A Note on Energy Shocks in a Risk-Centric Model

PAU BELDA*

April 2026

Very preliminary and incomplete.

Latest version here.

Abstract

Energy shocks are also uncertainty shocks that raise risk premia in financial markets. How should monetary policy respond to energy shocks, given this risk repricing? I study this question in a Risk-Centric New Keynesian model with a cost-push shock with stochastic volatility. The discretionary-optimal inflation–output trade-off is unchanged relative to the canonical New Keynesian model, but its implementation is not. Higher uncertainty raises risk premia, depresses asset prices, and dampens aggregate demand, delivering part of the contraction needed to stabilize inflation. Accordingly, the optimal policy response is less aggressive; ignoring the risk channel induces over-tightening and an unnecessarily deep recession. If the central bank smooths interest rates, risk premia can worsen the inflation–output trade-off, making policy gradualism more costly.

JEL Codes: E52, E44, G12, E31

Keywords: Monetary Policy, Risk Premia, Energy Shocks

*Research Economist at the Bank of England. I have benefited from conversations with Ricardo Caballero, Mridula Duggal, Michael McLeay, Ricardo Reis, Alp Simsek, Danny Walker and participants at the Bank of England's Macro Brownbag and LSE's Working group on Empirical Asset Pricing. I thank Hannah Copeland for sharing the energy shocks series for the Local Projections. 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.

1 Introduction

When supply tensions disrupt energy markets, costs for firms and households rise. But energy shocks might also increase uncertainty about the future path of cost pressures—through the duration of the disruption, the breadth of pass-through, or the risk of further escalation. This uncertainty is priced in financial markets, raising risk premia.

Figure 1 provides evidence. Using daily data for the U.S. and the U.K., I estimate local projections of option-implied equity risk premia on the absolute size of oil supply news shocks. The response is significantly positive: oil supply shocks of either sign raise risk compensation. The result is robust across alternative risk premia measures –as the excess bond premium– and direct risk measures –as the VIX–, and is stronger for larger shocks.[1] Energy shocks, in short, are also uncertainty shocks.

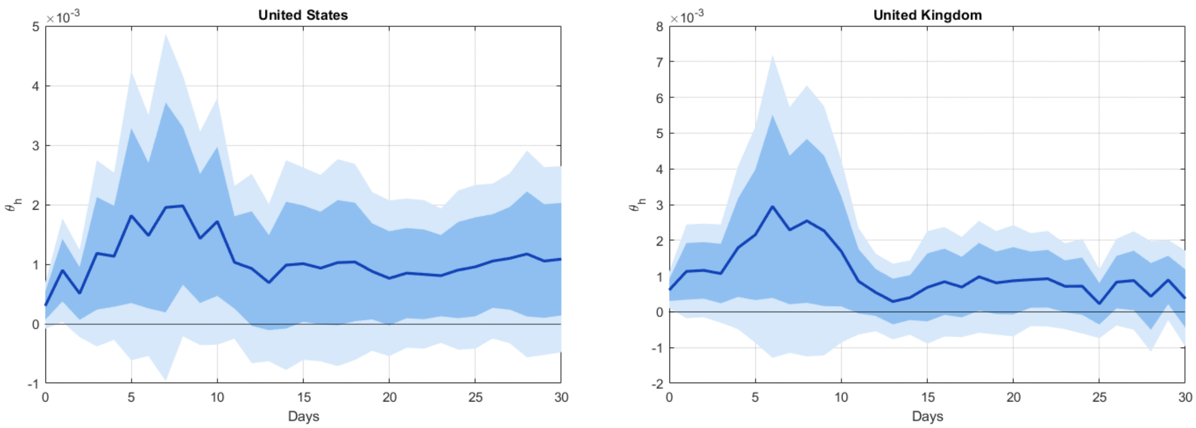

Figure 1: Response of the equity risk premium to oil supply news shocks. The figure reports local projections from the regression $erp_{t+h} = \alpha_h + \theta_h|z_t| + \Gamma_h' X_t + \varepsilon_{t+h}$ estimated separately for each horizon h, where erp_{t+h} denotes the option-implied equity risk premium at time $t+h$, z_t is the oil supply news shock, and X_t is a vector of controls. The controls are daily oil returns, stock returns, the 10-year minus 2-year government bond yield spread, and the central bank policy rate. Shocks follow the methodology in Kanzig (2021). Confidence bands are at the 68% and 90% levels. The sample is June 1999 to December 2025, determined by data availability. Data is at daily frequency. See Appendix A for details and robustness.

Higher risk premia tighten financial conditions and put downward pressure on aggregate spending. This naturally raises a set of monetary policy questions: how should the central bank respond to an energy shock when financial markets reprice risk? Does risk repricing alter the optimal inflation-output trade-off? And does it affect the implementation of the optimal allocation?

These questions are not fully addressed in standard New Keynesian (NK) models. In the canonical NK model, an energy shock is typically represented as a cost-push disturbance. In richer HANK environments, however, the same shock can also operate as a contractionary demand

[1] See Appendix A for details and robustness. This magnitude-based evidence is consistent with a broader literature linking energy shocks to uncertainty and risk compensation. Forni et al. (2025) show that oil supply news shocks increase uncertainty independently of sign and that both positive and negative shocks raise the excess bond premium. See also Kilian and Vigfusson (2011) and Elder and Serletis (2010) on the uncertainty channel. Kakhbod, Kermani, and Maciel (2026) further show that supply imbalances raise term premia and equity risk premia.

disturbance. In versions with imperfect information, the pass-through of the shock might be state-contingent, becoming more persistent when supply uncertainty is larger. In all cases, though, the shock is modeled as a first-moment disturbance.[2] In turn, the NK literature that emphasizes uncertainty often highlights precautionary pricing on the firm side and precautionary savings on the household side.[3] By contrast, this note suggests a complementary channel through asset pricing: the higher uncertainty associated with energy shocks raises risk premia, depressing asset prices, and weakening aggregate demand.

I study this channel in a streamlined version of the risk-centric model of Caballero and Simsek (2022). As in the canonical NK framework, prices are sticky and output is demand-determined. The model features a stylized TANK structure in which workers do not participate in asset markets and capitalists do not supply labor. The key departure is that capitalists' portfolio choice is delegated to wealth managers. Hence, capitalists' spending responds to asset prices but they do not price assets themselves. Instead, wealth managers maximize log wealth, which gives rise to a wealth-based stochastic discount factor. This separation changes how monetary policy works. The central bank does not steer aggregate demand in the first place by changing the intertemporal price of consumption—a channel for which direct empirical support is limited (Crawley , 2025). Instead, it operates through asset markets: the policy rate enters the stochastic discount factor, which determines asset valuations and financial conditions, and these in turn transmit to aggregate spending.

Relative to Caballero and Simsek (2022), I introduce a cost-push shock with stochastic volatility and shut down departures from rationality. The former captures energy shocks as both level and uncertainty shocks, generating time-varying risk premia. Abstracting from non-rational features isolates the risk channel, making it possible to see clearly how canonical NK prescriptions for cost-push shocks change when time-varying risk premia become part of the transmission mechanism.

While the goods supply side is still summarized by a standard New Keynesian Phillips Curve, the aggregate demand side is summarized by a distinct Risk-Centric IS curve. As in canonical NK models, the curve implies a negative relationship between the policy rate and the output gap. The key difference is the emergence of a time-varying risk-adjustment wedge that summarizes output and inflation risk premia and exerts an additional drag on aggregate demand and output.

The main result is that, under discretionary policy, the optimal inflation–output trade-off is the same as in the canonical NK benchmark. An energy cost-push shock still optimally calls for

[2] Through real-income losses, borrowing constraints, non-homothetic energy expenditure, labor-income redistribution, or unemployment risk, an inflationary shock can destroy demand. See Auclert et al. (2023), Chan, Diz, and Kanngiesser (2024), Bobasu, Dobrew, and Repele (2025), and Gnocato (2025). Closely related representative-agent analyses include Bodenstein, Guerrieri, and Kilian (2012), Natal (2012), and Plante (2014). For supply-chain uncertainty as an amplifier of pass-through, see (Merendino and Monacelli, 2026).

[3] For precautionary pricing and the mapping from uncertainty into cost-push-type inflation dynamics, see (Cho et al., 2021; Born and Pfeifer, 2014; Fernández-Villaverde et al., 2015; Mumtaz and Theodoridis, 2015; Oh, 2020). For precautionary saving and self-insurance in incomplete-markets and HANK environments, see (Kaplan, Moll, and Violante, 2018; Auclert, 2019; McKay, Nakamura, and Steinsson, 2016); for optimal policy in that class of models, especially (Acharya, Challe, and Dogra, 2023). For supply-chain uncertainty as an amplifier of cost-push pass-through, see (**?**). For related work on redistribution and risk premia in a HANK environment, see (Kekre and Lenel, 2022).

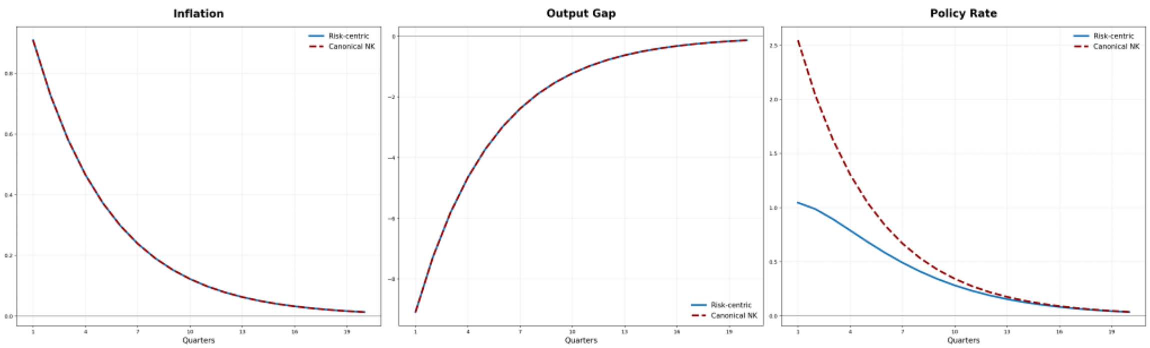

Figure 2: Responses of inflation, the output gap, and the policy rate one-standard-deviation shock to an energy shock under optimal discretionary policy in the Risk-Centric model and a canonical New Keynesian model.

higher inflation and lower output, but the additional risk generated by the shock does not alter the target allocation itself. What changes is the implementation of that allocation. As is standard, the central bank backs out the policy rate from the IS curve. In the risk-centric environment, however, the IS curve contains a risk wedge: higher uncertainty tightens financial conditions on its own, so part of the required contraction is delivered by private markets rather than by the policy rate. In log deviations from steady state, the optimal Risk-Centric rate satisfies

$$\hat{i}_t^{RC} = \hat{i}_t^{NK} - G^\star \hat{s}_t, \qquad G^\star \geq 0 \tag{1}$$

where $\hat{s}_t$ denotes the deviation of the conditional variance of cost-push innovations from steady state. Hence, following a positive volatility shock, the optimal policy rate lies strictly below its canonical NK counterpart: the same inflation–output mix is implemented with a lower rate because financial markets already provide part of the necessary tightening. Ignoring this risk channel leads the central bank to over-tighten and generate an unnecessarily deep recession. Figure 2 illustrates the analytical result.

This sharp separation between allocation and implementation breaks once the central bank smooths interest rates. With policy inertia, stochastic volatility is no longer neutral for allocations: by preventing the central bank from fully offsetting the volatility-induced tightening in private financial conditions, endogenous risk premia directly depress output and inflation even under optimal discretion. As a result, the inflation–output trade-off worsens relative to the canonical NK benchmark: achieving a similar inflation path may require a deeper recession, despite a less aggressive policy-rate path. The policy lesson is that interest-rate gradualism becomes more costly when risk premia respond endogenously to volatility. In the canonical NK model, inertia mainly redistributes stabilization over time; here, it also allows financial tightening to spill into allocations. Thus, when financial conditions deteriorate through higher risk premia, the case for gradualism is weaker than in the canonical benchmark.

