Trade Policy Uncertainty and Optimal Monetary Policy*

Lu Han and Yuko Imura
Bank of Canada

1 May 2026
Preliminary and Incomplete

Abstract

Trade policy uncertainty has become a salient macroeconomic risk, yet its transmission in open economies remains not fully understood analytically. In models with stochastic volatility, uncertainty shocks are higher-order objects that vanish from linearized approximations and are therefore typically studied only numerically. This paper provides an analytical characterization in a small open economy New Keynesian model with export-tariff uncertainty. Under complete markets, the equilibrium collapses to a tractable two-equation nonlinear system in producer-price inflation and the terms of trade, delivering closed-form third-order impulse responses. The solution reveals two opposing Jensen wedges—a New Keynesian Phillips curve (pricing) wedge and an asset-market (inverse-return) wedge—whose net effect on output depends on trade openness and trade elasticity and may change qualitatively depending on price stickiness and the monetary policy rule. Under a mild common-persistence condition, the full uncertainty response admits an exact decomposition into two standard first-order shocks—productivity and risk premium—with time-invariant weights, bridging higher-order and linear analyses. Under the Ramsey-optimal policy, strict stabilization of producer-price inflation is optimal and the divine coincidence extends to tariff uncertainty: allocations are invariant to the volatility state up to third order.

JEL classification: E32, E52, F41
Keywords: trade policy uncertainty, small open economy, stochastic volatility, third-order perturbation, Ramsey policy, Taylor rule

*We thank Gianluca Benigno, Giancarlo Corsetti, Erwan Gautier, Daniela Hauser, James MacGee, Ben Tomlin, and Edouard Tsague-Djeutem for their invaluable suggestions and feedback. The views expressed herein are those of the authors and not necessarily those of the Bank of Canada.

1 Introduction

Trade policy has become a prominent source of macroeconomic risk. In 2025–2026, the US administration introduced a series of import tariffs on its trading partners, as well as revisions to the Section 301 enforcement and exclusion lists for its trade with China.[1] This has raised uncertainty around future US trade policy and its potential impacts on US trading partner economies, with trade flows and investment decisions showing early signs of adjustments. These observations highlight that agents respond not only to the *realized* tariffs but also to *trade policy uncertainty*–changes in the perceived likelihood and magnitude of future tariff actions– even when current tariffs may be unchanged.

For trade-dependent economies such as Canada, uncertainty can also be compounded by institutional "cliff" risks around market access. The United States-Mexico-Canada Agreement (USMCA) is scheduled to undergo a joint review starting on July 1, 2026, and the governments have already initiated consultation processes ahead of the start of the review.[2] These actions and the impending trade negotiations can shift the perceived distribution of future tariffs and market access long before any changes to tariffs are announced or implemented.

This raises important questions critical for monetary policy making. When the perceived distribution of future tariffs widens without any contemporaneous change in actual tariffs, how does it affect current and future production and inflation, and through which channels—marginal cost, risk premia, and external relative prices—does the disturbance propagate? And what benchmark does optimal monetary policy set in an open economy with tariffs and time-varying risk?

Answering these questions is difficult with standard linear tools because tariff uncertainty is a higher-moment disturbance, and hence is not captured by first-order log-linear approximations. With time-varying exogenous risk, a volatility state affects equilibrium only through nonlinear expectations and conditional moments (Bloom, 2009; Fernández-Villaverde, Guerrón-Quintana, Rubio-Ramírez and Uribe, 2011; Benigno, Benigno and Nisticò, 2013). In our open economy environment, the key forward-looking restrictions are the New Keynesian Phillips curve (NKPC) and an Euler-consistent asset-pricing condition whose first-order counterpart is the familiar UIP relation. A change in tariff risk therefore shifts equilibrium through Jensen terms, and its transmission depends crucially on monetary policy.

This paper studies optimal monetary policy in the presence of tariff uncertainty using a tractable small open-economy New Keynesian model with tariffs that follow a stochastic-volatility process. We focus on uncertainty about foreign tariffs on the country's exports, and consider a volatility innovation that raises the conditional variance of future export tariffs while holding the tariff level

[1] See, for example, White House (2025*a*,*b*, 2026); USTR (2025*a*).

[2] See USTR (2025*b*); Villarreal, Fergusson and Meltzer (2026); Global Affairs Canada (2025).

fixed.[3] We solve the model using a third-order perturbation[4] and contrast complete international asset markets (complete risk sharing) with incomplete markets.

Our analysis delivers three key results. First, exploiting the fact that the complete-markets equilibrium collapses exactly to a two-equation nonlinear system in producer-price inflation (PPI) and the terms of trade (Lemma 1), we derive closed-form third-order responses of the economy to tariff uncertainty under Taylor rules (Proposition 1), and characterize how the sign and magnitude of the responses depend on trade openness, trade elasticity, nominal rigidity, and the Taylor-rule parameters. The closed-form solution reveals that a tariff-volatility shock operates through exactly two Jensen wedges: a New Keynesian Phillips curve (NKPC) pricing wedge and an asset-market wedge induced by the Euler-consistent inverse-return pricing object. These wedges are not sign-restricted in general. In our baseline calibration, the NKPC wedge is deflationary and, under a Taylor rule, induces monetary easing and a real depreciation that raises exports and output. On the other hand, the asset-market wedge pushes in the opposite direction, damping the real depreciation and partially offsetting the real expansion (Corollary 1). Because the two wedges enter the 2×2 system with different policy loadings, stronger inflation feedback, lower trade openness, or more flexible prices can reverse the sign of the output response (Figures 5–6).

Second, we establish a *two-channel spanning* result: under a mild common-persistence restriction, the third-order impulse response to tariff uncertainty is spanned by a constant-weight linear combination of two classical first-order shocks—a productivity shock and a UIP-wedge (risk-premium) shock—up to the $O(4)$ remainder (Proposition 2). A single first-order shock cannot match the joint sign pattern of the uncertainty response in general; two shocks are both necessary and sufficient in our setting. While the result relies on the tractable two-equation structure of the complete-markets model and the common-persistence assumption, it suggests a useful approach for policy analysis with first-order models: one can approximate the effects of a tariff-volatility shock by simulating two basis shocks and combining the resulting impulse responses with closed-form weights. In our baseline calibration, a one-standard-deviation tariff-volatility shock maps into a simultaneous 0.33% productivity shock and a 0.55% UIP-wedge shock, confirming that both the supply and external-finance channels carry quantitatively meaningful weights (Figure 7).

Third, we derive the Ramsey benchmark under stochastic volatility using a Quadratic–Cubic (QC) approximation. Under complete markets and producer-currency pricing, Ramsey policy implies strict PPI targeting, and divine coincidence extends to tariff uncertainty shocks: the Ramsey allocation is invariant to the volatility state up to third order (Proposition 3). This result has

[3] It is useful to distinguish two related disturbances. A tariff *level* shock is a surprise change in the current tariff that moves relative prices directly and can be captured by standard first-order (log-linear) analysis. A tariff *uncertainty* shock is a change in the conditional variance of future tariffs that leaves the tariff level unchanged on impact. Such shocks matter only through nonlinear expectations (Jensen terms) in forward-looking equilibrium conditions. In the model, tariff uncertainty takes the form of stochastic volatility in the export-tariff process: the uncertainty shock is an innovation to the volatility state $\sigma_{T,t}$.

[4] Our approach is related to Benigno, Benigno and Nisticò (2013), who develop approximation methods for models with time-varying exogenous risk. Under our log-volatility parameterization and perturbation scaling, a volatility innovation changes equilibrium only through conditional variances and has leading effects at third order (Lemma 2). Online Appendix C provides the coefficient-matching details.

a sharp policy implication: the sizable real effects of tariff uncertainty observed under Taylor rules are *policy-induced*, not a fundamental Ramsey trade-off. Under incomplete international asset markets, however, the exact bond-economy baseline activates an additional Euler/Jensen (precautionary-savings) wedge, and strict PPI targeting is no longer Ramsey optimal in general. In this case, uncertainty shocks can have real effects even under inflation stabilization (Online Appendix G). The incomplete-markets extension thus points to a genuine stabilization trade-off when international risk sharing is imperfect, and the two-channel spanning generalizes to three channels by adding a discount-factor (precautionary-savings) shock.

These theoretical results discipline our quantitative analysis and clarify the magnitudes of the impacts. In our baseline calibration, a one-standard-deviation tariff-volatility shock raises output by about 0.07% on impact under a PPI-based Taylor rule, alongside PPI deflation of about 0.99 percentage points and a 1.45 percentage-point cut in the policy rate (Figure 9). Under strict PPI stabilization, the same shock is neutral for real allocations in the complete-markets economy, confirming the Ramsey benchmark. The NKPC wedge contributes about +0.26% to impact output while the asset-market/Jensen wedge contributes about −0.19%, so the net expansion is materially smaller than the pricing channel alone (Figure 10). The closed-form solution also delivers sharp comparative statics across trade openness, nominal rigidity, and Taylor-rule feedback (Figures 5–6).

This paper contributes to several strands of the literature. First, it speaks to recent open-economy New Keynesian analyses of tariffs and trade wars (Bergin and Corsetti, 2023; Auray, Devereux and Eyquem, 2025; Monacelli, 2025), in which monetary policy shapes whether tariffs are contractionary or expansionary and governs the inflation–output trade-off. These studies focus on tariff *level* shocks and typically rely on numerical solutions; we complement them by studying tariff *uncertainty*—changes in the perceived distribution of future tariffs that leave the current tariff unchanged—and by providing closed-form characterizations of the transmission channels and their interaction with the policy rule. Our results highlight that expectations about future tariffs alone can influence equilibrium allocations, giving rise to a distinct role for monetary policy.

Second, it connects to the empirical trade-uncertainty literature. Handley and Limão (2017, 2022) show that trade policy uncertainty depresses investment and trade flows even absent actual tariff changes; Caldara, Iacoviello, Molligo, Prestipino and Raffo (2020) construct a comprehensive measure of trade policy uncertainty and document its adverse macroeconomic effects; and KalemliÖzcan, Soylu and Yildirim (2026) highlight an external-finance channel through which tariff uncertainty moves exchange rates via risk premia. Our two-wedge decomposition provides a structural counterpart to these empirical findings: the NKPC wedge captures the supply-side channel emphasized in the investment and trade literature, while the asset-market/Jensen wedge captures the external-finance and risk-premium channel, and the monetary policy rule governs their relative strength.

Third, it contributes to the broader literature on macroeconomic uncertainty and higher-order perturbation methods (Bloom, 2009; Fernández-Villaverde, Guerrón-Quintana, Rubio-Ramírez and Uribe, 2011; Benigno, Benigno and Nisticò, 2012, 2013; Basu and Bundick, 2017). Much of this work studies time-varying risk numerically; Basu and Bundick (2017) study demand uncertainty

in a closed-economy New Keynesian model and emphasize precautionary pricing. We contribute analytically by showing that—under complete international asset markets—the open-economy equilibrium reduces exactly to a two-equation nonlinear system, making closed-form third-order impulse responses and the constant-weight spanning representation tractable. This tractability extends to optimal policy: building on the open-economy optimal policy literature (Benigno and Benigno, 2003, 2008; Corsetti, Dedola and Leduc, 2010) and on QC approximations to nonlinear optimal policy (Gross and Hansen, 2021), our Ramsey analysis establishes that divine coincidence—a property known for first-order disturbances in the canonical SOE model of Galí and Monacelli (2005) and for closed-economy uncertainty shocks by Cho, Han, Oh and Rogantini Picco (2021)—carries over to tariff uncertainty in the open economy, a purely higher-order shock.

The remainder of the paper is organized as follows. Section 2 introduces the small open economy New Keynesian model, defines tariff uncertainty shocks (Subsection 2.8), and contrasts complete and incomplete international asset markets. Section 3 presents the closed-form characterization of uncertainty responses under a Taylor rule. Section 4 proves the two-channel spanning result. Section 5 derives the Ramsey benchmark and the associated divine coincidence result for uncertainty shocks, and provides quantitative illustrations and policy-rule comparisons. The online appendices collect the full nonlinear equilibrium system, the minimal nonlinear representation, and the proofs.

2 Model

Time is discrete. Our model economy builds on the canonical SOE New Keynesian framework (Galí and Monacelli, 2005), and is augmented with tariffs and stochastic volatility. Home is a small open economy and takes Foreign variables as exogenous. We assume that Home firms set their prices in the Home currency (producer-currency pricing, PCP), so exchange-rate movements are fully passed through to export prices in Foreign currency. Prices are set under Rotemberg adjustment costs (Rotemberg, 1982). Home exporters face a Foreign tariff T_t^* on their exports, while Home levies an import tariff $T_{F,t}$ on its imports from Foreign. Monetary policy sets a nominal interest rate following a Taylor rule in gross levels.

