the market claim, a one-period real safe bond, and a one-period nominal safe bond. Intermediate firms produce differentiated varieties using labor and set prices subject to Calvo frictions. A competitive final-goods firm aggregates those varieties into the unique consumption good. The central bank sets the nominal short rate.

The exogenous supply block is the scalar cost-push process (2) with stochastic conditional variance (3). Conditional on $\mathcal{I}_t$, the level innovation η_{t+1} is Gaussian with variance s_t. The volatility innovation ν^s_{t+1} is conditionally independent of η_{t+1} and has the bounded, time-invariant, conditionally symmetric distribution specified in Section 2.1. The support restriction (4) keeps the conditional variance strictly positive while preserving an affine volatility state.

Final good firm. The representative final-goods firm solves

$$\max_{\{Y_t(\nu)\}_{\nu\in[0,1]}} \left\{ Q_t Y_t - \int_0^1 P_t(\nu) Y_t(\nu)\, d\nu \right\}$$

subject to

$$Y_t = \left(\int_0^1 Y_t(\nu)^{\frac{\varepsilon-1}{\varepsilon}}\, d\nu \right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

Cost minimization implies the standard demand system

$$Y_t(\nu) = \left(\frac{P_t(\nu)}{Q_t} \right)^{-\varepsilon} Y_t,$$

with price index

$$Q_t = \left(\int_0^1 P_t(\nu)^{1-\varepsilon}\, d\nu \right)^{\frac{1}{1-\varepsilon}}.$$

Intermediate-goods firms. There is a continuum of monopolistically competitive intermediate-goods firms indexed by $\nu \in [0,1]$. Each firm produces a differentiated variety using labor only,

$$Y_t(\nu) = \mathcal{A}_t L_t(\nu)^{1-\alpha},$$

where $1-\alpha$ is labor's technological share. In the baseline, aggregate productivity is constant, so the only aggregate supply disturbance is the cost-push shock u_t. The corresponding flexible-price level of output is the natural-output benchmark y^n_t used in the main text.

Demand for each variety is

$$Y_t(\nu) = \left(\frac{P_t(\nu)}{Q_t}\right)^{-\varepsilon} Y_t.$$

Following standard practice in the New Keynesian literature (e.g. Galí (2008)), firms receive a constant subsidy at rate τ^s on their wage bill, financed through lump-sum taxation. Let $\mathcal{F}_t(\nu)$ denote the lump-sum tax on firm ν. Period profits are

$$\Pi_t(\nu) = P_t(\nu)Y_t(\nu) - (1-\tau^s)W_tL_t(\nu) - \mathcal{F}_t(\nu).$$

The lump-sum tax is taken as given by the individual firm and therefore does not affect marginal pricing or employment decisions.

Let

$$\mu \equiv \frac{\varepsilon}{\varepsilon - 1}$$

denote the desired gross monopoly markup. The subsidy is chosen so that

$$1 - \tau^s = \frac{1}{\mu} = \frac{\varepsilon - 1}{\varepsilon}, \qquad \tau^s = \frac{1}{\varepsilon}.$$

In the symmetric flexible-price allocation the firm's pricing condition is

$$Q_t = \mu(1-\tau^s)\frac{W_t}{MPL_t}.$$

Since $\mu(1-\tau^s) = 1$,

$$\frac{W_t}{Q_t} = MPL_t.$$

Thus the subsidy removes the steady-state monopoly wedge between the real wage and the marginal product of labor. Because τ^s is constant, its log deviation is zero and it does not alter the log-linear New Keynesian Phillips curve.

Whenever allowed to reoptimize, firm ν chooses its price to maximize the expected discounted value of profits subject to its demand schedule. Together with Calvo price setting, this gives the first-order supply block used in the main text,

$$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + \kappa x_t + u_t.$$

Define aggregate labor and aggregate lump-sum firm taxes as

$$L_t \equiv \int_0^1 L_t(\nu)\,d\nu, \qquad \mathcal{F}_t \equiv \int_0^1 \mathcal{F}_t(\nu)\,d\nu.$$

Using the zero-profit condition of the final-good firm,

$$\int_0^1 P_t(\nu)Y_t(\nu)\,d\nu = Q_tY_t,$$

aggregate after-tax profits are

$$\Pi_t \equiv \int_0^1 \Pi_t(\nu)\,d\nu = Q_tY_t - (1-\tau^s)W_tL_t - \mathcal{F}_t.$$

Fiscal authority. Let $\mathcal{T}_t$ denote the lump-sum transfer paid to workers. The government finances both this transfer and the employment subsidy through the lump-sum tax on firms. Its period-by-period balanced-budget condition is

$$\mathcal{F}_t = \mathcal{T}_t + \tau^s W_tL_t.$$

Substituting into aggregate profits gives

$$\begin{aligned}\Pi_t &= Q_tY_t - (1-\tau^s)W_tL_t - (\mathcal{T}_t + \tau^s W_tL_t)\\ &= Q_tY_t - W_tL_t - \mathcal{T}_t. \end{aligned} \tag{112}$$

The fiscal instruments are therefore purely redistributive at the aggregate level and absorb no real resources.

Workers. A representative worker solves

$$\max_{C_t^w, L_t}\left\{\log C_t^w - \chi_L\frac{L_t^{1+\varphi}}{1+\varphi}\right\}$$

subject to

$$Q_tC_t^w = W_tL_t + \mathcal{T}_t.$$

Because the transfer is lump-sum, the worker takes it as given when choosing labor. The labor-supply condition is

$$\frac{W_t}{Q_tC_t^w} = \chi_L L_t^{\varphi}.$$

Following Caballero and Simsek (2022), the transfer is chosen so that workers receive the technological labor share of aggregate output,

$$\mathcal{T}_t = (1-\alpha)Q_tY_t - W_tL_t.$$

Worker disposable income therefore satisfies

$$W_tL_t + \mathcal{T}_t = (1-\alpha)Q_tY_t,$$

and hence

$$C_t^w = (1-\alpha)Y_t. \tag{113}$$

Thus the worker consumption share is constant at all dates and in all states.

Combining (112) and (A) gives

$$\Pi_t = \alpha Q_tY_t. \tag{114}$$

Hence aggregate dividends accruing to capitalists are exactly the residual technological share αQ_tY_t.

Capitalists. Capitalists own the firms, so the profits in (114) are paid as dividends on the market claim. They also trade a one-period real safe bond and a one-period nominal safe bond directly.

Let

$$X_t \equiv R_t^p A_{t-1}$$

denote beginning-of-period real financial wealth, where A_{t-1} is saving carried from period $t-1$ and R_t^p is the

realized gross return on the capitalist's portfolio. A representative capitalist solves

$$V(X_t) = \max_{C_t^c, A_t, \omega_t, \omega_t^n} \{\log C_t^c + \beta \mathbb{E}_t[V(X_{t+1})]\}$$

subject to

$$C_t^c + A_t = X_t, \qquad X_{t+1} = A_t R_{t+1}^p,$$

where

$$R_{t+1}^p = R_t^f + \omega_t (R_{t+1} - R_t^f) + \omega_t^n \left(\frac{I_t}{Q_{t+1}/Q_t} - R_t^f \right). \tag{115}$$

Here R_{t+1} is the gross real return on the market claim, R_t^f is the one-period gross real safe return, and $I_t = \exp(i_t)$ is the gross nominal safe return.

Guessing

$$V(X) = a + b \log X$$

gives

$$b = \frac{1}{1-\beta}.$$

The consumption-saving first-order condition then implies

$$C_t^c = (1-\beta) X_t = (1-\beta) R_t^p A_{t-1}, \qquad A_t = \beta X_t. \tag{116}$$

Thus capitalists consume the constant fraction $1-\beta$ of financial wealth and save the remainder.

Conditional on saving A_t, portfolio choice satisfies

$$\mathbb{E}_t \left[\frac{R_{t+1} - R_t^f}{R_{t+1}^p} \right] = 0, \tag{117}$$

$$\mathbb{E}_t \left[\frac{I_t/(Q_{t+1}/Q_t) - R_t^f}{R_{t+1}^p} \right] = 0. \tag{118}$$

The capitalist consumption stochastic discount factor is

$$M_{t+1} \equiv \beta \frac{C_t^c}{C_{t+1}^c}.$$

Using (116) and $X_{t+1} = A_t R_{t+1}^p$ gives the exact identity

$$M_{t+1} = \frac{1}{R_{t+1}^p}. \tag{119}$$

Thus the inverse-wealth-return stochastic discount factor follows directly from capitalist optimization; no separate portfolio-management sector is required.