The note proceeds as follows. Section 2 describes the model and derives the equilibrium risk-centric IS curve. Section 3 characterizes optimal policy and unpacks its implementation. Section 4 characterizes optimal policy with interest rate inertia and policy with a Taylor rule. Section 5 concludes.

2 Model

The model is a streamlined version of Caballero and Simsek (2022). Relative to that framework, there are two changes. First, I shut down deviations from rationality. Second, I introduce stochastic volatility in the cost-push shock. Since the household side and the production block are otherwise unchanged, I keep the main-text presentation brief and relegate the full derivations to Appendix B.

2.1 Environment

Time is discrete. The economy is populated by households, wealth managers, monopolistically competitive intermediate-goods firms, a competitive final-goods firm, and a central bank. There are three financial assets: an equity claim on aggregate corporate profits (the market portfolio), a one-period real bond in zero net supply, and a one-period nominal bond in zero net supply. Monetary policy sets the nominal return on the nominal bond.

On the production side, a continuum of intermediate-goods firms produces differentiated varieties using labor. A competitive final-goods firm aggregates those varieties into the consumption good. Intermediate firms set prices subject to Calvo nominal rigidities.

On the demand side, there are three agents. Hand-to-mouth households (”workers”) consume all their current disposable income, supply labor and do not participate in asset markets. Asset holders (”capitalists”) consume and save out of wealth, do not participate in the labor market and delegate portfolio choice to wealth managers. Wealth managers are the financial markets players, allocating capitalists savings to the available assets to maximize their next period log wealth.

The key departure from Caballero and Simsek (2022) concerns a shock to firms’ costs. The economy is hit by a cost-push disturbance whose conditional variance is time-varying. As a result, an adverse realization is simultaneously a first-moment shock, raising current inflationary pressure, and a second-moment shock, increasing uncertainty about future cost conditions.[4] In this sense, an adverse energy shock is modeled as both an inflationary level shock and a risk shock. Formally, let the cost-push shock follow

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

and let the volatility state admit the local affine representation

$$s_{t+1} = (1-\rho_s)\bar{s} + \rho_s s_t + \iota\eta_{t+1} + \nu_{t+1}, \qquad \nu_{t+1} = \omega_\nu \xi_{t+1}, \quad \xi_{t+1} \mid \mathcal{I}_t \sim \mathcal{N}(0,1)$$

[4]Unlike Cho et al. (2021), who study a pure uncertainty shock by placing stochastic volatility on productivity, I place stochastic volatility directly on the cost-push wedge itself. See Section 3.2 for an extension with stochastic volatility on TFP shocks.

with $\bar{s} > 0$ and $\text{Corr}_t(\varepsilon_{t+1}, \xi_{t+1}) = 0$. When $\iota > 0$, adverse cost-push realizations increase both current inflationary pressure and future uncertainty.[5] Equations above are interpreted as a local approximation around a deterministic steady state with $\bar{s} > 0$; accordingly, all results are local and conditional on the equilibrium path remaining in an admissible subset of $\mathbb{R}_{++}$.

2.2 Equilibrium

In this section, I summarize the equilibrium conditions in goods, labor and asset markets for a given interest rate. The next section closes the model with an optimal interest rate rule.

New Keynesian Phillips Curve. Since the production block is unchanged relative to Caballero and Simsek (2022), the supply side retains the standard New Keynesian form:

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t^{\mathcal{F}}[\pi_{t+1}] + u_t \tag{2}$$

where π is inflation and x is the output gap. The output gap is given by

$$x_t \equiv y_t - y_t^p$$

where y_t^p is log potential output that evolves according to

$$y_t^p = y_{t-1}^p + z_t, \qquad z_t \sim \mathcal{N}(0, \sigma_z^2) \tag{3}$$

with z are TFP shocks. Relative to the textbook NKPC, the only difference is that the cost-push disturbance now has stochastic volatility and expectations are taken under the representative firm's subjective measure $\mathcal{F}$.

A Risk-Centric IS curve. Equilibrium in goods and asset markets implies the following aggregate-demand relation:

$$x_t = \mathbb{E}_t^{\mathcal{M}}[x_{t+1}] - \left(i_t - \mathbb{E}_t^{\mathcal{M}}[\pi_{t+1}] - r_t^n\right) - \Psi_t \tag{4}$$

where

$$\Psi_t \equiv \frac{1}{2}\text{Var}_t^{\mathcal{M}}(r_{t+1}) + \frac{1}{2}\text{Var}_t^{\mathcal{M}}(\pi_{t+1}) + \text{Cov}_t^{\mathcal{M}}(\pi_{t+1}, r_{t+1}) \tag{5}$$

The corresponding natural rate is

$$r_t^n = \rho + \mathbb{E}_t^{\mathcal{M}}[\Delta y_{t+1}^p] \tag{6}$$

where $\rho = -\log \beta$ is a time preference parameter. Equation (4) is the risk-centric analog of the canonical NK IS curve. In the equation, i is the one-period nominal safe rate set by the Central Bank, r^n is the natural rate, r is the return on the risky asset. Expectations are taken using the

[5] The Local Projection evidence in the Introduction is reduced-form and based on shock magnitudes, whereas the model adopts a signed structural specification in which adverse cost-push realizations raise future uncertainty. This is a convenient tractable choice: it preserves the affine structure needed for closed-form results while capturing the main mechanism of interest, namely that energy shocks can simultaneously affect current costs and future risk.

subjective probability of the representative wealth manager $\mathcal{M}$. The curve shows that the current output gap depends on the expected future output gap and on the ex-ante real policy rate relative to the natural rate, with an additional risk-adjustment wedge Ψ_t that includes inflation and output risk. See Appendix C for details on the derivation.

Remark 1 (Comparison with the canonical IS curve)**.** *Up to the usual slope coefficient σ^{-1}, the canonical NK IS curve can be written as*

$$x_t = \mathbb{E}_t[x_{t+1}] - \sigma^{-1}\Big(i_t - \mathbb{E}_t[\pi_{t+1}] - r_t^{n,NK}\Big), \qquad r_t^{n,NK} \equiv \rho + \sigma\,\mathbb{E}_t[\Delta y_{t+1}^p].$$

Equation (4) *differs along three dimensions. First, the expected policy rate operates through asset-price discounting and wealth effects rather than through the representative household's intertemporal substitution margin. Second, macroeconomic risk enters directly through Ψ_t: higher output and inflation risks tighten financial conditions for a given expected real policy rate. Third, the relevant expectations are those of wealth managers, under the subjective measure $\mathcal{M}$, rather than those of households.*

More broadly, the presence of an additional wedge in the aggregate-demand block is not, by itself, novel within the New Keynesian family. Once one departs from the benchmark representative-agent environment, generalized IS curves are standard. In HANK models with cyclical risk, for instance, the IS relation is modified because precautionary saving and self-insurance affect the mapping from the real rate to aggregate demand. In collateral and borrowing-constraint models, asset prices influence spending by relaxing or tightening borrowing capacity, so the demand block inherits additional state dependence from balance-sheet conditions. In open-economy and other richer NK environments, the natural benchmark itself also becomes more elaborate, reflecting terms of trade, external demand, capital accumulation, habits, or other endogenous states. For related examples of generalized IS curves within the NK tradition, see, among many others, Iacoviello (2005), Cúrdia and Woodford (2016), Bilbiie (2008), Galí and Monacelli (2005), Kaplan, Moll, and Violante (2018), Auclert, Rognlie, and Straub (2018), Acharya and Dogra (2020), and Acharya, Challe, and Dogra (2023).

It might be useful to distinguish between a reduced-form IS wedge and the present risk-centric structure. Under discretionary policy, the latter is allocation-equivalent to a generalized NK model with an additive wedge,

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

with $\mu_t = \Psi_t$: the inflation–output trade-off is still pinned down by the NKPC and the loss function, and the wedge only shifts the policy rate required to implement a given allocation. In that sense, the reduced-form logic of Proposition 1 could be replicated by a suitably chosen exogenous wedge process.

What the present framework adds is that the wedge is not primitive. The term Ψ_t is disciplined by asset pricing and pinned down by structural primitives. Hence, the model imposes sign and cross-equation restrictions on the wedge, links it directly to risk premia and asset valuations, and implies *additional regime dependence once one departs from the frictionless discretionary benchmark. What is distinctive in* (4) *is therefore not the mere presence of a wedge relative to the canonical IS curve, but the specific* risk-centric *microfoundation of that wedge.*

3 Optimal Monetary Policy

The central bank's policy problem can be stated as choosing sequences for the output gap, inflation, and the policy rate, $\{x_t, \pi_t, i_t\}_{t=0}^{\infty}$, so as to minimize

$$\min_{\{x_t,\pi_t,i_t\}_{t=0}^{\infty}} \mathbb{E}_0^{\mathcal{C}} \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \lambda x_t^2\right) \tag{7}$$

subject to the sequence of NKPCs (2) and IS curves (4), given the exogenous state processes for (u_t, s_t, y_t^p).[6]

I focus on the case without commitment.[7] Hence, the problem becomes intratemporal in the standard New Keynesian sense: the central bank re-optimizes period by period, taking private-sector expectations as given.[8] To keep the comparison with the canonical New Keynesian benchmark as clean as possible, assume from now on that all agents have full-information rational expectations:

$$\mathbb{E}_t^{\mathcal{F}}[\cdot] = \mathbb{E}_t^{\mathcal{M}}[\cdot] = \mathbb{E}_t^{\mathcal{C}}[\cdot] = \mathbb{E}_t[\cdot]$$

The equilibrium concept is therefore the standard discretionary rational-expectations equilibrium.

Definition 1 (Discretionary Rational Expectations Equilibrium)**.** *Given the exogenous state processes for* (u_t, s_t, y_t^p)*, a discretionary equilibrium is a collection of stochastic processes for the endogenous variables such that:*

1. *The private-sector equilibrium conditions hold, namely, the New Keynesian Phillips Curve* (2) *and the risk-centric IS curve* (4)*.*

2. *In each period, taking expectations as given, the triplet* (x_t, π_t, i_t) *solves the central bank's static constrained problem* $\min_{x_t,\pi_t,i_t} \pi_t^2 + \lambda x_t^2$*, subject to the private-sector equilibrium conditions.*

[6] This quadratic objective also admits a welfare interpretation in the present two-agent environment. In Bilbiie (2026), the second-order approximation to welfare in a TANK/THANK model contains, in addition to inflation and output, a term in cyclical consumption inequality. Here, however, consumption inequality is constant in log deviations, so there is no independent cyclical inequality wedge. Hence the welfare-based quadratic approximation collapses to the standard loss in inflation and the welfare-relevant output gap.

[7] Commitment is a useful normative benchmark, but would add history dependence that is orthogonal to the mechanism emphasized here.

[8] Since policy is derived from an optimization problem rather than from an exogenous interest-rate rule, determinacy in the Taylor-rule sense is not the relevant issue here. Under the usual boundedness condition $|\beta\rho_u| < 1$, the bounded linear discretionary equilibrium is unique.

3.1 Optimal allocation

The first-order condition of the Central Bank's problem is the familiar intratemporal trade-off

$$\lambda x_t + \kappa \pi_t = 0 \tag{8}$$

Combining (8) with the NKPC (2), and using the AR(1) law of motion for the cost-push shock, yields the bounded linear discretionary allocation

$$x_t^\star = -\frac{\kappa}{\kappa^2 + \lambda(1-\beta\rho_u)}\, u_t \;\equiv\; x_u^\star u_t, \qquad \pi_t^\star = \frac{\lambda}{\kappa^2 + \lambda(1-\beta\rho_u)}\, u_t \;\equiv\; \pi_u^\star u_t \tag{9}$$

These expressions are exactly the canonical NK discretionary allocations under rational expectations with a persistent cost-push shock. In particular, the volatility state s_t does not enter the optimal allocation. Hence the risk-centric features of the model do not alter the policymaker's inflation–output trade-off, nor the point chosen on the Phillips-curve frontier. They affect only the implementation problem, that is, the policy rate required to support the allocation (9).