For orientation, it is convenient to group the equilibrium conditions into five blocks (Online Appendix A records the full list in levels). The first block describes preferences and household optimality—utility (4), labor supply (5), and the Euler equation (6) (the Euler equation is redundant under complete markets but becomes essential under incomplete markets). The second block describes production and price setting—technology (7) and Rotemberg pricing, summarized by the nonlinear NKPC (8). The third block describes open-economy demand and goods-market clearing—the CES aggregator (2), export demand (3), the CPI index (9), and the resource constraint (10). The fourth block is the international asset-market closure—the Euler-consistent asset-pricing relation (11) together with either complete-markets risk sharing (12) or, under incomplete markets, the no-risk-sharing closure. The fifth block is monetary policy—either the Taylor rule (14) or the Ramsey problem studied in Section 5.

2.1 Key objects and notations

Let $P_{H,t}$ be the Home producer price index for domestically produced goods and let P_t be the Home CPI. Define

$$G_t \equiv \frac{P_t}{P_{H,t}}, \qquad \Pi_{H,t} \equiv \frac{P_{H,t}}{P_{H,t-1}}, \qquad \Pi_{C,t} \equiv \frac{P_t}{P_{t-1}}.$$

The wedge $G_t = P_t/P_{H,t}$ is the CPI–PPI wedge: it moves when import prices (and therefore the exchange rate, foreign prices, and import tariffs) move, and it is the object that makes CPI inflation differ from PPI inflation. $\Pi_{H,t}$ denotes PPI inflation, while $\Pi_{C,t}$ denotes CPI inflation. Let $\mathcal{E}_t$ be the nominal exchange rate (Home currency per unit of Foreign currency) and let P_t^* denote the Foreign CPI. The (world) terms of trade is

$$Q_t \equiv \frac{\mathcal{E}_t P_t^*}{P_{H,t}}. \tag{1}$$

In our normalization with constant Foreign prices, a rise in Q_t is a real depreciation for Home: Home goods become cheaper relative to Foreign consumption, making Home exports more competitive. Home consumption is a CES aggregate of Home and Foreign goods with import share $\alpha \in (0,1)$ and trade elasticity $\gamma > 0$:

$$C_t = \left[(1-\alpha)^{1/\gamma} C_{H,t}^{(\gamma-1)/\gamma} + \alpha^{1/\gamma} C_{F,t}^{(\gamma-1)/\gamma}\right]^{\gamma/(\gamma-1)}. \tag{2}$$

The foreign demand scale is set equal to the Home import share α for parsimony. Foreign demand for Home goods is isoelastic with the same elasticity and is scaled by α:

$$C_{H,t}^* = \alpha \left(\frac{T_t^*}{Q_t}\right)^{-\gamma} C_t^*. \tag{3}$$

2.2 Households

The representative Home household has CRRA utility

$$U(C_t, L_t) = \frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{L_t^{1+\varphi}}{1+\varphi}, \qquad \sigma > 0,\ \varphi \geq 0, \tag{4}$$

where σ controls intertemporal curvature (risk aversion / inverse intertemporal elasticity of substitution), φ controls how elastic labor supply is, and χ scales the disutility of work. The household faces the nominal budget constraint

$$P_t C_t + B_t = W_t L_t + R_{t-1} B_{t-1} + \mathcal{P}_t + T_t^{gov},$$

where B_t is a one-period nominal bond, $\mathcal{P}_t$ is firm profits (dividends), and T_t^{gov} is a lump-sum rebate of tariff revenues. Optimality delivers the labor supply condition and Euler equation

$$\begin{aligned} \lambda_t &\equiv C_t^{-\sigma}, \\ \frac{W_t}{P_t} &= \chi L_t^{\varphi} C_t^{\sigma}, \qquad (5) \\ 1 &= \beta R_t \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \right]. \qquad (6) \end{aligned}$$

2.3 Firms, production, and price setting

Home firms are monopolistically competitive and produce output using labor:

$$Y_t = A_t L_t, \qquad (7)$$

where A_t is aggregate productivity. Firms face a Dixit–Stiglitz demand aggregator with elasticity $\epsilon > 1$ (Dixit and Stiglitz, 1977) and adjust prices subject to Rotemberg quadratic adjustment costs (Rotemberg, 1982). We impose the optimal production subsidy $(1 - \tau_s) = (\epsilon - 1)/\epsilon$, so deterministic steady-state marginal cost equals one. The Rotemberg optimality condition can then be written as the nonlinear NKPC in PPI inflation:

$$\kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) + \beta\kappa \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \frac{\lambda_{t+1}}{\lambda_t} \frac{G_t}{G_{t+1}} \right], \qquad (8)$$

where $mc_t \equiv MC_t / P_{H,t}$ is real marginal cost in PPI units, and $\lambda_t = C_t^{-\sigma}$.

2.4 Price indices, goods market clearing, and tariffs

The CES demand system implies (i) the CPI price index and (ii) Home and Foreign demands for each good. With an import tariff $T_{F,t}$ applied to Foreign goods at the border, the CPI is

$$P_t = \left[(1 - \alpha) P_{H,t}^{1-\gamma} + \alpha \left(T_{F,t} \mathcal{E}_t P_{F,t}^* \right)^{1-\gamma} \right]^{1/(1-\gamma)}. \qquad (9)$$

Define $P_{F,t}^w \equiv \mathcal{E}_t P_{F,t}^*$ as the Home-currency price of Foreign goods. In the one-Foreign-good SOE environment used throughout, $P_t^* = P_{F,t}^*$ and the terms of trade definition (1) can equivalently be written $Q_t = P_{F,t}^w / P_{H,t}$. Under the SOE limit, we take Foreign inflation as exogenous and focus on cases where $P_{F,t}^*$ is constant; therefore, the ToT dynamics are driven by $(\mathcal{E}_t, \Pi_{H,t})$.

Goods market clearing for Home output is

$$Y_t = C_{H,t} + C_{H,t}^* + \frac{\kappa}{2} (\Pi_{H,t} - 1)^2 Y_t. \qquad (10)$$

2.5 International asset markets: complete vs incomplete

The international asset-market block combines an asset-pricing relation with a closure that determines the degree of international risk sharing. Let R_t and R_t^* denote gross one-period nominal interest rates at Home and Foreign. Under complete markets, the exact nonlinear closure used below is the Euler-consistent asset-pricing condition

$$1 = \frac{R_t}{R_t^*} e^{\psi_t} \mathbb{E}_t \left[\frac{\mathcal{E}_t}{\mathcal{E}_{t+1}} \right], \tag{11}$$

which, under constant Foreign prices, is equivalent to

$$1 = \beta R_t e^{\psi_t} \mathbb{E}_t \left[\frac{Q_t}{Q_{t+1} \Pi_{H,t+1}} \right].$$

The wedge ψ_t is normalized so that, at first order, it shifts the familiar log UIP restriction $r_t - r_t^* + \psi_t = \mathbb{E}_t[\Delta e_{t+1}]$. This normalization is useful for the decomposition result in Section 4, where a one-standard-deviation innovation to ψ_t serves as a familiar first-order "basis shock." Later, tariff uncertainty will generate an additional *endogenous* asset-market/Jensen correction under the Euler-consistent closure; to avoid confusion with the reduced-form ψ_t, we refer to that higher-order object as the asset-market uncertainty wedge.

Under complete international asset markets, international risk sharing holds:

$$\frac{C_t^{-\sigma}}{(C_t^*)^{-\sigma}} = \vartheta \cdot \frac{P_t}{\mathcal{E}_t P_t^*}, \tag{12}$$

for a constant $\vartheta > 0$. In the baseline analysis we normalize (C_t^*, P_t^*) to constants, so risk sharing pins Home consumption to the real exchange rate:

$$C_t = \left(\frac{G_t}{Q_t} \right)^{-1/\sigma}. \tag{13}$$

2.6 Monetary policy

We consider two monetary policy regimes. Under the Taylor rule, the nominal interest rate responds to PPI inflation and output:

$$\frac{R_t}{\bar{R}} = \Pi_{H,t}^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y}, \tag{14}$$

where $\bar{R}$ and $\bar{Y}$ denote deterministic steady-state values (in our stationary benchmark, $\bar{R} = 1/\beta$), and (ϕ_π, ϕ_y) satisfy the Taylor principle (roughly, $\phi_\pi > 1$) in our baseline. Under Ramsey policy, a planner chooses paths for allocations and inflation to maximize expected welfare subject to equilibrium constraints. Section 5 shows that the Ramsey allocation coincides with strict PPI targeting under complete markets.

2.7 Equilibrium

Given the exogenous processes $\{T_{F,t}, T_t^*, A_t, \psi_t\}$ (and exogenous Foreign variables in the SOE limit), an equilibrium is a collection of stochastic processes for allocations and prices that satisfy: (i) household optimality (labor supply and the Euler equation), (ii) the nonlinear Rotemberg NKPC for PPI inflation, (iii) goods market clearing, (iv) the international asset-market block (complete markets: Euler-consistent asset pricing plus risk sharing; incomplete markets: no risk sharing and, in quantitative work, a bond-economy closure), and (v) the monetary policy rule. Online Appendix A records the full nonlinear system in levels.

Under complete markets, risk sharing collapses the relevant intertemporal term in the NKPC and pins consumption to the real exchange rate, allowing a reduction to a minimal 2×2 nonlinear system in $(\Pi_{H,t}, Q_t)$ (Lemma 1). Under incomplete markets, the Euler equation becomes a non-redundant restriction and the minimal representation expands to three dynamic equations (Online Appendix G).

Lemma 1 (Exact minimal nonlinear system under complete markets)**.** *Maintain producer-currency pricing (PCP) and assume complete international asset markets with constant Foreign objects* (C_t^*, P_t^*, R_t^*) *and normalization* $R_t^* = \bar{R}$*. Then the complete-markets equilibrium is pinned down by the following* exact *nonlinear* 2×2 *system in* $(\Pi_{H,t}, Q_t)$*:*

$$\begin{aligned}
\kappa(\Pi_{H,t} - 1)\Pi_{H,t} &= (\epsilon - 1)(mc_t - 1) + \beta\kappa\mathbb{E}_t\left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1}\frac{Y_{t+1}}{Y_t}\frac{Q_t}{Q_{t+1}}\right], \\
1 &= \Pi_{H,t}^{\phi_\pi}\left(\frac{Y_t}{\bar{Y}}\right)^{\phi_y} e^{\psi_t}\mathbb{E}_t\left[\frac{Q_t}{Q_{t+1}\Pi_{H,t+1}}\right],
\end{aligned} \tag{15}$$

where output and real marginal cost are static functions of $(\Pi_{H,t}, Q_t)$ *and exogenous states:*

$$\begin{aligned}
\Phi_t &\equiv (1-\alpha) + \alpha\,(T_{F,t}Q_t)^{1-\gamma}, \qquad G_t \equiv \frac{P_t}{P_{H,t}} = \Phi_t^{1/(1-\gamma)}, \\
C_t &= \left(\frac{G_t}{Q_t}\right)^{-1/\sigma}, \\
C_{H,t} &= (1-\alpha)\left(\frac{P_{H,t}}{P_t}\right)^{-\gamma} C_t = (1-\alpha)G_t^{\gamma} C_t, \\
X_t \equiv C_{H,t}^* &= \alpha\left(\frac{T_t^*}{Q_t}\right)^{-\gamma} C_t^*, \\
Y_t &= \frac{C_{H,t} + X_t}{1 - \frac{\kappa}{2}(\Pi_{H,t} - 1)^2}, \\
mc_t &= \chi\, Q_t\, Y_t^{\varphi} A_t^{-(1+\varphi)}.
\end{aligned}$$

The reduction in Lemma 1 is possible for two reasons. First, under complete markets, risk sharing eliminates an independent intertemporal margin for Home consumption: the stochastic discount factor is tied to the terms of trade, so the NKPC discounting term collapses to Q_t/Q_{t+1}.

Second, the Euler-consistent asset-pricing relation (11), combined with the Taylor rule and the definition of Q_t (under constant Foreign prices), delivers (15). Everything else—output, marginal cost, absorption, and exports—is then recovered from a static block as a function of $(\Pi_{H,t}, Q_t)$ and exogenous states. The only approximations in our uncertainty analysis enter when we apply a third-order perturbation to this nonlinear system. The detailed derivation is in Online Appendix B.

2.8 Tariff Uncertainty Shocks

The focus of our analysis is *export-tariff uncertainty*: uncertainty about the future foreign tariff faced by Home exporters. The uncertainty experiment is deliberately "clean." On impact, agents learn that future tariff realizations are more dispersed, but the tariff level remains unchanged. We model tariff uncertainty as stochastic volatility in the export tariff process. Let $\tau_t \equiv \log T_t^*$ denote the log export tariff and let $\sigma_{T,t}$ denote its volatility state. We assume

$$\tau_t = \rho_\tau \tau_{t-1} + \bar{\sigma}_\tau e^{\sigma_{T,t-1}} \varepsilon_{\tau,t}, \tag{16}$$

$$\sigma_{T,t} = \rho_\sigma \sigma_{T,t-1} + \eta_\sigma \varepsilon_{\sigma,t}, \tag{17}$$

with $\mathbb{E}[\varepsilon_{\tau,t}] = \mathbb{E}[\varepsilon_{\sigma,t}] = 0$ and unit variances. Productivity and the UIP wedge follow

$$\begin{aligned} a_t &\equiv \log A_t = \rho_A a_{t-1} + \sigma_A \varepsilon_{A,t}, \\ \psi_t &= \rho_\psi \psi_{t-1} + \sigma_\psi \varepsilon_{\psi,t}. \end{aligned} \tag{18}$$

Assumption 1 (Innovation regularity)**.** The innovations $\{\varepsilon_{\tau,t}, \varepsilon_{\sigma,t}, \varepsilon_{A,t}, \varepsilon_{\psi,t}\}_{t\geq 0}$ are mutually independent across shocks and i.i.d. over time with mean zero and unit variance. We also assume symmetry (e.g. Gaussian shocks), so that third-order Jensen terms depend on conditional variances but not on conditional skewness.