In equilibrium the real and nominal safe bonds are in zero net supply, so aggregate capitalist wealth is entirely the market claim. Hence

$$R_{t+1}^p = R_{t+1}$$

in equilibrium and

$$M_{t+1} = \frac{1}{R_{t+1}}. \tag{120}$$

Goods-market and asset-market clearing. Since there are no government purchases or investment,

$$Y_t = C_t^w + C_t^c.$$

Together with (113), this implies

$$C_t^c = \alpha Y_t. \tag{121}$$

Therefore

$$C_t^w = (1-\alpha) Y_t, \qquad C_t^c = \alpha Y_t.$$

Let P_t denote the ex-dividend real price of the market claim. Since capitalists receive current dividends αY_t and own the market portfolio, equilibrium financial wealth satisfies

$$R_t^p A_{t-1} = \alpha Y_t + P_t. \tag{122}$$

Combining (116), (121), and (122) yields the output–asset-price relation used in the main text,

$$Y_t = \frac{1-\beta}{\alpha\beta} P_t. \tag{123}$$

B. Derivation of the Risk-Centric IS Curve

This appendix derives the aggregate-demand block used in the text. The key simplification relative to earlier versions of the model is that the output–asset-price relation implies an exact return identity, so no Campbell–Shiller approximation is required.

Exact output–wealth relation and market return. From Appendix A,

$$Y_t = \frac{1-\beta}{\alpha\beta} P_t. \tag{124}$$

Equivalently,

$$P_t = \frac{\alpha\beta}{1-\beta} Y_t.$$

The market claim pays the capitalist share αY_{t+1} as its dividend and has ex-dividend price P_{t+1}. Its gross real return is therefore

$$R_{t+1} = \frac{\alpha Y_{t+1} + P_{t+1}}{P_t}. \tag{125}$$

Substituting (124) into (125) gives

$$\begin{aligned} R_{t+1} &= \frac{\alpha Y_{t+1} + \frac{\alpha\beta}{1-\beta} Y_{t+1}}{\frac{\alpha\beta}{1-\beta} Y_t} \\ &= \frac{1}{\beta} \frac{Y_{t+1}}{Y_t}. \end{aligned} \tag{126}$$

Hence the log return satisfies the exact identity

$$r_{t+1} = \rho + y_{t+1} - y_t, \qquad \rho \equiv -\log \beta. \tag{127}$$

Taking conditional expectations,

$$\mathbb{E}_t r_{t+1} = \rho + \mathbb{E}_t \Delta y_{t+1}. \tag{128}$$

Under the flexible-price allocation the same relation gives

$$r^n_{t+1} = \rho + y^n_{t+1} - y^n_t, \qquad \bar{r}^n_t \equiv \mathbb{E}_t r^n_{t+1} = \rho + \mathbb{E}_t \Delta y^n_{t+1}. \tag{129}$$

Since $y_t = x_t + y^n_t$,

$$r_{t+1} - r^n_{t+1} = x_{t+1} - x_t,$$

and therefore

$$x_t = \mathbb{E}_t x_{t+1} - (\mathbb{E}_t r_{t+1} - \bar{r}^n_t). \tag{130}$$

Equation (130) is the primitive aggregate-demand relation: the output gap is governed by the difference between the expected return required on actual aggregate wealth and its natural-allocation counterpart.

Approximation convention. The demand and financial block above is exact. In particular, (127) follows algebraically from equilibrium goods and asset-market clearing, and the asset-pricing conditions below retain their nonlinear exponential form. The paper deliberately combines this exact financial block with a first-order New Keynesian supply block. This is not intended as a uniform higher-order perturbation of a fully nonlinear economy. Rather, the asymmetry isolates the risk-pricing channel while preserving the familiar inflation–output stabilization problem and analytical tractability. Keeping the Phillips curve log-linear shuts down additional higher-order supply-side mechanisms, such as precautionary price setting, which are conceptually distinct from the financial channel studied here.

Conditional second moments enter because the safe-bond Euler equations are nonlinear, not because selected second-order terms are retained from a perturbation expansion. In particular, if a return innovation z_{t+1} is conditionally Gaussian,

$$\log \mathbb{E}_t \exp(z_{t+1}) = \mathbb{E}_t z_{t+1} + \frac{1}{2}\mathrm{Var}_t(z_{t+1}), \tag{131}$$

exactly. More generally, the relevant conditional exponential moment is retained without approximation. The bounded volatility innovation ν^s_{t+1} has a time-invariant distribution; conditional on a fixed policy regime, any risk adjustment associated solely with that time-invariant component is absorbed into the regime-specific constant, while variation in s_t generates the time-varying priced-risk component studied in the text.

Asset pricing and real wealth risk. From (120), the equilibrium stochastic discount factor is

$$M_{t+1} = \frac{1}{R_{t+1}}.$$

Let $R^f_t = \exp(r^f_t)$ denote the gross one-period real safe return. Its Euler equation is therefore

$$1 = \mathbb{E}_t \left[\frac{R^f_t}{R_{t+1}} \right] = \mathbb{E}_t \exp(r^f_t - r_{t+1}). \tag{132}$$

Define the real wealth-return innovation

$$\varepsilon^R_{t+1} \equiv r_{t+1} - \mathbb{E}_t r_{t+1}.$$

Then

$$1 = \exp\left(r^f_t - \mathbb{E}_t r_{t+1}\right) \mathbb{E}_t \exp(-\varepsilon^R_{t+1}),$$
$$\mathbb{E}_t r_{t+1} = r^f_t + \log \mathbb{E}_t \exp(-\varepsilon^R_{t+1}). \tag{133}$$

This is the exact compensation associated with real aggregate wealth-return risk. Because $r_{t+1} = \rho + \Delta y_{t+1}$ exactly, the underlying risk is aggregate output risk.

In the conditionally Gaussian benchmark used for the closed-form mechanism in the main text, define

$$\mathcal{V}^y_t \equiv \mathrm{Var}_t(r_{t+1}). \tag{134}$$

Then

$$\mathbb{E}_t r_{t+1} = r^f_t + \frac{1}{2}\mathcal{V}^y_t. \tag{135}$$

Thus $\mathcal{V}^y_t$ denotes the conditional variance itself, rather than the compensation term $\frac{1}{2}\mathcal{V}^y_t$ that appears in the Euler equation.

The flexible-price allocation admits the same decomposition. Exactly,

$$\bar{r}^n_t = r^{f,n}_t + \log \mathbb{E}_t \exp\left[-\left(r^n_{t+1} - \mathbb{E}_t r^n_{t+1}\right)\right]. \tag{136}$$

Under conditional Gaussianity, define

$$\mathcal{V}^n_t \equiv \mathrm{Var}_t(r^n_{t+1}), \tag{137}$$

so the natural safe rate is

$$r^{f,n}_t = \bar{r}^n_t - \frac{1}{2}\mathcal{V}^n_t = \rho + \mathbb{E}_t \Delta y^n_{t+1} - \frac{1}{2}\mathcal{V}^n_t. \tag{138}$$

Nominal bonds and the aggregate risk adjustment. Let $I_t = \exp(i_t)$ denote the gross nominal safe return and $Q_{t+1}/Q_t = \exp(\pi_{t+1})$ gross inflation. The nominal-bond Euler equation is

$$1 = \mathbb{E}_t \left[\frac{I_t}{R_{t+1}(Q_{t+1}/Q_t)} \right]. \tag{139}$$

Define the nominal return on aggregate wealth,

$$r^N_{t+1} \equiv r_{t+1} + \pi_{t+1}, \tag{140}$$

and its innovation,

$$\varepsilon^N_{t+1} \equiv r^N_{t+1} - \mathbb{E}_t r^N_{t+1}. \tag{141}$$

Then (139) can be written as

$$\begin{aligned} 1 &= \mathbb{E}_t \exp(i_t - r^N_{t+1}) \\ &= \exp\left(i_t - \mathbb{E}_t r^N_{t+1}\right) \mathbb{E}_t \exp(-\varepsilon^N_{t+1}), \end{aligned} \tag{142}$$

which implies

$$\mathbb{E}_t r^N_{t+1} = i_t + \Psi_t, \tag{143}$$

where

$$\Psi_t \equiv \log \mathbb{E}_t \exp(-\varepsilon^N_{t+1}). \tag{144}$$

This is the aggregate risk adjustment that enters the Risk-Centric IS curve.