3.2 Implementability of optimal policy

The optimal allocation (9) prescribes a path for $(\pi_t^\star, x_t^\star)$. The central bank must set the nominal rate i_t so that this path is consistent with the IS curve. Inverting (4) yields the general implementability condition

$$i_t^\star = r_t^n + \mathbb{E}_t^{\mathcal{M}}[x_{t+1}^\star] + \mathbb{E}_t^{\mathcal{M}}[\pi_{t+1}^\star] - x_t^\star - \Psi_t^\star \tag{10}$$

Under the discretionary optimum, the output gap and inflation depend only on the cost-push state, as in (9). Since wealth managers correctly understand optimal policy,

$$\mathbb{E}_t^{\mathcal{M}}[x_{t+1}^\star] = x_u^\star \rho_u u_t, \qquad \mathbb{E}_t^{\mathcal{M}}[\pi_{t+1}^\star] = \pi_u^\star \rho_u u_t \tag{11}$$

To compute the risk premium, note that the output gap and the output–asset-price relation imply

$$y_t = x_t^\star + y_t^p, \qquad p_t = x_t^\star + y_t^p - m$$

Using the Campbell–Shiller approximation, the log return on aggregate wealth can be written as

$$r_{t+1} = \chi + \Delta x_{t+1}^\star + z_{t+1} \tag{12}$$

Hence, return risk has two components: one from the cost-push block, via the optimal response of the output gap, and one from innovations to potential-output growth. Because z_{t+1} is i.i.d. and independent of η_{t+1}, the conditional second moments simplify to

$$\mathrm{Var}_t^{\mathcal{M}}(r_{t+1}) = (x_u^\star)^2 s_t + \sigma_z^2 \tag{13}$$

$$\mathrm{Var}_t^{\mathcal{M}}(\pi_{t+1}) = (\pi_u^\star)^2 s_t \tag{14}$$

$$\mathrm{Cov}_t^{\mathcal{M}}(\pi_{t+1}, r_{t+1}) = x_u^\star \pi_u^\star s_t \tag{15}$$

Substituting (13)–(15) into (5) yields

$$\Psi_t^\star = G^\star s_t + \frac{1}{2}\sigma_z^2 \tag{16}$$

where

$$G^\star \equiv \frac{1}{2}\left(\frac{\lambda - \kappa}{\kappa^2 + \lambda(1 - \beta\rho_u)}\right)^2 \geq 0 \tag{17}$$

The term $G^\star s_t$ is the endogenous risk-pricing component associated with stochastic volatility in the cost-push shock. The term $\frac{1}{2}\sigma_z^2$ is a constant intercept shift induced by exogenous potential-output risk. Substituting the optimal expected allocations into (10) and simplifying gives the policy rule that implements the discretionary optimum.

Proposition 1 (Optimal policy rate with risk premia)**.** *Under the risk-centric discretionary optimum, the policy rate that implements the optimal allocation is (in deviations from the deterministic steady state):*

$$\hat{i}_t^{RC} = \hat{i}_t^{NK} - G^\star \hat{s}_t, \qquad G^\star \geq 0 \tag{18}$$

where

$$\hat{i}_t^{NK} = \frac{\kappa(1 - \rho_u) + \lambda\rho_u}{\kappa^2 + \lambda(1 - \beta\rho_u)} u_t$$

Proof. See Appendix D. □

Unlike in the canonical NK model, a rise in s_t raises risk premia, depressing asset valuations and compressing demand. Thus, the Risk-Centric implementing rate is weakly below the canonical NK counterpart when $\hat{s}_t \geq 0$.[9]

The stronger the uncertainty response triggered by the energy shock, the smaller the optimal policy tightening. In that sense, greater uncertainty calls for a more dovish implementation of the optimal allocation. Under discretion, this follows directly from the implementation equation and the chain rule:

$$\frac{d\hat{i}_t^{RC}}{d\iota} = \frac{\partial \hat{i}_t^{RC}}{\partial \hat{s}_t}\frac{d\hat{s}_t}{d\iota} = -\,G^\star \frac{d\hat{s}_t}{d\iota} < 0$$

[9] In levels, the optimal policy rate reads as

$$i_t^{RC} = \rho + \omega u_t - G^\star s_t - \frac{1}{2}\sigma_z^2$$

. Thus, constant variance terms appear from the original Caballero and Simsek (2022) structure. The distinct element of this note is the fact that some variances are time-varying.

Hence, when a larger ι makes the energy shock generate a stronger volatility response, the central bank sets a lower policy rate to implement the same inflation–output allocation.

A knife-edge case. When $\lambda = \kappa$, $G^\star = 0$ and the risk-centric rate collapses to the New Keynesian one.

Taylor rule representation. The discretionary optimum can also be written in a Taylor-style form. This is useful for comparing the Risk-Centric equilibrium with the canonical NK benchmark. The key difference is that, in the Risk-Centric economy, the implementing rate responds not only to inflation but also to volatility, because volatility endogenously tightens financial conditions through risk premia.

Using the optimal allocation in (9), the implementing rule in Proposition 1 can be rewritten as

$$\hat{i}_t^{RC} = \phi^\star \hat{\pi}_t^\star - G^\star \hat{s}_t, \tag{19}$$

where

$$\phi^\star \equiv \frac{\omega}{\pi_u^\star} = \rho_u + \frac{\kappa(1-\rho_u)}{\lambda}. \tag{20}$$

Hence, the discretionary optimum admits a Taylor-style representation: it responds to inflation with coefficient $\phi^\star$, but it also leans against the endogenous tightening generated by volatility through the term $-G^\star \hat{s}_t$. In the canonical NK benchmark, by contrast, the implementing rule is simply

$$\hat{i}_t^{NK} = \phi^\star \hat{\pi}_t^\star. \tag{21}$$

Thus, relative to the canonical model, the Risk-Centric economy features the same inflation coefficient, but an additional adjustment for risk.

The derivation is immediate from (18). Since the optimal allocation implies $\hat{\pi}_t^\star = \pi_u^\star u_t$, the level-shock term can be rewritten as $\omega u_t = (\omega/\pi_u^\star)\hat{\pi}_t^\star$. Defining $\phi^\star \equiv \omega/\pi_u^\star$ yields (19).[10]

Transmission mechanism. A way of summarizing the transmission mechanism under optimal policy is as follows. Once the central bank chooses the nominal policy rate required to implement the optimal allocation, the policy stance first affects the risk-adjusted real short rate through the Fisher relation. That real rate then pins down the expected return required by wealth managers on risky assets. In turn, the required return is capitalized into equilibrium asset valuations: a higher policy rate raises the relevant discount rate, lowers risky asset prices, and tightens financial conditions. Because aggregate demand is governed by wealth effects, lower asset valuations depress current spending and output relative to potential, whereas higher valuations stimulate demand.

Hence, in this environment, optimal monetary policy is transmitted not primarily through the

[10] Conceptually, this does not mean that optimal policy is obtained by imposing a Taylor rule ex ante. Rather, once the discretionary optimum has been solved for, the implementing interest-rate rule can be expressed in Taylor-like form.

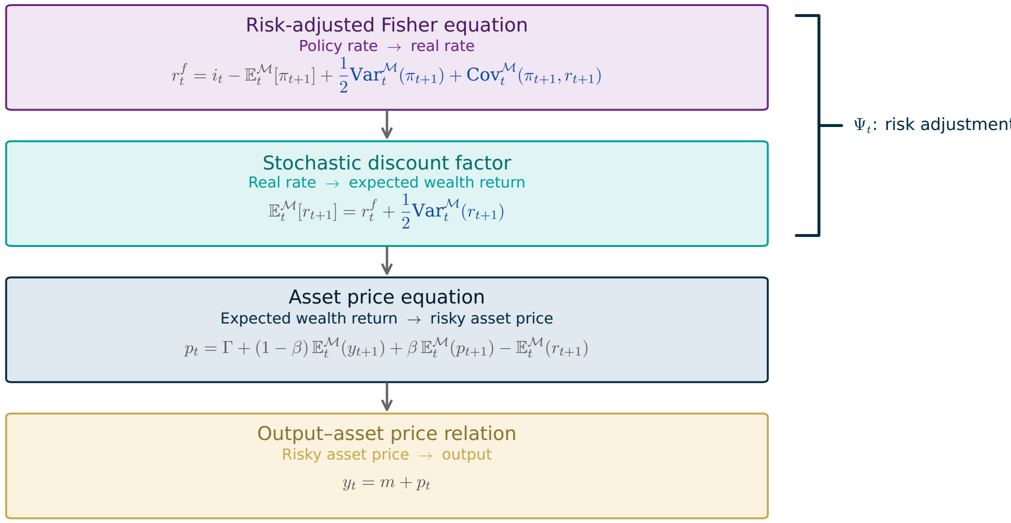

Figure 3: The transmission of monetary policy in the Risk-Centric model.

representative household's intertemporal substitution margin, but through a sequence running from the policy rate to the risk-adjusted real rate, from there to required wealth returns, then to asset prices, and finally to aggregate demand. Macro-financial risk matters at each step of this chain, since it affects both the risk-adjusted Fisher relation and the discounting of risky payoffs, thereby altering how a given policy-rate path translates into financial conditions and output. Figure 3 summarizes this mechanism.

This different transmission mechanism is illustrated by the optimal responses to the energy shock in Figure 4.[11] In the baseline simulation, when the shock hits, the central bank raises the policy rate, but by less than in a standard NK model, and the expected path of future policy rates shifts up more modestly, as reflected in the lower one-year-ahead expected rate and the flatter term structure of expected short rates. At the same time, the shock generates a persistent increase in the endogenous risk premium Ψ_t, which lowers risky asset prices relative to baseline. This asset-price decline tightens financial conditions and weakens demand, so part of the contractionary adjustment is delivered directly by financial markets. As a result, the real yield rises more gradually and displays a hump-shaped profile: in this calibration, financial markets absorb part of the tightening through repricing, not only through higher expected policy rates.

Relative to the canonical NK model, the paths of inflation and the output gap are unchanged, but the composition of the tightening differs markedly. In canonical NK, there is no independent risk-premium channel, so implementing the same stabilization requires a larger increase in both the current policy rate and expected future policy rates, which translates into a higher real rate. In the risk-centric economy, by contrast, the central bank implements the same inflation–output

[11] The figure is intended to illustrate the composition of the tightening in the baseline calibration, rather than to establish an additional general result beyond the analytical implementation equations.

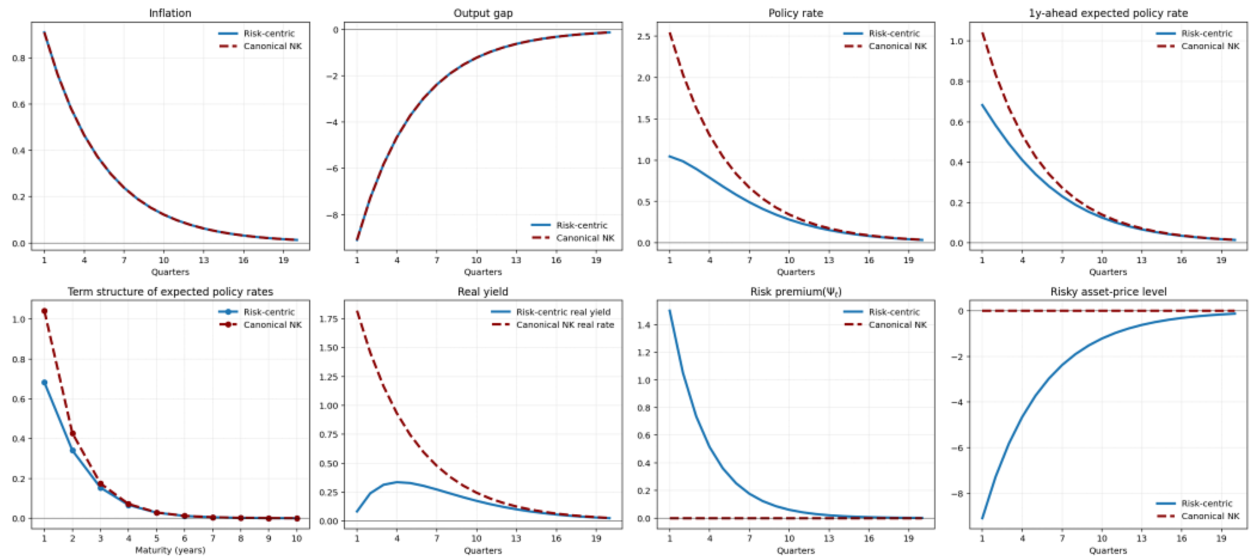

Figure 4: Responses to a one-standard-deviation shock to η_t under optimal discretionary policy in the Risk-Centric model and a canonical New Keynesian model.

allocation with a smaller policy-rate response because higher risk premia and lower asset prices already depress demand endogenously. The key difference, therefore, is not the allocation delivered by optimal policy, but the transmission mechanism through which that allocation is implemented.