The key timing for our uncertainty experiment is that the volatility state $\sigma_{T,t}$ is known at time t and governs the conditional variance of the next tariff innovation:

$$\mathrm{Var}_t(\tau_{t+1}) = \bar{\sigma}_\tau^2 e^{2\sigma_{T,t}}. \tag{19}$$

Because the NKPC and UIP blocks are forward-looking, changes in (19) affect time-t allocations even when the tariff level is unchanged.

Definition 1 (Tariff-uncertainty (volatility) shock)**.** The *tariff uncertainty shock* is the experiment where $\varepsilon_{\sigma,0} = 1$, $\varepsilon_{\sigma,t} = 0$ for $t \geq 1$, and all other innovations are set to zero. Starting from $\sigma_{T,-1} = 0$, this implies $\sigma_{T,t} = \rho_\sigma^t \eta_\sigma$ for $t \geq 0$ while the tariff level stays at $\tau_t \equiv 0$.

Lemma 2 (Volatility enters through conditional variances)**.** *Under Assumption 1 and the tariff-volatility experiment in Definition 1, the volatility state $\sigma_{T,t}$ affects equilibrium (up to third order) only through conditional second moments of next-period tariff innovations, i.e. through objects proportional to* $\mathrm{Var}_t(\tau_{t+1}) = \bar{\sigma}_\tau^2 e^{2\sigma_{T,t}}$.

Sketch. Under the volatility experiment ($\varepsilon_{\tau,t} \equiv 0$), the volatility state appears only in the *scale* of the next tariff innovation in (16), through the factor $\bar{\sigma}_\tau e^{\sigma_{T,t}}$. Since the forward-looking blocks (UIP and the Rotemberg NKPC) take expectations of nonlinear functions of next-period objects, a third-order expansion replaces those nonlinear expectations by (i) the expectation of the underlying linear term, and (ii) a Jensen correction proportional to conditional variance. Under symmetric innovations (Assumption 1), third-order terms linked to conditional skewness drop out, so only conditional variances remain. □

We characterize uncertainty shocks using a third-order perturbation expansion around the deterministic steady state. Throughout, $O(k)$ denotes terms of order k or higher in that expansion. Under the volatility experiment, the shock affects equilibrium only through conditional variances and related Jensen terms, as in Benigno, Benigno and Nisticò (2013) and the broader uncertainty-shock literature (Basu and Bundick, 2017). This structure underlies both our closed-form third-order responses in Section 3 and the spanning result in Section 4.

3 Tariff Uncertainty under a Taylor Rule

This section derives closed-form third-order impulse responses to a tariff-volatility shock when monetary policy follows a Taylor rule. This greatly facilitates the illustration of transmission mechanisms: *which wedges move, which relative prices move, and when output rises versus falls.* For exposition, we treat the import tariff $T_{F,t}$ as constant and focus on uncertainty in the export tariff T_t^*. Making $T_{F,t}$ time-varying affects only the static block (and therefore the coefficient objects) but does not change the logic of the two-equation reduction in Lemma 1.

Under complete markets, Lemma 1 implies that the equilibrium under a Taylor rule is pinned down by producer-price inflation and the terms of trade; output, exports, and the policy rate are recovered from static relationships and the policy rule. To compute the economy's responses to tariff uncertainty, we take this *exact* two-equation system and expand it around the deterministic steady state. A pure volatility innovation leaves the tariff level unchanged and raises only the conditional variances. Thus, its leading effects are third order: volatility matters only through Jensen corrections in the nonlinear expectation terms of the NKPC and asset-pricing conditions. These corrections can be collected into two endogenous wedges, one in the pricing block (an "NKPC uncertainty wedge") and the other in the asset-market block (an "asset-market/Jensen wedge"). Online Appendix C provides the coefficient-matching details underlying the closed forms reported below.

3.1 Closed-form uncertainty coefficients

The goal is to obtain closed-form impulse responses to a tariff-volatility shock under the Taylor rule. The solution proceeds in three steps: (i) rewrite the complete-markets equilibrium as an exact nonlinear system in logs, (ii) expand that system around the deterministic steady state to identify

the perturbation coefficients needed at each order, and (iii) solve a 2×2 linear system for the uncertainty coefficients. We summarize below the key objects and the result, and report the full coefficient-matching algebra in Online Appendix C.

Step 1: Exact log representation. Let $\pi_t \equiv \log \Pi_{H,t}$, $q_t \equiv \log Q_t$, $y_t \equiv \log(Y_t/\bar{Y})$, and define

$$\kappa_{mc} \equiv \frac{\epsilon - 1}{\kappa}. \tag{20}$$

Under complete markets, risk sharing ties the stochastic discount factor to the terms of trade (see Lemma 1). In logs, the minimal system becomes

$$(e^{\pi_t} - 1)e^{\pi_t} = \kappa_{mc}\big(e^{q_t + \varphi y_t - (1+\varphi)a_t} - 1\big) + \beta \mathbb{E}_t[(e^{\pi_{t+1}} - 1)e^{\pi_{t+1}}\, \Omega_{t+1}], \tag{21}$$

$$1 = \exp\big(\phi_\pi \pi_t + \phi_y y_t + \psi_t\big)\mathbb{E}_t\big[\exp\big(-(q_{t+1} - q_t) - \pi_{t+1}\big)\big], \tag{22}$$

where $\Omega_{t+1} \equiv \exp\big((y_{t+1} - y_t) + (q_t - q_{t+1})\big)$ is the NKPC discounting term under risk sharing, and $a_t \equiv \log A_t$.

Step 2: Perturbation expansion. Under complete markets, the static block provides a smooth map $y_t = y(\pi_t, q_t, \tau_t)$, where $\tau_t \equiv \log T_t^*$, with first-order derivatives $\theta_q \equiv \partial y/\partial q|_{ss} > 0$ and $\theta_\tau \equiv -\partial y/\partial \tau|_{ss} > 0$. In closed form (Online Appendix C.1),

$$\theta_q = (1-\alpha)\left[\frac{1}{\sigma} + \alpha\left(\gamma - \frac{1}{\sigma}\right)\right] + \alpha\gamma, \qquad \theta_\tau = \alpha\gamma.$$

The perturbation objects naturally separate into (i) *tariff-level exposures* and (ii) *volatility wedges*. Along deterministic tariff paths with no innovations, equilibrium decision rules admit the expansion

$$\pi_t = A_\pi \tau_t + \tfrac{1}{2} B_\pi \tau_t^2 + O(\tau_t^3), \qquad q_t = A_q \tau_t + \tfrac{1}{2} B_q \tau_t^2 + O(\tau_t^3), \tag{23}$$

where (A_π, A_q) are first-order tariff-exposure coefficients and (B_π, B_q) capture curvature in the tariff level.

By contrast, the volatility experiment in Definition 1 changes only the conditional distribution of τ_{t+1}. Since the reduced complete-markets equilibrium has exactly two forward-looking nonlinear equations (21)–(22), tariff uncertainty enters only through Jensen corrections in these two expectations. Taking logs of the asset-pricing condition and expanding the log-expectation term yields the familiar first-order UIP restriction plus an additive uncertainty wedge:

$$(\phi_\pi \pi_t + \phi_y y_t + \psi_t) = \mathbb{E}_t[(q_{t+1} - q_t) + \pi_{t+1}] + \Delta_t^{UIP} + O(4), \tag{24}$$

where Δ_t^{UIP} collects the baseline-subtracted conditional second-moment terms induced by the log-expectation in (22) (a negative variance component plus a curvature component). We retain the superscript "UIP" because this wedge shifts the linearized UIP relation, even though the

underlying nonlinear closure is Euler-consistent. Similarly, expanding the nonlinear NKPC produces the standard log-linear NKPC plus an additive *pricing wedge*:

$$\pi_t = \kappa_{mc}\big(q_t + \varphi y_t - (1+\varphi)a_t\big) + \beta\mathbb{E}_t[\pi_{t+1}] + \Delta_t^{NK} + O(4), \tag{25}$$

where Δ_t^{NK} collects the baseline-subtracted conditional second-moment terms generated by the expectation in (21). Under the volatility experiment, both wedges are proportional to the volatility state at third order,

$$\Delta_t^{NK} = u^{NK}\,\sigma_{T,t} + O(4), \qquad \Delta_t^{UIP} = u^{UIP}\,\sigma_{T,t} + O(4),$$

so the pair (u^{NK}, u^{UIP}) are sufficient statistics for the uncertainty impulse response.

We therefore define the *third-order uncertainty coefficients* (χ_π, χ_q) by

$$\pi_t = \bar{\pi} + \chi_\pi\,\sigma_{T,t} + O(4), \qquad q_t = \bar{q} + \chi_q\,\sigma_{T,t} + O(4), \tag{26}$$

where $(\bar{\pi}, \bar{q}) = O(\bar{\sigma}_\tau^2)$ are second-order stochastic-mean shifts. Under the volatility experiment in Definition 1, the impulse responses of the two forward-looking variables are geometric:

$$\Delta\pi_t = \chi_\pi\,\rho_\sigma^t\eta_\sigma, \qquad \Delta q_t = \chi_q\,\rho_\sigma^t\eta_\sigma, \qquad t \geq 0, \tag{27}$$

with impact effects $\Delta\pi_0 = \eta_\sigma\chi_\pi$ and $\Delta q_0 = \eta_\sigma\chi_q$.[5]

Step 3: Coefficient matching via a common 2×2 template. A key simplification is that all three coefficient layers solve a 2×2 linear system with the *same* matrix, evaluated at different persistence parameters. Define

$$\mathbf{M}(\rho) \equiv \begin{pmatrix} 1-\beta\rho & -\tilde{\kappa} \\ \phi_\pi - \rho & \phi_y\theta_q + 1 - \rho \end{pmatrix}, \qquad \tilde{\kappa} \equiv \kappa_{mc}(1+\varphi\theta_q), \tag{28}$$

where κ_{mc} is given in (20). At first order, (A_π, A_q) solve $\mathbf{M}(\rho_\tau)\mathbf{x} = \mathbf{b}^{(1)}$; at second order, (B_π, B_q) solve $\mathbf{M}(\rho_\tau^2)\mathbf{x} = \mathbf{b}^{(2)}$; at the volatility order, (χ_π, χ_q) solve

$$\mathbf{M}(\rho_\sigma)\begin{pmatrix}\chi_\pi \\ \chi_q\end{pmatrix} = \begin{pmatrix} u^{NK} \\ u^{UIP}\end{pmatrix} \equiv \bar{\sigma}_\tau^2 \begin{pmatrix} \beta\left(B_\pi + 3A_\pi^2 + 2A_\pi s_1\right) \\ (B_\pi + B_q) - (A_\pi + A_q)^2 \end{pmatrix}, \tag{29}$$

where $s_1 \equiv (\theta_q - 1)A_q - \theta_\tau$. When $\det\mathbf{M}(\rho_\sigma) \neq 0$, this system admits a unique closed-form solution. The right-hand side of (29) depends on the lower-order tariff coefficients (A_π, A_q, B_π, B_q), which themselves solve the same template at the appropriate persistence; explicit expressions for all layers are recorded in Online Appendix C.

[5]The notations for the uncertainty coefficients, (χ_π, χ_q), are not related to the preference parameter χ in the household utility (4).

We now state the closed-form coefficient layers and the implied impulse responses. Section 3.2 then interprets the economic content of the two wedges on the right-hand side of (29).