Since $r^N_{t+1} = r_{t+1} + \pi_{t+1}$, (143) is equivalently

$$\mathbb{E}_t r_{t+1} = i_t - \mathbb{E}_t \pi_{t+1} + \Psi_t. \tag{145}$$

Substituting into (130) gives

$$x_t = \mathbb{E}_t x_{t+1} - (i_t - \mathbb{E}_t \pi_{t+1} - \bar{r}^n_t) - \Psi_t, \tag{146}$$

which is equation (29) in the main text.

Thus no asset-price approximation is needed to obtain the Risk-Centric IS curve. The exact output–wealth equilibrium maps the expected return on aggregate wealth into aggregate demand, while the nominal-bond Euler equation maps the policy rate into that expected return through the compensation required for bearing nominal wealth risk.

Relation to conventional risk-premium measures. For completeness, define the nominal log market risk premium as

$$rp^N_t \equiv \log\left(\frac{\mathbb{E}_t R^N_{t+1}}{I_t}\right) = \log \mathbb{E}_t \exp(r^N_{t+1}) - i_t. \tag{147}$$

Using $r^N_{t+1} = \mathbb{E}_t r^N_{t+1} + \varepsilon^N_{t+1}$ and (143),

$$rp^N_t = \Psi_t + \log \mathbb{E}_t \exp(\varepsilon^N_{t+1}). \tag{148}$$

If the conditional distribution of ε^N_{t+1} is symmetric,

$$\mathbb{E}_t \exp(\varepsilon^N_{t+1}) = \mathbb{E}_t \exp(-\varepsilon^N_{t+1}),$$

so

$$rp^N_t = 2\Psi_t, \qquad \Psi_t = \frac{1}{2} rp^N_t. \tag{149}$$

Conditional Gaussianity gives the particularly transparent special case

$$\Psi_t = \frac{1}{2}\mathrm{Var}_t(r^N_{t+1}), \qquad rp^N_t = \mathrm{Var}_t(r^N_{t+1}). \tag{150}$$

Output risk, inflation risk, and the Fisher decomposition. It is also useful to decompose Ψ_t into the two conditional risks and their dependence. Write

$$\varepsilon^\pi_{t+1} \equiv \pi_{t+1} - \mathbb{E}_t \pi_{t+1}.$$

Combining the real- and nominal-bond Euler equations gives the exact relation

$$\begin{aligned} r^f_t =& i_t - \mathbb{E}_t \pi_{t+1} + \log \mathbb{E}_t \exp(-\varepsilon^\pi_{t+1}) \\ &- \Big\{ \log \mathbb{E}_t \exp(-\varepsilon^R_{t+1}) + \log \mathbb{E}_t \exp(-\varepsilon^\pi_{t+1}) - \log \mathbb{E}_t \exp[-(\varepsilon^R_{t+1} - \end{aligned} \tag{151}$$

The object in braces isolates the interaction between inflation and real wealth returns. Outside the Gaussian benchmark it is a cumulant interaction rather than a covariance, so I leave it unnamed.

Under joint conditional Gaussianity, define

$$\mathcal{V}^y_t \equiv \mathrm{Var}_t(r_{t+1}), \qquad \mathcal{V}^\pi_t \equiv \mathrm{Var}_t(\pi_{t+1}), \qquad \mathcal{H}_t \equiv -\mathrm{Cov}_t(r_{t+1}, \pi_{t+1}). \tag{152}$$

The interaction term in braces is then exactly $\mathcal{H}_t$. Hence

$$r^f_t = i_t - \mathbb{E}_t \pi_{t+1} + \frac{1}{2}\mathcal{V}^\pi_t - \mathcal{H}_t. \tag{153}$$

Likewise,

$$\Psi_t = \frac{1}{2}\mathrm{Var}_t(r_{t+1} + \pi_{t+1}) = \frac{1}{2}\mathcal{V}^y_t + \frac{1}{2}\mathcal{V}^\pi_t - \mathcal{H}_t. \tag{154}$$

Thus the main-text decomposition is literally a variance decomposition of nominal wealth. $\mathcal{V}^y_t$ is output risk, $\mathcal{V}^\pi_t$ is inflation risk, and $\mathcal{H}_t$ is the amount of those risks removed by their negative covariance. In conventional asset-pricing terminology $\mathcal{H}_t$ is also the inflation risk premium on nominal bonds, but the hedging interpretation is the one relevant for the policy mechanism.

C. Proof of Proposition 1

The target criterion and NKPC imply (35). The conditional expectations are

$$\mathbb{E}_t x^\star_{t+1} = x^\star_u \rho_u u_t, \qquad \mathbb{E}_t \pi^\star_{t+1} = \pi^\star_u \rho_u u_t,$$

where

$$x^\star_u = -\frac{\kappa}{D(\lambda)}, \qquad \pi^\star_u = \frac{\lambda}{D(\lambda)}.$$

Because the cost-push disturbance does not move natural output, the innovation in the nominal return on aggregate wealth is

$$r^{N,\star}_{t+1} - \mathbb{E}_t r^{N,\star}_{t+1} = (x^\star_u + \pi^\star_u)\,\eta_{t+1} = b(\lambda)\eta_{t+1}.$$

Conditional Gaussianity therefore gives

$$\Psi^\star_t = \frac{1}{2} b(\lambda)^2 s_t = \frac{1}{2}\left[\frac{\lambda - \kappa}{\kappa^2 + \lambda\chi_u}\right]^2 s_t, \tag{155}$$

which is (40).

Inverting the Risk-Centric IS curve and using the optimal allocation yields

$$i_t^{RC,\star} = \bar{r}_t^n + \left[(x_u^\star + \pi_u^\star)\rho_u - x_u^\star\right] u_t - \Psi_t^\star.$$

The coefficient on u_t is

$$(x_u^\star + \pi_u^\star)\rho_u - x_u^\star = \frac{\kappa(1-\rho_u) + \lambda\rho_u}{D(\lambda)} = \omega_u(\lambda),$$

so

$$i_t^{RC,\star} = i_t^{NK,\star} - \Psi_t^\star.$$

Holding the level of the cost-push shock fixed, a rise in s_t therefore raises the risk-premium wedge and lowers the implementing policy rate one-for-one. This proves the proposition. □

D. Interest-rate inertia: affine equilibrium and proof of Proposition ??

This appendix derives the affine equilibrium used in Section 3 and proves Proposition **??**. Work in deviations from the regime-specific Risk-Centric steady state. In addition to the assumptions in the main text, assume that the conditional law of ν_{t+1}^s is time invariant, independent of η_{t+1}, and has a finite moment-generating function in a neighborhood of the origin.

D.1. Affine equilibrium

The equilibrium conditions are

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t \pi_{t+1} + u_t, \tag{156}$$

$$x_t = \mathbb{E}_t(x_{t+1} + \pi_{t+1}) - \hat{i}_t - \hat{\Psi}_t^{\mathcal{I}}, \tag{157}$$

$$\lambda x_t + \kappa\pi_t = \delta(\hat{i}_t - \hat{i}_{t-1}). \tag{158}$$

Inherited-rate block. Conjecture coefficients (x_i, π_i, q_i) on $\hat{i}_{t-1}$ and define $L_i \equiv x_i + \pi_i$. Since $\hat{i}_t = q_i \hat{i}_{t-1}$ within this block, coefficient matching gives

$$\pi_i = \kappa x_i + \beta q_i \pi_i, \qquad x_i = (L_i - 1)q_i, \qquad \lambda x_i + \kappa\pi_i = \delta(q_i - 1). \tag{159}$$

Define

$$\mathcal{D}(q) \equiv (1-q)(1-\beta q) - \kappa q. \tag{160}$$

The first two equations in (159) imply

$$x_i = -\frac{q_i(1-\beta q_i)}{\mathcal{D}(q_i)}, \qquad \pi_i = -\frac{\kappa q_i}{\mathcal{D}(q_i)}, \qquad 1 - L_i = \frac{1-\beta q_i}{\mathcal{D}(q_i)}. \tag{161}$$

Substitution into the last equation in (159) gives the root condition

$$q_i\left[\lambda(1-\beta q_i) + \kappa^2\right] + \delta(q_i - 1)\mathcal{D}(q_i) = 0. \tag{162}$$

Let $\bar{q} \in (0,1)$ denote the unique root of $\mathcal{D}(q) = 0$. On the locally stable branch, $q_i \in (0, \bar{q})$ and therefore $\mathcal{D}(q_i) > 0$. To see existence and uniqueness, rewrite (162) on this interval as

$$\delta = R(q_i, \lambda) \equiv \frac{q_i\left[\lambda(1-\beta q_i) + \kappa^2\right]}{(1-q_i)\mathcal{D}(q_i)}. \tag{163}$$

The function R rises from zero to infinity on $(0, \bar{q})$. In particular,

$$\frac{R_q}{R} = \frac{1}{q} + \frac{1}{1-q} - \frac{\beta\lambda}{\lambda(1-\beta q) + \kappa^2} - \frac{\mathcal{D}'(q)}{\mathcal{D}(q)} > 0, \tag{164}$$

because $\mathcal{D}'(q) < 0$ on $(0,1)$ and

$$\frac{1}{q} + \frac{1}{1-q} - \frac{\beta\lambda}{\lambda(1-\beta q) + \kappa^2} = \frac{\kappa^2 + \lambda[1 - \beta q(2-q)]}{q(1-q)[\lambda(1-\beta q) + \kappa^2]} > 0. \tag{165}$$

Thus the stable root is unique. Moreover,

$$\frac{\partial q_i}{\partial \delta} = \frac{1}{R_q} > 0, \qquad \frac{\partial q_i}{\partial \lambda} = -\frac{R_\lambda}{R_q} < 0, \tag{166}$$

since $R_\lambda > 0$.