3.2.1 Implementability with time-varying uncertainty in potential output

Large energy shocks may raise not only current marginal costs, but also uncertainty about future supply conditions—for instance, through persistent supply-chain disruptions. To capture this possibility, let potential output continue to follow the random walk $y_t^p = y_{t-1}^p + z_t$, but assume

$$z_{t+1} \mid \mathcal{I}_t \sim \mathcal{N}(0, s_{z,t}) \tag{22}$$

with conditional variance

$$s_{z,t+1} = (1-\rho_{sz})\bar{s}_z + \rho_{sz} s_{z,t} + \iota_z \, \eta_{t+1} + \nu_{t+1}^z, \qquad \nu_{t+1}^z \mid \mathcal{I}_t \sim \mathcal{N}(0, \omega_{\nu_z}^2) \tag{23}$$

where z_{t+1} is conditionally independent of η_{t+1} and ν_{t+1}^z. Again, equations above are interpreted as a local approximation around a deterministic steady state $\bar{s}_z$.

This extension does not change the discretionary target allocation, because neither the Phillips-curve frontier nor the central bank's objective is affected. It changes only implementation. In particular, potential-output uncertainty adds a time-varying term to the risk premium:

$$\Psi_t^\star = G^\star s_t + \frac{1}{2} s_{z,t} \tag{24}$$

Hence, the policy rate that implements the discretionary optimum becomes, in deviations from steady state,

$$\hat{i}_t^{\star} = \hat{i}_t^{NK} - G^{\star}\hat{s}_t - \frac{1}{2}\hat{s}_{z,t} \tag{25}$$

where $\hat{s}_{z,t} \equiv s_{z,t} - \bar{s}_z$. Thus, time-varying supply uncertainty does not modify the optimal inflation–output trade-off; it only lowers the rate required to implement it. If an energy shock also increases uncertainty about future supply conditions, risk premia rise further, financial conditions tighten more, and the central bank must cut the policy rate by more to sustain the optimal allocation. In this sense, supply uncertainty reinforces the baseline result: optimal policy responds less aggressively to the shock.

4 Monetary Policy Inertia

The discretionary optimum above abstracts from a salient feature of monetary policy: central banks typically adjust policy rates gradually. Estimated policy rules display substantial interest-rate smoothing, and central-bank practice often emphasizes gradualism, uncertainty about transmission, and a preference to avoid sharp policy reversals; see, among others, Clarida, Galí, and Gertler (1999) and Woodford (2003). It is therefore natural to ask whether policy inertia modifies the previous results, especially the separation between allocations and implementation that insulated allocations from stochastic volatility.

To capture this feature in the simplest way, suppose the central bank acts under discretion but minimizes

$$\sum_{t=0}^{\infty} \beta^t \Big(\pi_t^2 + \lambda x_t^2 + \delta_i (i_t - i_{t-1})^2\Big), \qquad \delta_i > 0$$

subject to the risk-centric IS curve and the forward-looking Phillips curve. Unlike in the baseline discretionary problem, the lagged policy rate i_{t-1} is now a state variable.

The discretionary equilibrium can be expressed in affine form as

$$\hat{x}_t^{\mathcal{I}} = x_i\,\hat{i}_{t-1} + x_u\,u_t + x_s\,\hat{s}_t, \qquad \hat{\pi}_t^{\mathcal{I}} = \pi_i\,\hat{i}_{t-1} + \pi_u\,u_t + \pi_s\,\hat{s}_t \tag{26}$$

with the policy rate

$$\hat{i}_t^{\mathcal{I}} = q_i\,\hat{i}_{t-1} + q_u\,u_t + q_s\,\hat{s}_t \tag{27}$$

Appendix E derives the fixed-point system that characterizes these coefficients. Under the stable-root conditions reported there, the volatility loadings satisfy

$$x_s < 0, \qquad \pi_s < 0, \qquad q_s < 0$$

The time-varying risk premium satisfies

$$\Psi_t = \psi_s s_t + \frac{1}{2}\omega_\nu^2 (x_s + \pi_s)^2 + \frac{1}{2}\sigma_z^2 \tag{28}$$

where

$$\psi_s = \frac{1}{2}\big(x_u + \pi_u + \iota(x_s + \pi_s)\big)^2 \geq 0$$

Relative to the frictionless discretionary equilibrium in equation (16), the sensitivity of risk premia to volatility is now itself shaped by the inertia problem.[12] Thus, unlike under frictionless discretion, optimal allocations are no longer insulated from stochastic volatility. Because sharp movements in the policy rate are costly, the central bank does not fully offset the volatility-induced tightening on impact. Part of the adjustment is therefore absorbed by quantities, so both output and inflation become directly sensitive to $\hat{s}_t$.

Proposition 2 (Optimal policy with risk premia and policy inertia)**.** *Consider the discretionary equilibria with policy inertia in the NK and RC economies. For any fixed common state* $(\hat{i}_{t-1}, u_t, \hat{s}_t)$ *with* $\hat{s}_t \geq 0$*, the Risk-Centric equilibrium is weakly below the NK benchmark:*

$$\hat{x}_t^{RC,\mathcal{I}} \leq \hat{x}_t^{NK,\mathcal{I}}, \qquad \hat{\pi}_t^{RC,\mathcal{I}} \leq \hat{\pi}_t^{NK,\mathcal{I}}, \qquad \hat{i}_t^{RC,\mathcal{I}} \leq \hat{i}_t^{NK,\mathcal{I}}$$

Proof. See Appendix E. □

The proof shows that the coefficients on $(\hat{i}_{t-1}, u_t)$ coincide across the two economies,

$$(x_i^{RC}, x_u^{RC}, \pi_i^{RC}, \pi_u^{RC}, q_i^{RC}, q_u^{RC}) = (x_i^{NK}, x_u^{NK}, \pi_i^{NK}, \pi_u^{NK}, q_i^{NK}, q_u^{NK})$$

so the Risk-Centric equilibrium differs from the NK equilibrium only through the volatility block:

$$\hat{x}_t^{RC,\mathcal{I}} = \hat{x}_t^{NK,\mathcal{I}} + x_s\,\hat{s}_t \tag{29}$$

$$\hat{\pi}_t^{RC,\mathcal{I}} = \hat{\pi}_t^{NK,\mathcal{I}} + \pi_s\,\hat{s}_t \tag{30}$$

$$\hat{i}_t^{RC,\mathcal{I}} = \hat{i}_t^{NK,\mathcal{I}} + q_s\,\hat{s}_t \tag{31}$$

Importantly, these inequalities are state-by-state comparisons, not rankings of dynamic impulse responses. Along simulated transition paths, the lagged policy state generally differs across the two economies, so the pathwise ordering of $\hat{x}_t$ and $\hat{\pi}_t$ need not be monotone. Figure 5 illustrates this point in the baseline calibration: although the proposition delivers a comparison for a common state, the simulated impulse responses need not exhibit a uniform ordering at each horizon.

In the used (illustrative) calibration, the contraction in activity is initially larger in the Risk-Centric economy. The reason is that the energy shock raises stochastic volatility and, with it, risk premia, generating an additional tightening in private financial conditions. The central bank partly accommodates that tightening by setting a lower policy rate than in the NK benchmark, but the adjustment of the instrument is dampened by the inertia term in the loss function. In this calibration, the lower policy-rate path is not sufficient to fully offset the contractionary effect of higher risk premia on impact, so the output gap falls by more than in the canonical NK economy. At later

[12]See Appendix E for the derivation of equation (28).

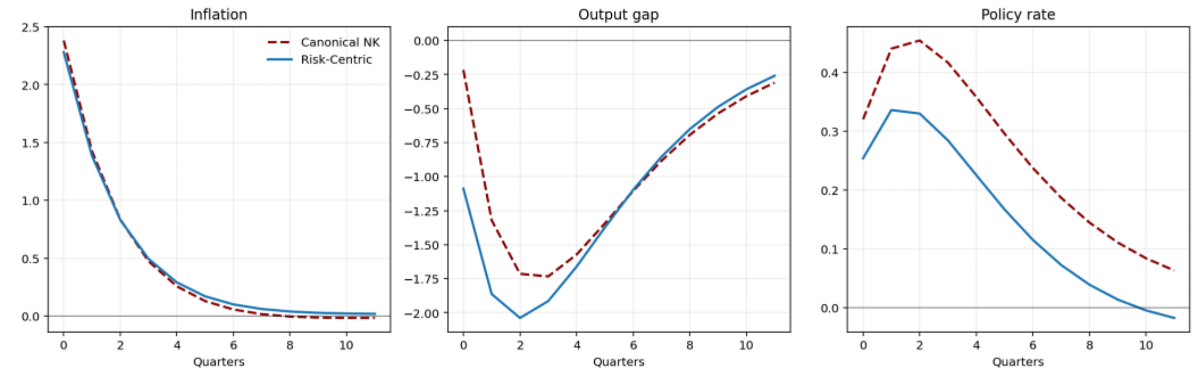

Figure 5: Responses of inflation, the output gap, and the policy rate to a one-standard-deviation shock to η_t under optimal under discretionary policy with interest-rate inertia.

horizons, however, as the volatility component dissipates and the cumulative effect of persistently lower rates works through demand, the Risk-Centric path catches up and may temporarily move above the NK benchmark.

Inflation behaves somewhat differently. Despite the lower policy-rate path in the Risk-Centric economy, the inflation response initially remains close to the NK benchmark because the additional contraction in activity generated by higher risk premia largely offsets the weaker monetary tightening. Only in the medium run, once the lagged effects of lower rates become more important and output in the Risk-Centric economy overtakes the NK path, does inflation become slightly more persistent.

The main message of Figure 5 is therefore that, in the baseline calibration, a similar inflation path to the canonical NK benchmark is associated in the Risk-Centric economy with both a deeper short-run recession and a lower policy-rate path. In that sense, endogenous risk premia can worsen the short-run inflation–output trade-off under policy inertia.[13]

Taylor-style representation with policy inertia As before, For interpretation, it is useful to rewrite the implementing policy rule in a Taylor-style form. This makes the comparison with both the frictionless Risk-Centric equilibrium and the canonical NK benchmark more transparent. In particular, using the inflation equation

$$\hat{\pi}_t^{\mathcal{I}} = \pi_i \,\hat{i}_{t-1} + \pi_u \,u_t + \pi_s \,\hat{s}_t$$

to eliminate the cost-push shock u_t from

$$\hat{i}_t^{\mathcal{I}} = q_i \,\hat{i}_{t-1} + q_u \,u_t + q_s \,\hat{s}_t,$$

[13]More broadly, policy inertia shifts stabilization away from the front-loaded adjustment of the frictionless case and toward a delayed and more distributed response over time, a standard implication of interest-rate smoothing and gradualism in monetary policy; see, for example, Clarida, Galí, and Gertler (1999) and Woodford (2003).

the implementing rule can be written as

$$\hat{i}_t^{\mathcal{I}} = \rho_i^\star \hat{i}_{t-1} + \phi_\pi^\star \hat{\pi}_t^{\mathcal{I}} - G_{\mathcal{I}}^\star \hat{s}_t, \tag{32}$$

where

$$\rho_i^\star \equiv q_i - \frac{q_u \pi_i}{\pi_u}, \qquad \phi_\pi^\star \equiv \frac{q_u}{\pi_u}, \qquad G_{\mathcal{I}}^\star \equiv \frac{q_u \pi_s}{\pi_u} - q_s.$$

Equation (32) is the natural counterpart of the frictionless representation,

$$\hat{i}_t^{\mathrm{RC}} = \phi_\pi^\star \hat{\pi}_t^{\mathrm{RC}} - G^\star \hat{s}_t.$$

Hence, in both the frictionless and inertial environments, the key departure from the canonical NK representation is the presence of an additional volatility term. In standard NK, once the optimal allocation is rewritten as a Taylor-style rule, the policy rate responds only to inflation and, with interest-rate smoothing, to its own lag. Here, by contrast, volatility enters separately because it moves risk premia and therefore shifts private financial conditions independently of the expected path of policy rates. The extra $\hat{s}_t$ term is thus the reduced-form footprint of the risk-premium channel in the implementing rule.