Proposition 1 (Closed-form uncertainty responses under a Taylor rule)**.** *Under complete markets, a tariff-volatility innovation is fully characterized by three coefficient layers: first-order tariff-exposure coefficients* (A_π, A_q)*, second-order tariff-curvature coefficients* (B_π, B_q)*, and third-order uncertainty coefficients* (χ_π, χ_q)*. Once these are known, all equilibrium impulse responses follow in closed form.*

Formally, maintain PCP and complete international asset markets as in Lemma 1, impose the PPI-based Taylor rule (14), and suppose $\det \mathbf{M}(\rho) \neq 0$ *for* $\rho \in \{\rho_\tau, \rho_\tau^2, \rho_\sigma\}$*.*[6]

First- and second-order tariff coefficients. *The first-order coefficients* (A_π, A_q) *in (23) solve* $\mathbf{M}(\rho_\tau)\mathbf{x} = \mathbf{b}^{(1)}$*, and the second-order curvature coefficients* (B_π, B_q) *solve* $\mathbf{M}(\rho_\tau^2)\mathbf{x} = \mathbf{b}^{(2)}$*, where the forcing vectors depend on static-block derivatives, tariff persistence, and the first-order coefficients. All expressions are available in closed form by Cramer's rule and are reported in Online Appendix C.*

Third-order uncertainty coefficients and geometric IRFs. *Let* (u^{NK}, u^{UIP}) *denote the NKPC pricing-wedge loading and the asset-market uncertainty-wedge loading defined in (25) and (24), respectively (equivalently, the right-hand side of (29)). Then*

$$\chi_\pi = \frac{u^{NK}(\phi_y\theta_q + 1 - \rho_\sigma) + \tilde{\kappa}\, u^{UIP}}{\det \mathbf{M}(\rho_\sigma)},$$
$$\chi_q = \frac{(1 - \beta\rho_\sigma)u^{UIP} - (\phi_\pi - \rho_\sigma)u^{NK}}{\det \mathbf{M}(\rho_\sigma)}.$$

A fully substituted representation in primitive parameters only (eliminating A_π, A_q, B_π, B_q*) is reported in Online Appendix C.4.4.*

Implied IRFs for remaining variables (Figure 3). *Let* $\hat{r}_t \equiv \log(R_t/\bar{R})$ *and let* $e_t \equiv \log \mathcal{E}_t$ *denote the log nominal exchange rate (Home currency per unit Foreign currency), with constant Foreign prices so that* $e_t - e_{t-1} = (q_t - q_{t-1}) + \pi_t$*. Then the third-order impulse responses implied by (27) satisfy*

$$\begin{aligned}
\Delta y_t &= \theta_q\, \Delta q_t + O(4), \\
\Delta \log C^*_{H,t} &= \gamma\, \Delta q_t + O(4), \\
\Delta \hat{r}_t &= \phi_\pi\, \Delta \pi_t + \phi_y\, \Delta y_t + O(4), \\
\Delta(e_t - e_{t-1}) &= \Delta \pi_t + (\Delta q_t - \Delta q_{t-1}), \qquad \Delta q_{-1} = 0.
\end{aligned}$$

3.2 Anatomy of the uncertainty wedges

Equation (29) makes clear that a volatility innovation enters the reduced system only through the two wedge loadings (u^{NK}, u^{UIP}), one per forward-looking equation. We now unpack these wedges

[6] Equivalently, the 2×2 coefficient-matching systems at each perturbation order are invertible.

directly: *what determines their magnitude and sign before they are propagated through the* 2×2 *system?*

The NKPC wedge is

$$u^{NK} = \beta \bar{\sigma}_\tau^2 \left(3A_\pi^2 + 2A_\pi s_1 + B_\pi\right),$$

which consists of three components with distinct economic roles. Here A_π is the first-order exposure of inflation to a deterministic tariff-level change and B_π is the curvature of that exposure (as in (23)). The term $s_1 \equiv (\theta_q - 1)A_q - \theta_\tau$ is the first-order loading of the NKPC discounting kernel Ω_{t+1} in (21) on the tariff state.

The first term $3\beta\bar{\sigma}_\tau^2 A_\pi^2$ is a pure convexity term from Rotemberg pricing: because price adjustment costs are convex in inflation, mean-preserving inflation risk raises expected pricing costs, so this contribution is always weakly positive. The term $2\beta\bar{\sigma}_\tau^2 A_\pi s_1$ is an inflation–discounting covariance term that reflects how the NKPC discounting kernel Ω_{t+1} weights inflation across tariff states. The term $\beta\bar{\sigma}_\tau^2 B_\pi$ is a curvature term inherited from the nonlinear deterministic response $\pi(\tau)$. With these three terms together, the sign of the net wedge determines the direction of pricing pressure: $u^{NK} > 0$ adds inflationary pressure to the NKPC while $u^{NK} < 0$ adds deflationary pressure.

The asset-market wedge

$$u^{UIP} = \bar{\sigma}_\tau^2 \left[(B_\pi + B_q) - (A_\pi + A_q)^2\right]$$

has two pieces. The first, $-\bar{\sigma}_\tau^2(A_\pi + A_q)^2$, is the variance component implied by the Euler-consistent inverse-return pricing object in (22): at first order, next-period depreciation $(q_{t+1} - q_t) + \pi_{t+1}$ loads on $(A_\pi + A_q)\tau_{t+1}$, so higher tariff risk mechanically lowers this component through Jensen's inequality. The second, $\bar{\sigma}_\tau^2(B_\pi + B_q)$, is a curvature correction inherited from the second-order tariff-level mappings in (23); it is negative when depreciation is locally concave in tariffs and positive when it is locally convex. Overall, u^{UIP} summarizes the net asset-market wedge in the complete-markets Euler closure: $u^{UIP} > 0$ pushes the economy toward a real depreciation while $u^{UIP} < 0$ pushes toward a real appreciation (in the standard parameter region where $\det \mathbf{M}(\rho_\sigma) > 0$).[7]

Figure 1 maps this decomposition to the calibration space. In the plotted ranges, the convexity piece of u^{NK} is positive while the covariance and curvature pieces are negative and dominate, so the net NKPC wedge is negative (figure 1, top row). For u^{UIP} (figure 1, bottom row), the variance piece is mechanically negative under the Euler-consistent closure, and the curvature piece is also negative near baseline, so the net asset-market wedge is unambiguously contractionary in the baseline calibration. To understand why the covariance and curvature components take negative signs near baseline, we next consider each in turn.

The covariance component $2\beta\bar{\sigma}_\tau^2 A_\pi s_1$ in u^{NK} reflects that the NKPC does not take a simple average of future inflation across tariff states: the expectation in the pricing block weights them by

[7] To see this formally, set $u^{NK} = 0$ in (29). The first row pins down $\chi_\pi = \frac{\tilde{\kappa}}{1-\beta\rho_\sigma}\chi_q$, and substituting into the asset-pricing row yields $\chi_q = \frac{(1-\beta\rho_\sigma)u^{UIP}}{\det \mathbf{M}(\rho_\sigma)}$. Since $1 - \beta\rho_\sigma > 0$ and a rise in $q_t = \log Q_t$ is a real depreciation in our convention, the sign of the single-channel asset-market response coincides with the sign of u^{UIP} whenever $\det \mathbf{M}(\rho_\sigma) > 0$.

the discounting kernel Ω_{t+1}. When the inflationary tariff states are precisely the states in which Ω_{t+1} is low (for example because tariffs depress activity and export growth), firms put less weight on those inflationary outcomes when setting prices today. This composition effect makes tariff *risk* act like an additional deflationary pricing wedge.

The curvature components $\beta\bar{\sigma}_\tau^2 B_\pi$ in u^{NK} and $\bar{\sigma}_\tau^2(B_\pi + B_q)$ in u^{UIP} come from the fact that the equilibrium responses $\pi(\tau)$ and $q(\tau)$—and hence their sum $\pi(\tau) + q(\tau)$, which governs log nominal depreciation inside the asset-pricing expectation—are not linear in the tariff level. When these mappings are locally concave, more tariff risk (while holding the mean tariff fixed) lowers average inflation and average depreciation by Jensen's inequality, so the curvature terms enter with a negative sign.

Two general-equilibrium forces push toward concavity. First, an export tariff erodes the export base: as τ rises, exports fall, so the tariff bites on a smaller quantity and the marginal effect of further tariff increases on market clearing and marginal costs shrinks, making the economy effectively less exposed at the margin. Second, Rotemberg adjustment costs and Taylor-rule stabilization dampen incremental movements in inflation as the distortion grows; because (π, q) are jointly determined by the NKPC and asset-pricing blocks, this dampening also bends the $q(\tau)$ mapping and, through it, the $\pi(\tau)+q(\tau)$ composite. In the baseline calibration, these forces produce $B_\pi < 0$ and $B_\pi + B_q < 0$, and Figure 1 confirms that both curvature components remain negative across the plotted parameter ranges.

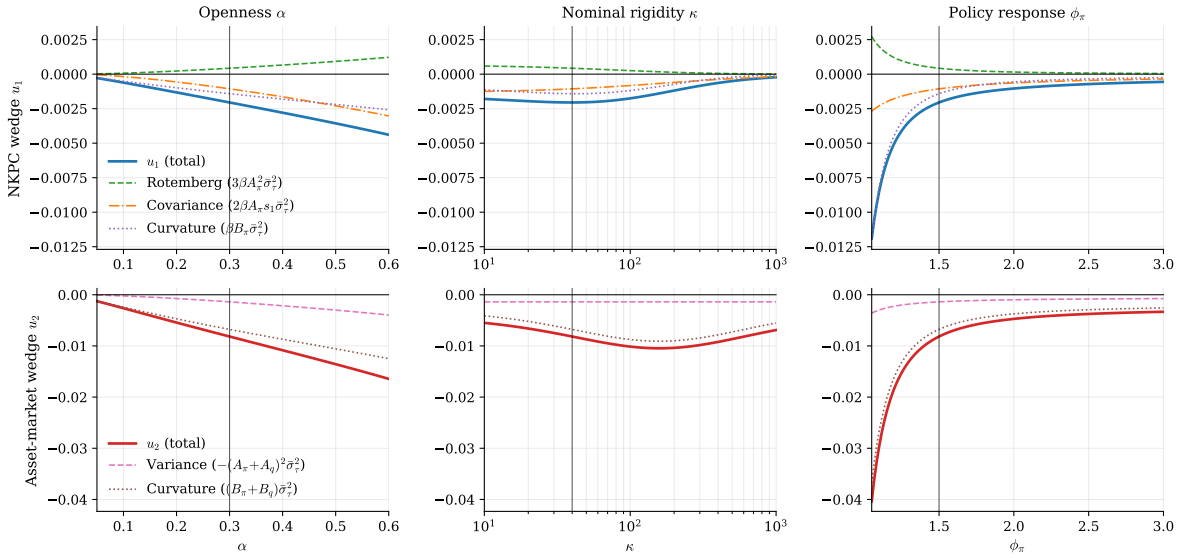

Figure 1: Decomposition of the NKPC and asset-market uncertainty wedges.

Notes: Top row: the NKPC wedge u^{NK} (solid blue) and its three components—Rotemberg convexity $3\beta A_\pi^2\bar{\sigma}_\tau^2$ (dashed green), inflation–SDF covariance $2\beta A_\pi s_1\bar{\sigma}_\tau^2$ (dash-dotted orange), and curvature $\beta B_\pi\bar{\sigma}_\tau^2$ (dotted purple)—as a function of openness α, nominal rigidity κ, and the Taylor inflation coefficient ϕ_π. Bottom row: the asset-market wedge u^{UIP} (solid red) and its two components—the inverse-return variance term $-(A_\pi + A_q)^2\bar{\sigma}_\tau^2$ (dashed pink) and curvature $(B_\pi + B_q)\bar{\sigma}_\tau^2$ (dotted brown). Vertical lines mark the baseline calibration. All other parameters are held at their baseline values.

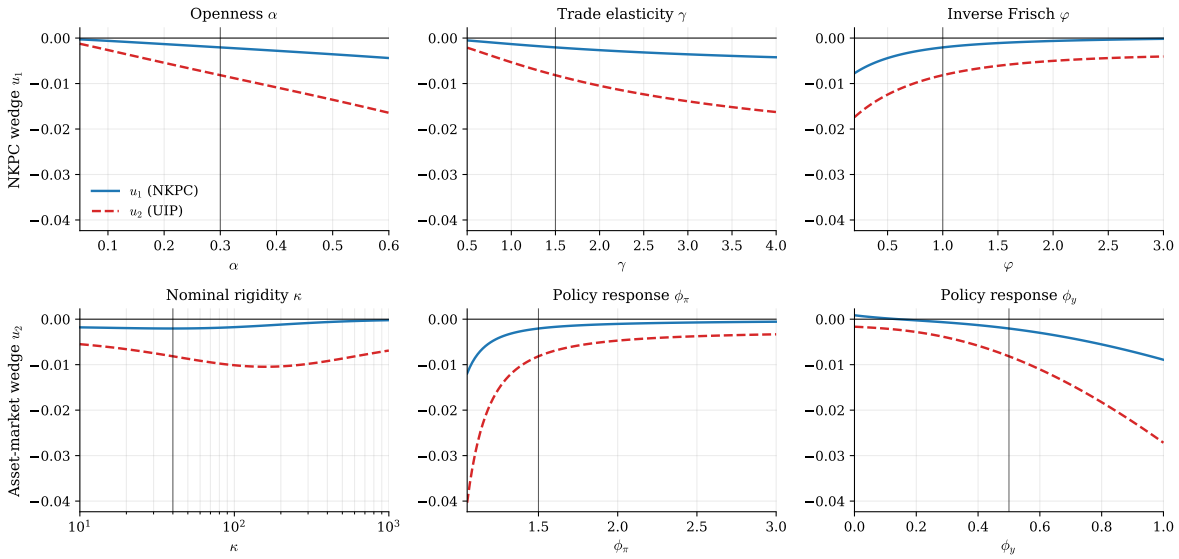

Figure 2: Sensitivity of the NKPC and asset-market uncertainty wedges to structural parameters.