Forcing blocks. For a generic AR(1) forcing variable w_t satisfying $\mathbb{E}_t w_{t+1} = \varrho w_t$, suppose that it enters the Phillips curve with coefficient c_π and the IS curve with coefficient c_x. Define

$$\Delta(\varrho) \equiv (1 - L_i)\left[\kappa^2 + \lambda(1 - \beta\varrho)\right] - \beta\left[\kappa(1-\varrho) + \lambda\varrho\right]\pi_i + \delta\left[1 - \varrho(1 + \beta + \kappa) + \beta\varrho^2\right]. \tag{167}$$

Assume

$$\Delta(\rho_u) > 0, \qquad \Delta(\rho_s) > 0. \tag{168}$$

Coefficient matching in (156)–(158) gives

$$q_w = \frac{c_x[\kappa^2 + \lambda(1-\beta\varrho)] + c_\pi[\kappa(1-\varrho) + \lambda\varrho]}{\Delta(\varrho)}, \tag{169}$$

$$x_w = \frac{c_\pi[\delta\varrho - \kappa(1-L_i)] + c_x[\delta(1-\beta\varrho) - \beta\kappa\pi_i]}{\Delta(\varrho)}, \tag{170}$$

$$\pi_w = \frac{c_\pi[\lambda(1-L_i) + \delta(1-\varrho)] + c_x[\delta\kappa + \beta\lambda\pi_i]}{\Delta(\varrho)}. \tag{171}$$

For the cost-push shock, $(w_t, \varrho, c_\pi, c_x) = (u_t, \rho_u, 1, 0)$, so

$$q_u = \frac{\kappa(1-\rho_u) + \lambda\rho_u}{\Delta(\rho_u)} > 0, \tag{172}$$

$$x_u = \frac{\delta\rho_u - \kappa(1-L_i)}{\Delta(\rho_u)}, \tag{173}$$

$$\pi_u = \frac{\lambda(1-L_i) + \delta(1-\rho_u)}{\Delta(\rho_u)} > 0. \tag{174}$$

For the risk-premium state, $(w_t, \varrho, c_\pi, c_x) =$

$(\hat{\Psi}_t^{\mathcal{I}}, \rho_s, 0, -1)$, so

$$q_\Psi = -\frac{\kappa^2 + \lambda(1-\beta\rho_s)}{\Delta(\rho_s)} < 0, \tag{175}$$

$$x_\Psi = -\frac{\delta(1-\beta\rho_s) - \beta\kappa\pi_i}{\Delta(\rho_s)} < 0, \tag{176}$$

$$\pi_\Psi = -\frac{\delta\kappa + \beta\lambda\pi_i}{\Delta(\rho_s)} < 0. \tag{177}$$

The last sign follows from the root condition, which implies

$$\delta\kappa + \beta\lambda\pi_i = \frac{\kappa q_i[\kappa^2 + \lambda(1-\beta)]}{(1-q_i)\mathcal{D}(q_i)} > 0. \tag{178}$$

Equations (161) and (172)–(177) establish the affine representation in (59)–(61).

D.2. The equilibrium risk premium

Define $L_\Psi^{\mathcal{I}} \equiv x_\Psi + \pi_\Psi$ and conjecture

$$\hat{\Psi}_t^{\mathcal{I}} = G_u^{\mathcal{I}} \hat{s}_t. \tag{179}$$

Then

$$\hat{\Psi}_{t+1}^{\mathcal{I}} = \rho_s \hat{\Psi}_t^{\mathcal{I}} + G_u^{\mathcal{I}} \nu_{t+1}^s. \tag{180}$$

The innovation in the nominal return on aggregate wealth is therefore

$$r_{t+1}^{N,\mathcal{I}} - \mathbb{E}_t r_{t+1}^{N,\mathcal{I}} = L_u^{\mathcal{I}} \eta_{t+1} + L_\Psi^{\mathcal{I}} G_u^{\mathcal{I}} \nu_{t+1}^s. \tag{181}$$

Conditional independence and conditional Gaussianity of η_{t+1} imply

$$\Psi_t^{\mathcal{I}} = \frac{1}{2}\left(L_u^{\mathcal{I}}\right)^2 s_t + K_\nu\left(-L_\Psi^{\mathcal{I}} G_u^{\mathcal{I}}\right), \qquad K_\nu(a) \equiv \log \mathbb{E}_t e^{a\nu_{t+1}^s}. \tag{182}$$

The second term is constant over time because the conditional law of ν_{t+1}^s is time invariant. Subtracting the regime-specific Risk-Centric steady state therefore gives

$$\hat{\Psi}_t^{\mathcal{I}} = \frac{1}{2}\left(L_u^{\mathcal{I}}\right)^2 \hat{s}_t. \tag{183}$$

Thus the conjecture is internally consistent with

$$G_u^{\mathcal{I}}(\lambda, \delta) = \frac{1}{2}\left(L_u^{\mathcal{I}}\right)^2. \tag{184}$$

D.3. Proof of Proposition ??

Adding (173) and (174) gives

$$L_u^{\mathcal{I}} = \frac{\delta + (\lambda - \kappa)(1 - L_i)}{\Delta(\rho_u)}. \tag{185}$$

Using (161) and the root condition (162), one obtains the identity

$$\delta + (\lambda - \kappa)(1 - L_i) = \frac{[(\kappa+\delta)\lambda - \kappa^2](1-\beta q_i)}{[\kappa + \delta(1-q_i)]\mathcal{D}(q_i)}. \tag{186}$$

Hence

$$L_u^{\mathcal{I}} = \mathcal{E}_u^{\mathcal{I}}(\lambda, \delta)\left[\lambda - \frac{\kappa^2}{\kappa + \delta}\right], \tag{187}$$

where

$$\mathcal{E}_u^{\mathcal{I}}(\lambda, \delta) \equiv \frac{(\kappa+\delta)(1-\beta q_i)}{[\kappa + \delta(1-q_i)]\mathcal{D}(q_i)\Delta(\rho_u)} > 0. \tag{188}$$

This proves (**??**)–(70). Since $G_u^{\mathcal{I}} = \frac{1}{2}(L_u^{\mathcal{I}})^2$, the loading is non-negative and vanishes if and only if

$$\lambda = \lambda_{\mathcal{I}}^H(\delta) \equiv \frac{\kappa^2}{\kappa + \delta}. \tag{189}$$

At the hedge,

$$\left.\frac{\partial^2 G_u^{\mathcal{I}}}{\partial \lambda^2}\right|_{\lambda = \lambda_{\mathcal{I}}^H(\delta)} = \left[\mathcal{E}_u^{\mathcal{I}}(\lambda_{\mathcal{I}}^H(\delta), \delta)\right]^2 > 0, \tag{190}$$

which establishes the local U-shape and the unique global minimum. Differentiating the hedge gives (73).