At the same time, unlike in the frictionless case, the sign of the volatility coefficient under inertia is not immediate. Since

$$G_{\mathcal{I}}^\star = \frac{q_u \pi_s}{\pi_u} - q_s,$$

volatility affects the policy rate both indirectly, through its effect on equilibrium inflation, and directly, through the policy rule itself. These two channels need not reinforce each other, so the overall sign of $G_{\mathcal{I}}^\star$ is in general ambiguous.

4.1 Equilibrium with a Taylor rule

The exposure of allocations to stochastic volatility is not specific to the discretionary optimum with policy inertia. It also arises when monetary policy is delegated to a simple interest-rate rule. This is useful because it shows that the key mechanism is more general: once monetary policy does not fully neutralize the endogenous tightening in financial conditions, stochastic volatility feeds into equilibrium allocations through risk premia.

Suppose the central bank follows the Taylor rule

$$i_t = \phi \pi_t, \qquad \phi > 1$$

and that this reaction function is common knowledge. Then the equilibrium in deviations from the deterministic steady state can be written as

$$\hat{x}_t^{\mathcal{T}} = x_u\, u_t + x_s\, \hat{s}_t, \qquad \hat{\pi}_t^{\mathcal{T}} = \pi_u\, u_t + \pi_s\, \hat{s}_t \tag{33}$$

The coefficients on the level shock u_t coincide with those in the canonical NK model under the same Taylor rule. Hence, the key departure from canonical NK is the presence of the volatility terms x_s and π_s, which capture the effect of stochastic volatility on equilibrium allocations through risk premia. Closed-form expressions for all coefficients are reported in Appendix F.

The associated risk-adjustment term is

$$\Psi_t^{\mathcal{T}} = \frac{1}{2}\big[(x_u + \iota x_s) + (\pi_u + \iota \pi_s)\big]^2 s_t + \frac{1}{2}\omega_\nu^2 (x_s + \pi_s)^2 + \frac{1}{2}\sigma_z^2 \tag{34}$$

which is non-negative for all $s_t \geq 0$. Hence, even under a simple Taylor rule, stochastic volatility raises required returns and tightens private financial conditions.

Under the same rule, the canonical NK benchmark satisfies

$$\hat{x}_t^{\mathcal{T},NK} = x_u\, u_t, \qquad \hat{\pi}_t^{\mathcal{T},NK} = \pi_u\, u_t, \qquad \hat{i}_t^{\mathcal{T},NK} = \phi\, \hat{\pi}_t^{\mathcal{T},NK} \tag{35}$$

Therefore, for $\hat{s}_t > 0$,

$$\hat{x}_t^{\mathcal{T}} < \hat{x}_t^{\mathcal{T},NK}, \qquad \hat{\pi}_t^{\mathcal{T}} < \hat{\pi}_t^{\mathcal{T},NK}, \qquad \hat{i}_t^{\mathcal{T}} < \hat{i}_t^{\mathcal{T},NK} \tag{36}$$

The intuition is the same as under policy inertia. The endogenous wedge $\Psi_t^{\mathcal{T}}$ depresses asset prices and demand, so equilibrium output and inflation are lower than in the canonical NK economy.

5 Conclusions

When monetary policy works primarily through asset prices, rather than through household intertemporal substitution, the endogenous response of risk premia to macroeconomic uncertainty becomes relevant for the conduct of policy. An energy shock that raises both current cost pressures and uncertainty about their future path tightens financial conditions through risk pricing. That endogenous tightening substitutes, at least in part, for the adjustment in the policy rate.

The central message of the paper is that endogenous risk premia alter the implementation of optimal policy even when they do not alter the benchmark allocation. Under frictionless discretion, there is a sharp separation between allocation and implementation: the optimal inflation–output trade-off is the same as in the canonical New Keynesian benchmark, but the policy rate required to support that allocation is lower. The reason is that, in the risk-centric model, an increase in cost-push uncertainty raises risk premia, depresses asset prices, and contracts demand even if the policy rate itself does not move. Hence, a central bank that ignores this channel and follows the canonical NK prescription will over-tighten, generating an unnecessarily large contraction in activity.

The analysis with policy inertia qualifies and enriches this benchmark. Once sharp policy-rate movements are costly, the central bank does not fully offset the volatility-induced tightening on impact, and the separation between allocation and implementation breaks down: stochastic volatility now affects output and inflation directly, not only the implementing rate. In that case, achieving an inflation path similar to the canonical NK benchmark may require both a deeper recession and a lower policy-rate path. Thus, endogenous risk premia can worsen the short-run inflation–output trade-off.

Several extensions are left for future work. In the current model, whether financial tightening comes from higher expected policy rates or risk premia is irrelevant; what matters is the overall magnitude, as the price level of the market portfolio is the sufficient statistic for financial tightening. One interesting venue is to study time-varying risk premia in a multi-asset environment in which asset prices affect demand through channels beyond wealth effects. For instance, embedding the mechanism in settings with endogenous borrowing constraints could break the equivalence between expected policy rates and risk premia, delivering additional insights on the conduct of monetary policy with time-varying risk premia. Furthermore, while the model has focused on risk pricing in financial markets, exploring its potential interactions with other risk-related margins, such as precautionary savings by asset-holders, might be interesting. Finally, a quantitative version of the framework could be embedded in richer estimated models to assess the empirical importance of the mechanism and to trace out a concrete optimal path of policy in response to energy shocks.

References

Acharya, Sushant, and Keshav Dogra. 2020. "Understanding HANK: Insights from a PRANK." *Econometrica* 88 (3): 1113–1158.

Acharya, Sushant, Edouard Challe, and Keshav Dogra. 2023. "Optimal Monetary Policy According to HANK." *American Economic Review* 113 (7): 1741–1782.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2018. "The Intertemporal Keynesian Cross." *NBER Working Paper* No. 25020.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–2367.

Auclert, Adrien, Hugo Monnery, Matthew Rognlie, and Ludwig Straub. 2023. "Managing an Energy Shock: Fiscal and Monetary Policy." *NBER Working Paper* 31543. National Bureau of Economic Research.

Bilbiie, Florin O. 2008. "Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic." *Journal of Economic Theory* 140 (1): 162–196.

Bilbiie, Florin O. 2020. "The New Keynesian Cross." *Journal of Monetary Economics* 114: 90–108.

Bobasu, Alina, Michael Dobrew, and Amalia Repele. 2025. "Energy Price Shocks, Monetary Policy and Inequality." *European Economic Review* 175: 104986.

Bodenstein, Martin, Luca Guerrieri, and Lutz Kilian. 2012. "Monetary Policy Responses to Oil Price Fluctuations." *IMF Economic Review* 60: 470–504.

Born, Benjamin, and Johannes Pfeifer. 2014. "Policy Risk and the Business Cycle." *Journal of Monetary Economics* 68: 68–85.

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

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–398.

Chan, Jenny, Sebastian Diz, and Derrick Kanngiesser. 2024. "Energy Prices and Household Heterogeneity: Monetary Policy in a Gas-TANK." *Journal of Monetary Economics* 147 (S): 103620.

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

Clarida, Richard, Jordi Galí, and Mark Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–1707.

Cogley, Timothy, Christian Matthes, and Argia M. Sbordone. 2015. "Optimized Taylor Rules for Disinflation When Agents Are Learning." *Journal of Monetary Economics* 72: 131–147.

Crawley, S. Edmund. 2025. "Do Households Substitute Intertemporally? 10 Structural Shocks That Suggest Not" *Finance and Economics Discussion Series Working Paper.*

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

Dávila, Eduardo, and Andreas Schaab. 2023. "Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy." *NBER Working Paper* 30961. Cambridge, MA: National Bureau of Economic Research.

Elder, John, and Apostolos Serletis. 2010. "Oil Price Uncertainty." *Journal of Money, Credit and Banking* 42 (6): 1137–1159.

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.

Filardo, Andrew, Paul Hubert, and Phurichai Rungcharoenkitkul. 2019. "The Reaction Function Channel of Monetary Policy and the Financial Cycle." *BIS Working Papers* 816.

Forni, Mario, Alessandro Franconi, Luca Gambetti, and Luca Sala. 2025. "Asymmetric Transmission of Oil Supply News." *Quantitative Economics* 16: 947–979.

Galí, Jordi, and Tommaso Monacelli. 2005. "Monetary Policy and Exchange Rate Volatility in a Small Open Economy." *Review of Economic Studies* 72 (3): 707–734.

Gnocato, Nicolò. 2025. "Energy Price Shocks, Unemployment, and Monetary Policy." *Journal of Monetary Economics* 151: 103734.

Iacoviello, Matteo. 2005. "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle." *American Economic Review* 95 (3): 739–764.

Kakhbod, Ali, Amir Kermani, and Bernardo Maciel. 2026. "In the Fed's Mind." *NBER Working Paper* 35016.

Känzig, Diego R. "The macroeconomic effects of oil supply news: Evidence from OPEC announcements." *American Economic Review* 111.4 (2021): 1092-1125.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3): 697–743.

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

Kilian, Lutz, and Robert J. Vigfusson. 2011. "Nonlinearities in the Oil Price–Output Relationship." *Macroeconomic Dynamics* 15 (S3): 337–363.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review* 106 (10): 3133–3158.

Merendino, Alfonso, and Tommaso Monacelli. 2026. "Supply Chain Uncertainty, Energy Prices, and Inflation." Working paper.

Monacelli, Tommaso. 2006. "Optimal Monetary Policy with Collateralized Household Debt and Borrowing Constraints." *NBER Working Paper* No. 12470.

Mumtaz, Haroon, and Konstantinos Theodoridis. 2015. "The International Transmission of Volatility Shocks: An Empirical Analysis." *Journal of the European Economic Association* 13 (3): 512–533.

Natal, Jean-Marc. 2012. "Monetary Policy Response to Oil Price Shocks." *Journal of Money, Credit and Banking* 44 (1): 53–101.

Oh, Joonseok. 2020. "The Propagation of Uncertainty Shocks: Rotemberg versus Calvo." *International Economic Review* 61 (3): 1097–1113.

Plante, Michael. 2014. "How Should Monetary Policy Respond to Changes in the Relative Price of Oil? Considering Supply and Demand Shocks." *Journal of Economic Dynamics and Control* 44: 1–19.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton: Princeton University Press.

Appendix A Local projections on the size of energy shocks

This appendix describes the local projections used to provide motivating empirical evidence for the risk channel of energy shocks. The objective is to document that larger oil and gas supply news shocks are associated with a rise in market-implied compensation for risk. The dependent variable, denoted by y_t, is either the option-implied equity risk premium or, in alternative specifications, implied equity-market volatility. The shock series, denoted by z_t, is the daily oil or gas supply news shock. For each horizon $h = 0, 1, \ldots, H$, the regressions are estimated by OLS separately horizon by horizon.