Notes: Each panel varies one parameter while holding the others at the baseline calibration in Table A.1 (vertical line). The solid blue line plots the NKPC wedge u^{NK}; the dashed red line plots the asset-market wedge u^{UIP}. Both are the wedge loadings from (29), which multiply the volatility state $\sigma_{T,t}$ in the uncertainty-coefficient system.

Figure 2 extends this analysis by plotting the sensitivity of u^{NK} and u^{UIP} jointly across six key parameters. In the plotted ranges around the baseline, both wedges are negative over most values. This pattern is consistent with the baseline deflationary response, but it is not sufficient on its own: by (29), the sign of outcomes also depends on the propagation matrix $\mathbf{M}(\rho_\sigma)^{-1}$. Trade exposure (α, γ) amplifies both wedges in these sweeps, because higher openness and trade elasticity increase tariff exposures (A_π, A_q) that enter the Jensen terms. Nominal rigidity and labor-supply curvature reshape relative wedge importance rather than scaling both uniformly: in the plotted calibration range, higher κ (more rigid prices in Rotemberg models) reduces $|u^{NK}|$ and increases $|u^{UIP}|$, while higher φ attenuates both wedges.

Taylor-rule coefficients are also quantitatively important. In the plotted range, lower ϕ_π and higher ϕ_y are associated with larger wedge magnitudes, especially $|u^{UIP}|$. Mechanically, policy parameters enter the lower-order systems for (A_π, A_q) and (B_π, B_q), so they alter both variance and curvature components before the final propagation through $\mathbf{M}(\rho_\sigma)^{-1}$. The equilibrium outcomes (χ_π, χ_q) therefore reflect two layers: variation in wedge loadings and variation in the propagation matrix $\mathbf{M}(\rho_\sigma)^{-1}$.

3.3 Transmission mechanism at baseline

We now trace how the two wedges propagate through the full system. Figure 3 plots the baseline impulse responses under the PPI-based Taylor rule. All responses are driven by the two uncertainty wedges (u^{NK}, u^{UIP}) defined in (25) and (24). The output response is not sign-restricted: it depends on how these wedges translate into the terms of trade under the policy rule.

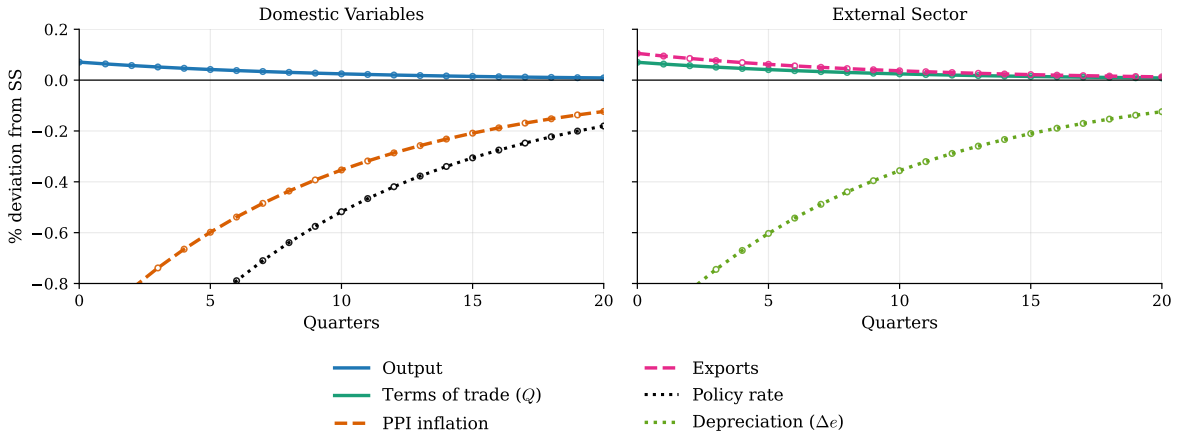

Figure 3: Tariff uncertainty mechanism under the PPI-based Taylor rule.

Notes: The figure plots third-order impulse responses to a one-standard-deviation innovation in tariff volatility at $t = 0$ (Definition 1), with all other shocks shut down, using the baseline calibration in Table A.1. Solid lines are the closed-form predictions; circle markers are the simulated Dynare responses (order 3). The volatility innovation generates two higher-order wedges: an asset-market/Jensen wedge in the asset-market block and an NKPC uncertainty wedge in the pricing block. In the baseline calibration both wedge loadings are negative and, under the Taylor rule, contribute to PPI deflation and monetary easing, but they push relative prices in different directions; the output and trade responses reflect the net effect.

The key object is the terms of trade, q_t. Under PCP, the static block implies $\Delta \log X_t = \gamma \Delta q_t + O(4)$ and $\Delta y_t = \theta_q \Delta q_t + O(4)$ with $\theta_q > 0$, so the real-side response is governed by the sign of Δq_t. The closed-form solution in Proposition 1 implies

$$\chi_q = \frac{(1 - \beta\rho_\sigma)u^{UIP} - (\phi_\pi - \rho_\sigma)u^{NK}}{D_\sigma}, \qquad D_\sigma \equiv \det \mathbf{M}(\rho_\sigma).$$

This expression isolates how policy loads the two wedges. When $D_\sigma > 0$ and $\phi_\pi > \rho_\sigma$, a deflationary NKPC wedge ($u^{NK} < 0$) raises χ_q (a real depreciation), while a negative asset-market/Jensen wedge ($u^{UIP} < 0$) lowers χ_q (a real appreciation). The net sign is not restricted: it depends on the relative magnitudes of the two wedges and on how the Taylor rule enters D_σ.

In the baseline calibration, both wedges are negative and the NKPC component dominates on impact, so $q_0 > 0$ and output rises even though the nominal exchange rate appreciates. The nominal appreciation follows from $\Delta e_t = (q_t - q_{t-1}) + \pi_t$ and the fact that PPI deflation is sizable on impact: the jump in q_t partially offsets $\pi_0 < 0$ at $t = 0$, but as q_t reverts, $(q_1 - q_0)$ reinforces $\pi_1 < 0$, making the depreciation temporarily more negative at $t = 1$ before it converges back to steady state.

The result is "output up, inflation down" as seen in Figure 3. Online Appendix Figure I.1 reports the full set of model variables from the same closed-form solution.

3.4 Special cases

The baseline mechanism in Figure 3 combines the two uncertainty wedges, a policy rule, and persistent dynamics all at once. To build intuitions for how each ingredient contributes, we isolate

their independent impact one at a time, starting from the simplest case and adding complexity. Figure 4 overlays the impulse responses for each special case against the baseline. Table 1 summarizes the key formulas. Derivations are reported in Online Appendix C.4.3.

a. Isolating each channel (Corollary 1). The clearest way to understand the baseline response is to ask: *what would happen if only one of the two wedges was active?*

Corollary 1 (Single-channel limits)**.** *If the NKPC uncertainty forcing is shut down ($u^{NK} = 0$), then*

$$\chi_\pi = \frac{\tilde{\kappa}}{D_\sigma} u^{UIP}, \qquad \chi_q = \frac{1 - \beta\rho_\sigma}{D_\sigma} u^{UIP}, \qquad \frac{\chi_\pi}{\chi_q} = \frac{\tilde{\kappa}}{1 - \beta\rho_\sigma}.$$

If the asset-market/Jensen forcing is shut down ($u^{UIP} = 0$), then

$$\chi_\pi = \frac{1 - \rho_\sigma + \phi_y\theta_q}{D_\sigma} u^{NK}, \qquad \chi_q = -\frac{\phi_\pi - \rho_\sigma}{D_\sigma} u^{NK}, \qquad \frac{\chi_q}{\chi_\pi} = -\frac{\phi_\pi - \rho_\sigma}{1 - \rho_\sigma + \phi_y\theta_q}.$$

Moreover, if $\phi_\pi > \rho_\sigma$ and $u^{NK}(\kappa), u^{UIP}(\kappa)$ remain finite as $\kappa \to 0^+$, with $u^{UIP}(\kappa) \to u^{UIP,0}$, then the flexible-price limits are

$$\lim_{\kappa\to 0^+} \chi_\pi^{NK} = 0, \qquad \lim_{\kappa\to 0^+} \chi_q^{NK} = 0,$$

$$\lim_{\kappa\to 0^+} \chi_q^{UIP} = 0, \qquad \lim_{\kappa\to 0^+} \chi_\pi^{UIP} = \frac{u^{UIP,0}}{\phi_\pi - \rho_\sigma}.$$

In the NKPC-only case ($u^{UIP} = 0$, green dashed line in Figure 4), only the pricing wedge is active. The NKPC uncertainty wedge shifts the Phillips curve and induces deflationary pressure in our baseline calibration. Through the Taylor rule, this generates monetary easing and a real depreciation, so terms of trade, exports, and output rise. Formally, the sign condition is $\text{sign}(\chi_q^{NK}) = \text{sign}(-(\phi_\pi - \rho_\sigma)u^{NK}/D_\sigma)$, so the direction is calibration-dependent; in the baseline it is positive and matches the "output up, inflation down" pattern.

Figure 4 also reports a low-rigidity NKPC-only variant (cyan dashed), setting $\kappa = 4$ while keeping $u^{UIP} = 0$. Relative to the baseline NKPC-only path ($\kappa = 40$), impact responses are much smaller: output falls from about $+0.19\%$ to about $+0.02\%$, and PPI deflation from about -0.19 pp to about -0.02 pp. The key mechanism is attenuation, not sign reversal. Lower κ (more flexible prices in the Rotemberg setup) changes both the wedge loading u^{NK} and the propagation matrix through $\tilde{\kappa} = \kappa_{mc}(1 + \varphi\theta_q)$ in (29). In the NKPC-only mapping, $\chi_q^{NK} = -(\phi_\pi - \rho_\sigma)u^{NK}/D_\sigma$ and $\chi_\pi^{NK} = (1 - \rho_\sigma + \phi_y\theta_q)u^{NK}/D_\sigma$, so for a given u^{NK} a larger $\tilde{\kappa}$ raises D_σ and shrinks both responses. In our calibration, u^{NK} itself also becomes small in magnitude as prices become more flexible (Figure 2), reinforcing the attenuation. Corollary 1 gives the limiting benchmark: as $\kappa \to 0^+$, the NKPC-only channel vanishes, while the asset-only channel delivers no terms-of-trade response and a finite inflation limit.

In the asset-only case ($u^{NK} = 0$, red dash-dotted line), only the asset-market wedge is active. The Jensen correction enters the Euler-consistent asset-pricing condition as a risk-premium-type disturbance. In the baseline calibration this implies nominal appreciation and lower terms of trade, so

output and exports contract. PPI inflation still falls because the terms-of-trade movement feeds into marginal cost, and the policy rate falls. Again, the sign is disciplined by $\chi_q^{UIP} = (1-\beta\rho_\sigma)u^{UIP}/D_\sigma$: in the baseline it is negative, opposite to the NKPC-only channel.

Figure 4 now also reports a low-rigidity asset-only variant (brown dotted), setting $\kappa = 4$ with $u^{NK} = 0$. Relative to baseline asset-only dynamics ($\kappa = 40$), terms-of-trade and real responses collapse toward zero (output from about -0.14% to about -0.01%; terms of trade from about -0.14% to about -0.01%), while inflation remains nonzero though smaller in magnitude (about -0.52 pp to about -0.38 pp). This is exactly the flexible-price limit in Corollary 1: as $\kappa \to 0^+$, $\chi_q^{UIP} \to 0$ but χ_π^{UIP} converges to a finite object pinned by the asset-market wedge and policy feedback.

In our calibration, the NKPC channel dominates on the real side (the output response of $+0.19\%$ under the NKPC-only case versus -0.14% under the asset-only case), and output and exports rise. On the nominal side, the two channels reinforce each other, resulting in a sizable deflation.

b. Removing persistence (Corollary 2).

Corollary 2 (Transitory volatility state: $\rho_\sigma = 0$)**.** *Under Proposition 1, if $\rho_\sigma = 0$ then*

$$\chi_\pi = \frac{u^{NK}(1+\phi_y\theta_q) + \tilde{\kappa}u^{UIP}}{(1+\phi_y\theta_q) + \tilde{\kappa}\phi_\pi}, \qquad \chi_q = \frac{u^{UIP} - \phi_\pi u^{NK}}{(1+\phi_y\theta_q) + \tilde{\kappa}\phi_\pi}.$$

Uncertainty effects are impact-only: $\Delta\pi_0 = \eta_\sigma\chi_\pi$, $\Delta q_0 = \eta_\sigma\chi_q$, and $\Delta\pi_t = \Delta q_t = 0$ for $t \geq 1$.

When the volatility state has no persistence, agents know that the tariff variance returns to its unconditional level after the impact period. Consistent with Corollary 2, $\Delta\pi_t = \Delta q_t = 0$ for $t \geq 1$, so all variables that are contemporaneous functions of (π_t, q_t) are impact-only (orange line in Figure 4). Variables with lagged terms-of-trade components (such as the depreciation rate $\Delta e_t = (q_t - q_{t-1}) + \pi_t$) can still display a one-period mechanical echo at $t = 1$ through $-q_0$. Relative to the baseline, persistence is quantitatively central: in our calibration, impact output is about $+0.05\%$ with persistent volatility but about -0.12% when $\rho_\sigma = 0$, because the two impact channels offset differently once all future-state propagation is shut down.

c. Shutting down output-gap feedback (Corollary 3).