It remains to establish the global comparative static with respect to inertia. Use (163) to parameterize the stable branch by q rather than δ. Let

$$\widetilde{\Delta}_u(q, \lambda) \equiv \Delta(\rho_u)|_{\delta = R(q,\lambda)}, \tag{191}$$

and let $\widetilde{L}_u(q, \lambda)$ denote (185) after the same substitution. Direct differentiation and collection of terms gives

$$\frac{\partial \widetilde{L}_u}{\partial q} = \frac{\kappa(1-\rho_u) + \lambda\rho_u}{\mathcal{Z}(q,\lambda)^2}\Bigg\{\kappa^2\left[(1-\beta q)^2 + \kappa + \beta(1-\rho_u)(1-\beta q^2)\right] + \lambda\left[(1-\beta\rho_u)(1-\beta q)^2 + \kappa(1-\beta)\right] \tag{192}$$

where

$$\mathcal{Z}(q, \lambda) \equiv (1-q)\mathcal{D}(q)\widetilde{\Delta}_u(q, \lambda). \tag{193}$$

Every term in braces is strictly positive for $0 < \beta, \rho_u < 1$, $\kappa > 0$, and $0 < q < \bar{q}$. Since (166) gives $\partial q_i / \partial \delta > 0$, the chain rule yields

$$\frac{\partial L_u^{\mathcal{I}}}{\partial \delta} = \frac{\partial \widetilde{L}_u}{\partial q}\frac{\partial q_i}{\partial \delta} > 0, \tag{194}$$

proving (**??**).

Finally,

$$\frac{\partial G_u^{\mathcal{I}}}{\partial \delta} = L_u^{\mathcal{I}} \frac{\partial L_u^{\mathcal{I}}}{\partial \delta}, \tag{195}$$

so the sign of the derivative is the sign of $L_u^{\mathcal{I}}$ and therefore the sign of $\lambda - \lambda_{\mathcal{I}}^H(\delta)$. This proves (72). For $0 < \lambda < \kappa$, the unique crossing occurs at

$$\delta^H(\lambda) = \frac{\kappa^2}{\lambda} - \kappa > 0, \tag{196}$$

so $G_u^{\mathcal{I}}$ falls before the crossing and rises after it. For

$\lambda \geq \kappa$, the economy is on or above the positive-exposure side already at $\delta = 0$, so $G_u^{\mathcal{I}}$ rises for every $\delta > 0$.

As $\delta \downarrow 0$, (163) implies $q_i \to 0$. Equation (188) then gives

$$\mathcal{E}_u^{\mathcal{I}}(\lambda, 0) = \frac{1}{\kappa^2 + \lambda(1 - \beta\rho_u)}, \tag{197}$$

so

$$L_u^{\mathcal{I}}(\lambda, 0) = \frac{\lambda - \kappa}{\kappa^2 + \lambda(1 - \beta\rho_u)} = b(\lambda), \tag{198}$$

and the no-inertia risk-premium curve of Section 2.3 is recovered. □

D.4. Proof of Proposition 2

This appendix proves the comparative statics underlying Proposition 2. All derivatives are impact comparative statics holding the primitive states $(u_t, \hat{s}_t)$ fixed and evaluated from the regime-specific steady state, so that $\hat{i}_{t-1} = 0$.

Let

$$q \equiv q_i.$$

On the locally stable branch,

$$0 < q < \bar{q} < 1, \qquad D(q) > 0,$$

where

$$D(q) \equiv (1 - q)(1 - \beta q) - \kappa q. \tag{199}$$

The inherited-rate block can be parameterized by

$$\delta = R(q, \lambda) \equiv \frac{q\left[\lambda(1 - \beta q) + \kappa^2\right]}{(1 - q)D(q)}. \tag{200}$$

As established in Appendix D,

$$R_q(q, \lambda) > 0,$$

and hence

$$q_\delta \equiv \frac{\partial q_i}{\partial \delta} = \frac{1}{R_q(q, \lambda)} > 0. \tag{201}$$

It is therefore sufficient to establish the signs of the relevant coefficients with respect to q.

For compactness, define

$$A_0 \equiv \kappa^2 + \lambda(1 - \beta) > 0, \tag{202}$$

$$A(\varrho) \equiv \kappa^2 + \lambda(1 - \beta\varrho) > 0, \tag{203}$$

$$B(\varrho) \equiv \kappa(1 - \varrho) + \lambda\varrho > 0, \tag{204}$$

$$\mathcal{D}(\varrho, q) \equiv \left[(1 - q)D(q)\Delta(\varrho)\right]^2 > 0. \tag{205}$$

Also define

$$\mathcal{Q}(q, \varrho) \equiv (1 - \varrho)\mathcal{Q}_0(q) + \varrho\mathcal{Q}_1(q), \tag{206}$$

where

$$\mathcal{Q}_0(q) \equiv \kappa^2\left(2\beta q - \beta - 1\right) - \lambda(1 - \beta q)^2, \tag{207}$$

$$\mathcal{Q}_1(q) \equiv -A_0\left[(1 - \beta q)^2 + \beta\kappa q^2\right]. \tag{208}$$

Because $0 < q < 1$ and $0 < \beta < 1$,

$$2\beta q - \beta - 1 < \beta - 1 < 0,$$

so that $\mathcal{Q}_0(q) < 0$. Likewise, $\mathcal{Q}_1(q) < 0$. Since $0 < \varrho < 1$,

$$\mathcal{Q}(q, \varrho) < 0. \tag{209}$$

The standard gradualism channel. Consider first the cost-push coefficients x_u and π_u. Substituting (200) into the closed-form expressions for x_u and π_u derived in Appendix D, and differentiating with respect to q, gives

$$\frac{\partial x_u}{\partial q} = -\frac{B(\rho_u)\,\mathcal{Q}(q, \rho_u)}{\mathcal{D}(\rho_u, q)} > 0, \tag{210}$$

and

$$\frac{\partial \pi_u}{\partial q} = \frac{\kappa\left(1 - \beta\rho_u q^2\right) A_0 B(\rho_u)}{\mathcal{D}(\rho_u, q)} > 0. \tag{211}$$

The signs follow from (202)–(209) and $1 - \beta\rho_u q^2 > 0$.

Using (201),

$$x_{u,\delta} = \frac{\partial x_u}{\partial q} q_\delta > 0, \tag{212}$$

$$\pi_{u,\delta} = \frac{\partial \pi_u}{\partial q} q_\delta > 0. \tag{213}$$

Thus, absent time-varying risk premia, greater gradualism reduces the contemporaneous output contraction and increases inflation following an inflationary cost-push shock.

The absorption effect. Now consider the response to the risk-premium state. The affine solution gives

$$x_\Psi < 0, \qquad \pi_\Psi < 0.$$

Using the same stable-branch parameterization, direct differentiation gives

$$\frac{\partial x_\Psi}{\partial q} = \frac{A(\rho_s)\,\mathcal{Q}(q, \rho_s)}{\mathcal{D}(\rho_s, q)} < 0, \tag{214}$$

and

$$\frac{\partial \pi_\Psi}{\partial q} = -\frac{\kappa\left(1 - \beta\rho_s q^2\right) A_0 A(\rho_s)}{\mathcal{D}(\rho_s, q)} < 0. \tag{215}$$

Hence,

$$x_{\Psi,\delta} = \frac{\partial x_\Psi}{\partial q} q_\delta < 0, \tag{216}$$

$$\pi_{\Psi,\delta} = \frac{\partial \pi_\Psi}{\partial q} q_\delta < 0. \tag{217}$$

Greater gradualism therefore increases the amount of a given financial tightening that is absorbed through lower output and inflation.

The repricing effect. The time-varying risk premium is

$$\hat{\Psi}_t^{\mathcal{I}} = G_u^{\mathcal{I}}(\lambda, \delta)\hat{s}_t. \tag{218}$$

From the risk-premium comparative static established in the previous subsection,

$$\operatorname{sgn} G_{u,\delta}^{\mathcal{I}} = \operatorname{sgn}\left[\lambda - \lambda_{\mathcal{I}}^{H}(\delta)\right]. \tag{219}$$

Hence, on the inflation-risk side of the hedge,

$$\lambda > \lambda_{\mathcal{I}}^{H}(\delta),$$

we have

$$G_u^{\mathcal{I}} > 0, \qquad G_{u,\delta}^{\mathcal{I}} > 0. \tag{220}$$

Recall the risk-feedback loadings

$$\mathcal{R}_x(\lambda, \delta) = x_{\Psi,\delta} G_u^{\mathcal{I}} + x_\Psi G_{u,\delta}^{\mathcal{I}}, \tag{221}$$

$$\mathcal{R}_\pi(\lambda, \delta) = \pi_{\Psi,\delta} G_u^{\mathcal{I}} + \pi_\Psi G_{u,\delta}^{\mathcal{I}}. \tag{222}$$

Combining (216)–(217) with (220) and $x_\Psi < 0$, $\pi_\Psi < 0$ gives

$$x_{\Psi,\delta} G_u^{\mathcal{I}} < 0, \qquad x_\Psi G_{u,\delta}^{\mathcal{I}} < 0, \tag{223}$$

and

$$\pi_{\Psi,\delta} G_u^{\mathcal{I}} < 0, \qquad \pi_\Psi G_{u,\delta}^{\mathcal{I}} < 0. \tag{224}$$

Therefore

$$\mathcal{R}_x(\lambda, \delta) < 0, \qquad \mathcal{R}_\pi(\lambda, \delta) < 0 \tag{225}$$

throughout the inflation-risk region.