Let X_t denote the vector of controls included in each regression. In the baseline implementation, X_t contains lags of the dependent variable and lags of the daily financial controls described in the main text. Standard errors are computed using a heteroskedasticity- and autocorrelation-consistent Newey–West covariance estimator.

Baseline specification: linear in shock size. The baseline specification relates the future outcome to the absolute size of the energy shock:

$$y_{t+h} = \alpha_h + \theta_h |z_t| + \Gamma_h' X_t + \varepsilon_{t+h} \tag{37}$$

This specification imposes sign symmetry: positive and negative shocks of the same absolute magnitude are assumed to have the same effect on y_{t+h}. The coefficient θ_h therefore captures the response of the outcome to shock size, regardless of sign. A positive estimate of θ_h means that larger energy shocks are associated with a higher option-implied equity risk premium (or higher implied volatility) at horizon h. Controls included lagged oil returns, stock returns, the gap between the 10y and 2y government bond yields, the central bank policy rate and lagged values of the dependent variable. The results use 10 lags, but are robust to the number of lags and removing controls.

The main text uses the option-implied equity risk premium as the dependent variable. Figure shows that the result goes through using an alternative measure of risk premium more directly linked to the corporate borrowing costs –the excess bond premium– and to use a direct measure of financial uncertainty –the VIX. This exercise is for the U.S.

It is useful to stress the interpretation of (37). Because the regressor is $|z_t|$, the coefficient θ_h is the slope of the local projection with respect to the magnitude of the shock. In that sense, (37) is designed to test whether larger energy disturbances, irrespective of sign, are followed by tighter market pricing of risk.

Nonlinear specification: small versus large shocks. To allow the response to differ between small and large shocks, I augment the baseline regression with a piecewise-linear tail term:

$$y_{t+h} = \alpha_h + \theta_h |z_t| + \gamma_h \big(|z_t| - b\big)_+ + \Gamma_h' X_t + \varepsilon_{t+h} \tag{38}$$

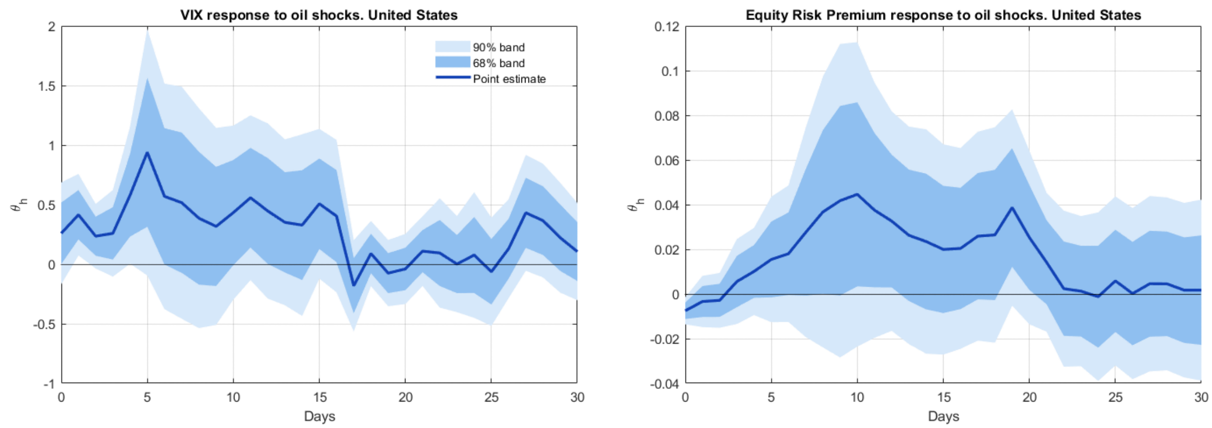

Figure 6: Daily response of the VIX (left hand side) and the excess bond premium (right hand side) to oil supply news shocks for the U.S.

where $(x)_+ \equiv \max\{x, 0\}$ and $b > 0$ is a threshold chosen from the empirical distribution of non-zero shock magnitudes. Equation (38) allows the slope of the response to change once the absolute size of the shock exceeds b. Thus, γ_h measures whether the response becomes steeper once the shock is sufficiently large. If $\gamma_h > 0$, large shocks raise the option-implied equity risk premium more strongly than small shocks. If $\gamma_h = 0$, the response is linear in shock size. Since (38) is still written in terms of $|z_t|$, it continues to impose the same response for positive and negative shocks of equal absolute magnitude. The left panel in figure 7 plots the marginal effects for small shocks (θ_h) and large shocks ($\theta_h + \gamma_h$), with larger shocks trigger a larger effect.

Signed-tail robustness: large positive versus large negative shocks. A natural concern is that the magnitude-based specification may conceal differences between large positive and large negative shocks. To assess this possibility, I estimate the following robustness regression:

$$y_{t+h} = \alpha_h + \beta_h |z_t| + \gamma_h^+ (z_t - b)_+ + \gamma_h^- (-z_t - b)_+ + \Gamma_h' X_t + \varepsilon_{t+h} \tag{39}$$

Hence, γ_h^+ and γ_h^- capture whether the response to large positive and large negative shocks departs from the common magnitude effect. If $\gamma_h^+ = \gamma_h^-$, then large positive and large negative shocks have the same incremental effect once one conditions on their absolute size. If instead $\gamma_h^+ \neq \gamma_h^-$, the data point to sign asymmetries in the tails. The right panel in figure 7 reports the effects for large positive vs large negative shocks. Both trigger a higher risk premia, but negative shocks seems to have a larger effect.

Appendix B Model Setup

This environment is unchanged relative to Caballero and Simsek (2022); see that paper for additional detail. The economy is populated by five types of agents: workers, capitalists, wealth

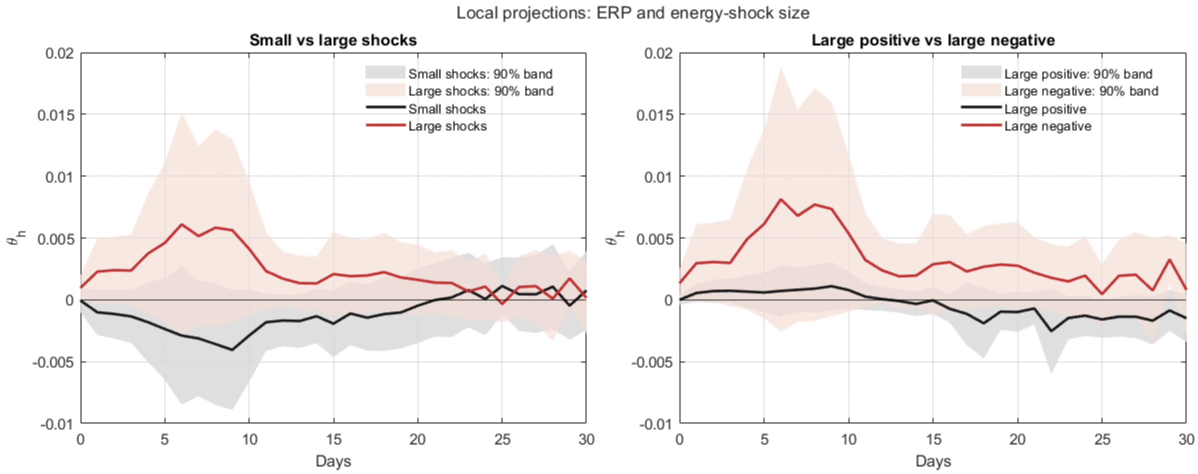

Figure 7: Nonlinear responses of the equity risk premium to oil shocks for the U.S.

managers, monopolistically competitive intermediate-goods firms, and a competitive final-goods firm, together with a central bank that sets the nominal short rate. Workers consume the final good, supply labor, and do not participate in asset markets. Capitalists also consume the final good, save out of wealth, and delegate portfolio choice to wealth managers. Wealth managers allocate capitalist savings across the market portfolio, 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. The final-goods firm aggregates those varieties into the unique consumption good.

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 for intermediate varieties

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

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, with Cobb–Douglas technology

$$Y_t(\nu) = A_t L_t(\nu)^{1-\alpha},$$

where A_t is aggregate productivity and $1-\alpha$ is labor's share in production. Demand for each variety is inherited from the final-good firm's problem:

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

Period profits are

$$\Pi_t(\nu) = P_t(\nu)Y_t(\nu) - W_t L_t(\nu) - T_t,$$

where T_t is the lump-sum tax term used to redistribute part of profits to workers, as in Caballero and Simsek (2022). Whenever allowed to reoptimize, firm ν chooses its price to maximize the expected discounted value of profits subject to the demand schedule above. With Calvo price setting, this delivers the Phillips-curve block used in the text. Aggregating across firms yields total profits received by capitalists:

$$\Pi_t \equiv \int_0^1 \Pi_t(\nu)\, d\nu = \alpha Q_t Y_t.$$

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

$$Q_t C_t^w = W_t L_t + T_t,$$

where W_t is the nominal wage and T_t is a lump-sum transfer. The first-order condition for labor supply is

$$\frac{W_t}{Q_t C_t^w} = \chi_L L_t^{\varphi},$$

or

$$L_t = \left(\frac{W_t}{\chi_L Q_t C_t^w}\right)^{1/\varphi}.$$

Following Caballero and Simsek (2022), transfers are chosen so that workers receive the technological labor share of output. Thus worker disposable income satisfies

$$W_t L_t + T_t = (1-\alpha) Q_t Y_t,$$

and therefore

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

Capitalists. Let

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

denote beginning-of-period real wealth, where A_{t-1} is wealth carried from period $t-1$ and R_t^p is the realized gross return on the managed portfolio between $t-1$ and t. A representative asset-holder household solves

$$V_t(X_t) = \max_{C_t^c, A_t} \left\{ \log C_t^c + \beta \mathbb{E}_t \left[V_{t+1}(R_{t+1}^p A_t) \right] \right\}$$

subject to

$$C_t^c + A_t = X_t.$$

Guess

$$V_t(X) = a_t + b \log X.$$

The Bellman equation implies $b = 1/(1-\beta)$, and the first-order condition for savings is

$$\frac{1}{C_t^c} = \frac{\beta b}{A_t}.$$

Combining this with the budget constraint yields the policy rules

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

Hence capitalists consume a constant fraction $1-\beta$ of beginning-of-period wealth and save the rest.

Wealth managers. Given capitalist savings A_t, wealth managers choose the portfolio shares invested in the market portfolio and in the nominal bond. Let ω_t denote the share invested in the market portfolio and ω_t^n the share invested in the nominal bond; the residual is invested in the real safe bond. The wealth manager solves

$$\max_{\omega_t, \omega_t^n} \mathbb{E}_t^{\mathcal{M}} \left[\log\left(A_t R_{t+1}^p \right) \right]$$

where

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

Since A_t is predetermined at the portfolio-allocation stage, this is equivalent to maximizing $\mathbb{E}_t^{\mathcal{M}}[\log R_{t+1}^p]$. The first-order conditions are

$$\mathbb{E}_t^{\mathcal{M}} \left[\frac{R_{t+1} - R_t^f}{R_{t+1}^p} \right] = 0, \qquad \mathbb{E}_t^{\mathcal{M}} \left[\frac{\frac{R_t^n}{Q_{t+1}/Q_t} - R_t^f}{R_{t+1}^p} \right] = 0$$

Defining the wealth-based stochastic discount factor as

$$\mathcal{M}_{t+1} \equiv \frac{1}{R^p_{t+1}}$$

these conditions can be written as

$$\mathbb{E}^{\mathcal{M}}_t\Big[\mathcal{M}_{t+1}(R_{t+1} - R^f_t)\Big] = 0$$

$$\mathbb{E}^{\mathcal{M}}_t\left[\mathcal{M}_{t+1}\left(\frac{R^n_t}{Q_{t+1}/Q_t} - R^f_t\right)\right] = 0$$

Because portfolio choice is delegated, asset prices are pinned down by the wealth managers' portfolio problem. Under the log-utility assumption used here, the induced wealth-based SDF is numerically identical to the asset holders' consumption kernel. The novelty relative to canonical NK is therefore not an algebraic difference between wealth- and consumption-based kernels per se, but that the relevant kernel is tied to the wealth of the asset-holding block rather than to aggregate household consumption.