Corollary 3 (Pure inflation Taylor rule: $\phi_y = 0$)**.** *Under Proposition 1, suppose policy responds only to PPI inflation, so $\phi_\pi > 0$ and $\phi_y = 0$. Then*

$$\chi_\pi = \frac{(1-\rho_\sigma)u^{NK} + \tilde{\kappa}u^{UIP}}{(1-\beta\rho_\sigma)(1-\rho_\sigma) + \tilde{\kappa}(\phi_\pi - \rho_\sigma)}, \qquad \chi_q = \frac{(1-\beta\rho_\sigma)u^{UIP} - (\phi_\pi - \rho_\sigma)u^{NK}}{(1-\beta\rho_\sigma)(1-\rho_\sigma) + \tilde{\kappa}(\phi_\pi - \rho_\sigma)}.$$

Removing output-gap feedback changes the propagation matrix and therefore the mapping from wedges to outcomes. In our baseline calibration this weakens the real-side responses (purple dotted

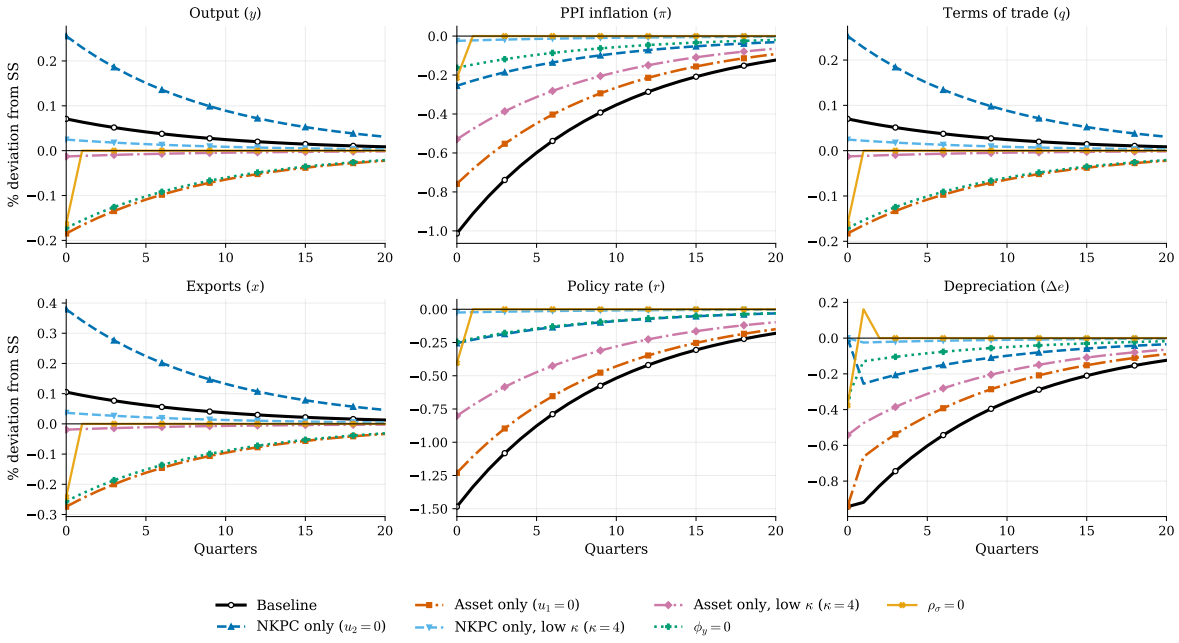

Figure 4: Impulse responses under special cases.

Notes: Each panel overlays the closed-form third-order impulse responses to a one-standard-deviation tariff-volatility innovation (Definition 1) under the baseline calibration and special-case variants. Legend order is: baseline, NKPC-only, asset-only, NKPC-only low-κ, asset-only low-κ, $\phi_y = 0$, and $\rho_\sigma = 0$. "NKPC only" shuts down the asset-market/Jensen forcing ($u^{UIP} = 0$); "asset only" shuts down the NKPC forcing ($u^{NK} = 0$); "$\rho_\sigma = 0$" removes volatility persistence; "$\phi_y = 0$" removes output-gap feedback from the Taylor rule. The additional lines "NKPC only, low κ" and "asset only, low κ" apply the respective single-channel cases at $\kappa = 4$ (more flexible prices). Distinct marker symbols are used to improve readability when lines overlap. All responses are percent deviations from steady state.

line in Figure 4): with $\phi_y = 0$, the policy rate reacts only to inflation, which dampens the induced terms-of-trade movement relative to the baseline rule and reduces output/export responses.

Table 1: Special-case map for uncertainty transmission under a Taylor rule

Case	Closed-form implication	Economic takeaway
$u^{NK} = 0$ or $u^{UIP} = 0$	Channel-specific slopes χ_π/χ_q or χ_q/χ_π in closed form	Under baseline signs, the NKPC channel implies real depreciation and expansion, while the asset-market channel implies appreciation and contraction.
$u^{UIP} = 0,\ \kappa = 4$	NKPC-only mapping evaluated at lower Rotemberg cost	More flexible prices strongly attenuate the NKPC-only real and inflation responses in the baseline calibration.
$u^{NK} = 0,\ \kappa = 4$	Asset-only mapping evaluated at lower Rotemberg cost	Terms-of-trade and real responses shrink toward zero, while inflation remains nonzero (finite flexible-price asset-only limit).
$\rho_\sigma = 0$	Impact-only responses ($\Delta\pi_t = \Delta q_t = 0$ for $t \geq 1$)	Eliminating persistence removes geometric propagation; relative to baseline, this materially attenuates real responses and can change the output sign.
$\phi_y = 0$	Denominator simplifies to $(1 - \beta\rho_\sigma)(1 - \rho_\sigma) + \tilde{\kappa}(\phi_\pi - \rho_\sigma)$	In the baseline calibration, removing output-gap feedback weakens the real-side response.

Notes: Summary of Corollaries 1, 2, and 3. Here u^{NK} is the NKPC uncertainty forcing and u^{UIP} is the asset-market/Jensen forcing from Proposition 1.

3.5 Comparative statics

The special cases above isolate individual ingredients by shutting them down one at a time. We now vary parameters continuously to trace the full landscape of impact responses, reflecting both changes in the wedge loadings (Section 3.2) and in the propagation matrix $\mathbf{M}(\rho_\sigma)^{-1}$. Figures 5–6 show the results.

Three patterns stand out. First, openness and trade elasticity (α, γ) amplify both output and inflation responses monotonically—a direct consequence of the wedge amplification documented in Figure 2. Second, policy parameters (ϕ_π, ϕ_y) govern amplification and can flip the sign of the output response, consistent with Corollary 3 and with the near-singularity behavior of the

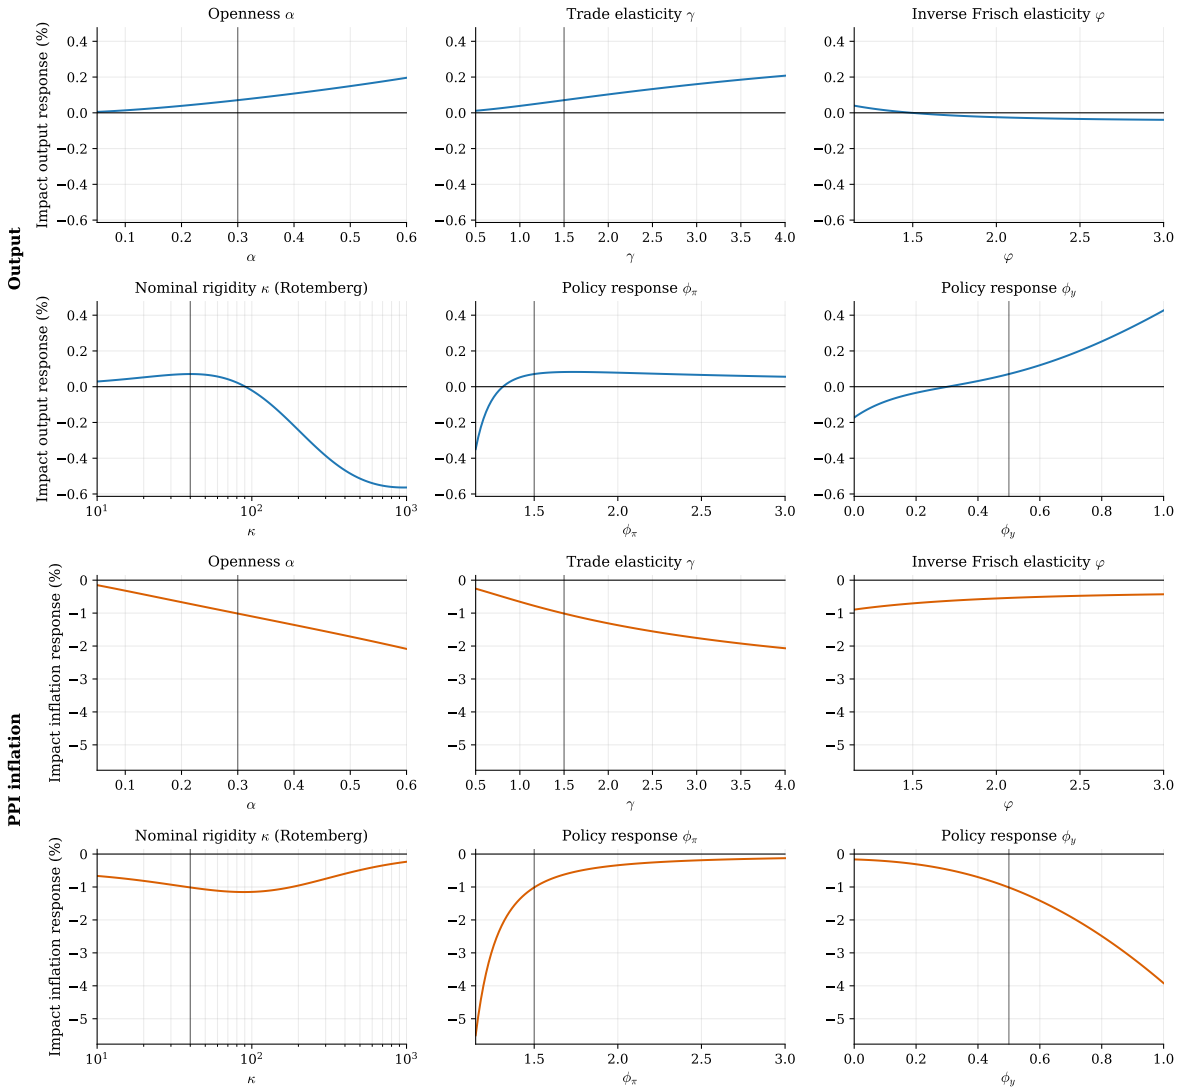

Figure 5: Comparative statics: impact responses to a tariff-uncertainty shock.

Notes: Each panel varies one parameter while holding the others at the baseline calibration in Table A.1 (vertical line). Top two rows: period-0 response of output (blue); bottom two rows: period-0 response of PPI inflation (orange). Both are percent deviations from steady state implied by the closed-form third-order solution in Proposition 1; the κ axis is on a log scale. Trade exposure (α, γ) amplifies both responses monotonically. Nominal rigidity κ and labor-supply curvature φ reshape (and can flip) the output response, while in these sweeps the inflation response is uniformly negative but attenuates toward zero for large φ or κ. Policy feedback (ϕ_π, ϕ_y) generates large swings in both responses, reflecting the dual role of these parameters in determining the wedge loadings and the general-equilibrium mapping.