Reversal thresholds. At impact,

$$\frac{dx_t^{\mathcal{I}}}{d\delta} = x_{u,\delta} u_t + \mathcal{R}_x(\lambda, \delta)\hat{s}_t, \tag{226}$$

$$\frac{d\pi_t^{\mathcal{I}}}{d\delta} = \pi_{u,\delta} u_t + \mathcal{R}_\pi(\lambda, \delta)\hat{s}_t. \tag{227}$$

For an inflationary cost-push shock $u_t > 0$, the first terms are positive by (212) and (213), whereas the coefficients multiplying $\hat{s}_t$ are negative on the inflation-risk side. It follows immediately that

$$\frac{dx_t^{\mathcal{I}}}{d\delta} < 0 \quad \Longleftrightarrow \quad \hat{s}_t > \frac{x_{u,\delta} u_t}{-\mathcal{R}_x(\lambda, \delta)} \equiv \hat{s}_t^{x}, \tag{228}$$

and

$$\frac{d\pi_t^{\mathcal{I}}}{d\delta} < 0 \quad \Longleftrightarrow \quad \hat{s}_t > \frac{\pi_{u,\delta} u_t}{-\mathcal{R}_\pi(\lambda, \delta)} \equiv \hat{s}_t^{\pi}. \tag{229}$$

This establishes the two reversal conditions in Proposition 2. □

E. Forward-Looking Interest-Rate Smoothing

Section 3 studies a deliberately myopic formulation of interest-rate gradualism. This subsection shows that the central mechanism does not depend on that assumption. I restore the continuation value associated with the inherited policy rate and solve for a Markov-perfect discretionary equilibrium.

Let ω_t collect the exogenous states of the economy. The central bank solves

$$V(i_{t-1}, \omega_t) = \min_{x_t, \pi_t, i_t} \left\{ \frac{1}{2}\left[\pi_t^2 + \lambda x_t^2 + \delta(i_t - i_{t-1})^2\right] + \beta E_t V(i_t, \omega_{t+1}) \right\} \tag{230}$$

subject to the New Keynesian Phillips curve and the Risk-Centric IS curve,

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t, \tag{231}$$

$$x_t = E_t(x_{t+1} + \pi_{t+1}) - i_t - \Psi_t, \tag{232}$$

where variables are expressed as deviations from their regime-specific steady states and the natural-rate component has been suppressed for notational convenience.

The difference relative to Section 3 is that the current policy rate becomes tomorrow's inherited rate. A current rate adjustment therefore affects aggregate demand both directly and through the future policy path. This is the history-dependence mechanism emphasized by Woodford (2003). The formulation in Section 3 shuts down this continuation-value channel in order to isolate the direct cost of moving the policy rate. The Bellman problem in (230) restores it.

E.1. The inherited-rate block

Consider first the equilibrium response to the inherited policy rate and conjecture

$$i_t = q i_{t-1} + \cdots. \tag{233}$$

Define

$$Q \equiv 1 - \beta q, \qquad D \equiv (1 - q)(1 - \beta q) - \kappa q. \tag{234}$$

Along the continuation equilibrium, the total effect of a current rate choice on current output and inflation is

$$\Xi_i \equiv \frac{dx_t}{di_t} = -\frac{Q}{D}, \qquad \Pi_i \equiv \frac{d\pi_t}{di_t} = -\frac{\kappa}{D}. \tag{235}$$

On the regular stable branch,

$$|\Xi_i| > 1, \qquad |\Pi_i| > \kappa. \tag{236}$$

Thus a current interest-rate movement is more powerful than in the myopic benchmark: it affects demand contemporaneously and also shifts expected future rates.

Using the envelope condition, the optimality condition for the current policy rate becomes

$$(-\lambda \Xi_i) x_t + (-\Pi_i)\pi_t = \delta(i_t - i_{t-1}) + \beta\delta\left(i_t - E_t i_{t+1}\right). \tag{237}$$

The final term is absent from the myopic problem. It is the continuation-value cost associated with choosing today's rate, which becomes the inherited policy rate next period.

In the myopic benchmark, $\Xi_i = -1$ and $\Pi_i = -\kappa$, so (237) reduces to the target criterion used in Section 3,

$$\lambda x_t + \kappa \pi_t = \delta(i_t - i_{t-1}). \tag{238}$$

The forward-looking equilibrium value of q is determined by

$$q\left[\kappa^2 + \lambda Q^2\right] = \delta(1-q)QD^2. \tag{239}$$

This scalar equation replaces the corresponding inherited-rate condition in the myopic problem.

E.2. Risk-premium disturbances

The main question is whether introducing the continuation value restores the separation result of Section 2. It does not.

Consider a pure uncertainty disturbance and let the associated time-varying component of the risk premium satisfy

$$E_t \widehat{\Psi}_{t+1} = \rho_s \widehat{\Psi}_t, \qquad 0 < \rho_s < 1. \tag{240}$$

With all other states set to zero, conjecture

$$x_t = x_\Psi^V \widehat{\Psi}_t, \qquad \pi_t = \pi_\Psi^V \widehat{\Psi}_t, \qquad i_t = q_\Psi^V \widehat{\Psi}_t, \tag{241}$$

where the superscript V distinguishes the forward-looking solution from the myopic equilibrium of Section 3.

Define

$$T \equiv 1 - \beta\rho_s > 0, \qquad K \equiv 1 + \beta(1 - q - \rho_s) > 0. \tag{242}$$

Under the regular stable solution, the coefficient system has determinant $\Delta_V(\rho_s) > 0$, and Cramer's rule gives

$$x_\Psi^V = -\frac{\delta DKT + \beta\kappa^2 q/D}{\Delta_V(\rho_s)} < 0, \tag{243}$$

$$\pi_\Psi^V = -\frac{\kappa}{\Delta_V(\rho_s)}\left(\delta DK - \frac{\beta\lambda qQ}{D}\right) < 0, \tag{244}$$

$$q_\Psi^V = -\frac{\kappa^2 + \lambda QT}{\Delta_V(\rho_s)} < 0. \tag{245}$$

The only non-immediate sign in (243)–(245) is that of the expression in parentheses in (244). Using (239),

$$\delta DK - \frac{\beta\lambda qQ}{D} = \frac{qK\kappa^2 + \lambda Q^2 T}{D(1-q)Q} > 0, \tag{246}$$

which establishes $\pi_\Psi^V < 0$.

Proposition 4 (Forward-looking gradualism and priced risk)**.** *Suppose $\delta > 0$ and the economy is on the regular stable Markov branch. Restoring the continuation value associated with the inherited policy rate does not restore the allocation–implementation separation. A rise in the risk-premium wedge satisfies*

$$x_\Psi^V < 0, \qquad \pi_\Psi^V < 0, \qquad q_\Psi^V < 0. \tag{247}$$

Hence higher priced uncertainty tightens private financial conditions and induces the central bank to lower the policy rate, but the rate adjustment is incomplete: output and inflation also decline.

The proposition is the forward-looking counterpart of the sign result in Section 3. Interest-rate smoothing still makes the policy instrument part of the stabilization trade-off. Because fully neutralizing a higher risk premium requires a costly movement in the policy rate, uncertainty is no longer absorbed solely through implementation. The continuation value strengthens the effect of current policy on private demand, but does not eliminate this basic force.

E.3. Monetary policy and return exposure

The second mechanism in Section 3 also survives. Let x_u^V and π_u^V denote the forward-looking equilibrium responses to the cost-push innovation. Nominal aggregate wealth continues to inherit the exposure

$$L_u^V \equiv x_u^V + \pi_u^V, \tag{248}$$

so that the cost-push component of the equilibrium risk-premium loading remains

$$G_u^V = \frac{1}{2}\left(L_u^V\right)^2. \tag{249}$$

Thus the continuation value changes the equilibrium coefficients and therefore the *amount* of risk priced by financial markets, but not the underlying mapping

$$\text{policy} \longrightarrow (x_u, \pi_u) \longrightarrow L_u \longrightarrow G_u \longrightarrow \Psi_t. \tag{250}$$

Unlike in the myopic benchmark, however, the forward-looking solution does not in general deliver the same simple closed-form zero-risk locus. The continuation-value channel changes the endogenous exposure of nominal wealth and therefore the precise location at which output and inflation risk hedge one another. This is why the main text uses the myopic formulation to characterize the interaction between λ, δ, and priced risk analytically.