Appendix C Derivation of the IS curve

To derive (4), start from goods-market clearing. Aggregate consumption is the sum of workers' labor income and capitalists' wealth-based demand:

$$C_t = (1-\beta)\, R^p_t A_{t-1} + (1-\alpha)Y_t \tag{40}$$

where R^p_t is the gross return on wealth, A_{t-1} is beginning-of-period wealth at market value, β is the discount factor of asset holders, and $(1-\alpha)$ is the competitive labor share. Since only equity is held in equilibrium,

$$R^p_t A_{t-1} = R_t P_{t-1} = \alpha Y_t + P_t \tag{41}$$

that is, capitalists receive the corporate share of income and own the aggregate stock of wealth. Imposing goods-market clearing, $C_t = Y_t$, then yields

$$Y_t = \frac{1-\beta}{\alpha\beta}\, P_t \tag{42}$$

Log-linearizing around the steady state gives the output–asset-price relation

$$y_t = m + p_t, \qquad m \equiv \log\left(\frac{1-\beta}{\alpha\beta}\right) \tag{43}$$

This is the Caballero and Simsek (2022) output–asset-price mapping.

On the asset-pricing side, it is useful to begin with a Campbell–Shiller log-linearization around the steady-state price-dividend ratio.[14] Let $r_{t+1} \equiv \log R_{t+1}$. Then

$$r_{t+1} \approx \Gamma + (1-\beta)y_{t+1} + \beta p_{t+1} - p_t, \tag{44}$$

or, equivalently,

$$p_t \approx \Gamma + (1-\beta)\mathbb{E}_t^{\mathcal{M}}[y_{t+1}] + \beta\mathbb{E}_t^{\mathcal{M}}[p_{t+1}] - \mathbb{E}_t^{\mathcal{M}}[r_{t+1}], \tag{45}$$

where

$$\Gamma \equiv -\beta \log \beta - (1-\beta)\log(1-\beta) + (1-\beta)\log \alpha.$$

In the linear-Gaussian equilibrium considered here, conditional on $\mathcal{I}_t$, next-period endogenous variables are affine in the Gaussian innovations. It follows that, under the objective probability measure, r_{t+1} is conditionally normal

$$r_{t+1} \mid \mathcal{I}_t \sim \mathcal{N}\big(\mathbb{E}_t^{\mathcal{M}}[r_{t+1}],\, \mathrm{Var}_t^{\mathcal{M}}(r_{t+1})\big),$$

In the Rational Expectations Equilibrium studied in the main text, the wealth managers' measure $\mathcal{M}$ coincides with the objective one.

The equilibrium stochastic discount factor is the inverse of the return on aggregate wealth, $1/R_{t+1}$. Hence the Euler equations for real and nominal safe bonds are

$$1 = \mathbb{E}_t^{\mathcal{M}}\left[\frac{R_t^r}{R_{t+1}}\right], \qquad 1 = \mathbb{E}_t^{\mathcal{M}}\left[\frac{R_t^n}{R_{t+1}(Q_{t+1}/Q_t)}\right], \tag{46}$$

where R_t^r and R_t^n denote the gross returns on real and nominal safe bonds, and Q_t is the consumption price level. Using the conditional normality of r_{t+1}, the real-bond Euler equation implies

$$\mathbb{E}_t^{\mathcal{M}}[r_{t+1}] = r_t^f + \frac{1}{2}\mathrm{Var}_t^{\mathcal{M}}(r_{t+1}), \tag{47}$$

where $r_t^f \equiv \log R_t^r$ is the log real safe rate. Combining the real and nominal bond Euler equations then yields the risk-adjusted Fisher equation

$$r_t^f = i_t - \mathbb{E}_t^{\mathcal{M}}[\pi_{t+1}] + \frac{1}{2}\mathrm{Var}_t^{\mathcal{M}}(\pi_{t+1}) + \mathrm{Cov}_t^{\mathcal{M}}(\pi_{t+1}, r_{t+1}), \tag{48}$$

where $i_t \equiv \log R_t^n$ and $\pi_{t+1} \equiv \log(Q_{t+1}/Q_t)$. Thus, the real safe rate equals the nominal rate net of expected inflation, corrected for inflation risk and its covariance with wealth returns.

Substituting (47) and (48) into the Campbell–Shiller relation (45) yields

$$p_t = \Gamma + (1-\beta)\mathbb{E}_t^{\mathcal{M}}[y_{t+1}] + \beta\mathbb{E}_t^{\mathcal{M}}[p_{t+1}] - \big(i_t - \mathbb{E}_t^{\mathcal{M}}[\pi_{t+1}]\big) - \Psi_t, \tag{49}$$

with Ψ_t defined in (5).

[14] In steady state, $P_t/Y_t = \alpha\beta/(1-\beta)$.

Finally, substitute the output–asset-price relation (43) into (49), and write output as the sum of potential output and the output gap, $y_t = x_t + y_t^p$. Rearranging gives

$$x_t = \mathbb{E}_t^{\mathcal{M}}[x_{t+1}] - \Big[\big(i_t - \mathbb{E}_t^{\mathcal{M}}[\pi_{t+1}]\big) + \Psi_t - \Big(\rho + \mathbb{E}_t^{\mathcal{M}}[\Delta y_{t+1}^p]\Big)\Big] \tag{50}$$

which is exactly (4)–(6).

Appendix D Proof of Proposition 1

Under the optimal allocation (9), the output gap and inflation depend only on u_t. Define

$$D \equiv \kappa^2 + \lambda(1 - \beta\rho_u), \qquad x_u^\star \equiv -\frac{\kappa}{D}, \qquad \pi_u^\star \equiv \frac{\lambda}{D}$$

Then

$$x_{t+1}^\star = x_u^\star u_{t+1} = x_u^\star(\rho_u u_t + \eta_{t+1}), \tag{51}$$

$$\pi_{t+1}^\star = \pi_u^\star u_{t+1} = \pi_u^\star(\rho_u u_t + \eta_{t+1}). \tag{52}$$

Using the Campbell–Shiller approximation and

$$y_t = x_t^\star + y_t^p, \qquad p_t = x_t^\star + y_t^p - m,$$

the return on aggregate wealth is

$$\begin{aligned}
r_{t+1} &= \Gamma + (1-\beta)y_{t+1} + \beta p_{t+1} - p_t \\
&= \Gamma + (1-\beta)(x_{t+1}^\star + y_{t+1}^p) + \beta(x_{t+1}^\star + y_{t+1}^p - m) - (x_t^\star + y_t^p - m) \\
&= \chi + \Delta x_{t+1}^\star + \Delta y_{t+1}^p \\
&= \rho + x_u^\star\big[(\rho_u - 1)u_t + \eta_{t+1}\big] + z_{t+1}
\end{aligned} \tag{53}$$

where $\chi \equiv \Gamma + (1-\beta)m = -\log\beta = \rho$ and the last equality uses (3).

Since z_{t+1} is i.i.d. with variance σ_z^2 and independent of η_{t+1},

$$\mathrm{Var}_t^{\mathcal{M}}(r_{t+1}) = (x_u^\star)^2 s_t + \sigma_z^2 \tag{54}$$

$$\mathrm{Var}_t^{\mathcal{M}}(\pi_{t+1}) = (\pi_u^\star)^2 s_t \tag{55}$$

$$\mathrm{Cov}_t^{\mathcal{M}}(\pi_{t+1}, r_{t+1}) = x_u^\star \pi_u^\star s_t \tag{56}$$

Substituting (54)–(56) into the definition of Ψ_t gives

$$\begin{aligned}\Psi_t^\star &= \frac{1}{2}\mathrm{Var}_t^{\mathcal{M}}(r_{t+1}) + \frac{1}{2}\mathrm{Var}_t^{\mathcal{M}}(\pi_{t+1}) + \mathrm{Cov}_t^{\mathcal{M}}(\pi_{t+1}, r_{t+1}) \\ &= \frac{1}{2}(x_u^\star)^2 s_t + \frac{1}{2}(\pi_u^\star)^2 s_t + x_u^\star \pi_u^\star s_t + \frac{1}{2}\sigma_z^2 \\ &= \frac{1}{2}(x_u^\star + \pi_u^\star)^2 s_t + \frac{1}{2}\sigma_z^2 \end{aligned} \tag{57}$$

Now,

$$x_u^\star + \pi_u^\star = \frac{\lambda - \kappa}{\kappa^2 + \lambda(1 - \beta\rho_u)}$$

so

$$G^\star = \frac{1}{2}\left(\frac{\lambda - \kappa}{\kappa^2 + \lambda(1 - \beta\rho_u)}\right)^2$$

which establishes (17). Hence

$$\Psi_t^\star = G^\star s_t + \frac{1}{2}\sigma_z^2$$

The natural benchmark is

$$r_t^n = \rho + \mathbb{E}_t^{\mathcal{M}}[\Delta y_{t+1}^p].$$

Using (3) and $\mathbb{E}_t^{\mathcal{M}}[z_{t+1}] = 0$,

$$r_t^n = \rho$$

The IS curve (4) then implies

$$\begin{aligned} i_t^\star &= r_t^n + \mathbb{E}_t^{\mathcal{M}}[x_{t+1}^\star] + \mathbb{E}_t^{\mathcal{M}}[\pi_{t+1}^\star] - x_t^\star - \Psi_t^\star \\ &= \rho + (x_u^\star + \pi_u^\star)\rho_u u_t - x_u^\star u_t - G^\star s_t - \frac{1}{2}\sigma_z^2 \\ &= \rho + \left[(x_u^\star + \pi_u^\star)\rho_u - x_u^\star\right] u_t - G^\star s_t - \frac{1}{2}\sigma_z^2 \end{aligned} \tag{58}$$

Computing the coefficient on u_t,

$$\begin{aligned}(x_u^\star + \pi_u^\star)\rho_u - x_u^\star &= \frac{(\lambda - \kappa)\rho_u + \kappa}{\kappa^2 + \lambda(1 - \beta\rho_u)} \\ &= \frac{\kappa(1 - \rho_u) + \lambda\rho_u}{\kappa^2 + \lambda(1 - \beta\rho_u)} \equiv \omega \end{aligned}$$

Therefore,

$$i_t^\star = \rho + \omega u_t - G^\star s_t - \frac{1}{2}\sigma_z^2$$

In deviations from the steady state, it boils down to 18 (In the main text, $i^\star$ is referred to as i^{RC}). □

Appendix E Optimal policy with interest rate inertia

Proof of Proposition 2. Work in deviations from the deterministic steady state and write the implementability object as

$$\hat{\Xi}_t = \mathbb{E}_t[\hat{x}_{t+1} + \hat{\pi}_{t+1}] - \hat{\Psi}_t.$$

As in the main text, define

$$\Omega_F \equiv \kappa^2 + \lambda, \qquad A \equiv \frac{\delta_i}{\Omega_F + \delta_i}, \qquad B \equiv \frac{\kappa}{\Omega_F + \delta_i}, \qquad M \equiv \frac{\Omega_F}{\Omega_F + \delta_i} = 1 - A.$$

Step 1: Conditional policy problem. Let

$$m_t \equiv u_t + \beta \mathbb{E}_t[\pi_{t+1}].$$

Using the IS relation $i_t = \Xi_t - x_t$ and the NKPC $\pi_t = \kappa x_t + m_t$, the period loss is

$$(\kappa x_t + m_t)^2 + \lambda x_t^2 + \delta_i(\Xi_t - i_{t-1} - x_t)^2.$$

The first-order condition implies

$$x_t = A(\Xi_t - i_{t-1}) - B\, m_t,$$

hence

$$i_t = M\Xi_t + A i_{t-1} + B\, m_t, \qquad \pi_t = \kappa x_t + m_t.$$

Step 2: Affine equilibrium. Conjecture

$$\hat{\Xi}_t = \xi_i \hat{i}_{t-1} + \xi_u u_t + \xi_s \hat{s}_t, \qquad \hat{i}_t = q_i \hat{i}_{t-1} + q_u u_t + q_s \hat{s}_t,$$

$$\hat{x}_t = x_i \hat{i}_{t-1} + x_u u_t + x_s \hat{s}_t, \qquad \hat{\pi}_t = \pi_i \hat{i}_{t-1} + \pi_u u_t + \pi_s \hat{s}_t.$$