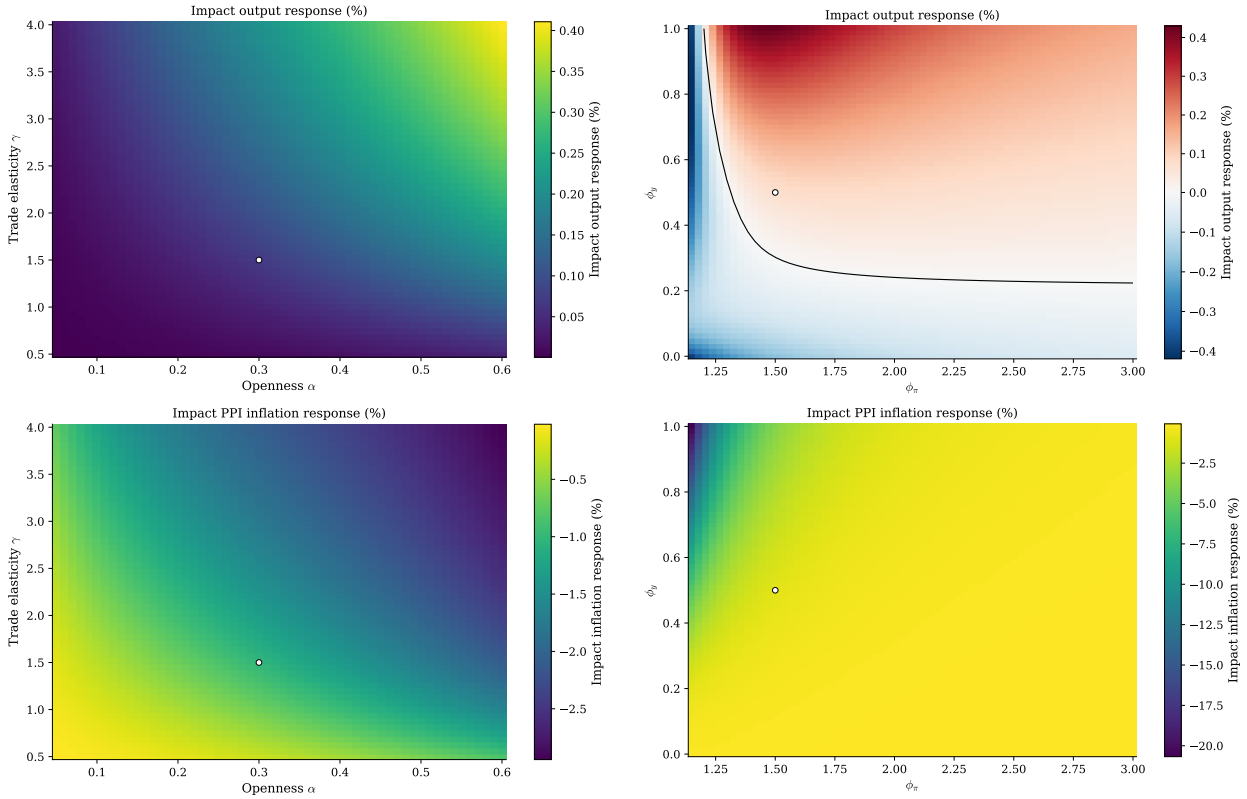

(a) Vary openness α and trade elasticity γ.

(b) Vary Taylor coefficients (ϕ_π, ϕ_y).

Figure 6: Comparative statics: impact responses across parameter pairs.

Notes: In each panel, the top heatmap reports the period-0 output response and the bottom heatmap reports the period-0 PPI inflation response (both in percent deviations from steady state) from Proposition 1; the marker indicates the baseline calibration in Table A.1. In panel (a), both output and inflation responses grow monotonically with openness and trade elasticity; in the plotted range the inflation response is uniformly negative, driven by the NKPC Jensen wedge. In panel (b), the output response can change sign (black contour), reflecting shifts in how policy maps the NKPC and asset-market/Jensen wedges into equilibrium relative prices and spending. In this sweep the inflation response is uniformly negative but its magnitude is highly sensitive to (ϕ_π, ϕ_y): the deflation grows large near the Taylor-principle boundary where $\det \mathbf{M}(\rho_\sigma)$ is small, consistent with the amplification mechanism discussed in the text.

propagation matrix when $\det \mathbf{M}(\rho_\sigma)$ becomes small. Third, nominal rigidity κ and labor-supply curvature φ reshape the *relative* importance of the two wedges (Figure 2), generating hump-shaped or sign-changing output responses even when inflation remains deflationary. Hence, the baseline "output up, inflation down" result is one point in a parameter-dependent map; the sensitivity figures make explicit where it persists and where it reverses.

4 A Two-Channel Spanning Result

Section 3 characterizes the primitive forces behind the impact of tariff uncertainty: two Jensen-correction wedges whose magnitudes and signs depend on structural and policy parameters in economically transparent ways. Those results describe what *drives* uncertainty effects—the nonlinear expectation terms that shift the NKPC and the asset-market block. They do not, however, answer a distinct and practically important question: what do uncertainty effects *look like* when viewed through the lens of a standard linearized model?

This question matters because most policy institutions work with first-order DSGE models populated by a handful of classical disturbances such as productivity shocks, markup shocks, risk-premium shocks, and these models cannot capture uncertainty effects by construction. Knowing that two wedges drive the mechanism (Section 3) does not by itself guarantee that two first-order shocks can replicate the *full equilibrium response* across all endogenous variables and all horizons. Such a representation requires that the decay structure of third-order uncertainty responses aligns with first-order dynamics, and that the spanning weights are time-invariant—neither of which follows from the wedge characterization alone. Our next result establishes that both properties hold under a common-persistence restriction in our model, yielding a closed-form link between the higher-order uncertainty mechanism and first-order impulse responses.

Assumption 2 (Common persistence)**.** All exogenous processes are AR(1) with a common persistence $\rho \in (0,1)$; in particular, $\rho_\tau = \rho_\sigma = \rho_A = \rho_\psi = \rho$ in (16)–(18). When persistence differs across shocks, a representation like (30) typically requires time-varying weights.

Proposition 2 (Two-channel decomposition)**.** *Consider the volatility experiment in Definition 1 under complete markets and the Taylor rule (14). Let* $\mathrm{IRF}_t^{\mathrm{unc}}(x)$ *denote the third-order impulse response of any endogenous variable* x_t *(relative to the baseline with* $\varepsilon_{\sigma,0} = 0$*). Let* $\mathrm{IRF}_t^{\mathrm{Prod}}(x)$ *and* $\mathrm{IRF}_t^{\psi}(x)$ *denote the first-order impulse responses to one-standard-deviation innovations in productivity and the UIP wedge* ψ_t*, respectively.*

Under Assumption 2, there exist constants $(\lambda_A, \lambda_\psi)$ *such that for every horizon* $t \geq 0$*,*

$$\mathrm{IRF}_t^{\mathrm{unc}}(x) = \lambda_A \, \mathrm{IRF}_t^{\mathrm{Prod}}(x) + \lambda_\psi \, \mathrm{IRF}_t^{\psi}(x) + O(4). \tag{30}$$

Moreover, the weights are time-invariant and can be computed in closed form from impact responses

in the linearized system. Let

$$v^{\text{Prod}} \equiv (\text{IRF}_0^{\text{Prod}}(\pi), \text{IRF}_0^{\text{Prod}}(q))^\top, \qquad v^{\psi} \equiv (\text{IRF}_0^{\psi}(\pi), \text{IRF}_0^{\psi}(q))^\top$$

denote the impact responses of (π_t, q_t) *in the* linearized *(first-order) system. Matching impact responses of* (π_t, q_t) *pins down the spanning weights by the* 2×2 *system*

$$\begin{pmatrix} v_\pi^{\text{Prod}} & v_\pi^{\psi} \\ v_q^{\text{Prod}} & v_q^{\psi} \end{pmatrix} \begin{pmatrix} \lambda_A \\ \lambda_\psi \end{pmatrix} = \begin{pmatrix} \eta_\sigma \chi_\pi \\ \eta_\sigma \chi_q \end{pmatrix}. \tag{31}$$

Whenever $\det\left[v^{\text{Prod}}, v^{\psi}\right] \neq 0$ *(i.e. the two basis shocks move* (π_t, q_t) *in linearly independent directions), solving* (31) *yields*

$$\begin{aligned} \lambda_A &= \eta_\sigma \frac{\tilde{\kappa}\chi_q - (1-\beta\rho)\chi_\pi}{\sigma_A c_a}, \\ \lambda_\psi &= \eta_\sigma \frac{-\big(\phi_y\theta_q + 1 - \rho\big)\chi_q - (\phi_\pi - \rho)\chi_\pi}{\sigma_\psi}, \end{aligned} \tag{32}$$

where $\tilde{\kappa} \equiv \kappa_{mc}(1+\varphi\theta_q)$ *is as in Proposition 1,* $c_a \equiv \kappa_{mc}(1+\varphi)$ *is the coefficient on productivity in the linearized NKPC, and* (σ_A, σ_ψ) *are the innovation scales in* (16)–(18)*. In the quantitative illustrations we normalize* $\sigma_A = \sigma_\psi = 0.01$ *(1% innovations).*

Figure 7 confirms Proposition 2 visually. Once persistence is aligned, the two-shock span lies essentially on top of the uncertainty impulse responses across all variables and horizons. No single first-order shock can replicate the uncertainty fingerprint: a productivity shock alone matches the real side but cannot reproduce the inflation and terms-of-trade responses, because tariff uncertainty moves the economy along two linearly independent directions simultaneously.

The spanning result offers two advantages beyond the wedge characterization. First, it converts a higher-order phenomenon into the language of classical first-order shocks; the two-dimensional structure is genuinely necessary, as no single first-order shock can replicate the joint sign pattern (Table 2). Second, it suggests a practical approach for approximating uncertainty effects in first-order models: simulate the two basis shocks and combine the resulting impulse responses using the constant weights $(\lambda_A, \lambda_\psi)$. The accuracy of this approximation in richer models is an open question. In the baseline calibration, the estimated weights are $\lambda_A \approx 0.33$ and $\lambda_\psi \approx 0.55$ (Table 2). Interpreted literally, a one-standard-deviation tariff-volatility shock produces the same *linear* impulse responses as simultaneously hitting the first-order system with a 0.33% productivity shock and a 0.55% UIP-wedge shock (all with the same persistence), even though the primitive disturbance is purely a change in conditional variance.

Table 2 reports the best one-shock approximation and the two-shock spanning fit in the baseline calibration (Table A.1). To make the "best" single shock a meaningful benchmark for real and nominal variables *simultaneously*, we evaluate fits using a sign-aware objective that weights each variable equally (so a proxy that flips the sign of output or the terms of trade is penalized even if it

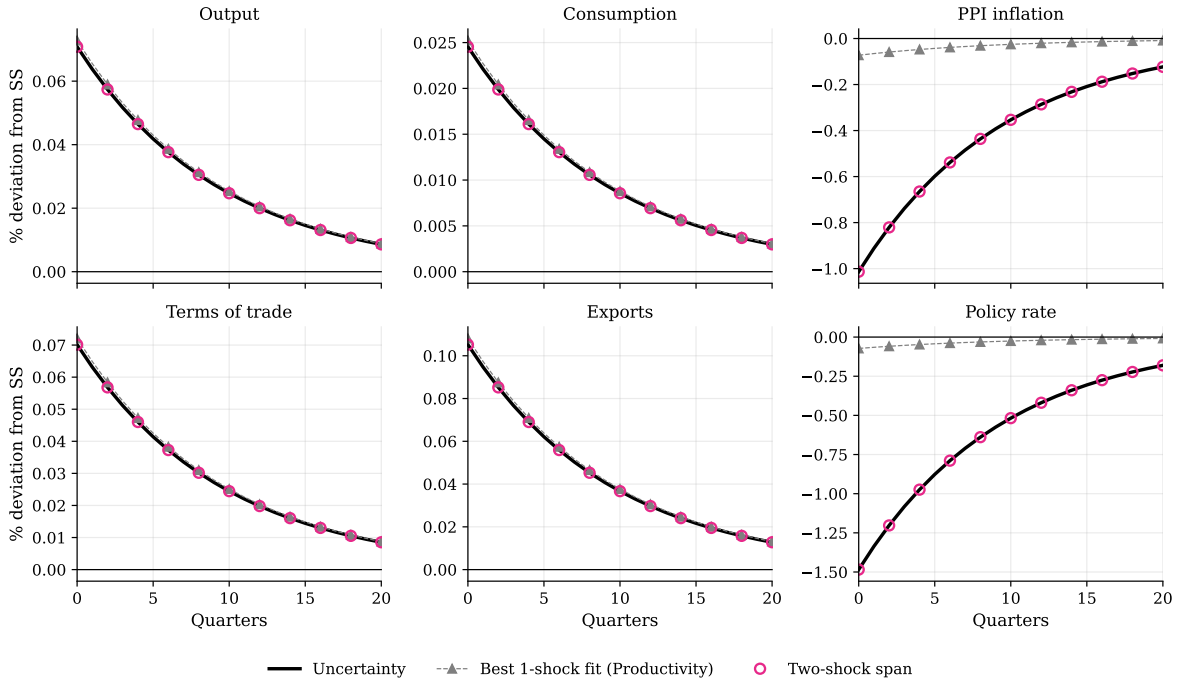

Figure 7: Two-channel spanning.

Notes: The experiment is a one-standard-deviation volatility innovation at $t = 0$ (Definition 1) under the baseline calibration in Table A.1. The solid line plots the third-order uncertainty impulse responses; the dotted line plots the best one-shock fit (productivity, under the balanced objective in the notes to Table 2); and the dashed line plots the fitted constant-weight combination of two first-order basis shocks (productivity and a UIP wedge). All variables are reported as percent deviations from the deterministic steady state (log points multiplied by 100). Imposing common persistence by setting the basis-shock persistence to $\rho = \rho_\sigma$ makes the fitted weights time-invariant.

matches inflation). Under this criterion, productivity is the best single-shock proxy: it matches the sign pattern of the uncertainty response but cannot match the joint magnitudes of inflation and the terms of trade. Adding the UIP wedge delivers a two-shock fit that is exact up to the $O(4)$ remainder (and numerically indistinguishable from perfect fit in Table 2).

Table 2: How many first-order shocks are needed to match uncertainty IRFs?