In the baseline calibration, restoring the continuation value makes monetary policy more powerful through expectations. The resulting responses satisfy

$$x_u^V < x_u^M < 0, \qquad 0 < \pi_u^V < \pi_u^M, \tag{251}$$

and therefore

$$\left|x_u^V + \pi_u^V\right| > \left|x_u^M + \pi_u^M\right|. \tag{252}$$

Thus the equilibrium risk premium is larger in the forward-looking solution in that calibration. This ordering is not general: the continuation value changes return exposure and can therefore either raise or lower the equilibrium premium.

The purpose of this extension is not to reproduce every closed-form comparative static of Section 3, but to establish that its central mechanism is not an artifact of policy myopia. Once current policy affects the future rate path, rate movements become more powerful through ex-

pectations. Nevertheless, costly adjustment still prevents the central bank from offsetting movements in financial conditions one-for-one, and monetary policy continues to shape the risk premium through its effect on the exposure of nominal wealth to aggregate shocks.

F. Rogoff delegation with priced aggregate risk

This appendix derives the inflation bias and proves Proposition **??**. Throughout, the socially desired output gap is $\bar{x} > 0$. Society has output-stabilization weight λ^S, while monetary policy is delegated to a central banker with weight λ^B. A lower λ^B corresponds to a more inflation-conservative central banker.

F.1. Discretionary equilibrium and the inflation bias

The delegated central banker minimizes

$$\mathcal{L}_t^B = \frac{1}{2}\left[\pi_t^2 + \lambda^B (x_t - \bar{x})^2\right] \tag{253}$$

subject to the New Keynesian Phillips curve,

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + u_t. \tag{254}$$

Under discretion, private expectations are taken as given. The first-order condition is therefore

$$\lambda^B (x_t - \bar{x}) + \kappa \pi_t = 0. \tag{255}$$

Conjecture a bounded affine equilibrium

$$x_t = \bar{x}^B + x_u^B u_t, \qquad \pi_t = \bar{\pi}^B + \pi_u^B u_t. \tag{256}$$

Using $\mathbb{E}_t u_{t+1} = \rho_u u_t$, the constant terms in (254) and (255) satisfy

$$(1-\beta)\bar{\pi}^B = \kappa \bar{x}^B, \tag{257}$$

$$\lambda^B (\bar{x}^B - \bar{x}) + \kappa \bar{\pi}^B = 0. \tag{258}$$

Hence

$$\bar{x}^B = \frac{(1-\beta)\lambda^B}{\kappa^2 + (1-\beta)\lambda^B}\,\bar{x}, \tag{259}$$

$$\bar{\pi}^B = \frac{\kappa \lambda^B}{\kappa^2 + (1-\beta)\lambda^B}\,\bar{x}. \tag{260}$$

Since $\bar{x} > 0$ and $\lambda^B > 0$, $\bar{\pi}^B > 0$. Moreover,

$$\frac{\partial \bar{\pi}^B}{\partial \lambda^B} = \frac{\kappa^3}{[\kappa^2 + (1-\beta)\lambda^B]^2}\,\bar{x} > 0. \tag{261}$$

This proves Proposition **??**(i).

The coefficients on the cost-push shock satisfy

$$(1 - \beta\rho_u)\pi_u^B = \kappa x_u^B + 1, \tag{262}$$

$$\lambda^B x_u^B + \kappa \pi_u^B = 0. \tag{263}$$

Defining $\chi_u \equiv 1 - \beta\rho_u$ as in (34),

$$x_u^B = -\frac{\kappa}{\kappa^2 + \lambda^B \chi_u}, \qquad \pi_u^B = \frac{\lambda^B}{\kappa^2 + \lambda^B \chi_u}. \tag{264}$$

Thus the distortion $\bar{x}$ changes the equilibrium means but not the marginal stabilization response to u_t.

F.2. The certainty-equivalent delegation problem

Society evaluates outcomes according to

$$\mathcal{L}_t^S = \frac{1}{2}\left[\pi_t^2 + \lambda^S (x_t - \bar{x})^2\right]. \tag{265}$$

Let $\sigma_u^2 \equiv \mathbb{E}[u_t^2] > 0$ denote the unconditional variance of the stationary cost-push process. Under the stationary distribution, minimizing expected discounted welfare is equivalent, up to the positive factor $1/(1-\beta)$, to minimizing expected per-period loss. Using (259)–(264), the certainty-equivalent delegation objective is

$$W^{CE}(\lambda^B; \lambda^S) = \frac{1}{2}\left[(\lambda^B)^2 + \lambda^S \kappa^2\right]\left\{\frac{\kappa^2 \bar{x}^2}{[\kappa^2 + (1-\beta)\lambda^B]^2} + \frac{\sigma_u^2}{[\kappa^2 + \chi_u \lambda^B]^2}\right\}. \tag{266}$$

Differentiating gives

$$\frac{\partial W^{CE}}{\partial \lambda^B} = \frac{\kappa^4 \bar{x}^2 [\lambda^B - (1-\beta)\lambda^S]}{[\kappa^2 + (1-\beta)\lambda^B]^3} + \frac{\sigma_u^2 \kappa^2 [\lambda^B - \chi_u \lambda^S]}{[\kappa^2 + \chi_u \lambda^B]^3}. \tag{267}$$

Since $0 < \rho_u < 1$,

$$0 < 1 - \beta < \chi_u = 1 - \beta\rho_u < 1. \tag{268}$$

It follows from (267) that

$$\frac{\partial W^{CE}}{\partial \lambda^B} < 0 \quad \text{for} \quad \lambda^B \le (1-\beta)\lambda^S, \tag{269}$$

and

$$\frac{\partial W^{CE}}{\partial \lambda^B} > 0 \quad \text{for} \quad \lambda^B \ge \chi_u \lambda^S. \tag{270}$$

Therefore every global minimizer $\lambda^R \in \arg\min_{\lambda^B > 0} W^{CE}(\lambda^B; \lambda^S)$ lies in the interval

$$(1-\beta)\lambda^S < \lambda^R < \chi_u \lambda^S < \lambda^S. \tag{271}$$

In particular,

$$0 < \lambda^R < \lambda^S, \tag{272}$$

which is the Rogoff conservative-central-banker result in the present parameterization. Because a larger λ means greater concern for output stabilization, $\lambda^R < \lambda^S$ means that society delegates policy to a more inflation-oriented central banker.

F.3. Why aggregate risk pricing does not change delegation when $\delta = 0$

Now restore the Risk-Centric IS curve. For any feasible allocation (x_t, π_t), Section 2.3 defines the certainty-equivalent implementing rate

$$i_t^{NK}(x, \pi) = \bar{r}_t^n + \mathbb{E}_t \pi_{t+1} + \mathbb{E}_t x_{t+1} - x_t. \tag{273}$$

The Risk-Centric IS curve then implies

$$i_t^{RP}(x, \pi) = i_t^{NK}(x, \pi) - \Psi_t. \tag{274}$$

When $\delta = 0$, the policy rate does not enter either the delegated central bank's loss (253) or society's loss (265). Moreover, Ψ_t does not enter the Phillips curve. Hence, for every λ^B, (274) implements exactly the same (x_t, π_t) allocation as in the certainty-equivalent economy:

$$x_t^{RP}(\lambda^B) = x_t^{CE}(\lambda^B), \qquad \pi_t^{RP}(\lambda^B) = \pi_t^{CE}(\lambda^B). \tag{275}$$

It follows immediately that society's delegation objective is identical pointwise in the delegated mandate,

$$W^{RP}(\lambda^B; \lambda^S) = W^{CE}(\lambda^B; \lambda^S) \qquad \text{for all } \lambda^B > 0. \tag{276}$$

Therefore

$$\arg\min_{\lambda^B > 0} W^{RP}(\lambda^B; \lambda^S) = \arg\min_{\lambda^B > 0} W^{CE}(\lambda^B; \lambda^S), \tag{277}$$

and any optimal mandate satisfies $\lambda_{RP}^{B,*} = \lambda^R < \lambda^S$.