Since

$$\mathbb{E}_t[\hat{\pi}_{t+1}] = \pi_i \hat{i}_t + \rho_u \pi_u u_t + \rho_s \pi_s \hat{s}_t,$$

we have

$$m_t = \beta \pi_i \hat{i}_t + (1 + \beta\rho_u\pi_u)u_t + \beta\rho_s\pi_s\hat{s}_t.$$

Substituting into the policy rule yields

$$q_i = \frac{M\xi_i + A}{1 - B\beta\pi_i}, \qquad q_u = \frac{M\xi_u + B(1 + \beta\rho_u\pi_u)}{1 - B\beta\pi_i}, \qquad q_s = \frac{M\xi_s + B\beta\rho_s\pi_s}{1 - B\beta\pi_i}. \tag{59}$$

Because $x_t = \Xi_t - i_t$,

$$x_i = \xi_i - q_i, \qquad x_u = \xi_u - q_u, \qquad x_s = \xi_s - q_s,$$

and the NKPC implies

$$\pi_i = \kappa x_i + \beta \pi_i q_i, \qquad \pi_u = \kappa x_u + 1 + \beta \pi_i q_u + \beta \rho_u \pi_u, \qquad \pi_s = \kappa x_s + \beta \pi_i q_s + \beta \rho_s \pi_s.$$

Step 3: The common $(\hat{\imath}_{t-1}, u_t)$-block. The wedge $\hat{\Psi}_t$ loads only on $\hat{s}_t$. Therefore the coefficient-matching equations for the $(\hat{\imath}_{t-1}, u_t)$-coefficients are identical in the NK economy ($\hat{\Psi}_t \equiv 0$) and in the RC economy. Hence the common block $(x_i, x_u, \pi_i, \pi_u, q_i, q_u)$ is the same in both economies.

For completeness, the stable affine equilibrium selects the unique economically relevant root $q_i \in (0,1)$ of

$$\beta \delta_i q_i^3 - \big[\beta\lambda + \delta_i(1 + 2\beta + \kappa)\big] q_i^2 + \big[\kappa^2 + \lambda + \delta_i(2 + \beta + \kappa)\big] q_i - \delta_i = 0.$$

Since $P(0) = -\delta_i < 0$ and $P(1) = \kappa^2 + (1-\beta)\lambda > 0$, there is a root in $(0,1)$, and the stable-root restriction selects that root. Using the identities above,

$$\pi_i = -\frac{\delta_i \kappa (1 - q_i)}{\kappa^2 + \lambda(1 - \beta q_i)} < 0, \qquad L_i \equiv x_i + \pi_i = -\frac{\delta_i (1 - q_i)\big(1 - \beta q_i + \kappa\big)}{\kappa^2 + \lambda(1 - \beta q_i)} < 0,$$

Step 4: Proof that $\psi_s > 0$. Because i_t is known at time t, the only innovations from t to $t+1$ come from η_{t+1} and ν_{t+1}. Using the affine conjecture and the shock laws,

$$(\hat{x}_{t+1} + \hat{\pi}_{t+1}) - \mathbb{E}_t[\hat{x}_{t+1} + \hat{\pi}_{t+1}] = (L_u + \iota L_s)\eta_{t+1} + L_s \nu_{t+1},$$

where

$$L_u \equiv x_u + \pi_u, \qquad L_s \equiv x_s + \pi_s.$$

Hence

$$\mathrm{Var}_t(\hat{x}_{t+1} + \hat{\pi}_{t+1}) = (L_u + \iota L_s)^2 s_t + L_s^2 \omega_\nu^2.$$

The second term is constant and therefore disappears in deviations from the deterministic steady state. Thus the wedge is affine in $\hat{s}_t$:

$$\hat{\Psi}_t = \frac{1}{2}\big(L_u + \iota L_s\big)^2 \hat{s}_t.$$

So

$$\psi_s = \frac{1}{2}\big(L_u + \iota L_s\big)^2 \geq 0,$$

and under the nondegeneracy condition $L_u + \iota L_s \neq 0$, we have $\psi_s > 0$.

Step 5: Volatility block. Let

$$\mathcal{D}_s \equiv \Omega_F + \delta_i - \beta\kappa\pi_i - \beta\rho_s(\delta_i + \lambda).$$

From (59) and the identities $x_s = \xi_s - q_s$, $\pi_s = \kappa x_s + \beta \pi_i q_s + \beta \rho_s \pi_s$, one obtains

$$q_s = \Gamma_s \xi_s, \qquad x_s = X_s \xi_s, \qquad \pi_s = P_s \xi_s,$$

where

$$\Gamma_s \equiv \frac{\kappa^2 + \lambda(1-\beta\rho_s)}{\mathcal{D}_s}, \qquad X_s \equiv \frac{\delta_i(1-\beta\rho_s) - \beta\kappa\pi_i}{\mathcal{D}_s}, \qquad P_s \equiv \frac{\delta_i \kappa + \beta\lambda\pi_i}{\mathcal{D}_s}.$$

Letting $L_s = x_s + \pi_s$, the fixed-point equation for Ξ_t gives

$$\xi_s = L_i q_s + \rho_s L_s - \psi_s.$$

Substituting the formulas above,

$$\xi_s = L_i \Gamma_s \xi_s + \rho_s (X_s + P_s)\xi_s - \psi_s,$$

so

$$\big[1 - L_i \Gamma_s - \rho_s(X_s + P_s)\big]\xi_s = -\psi_s.$$

Multiplying by $\mathcal{D}_s$, this becomes

$$\Delta_s \, \xi_s = -\mathcal{D}_s \, \psi_s,$$

where, after simplification,

$$\Delta_s = (1 - L_i)\big[\kappa^2 + \lambda(1-\beta\rho_s)\big] - \beta\big[\kappa(1-\rho_s) + \lambda\rho_s\big]\pi_i + \delta_i\big[1 - \rho_s(1+\beta+\kappa) + \beta\rho_s^2\big].$$

Therefore

$$\xi_s = -\frac{\mathcal{D}_s}{\Delta_s}\, \psi_s,$$

and

$$q_s = -\frac{\kappa^2 + \lambda(1-\beta\rho_s)}{\Delta_s}\, \psi_s, \qquad x_s = -\frac{\delta_i(1-\beta\rho_s) - \beta\kappa\pi_i}{\Delta_s}\, \psi_s, \qquad \pi_s = -\frac{\delta_i\kappa + \beta\lambda\pi_i}{\Delta_s}\, \psi_s.$$

Step 6: Signs and comparison. Assume $\Delta_s > 0$ and $L_u + \iota L_s \neq 0$, so $\psi_s > 0$.

First,

$$\kappa^2 + \lambda(1-\beta\rho_s) > 0,$$

hence $q_s < 0$.

Second, since $\pi_i < 0$,

$$\delta_i(1-\beta\rho_s) - \beta\kappa\pi_i > 0,$$

so $x_s < 0$.

Third,

$$\delta_i\kappa + \beta\lambda\pi_i = \delta_i\kappa\left[1 - \frac{\beta\lambda(1-q_i)}{\kappa^2+\lambda(1-\beta q_i)}\right]$$
$$= \delta_i\kappa\,\frac{\kappa^2+\lambda(1-\beta)}{\kappa^2+\lambda(1-\beta q_i)} > 0,$$

so $\pi_s < 0$.

Because the common $(\hat{i}_{t-1}, u_t)$-block is identical across economies, equations (29)–(31) follow immediately. The state-by-state inequalities then follow from $x_s < 0$, $\pi_s < 0$, and $q_s < 0$ when $\hat{s}_t \geq 0$. This same sign result is what underlies the discussion in the text that, with policy inertia, stochastic volatility affects allocations, whereas under frictionless discretion it only affects implementation. □

Appendix F Taylor-rule equilibrium

Under the Taylor rule $i_t = \phi\pi_t$, conjecture the affine equilibrium in deviations from the deterministic steady state:

$$\hat{x}_t^{\mathcal{T}} = x_u\,u_t + x_s\,\hat{s}_t, \qquad \hat{\pi}_t^{\mathcal{T}} = \pi_u\,u_t + \pi_s\,\hat{s}_t.$$

Because firms satisfy FIRE, the Phillips curve under the Taylor rule is the forward-looking NKPC

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

Substituting the affine conjecture into this equation gives

$$\pi_u = \kappa x_u + \beta\rho_u\pi_u + 1, \qquad \pi_s = \kappa x_s + \beta\rho_s\pi_s$$

Hence

$$\pi_u = \frac{\kappa x_u + 1}{1-\beta\rho_u}, \qquad \pi_s = \frac{\kappa}{1-\beta\rho_s}x_s \equiv \alpha_s x_s, \qquad \alpha_s \equiv \frac{\kappa}{1-\beta\rho_s}.$$

Substituting the affine conjecture into the IS curve under the Taylor rule and collecting terms in u_t yields

$$x_u = \rho_u x_u - \phi\pi_u + \rho_u\pi_u$$

so that

$$x_u(1-\rho_u) = (\rho_u - \phi)\pi_u = (\rho_u-\phi)\frac{\kappa x_u + 1}{1-\beta\rho_u}$$

Solving,

$$x_u = \frac{\rho_u - \phi}{D_u}, \qquad \pi_u = \frac{1-\rho_u}{D_u}, \qquad D_u \equiv (1-\beta\rho_u)(1-\rho_u) + \kappa(\phi-\rho_u)$$

Collecting terms in $\hat{s}_t$ yields

$$x_s = \rho_s x_s - \phi\pi_s + \rho_s\pi_s - \psi_s$$

where

$$\psi_s \equiv \frac{1}{2}\big[(x_u + \iota x_s) + (\pi_u + \iota \pi_s)\big]^2$$

Using $\pi_s = \alpha_s x_s$, this becomes

$$0 = \frac{1}{2}\Big[(x_u + \pi_u) + \iota(1+\alpha_s)x_s\Big]^2 + \frac{D_s}{1-\beta\rho_s}x_s$$

where

$$D_s \equiv (1-\beta\rho_s)(1-\rho_s) + \kappa(\phi - \rho_s)$$

Expanding gives (**??**), with coefficients

$$g_0 = \frac{1}{2}(x_u+\pi_u)^2, \qquad g_1 = \iota(1+\alpha_s)(x_u+\pi_u), \qquad g_2 = \frac{1}{2}\,\iota^2(1+\alpha_s)^2$$

The risk-adjustment term is computed from the conditional second moments of r_{t+1} and π_{t+1}, which depend on η_{t+1} and ν_{t+1}:

$$\begin{aligned}
\mathrm{Var}_t(r_{t+1}) &= (x_u + \iota x_s)^2\, s_t + \omega_\nu^2 x_s^2 \\
\mathrm{Var}_t(\pi_{t+1}) &= (\pi_u + \iota \pi_s)^2\, s_t + \omega_\nu^2 \pi_s^2 \\
\mathrm{Cov}_t(\pi_{t+1}, r_{t+1}) &= (x_u + \iota x_s)(\pi_u + \iota\pi_s)\, s_t + \omega_\nu^2 x_s \pi_s
\end{aligned}$$

Therefore,

$$\begin{aligned}
\Psi_t &= \frac{1}{2}\mathrm{Var}_t(r_{t+1}) + \frac{1}{2}\mathrm{Var}_t(\pi_{t+1}) + \mathrm{Cov}_t(\pi_{t+1}, r_{t+1}) \\
&= \frac{1}{2}\big[(x_u + \iota x_s) + (\pi_u + \iota\pi_s)\big]^2 s_t + \frac{1}{2}\,\omega_\nu^2(x_s+\pi_s)^2
\end{aligned}$$

which establishes (34).