Finally, since the mean distortion does not alter the shock loadings in (264), nominal-wealth exposure to the cost-push innovation is

$$b(\lambda^B) = x_u^B + \pi_u^B = \frac{\lambda^B - \kappa}{\kappa^2 + \lambda^B \chi_u}. \tag{278}$$

Under conditional Gaussianity, the associated risk-premium wedge is

$$\Psi_t^\star(\lambda^B) = \frac{1}{2} \left[\frac{\lambda^B - \kappa}{\kappa^2 + \lambda^B \chi_u} \right]^2 s_t. \tag{279}$$

The only effect of priced aggregate risk is therefore on implementation,

$$i_t^{RP}(\lambda^B) = i_t^{NK}(\lambda^B) - \Psi_t^\star(\lambda^B), \tag{280}$$

which proves Proposition **??**(ii).

F.4. Proof of Proposition 3

The delegation objective is

$$W^{RP}(\lambda; V_s, \delta) = W^{CE}(\lambda; \delta) + V_s \, \Omega(\lambda, \delta), \tag{281}$$

where, suppressing the superscript B to lighten notation,

$$\Omega(\lambda, \delta) = \frac{1}{8} J_\Psi(\lambda, \delta) \left[\mathcal{E}_u^{\mathcal{I}}(\lambda, \delta) \right]^4 [\lambda - \lambda_H(\delta)]^4. \tag{282}$$

On the maintained stable branch, $J_\Psi(\lambda, \delta) > 0$ and $\mathcal{E}_u^{\mathcal{I}}(\lambda, \delta) > 0$. Hence

$$\Omega(\lambda, \delta) \geq 0, \qquad \Omega(\lambda, \delta) = 0 \iff \lambda = \lambda_H(\delta). \tag{283}$$

We first show that the zero at $\lambda_H(\delta)$ is globally isolated.

Recall that the inherited-rate coefficient q_i solves

$$q_i \left[\lambda(1 - \beta q_i) + \kappa^2 \right] + \delta(q_i - 1) D(q_i) = 0, \tag{284}$$

with $q_i \in (0, 1)$ and $D(q_i) > 0$ on the stable branch. Dividing (284) by λ and taking $\lambda \to \infty$ gives

$$q_i(1 - \beta q_i) \to 0. \tag{285}$$

Since $0 < q_i < 1$ and $\beta < 1$, it follows that

$$q_i \to 0, \qquad x_i \to 0, \qquad \pi_i \to 0, \qquad L_i \to 0. \tag{286}$$

Using the cost-push exposure derived in Section 3,

$$L_u^{\mathcal{I}} = \frac{\delta + (\lambda - \kappa)(1 - L_i)}{\Delta(\rho_u)}, \tag{287}$$

and

$$\frac{\Delta(\rho_u)}{\lambda} \longrightarrow 1 - \beta \rho_u \equiv \chi_u > 0, \tag{288}$$

we obtain

$$L_u^{\mathcal{I}} \longrightarrow \frac{1}{\chi_u}. \tag{289}$$

Because

$$L_u^{\mathcal{I}} = \mathcal{E}_u^{\mathcal{I}}(\lambda, \delta) \left[\lambda - \lambda_H(\delta) \right], \tag{290}$$

this implies

$$\left[\mathcal{E}_u^{\mathcal{I}}(\lambda, \delta) \right]^4 [\lambda - \lambda_H(\delta)]^4 \longrightarrow \frac{1}{\chi_u^4}. \tag{291}$$

It remains to characterize the welfare loading J_Ψ. As $\lambda \to \infty$, the priced-uncertainty coefficients satisfy

$$q_\Psi \to -1, \qquad x_\Psi \to 0, \qquad \pi_\Psi \to 0, \tag{292}$$

while (286) implies that the inherited-rate propagation vanishes. Let the uncertainty state follow

$$\hat{s}_t = \rho_s \hat{s}_{t-1} + \nu_t^s, \qquad |\rho_s| < 1. \tag{293}$$

In the limit, a unit risk-premium disturbance is therefore absorbed entirely by the policy rate,

$$\hat{i}_t = -\hat{s}_t, \tag{294}$$

so that

$$\Delta \hat{i}_t = -(\hat{s}_t - \hat{s}_{t-1}). \tag{295}$$

Normalizing by $V_s = \mathrm{Var}(\hat{s}_t)$,

$$\frac{\mathrm{Var}(\Delta \hat{i}_t)}{V_s} = 2(1 - \rho_s). \tag{296}$$

The output and inflation contributions vanish in this limit,

and hence

$$J_\Psi(\lambda,\delta) \longrightarrow 2\delta(1-\rho_s) > 0. \tag{297}$$

Combining (291) and (297),

$$\Omega(\lambda,\delta) \longrightarrow \frac{\delta(1-\rho_s)}{4\chi_u^4} > 0 \qquad \text{as} \qquad \lambda \to \infty. \tag{298}$$

Thus, for every $\varepsilon > 0$, continuity of Ω, the unique-zero property (283), and the positive tail limit (298) imply that there exists a constant $m_\varepsilon > 0$ such that

$$|\lambda - \lambda_H(\delta)| \geq \varepsilon \quad \Longrightarrow \quad \Omega(\lambda,\delta) \geq m_\varepsilon. \tag{299}$$

For any closed set $\mathcal{A} \subset [0,\infty)$ bounded away from $\lambda_H(\delta)$, (299) implies

$$m_{\mathcal{A}} \equiv \inf_{\lambda\in\mathcal{A}} \Omega(\lambda,\delta) > 0. \tag{300}$$

Define

$$V^\star(\mathcal{A}) \equiv \sup_{\lambda\in\mathcal{A}} \frac{\left[W^{CE}(\lambda_H(\delta);\delta) - W^{CE}(\lambda;\delta)\right]_+}{\Omega(\lambda,\delta)}. \tag{301}$$

Because $W^{CE}(\lambda;\delta) \geq 0$ and (300) holds,

$$0 \leq V^\star(\mathcal{A}) \leq \frac{W^{CE}(\lambda_H(\delta);\delta)}{m_{\mathcal{A}}} < \infty. \tag{302}$$

Now suppose $V_s > V^\star(\mathcal{A})$. For any $\lambda \in \mathcal{A}$ such that $W^{CE}(\lambda;\delta) < W^{CE}(\lambda_H;\delta)$, the definition of $V^\star(\mathcal{A})$ gives

$$V_s\Omega(\lambda,\delta) > W^{CE}(\lambda_H(\delta);\delta) - W^{CE}(\lambda;\delta). \tag{303}$$

Hence

$$W^{RP}(\lambda;V_s,\delta) > W^{CE}(\lambda_H(\delta);\delta) = W^{RP}(\lambda_H(\delta);V_s,\delta)\,. \tag{304}$$

If instead $W^{CE}(\lambda;\delta) \geq W^{CE}(\lambda_H;\delta)$, then $\Omega(\lambda,\delta) > 0$ on $\mathcal{A}$ implies the same strict inequality for every $V_s > 0$. Therefore, whenever $V_s > V^\star(\mathcal{A})$, every mandate in $\mathcal{A}$ is dominated by the hedge and no global minimizer can belong to $\mathcal{A}$.

The three cases follow by choosing the excluded set appropriately. If

$$\lambda_R^B(\delta) < \lambda^S < \lambda_H(\delta),$$

take $\mathcal{A} = \mathcal{A}_R = [0,\lambda^S]$. Since the hedge lies outside this set, $V_R^\star = V^\star(\mathcal{A}_R) < \infty$, and $V_s > V_R^\star$ implies

$$\lambda_{RP}^{B,*}(V_s,\delta) > \lambda^S.$$

If

$$\lambda_R^B(\delta) < \lambda_H(\delta) < \lambda^S,$$

take

$$\mathcal{A} = \mathcal{A}_A = [0,\lambda_R^B(\delta)] \cup [\lambda^S,\infty).$$

Then $V_s > V_A^\star$ implies

$$\lambda_R^B(\delta) < \lambda_{RP}^{B,*}(V_s,\delta) < \lambda^S.$$

Finally, if

$$\lambda_H(\delta) < \lambda_R^B(\delta) < \lambda^S,$$

take

$$\mathcal{A} = \mathcal{A}_C = [\lambda_R^B(\delta),\infty).$$

Then $V_s > V_C^\star$ implies

$$\lambda_{RP}^{B,*}(V_s,\delta) < \lambda_R^B(\delta).$$

This proves all three finite-threshold results. □