Approximation	Classical shocks	Spanning weights	R^2
Best 1-shock fit	Productivity	$\lambda_A = 0.064$	0.705
Two-shock span	Productivity + UIP wedge	$\lambda_A = 0.334,\ \lambda_\psi = 0.553$	1.000

Notes: The table uses the volatility experiment in Definition 1 under complete markets and the PPI-based Taylor rule (14). For each candidate fit we choose weights to minimize squared error in stacked impulse responses of output, consumption, PPI inflation, the terms of trade, exports, and the policy rate over horizons $t = 0, \ldots, 20$, where uncertainty IRFs are third order and basis IRFs are first order. We match IRFs in levels (not absolute values), and weights are unrestricted (can be negative), so sign reversals contribute to squared error rather than being "fit away." Basis innovations are normalized to have 1% standard deviation, and we impose the common-persistence restriction by setting $\rho_A = \rho_\psi = \rho_\tau = \rho_\sigma$. To balance the comparison across variables, we normalize each variable's IRF by its own RMS magnitude under the uncertainty shock before stacking, so each variable contributes equally to the objective. The reported R^2 is computed on this normalized stacked vector as $R^2 \equiv 1 - \text{SSE/SST}$.

Even though the spanning representation is exact up to the $O(4)$ remainder under the common-persistence restriction, the weights are not structural constants. Because $(\lambda_A, \lambda_\psi)$ solve the impact-matching system (31), they depend both on the uncertainty coefficients (χ_π, χ_q)—which encode the wedge magnitudes characterized in Section 3—and on the linearized general-equilibrium mapping from the two basis shocks into (π_t, q_t). Figure 8 shows that trade openness tends to scale up both weights, while nominal rigidity and policy feedback reshape the *relative* importance of the two basis directions. The sensitivity patterns parallel those of the underlying wedges (Figure 2), but the weights also reflect how the first-order model translates each basis shock into (π_t, q_t), so the weight comparative statics are not simply a rescaling of the wedge comparative statics.

The comparative statics are economically transparent. In the openness and trade-elasticity panels, higher α and higher γ amplify the underlying uncertainty wedges (as documented in Section 3.2) and simultaneously make external relative prices more potent in the first-order model, so both λ_A and λ_ψ rise. At very low openness the UIP/risk-premium direction carries most of the adjustment, while at higher openness the supply (productivity-like) direction becomes relatively more important.

In the nominal-rigidity panel, as κ rises (prices become more rigid), the Phillips-curve slope falls, so inflation becomes less responsive to a given marginal-cost movement. Matching the deflationary component of the uncertainty response therefore requires larger movements in marginal cost, which tilts the spanning mixture toward the productivity (supply) basis shock: λ_A rises sharply and displays a hump shape. The UIP weight λ_ψ rises for moderate rigidity but falls at high κ, reflecting that exchange-rate movements alone cannot reproduce the inflation response when prices are very rigid.

In the policy-aggressiveness panel, more aggressive inflation feedback in the Taylor rule reduces

both spanning weights. Stronger stabilization shrinks the general-equilibrium mapping from wedges into (π_t, q_t), so tariff uncertainty becomes closer to neutral in equilibrium and is equivalent to smaller classical shocks.

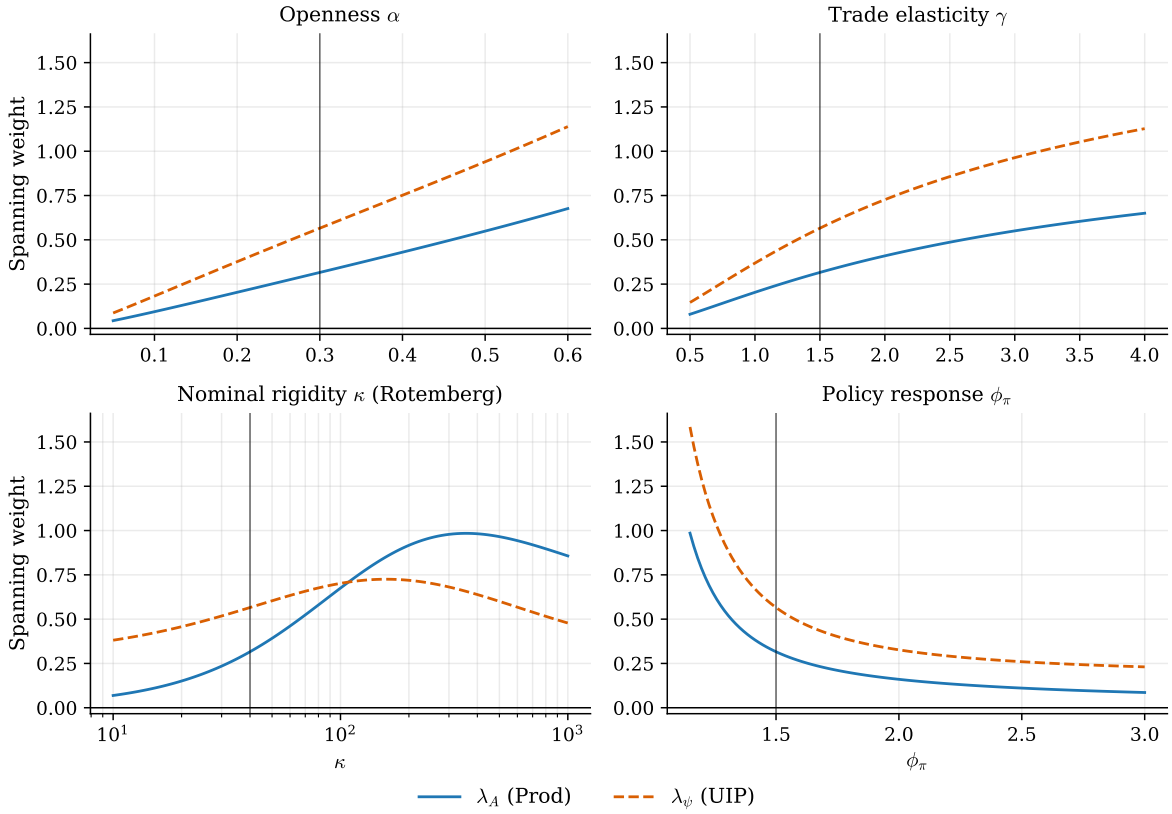

Figure 8: Sensitivity of spanning weights.

Notes: The figure reports the constant weights $(\lambda_A, \lambda_\psi)$ in Proposition 2 under the volatility experiment in Definition 1, varying one parameter at a time around the baseline calibration in Table A.1 (vertical lines). Basis shocks are normalized to 1% innovations, so weights are interpreted as percent-sized basis shocks (e.g. $\lambda_\psi = 0.30$ means a 0.30% UIP-wedge innovation), and persistence is set to $\rho = \rho_\sigma$ to satisfy Assumption 2.

5 Ramsey Policy and Divine Coincidence for Uncertainty Shocks

Sections 3–4 deliver a positive characterization of tariff uncertainty under a Taylor rule: two Jensen-correction wedges drive the mechanism (Section 3), and the resulting equilibrium responses are spanned by two classical first-order shocks up to the $O(4)$ remainder (Section 4). Both sets of results take the interest-rate rule as given. We now turn to the natural normative question: *what is the optimal monetary policy response to tariff uncertainty, and do the sizable real effects documented above reflect an inherent policy tradeoff or an artifact of the Taylor rule?*

To state the Ramsey benchmark precisely, define the flexible-price benchmark allocation and the associated gaps. Let $\{Y_t^{\text{eff}}, C_t^{\text{eff}}, L_t^{\text{eff}}\}$ denote the allocation that would arise under flexible prices

in the complete-markets economy for the same exogenous states. Define the (log) output gap as $\tilde{y}_t \equiv \log(Y_t/Y_t^{\text{eff}})$, and define $\pi_t \equiv \log(\Pi_{H,t})$ as (log) PPI inflation.

Proposition 3 (Ramsey policy under complete markets)**.** *Under complete international asset markets and PCP, assume there is no cost-push wedge. Then Ramsey-optimal monetary policy is strict PPI inflation targeting, $\Pi_{H,t} = 1$ for all t (equivalently, $\pi_t = 0$). This implies a zero output gap (the "divine coincidence": stabilizing PPI inflation also stabilizes the welfare-relevant output gap in this benchmark). Moreover, under the volatility experiment in Definition 1, the Ramsey allocation is invariant to the volatility state in (17), so tariff uncertainty shocks do not move real allocations (up to third order), though the implementing policy rate generally responds through the asset-market/Jensen term.*

Why does the divine coincidence extend to uncertainty shocks? The logic is straightforward (Online Appendix F provides the formal proof). Under strict PPI targeting, $\Pi_{H,t} = 1$ eliminates Rotemberg costs and pins marginal cost—and hence the output gap—at zero, state by state.[8] Since the benchmark flexible-price allocation depends on exogenous *levels* (tariffs, productivity) but not on the volatility state, the Ramsey allocation under strict PPI targeting is invariant to tariff uncertainty. The policy rate still adjusts to satisfy the Euler-consistent asset-pricing relation in the presence of Jensen terms, but PPI inflation and the output gap remain at zero.

Proposition 3 provides a sharp benchmark: in the complete-markets PCP economy, any sizable real effects of tariff uncertainty under a Taylor rule are *policy-induced*—they reflect the interaction between nonlinear expectations and a partial stabilization rule, not an inherent Ramsey tradeoff. Tariff uncertainty affects welfare only through benchmark-allocation risk, which is independent of policy.

Once we depart from the complete market assumption, under incomplete markets, the Euler equation becomes an additional binding constraint and a precautionary-savings channel appears at third order. In this case, strict PPI targeting need not remain Ramsey optimal (Online Appendix G); the quantitative illustrations below focus on the complete-markets benchmark.

5.1 Policy Rules and Quantitative Illustrations

Proposition 3 turns the positive results into a quantitative question: *how far do implementable interest-rate rules move the economy away from the Ramsey benchmark?* We illustrate with third-order simulated impulse responses under the baseline calibration in Online Appendix Table A.1, contrasting strict PPI targeting with PPI-based and CPI-based Taylor rules. We then use the closed-form representation to decompose impact effects into their NKPC and asset-market uncertainty-wedge components.

Figure 9 compares strict PPI targeting to PPI- and CPI-based Taylor rules under the tariff-volatility experiment. Under strict PPI targeting, the volatility innovation is neutral for inflation

[8] In complete markets, risk sharing pins consumption to relative prices, so the welfare loss is locally increasing in $(\pi_t, \tilde{y}_t)$ alone and there is no independent Euler stabilization motive.

and the output gap in the complete-markets economy. Under Taylor rules, the same mean-preserving innovation moves inflation and relative prices through the two uncertainty wedges in (29), generating non-trivial responses in quantities and trade flows. The Taylor-rule economy displays a characteristic pattern: PPI deflation, monetary easing, and a real depreciation (a rise in the terms of trade Q_t) that raises exports and output. This expansion is driven by how Jensen terms in the NKPC and the Euler-consistent asset-pricing block shift the equilibrium relationship between inflation, relative prices, and the policy rate.

Figure 10 decomposes the transmission mechanism into the NKPC wedge and the asset-market wedge. The decomposition exploits the linearity of the 2×2 system (29) in the forcing terms (u^{NK}, u^{UIP}): we solve (29) twice—once setting $u^{UIP} = 0$ (shutting down the asset-market/Jensen wedge) and once setting $u^{NK} = 0$ (shutting down the NKPC wedge)—to obtain counterfactual uncertainty coefficients $(\chi_\pi^{NK}, \chi_q^{NK})$ and $(\chi_\pi^{UIP}, \chi_q^{UIP})$, which we translate into impact responses of output ($y_0 = \theta_q \chi_q \eta_\sigma$), PPI inflation ($\pi_0 = \chi_\pi \eta_\sigma$), the terms of trade ($q_0 = \chi_q \eta_\sigma$), and the policy rate ($r_0 = \phi_\pi \pi_0 + \phi_y y_0$) through the static block and the Taylor rule. Linearity guarantees $(\chi_\pi, \chi_q) = (\chi_\pi^{NK} + \chi_\pi^{UIP},\ \chi_q^{NK} + \chi_q^{UIP})$, so the two sets of bars sum to the total (the dots), up to $O(4)$ remainders.

In the baseline calibration, the NKPC wedge contributes about $+0.19\%$ to both the impact terms-of-trade movement and output, while the asset-market wedge subtracts about -0.14%, leaving a net impact increase of about $+0.05\%$. For inflation and the policy rate, by contrast, the wedges reinforce each other: together they generate a sizable deflation and policy easing on impact of the shock. Economically, the NKPC wedge is a pricing-pressure disturbance: it shifts the Phillips curve through discounted Rotemberg costs and, under the Taylor rule, induces monetary easing and a real depreciation that raises competitiveness and output. The asset-market wedge is an external-finance disturbance: it shifts the Euler-consistent asset-pricing restriction through conditional second moments (a hedge-value component plus curvature), which in our baseline implies nominal appreciation and an offsetting force toward a smaller real depreciation. Hence the two wedges partially offset in the terms of trade and output but reinforce in inflation.

Figure 9 also compares a PPI-based Taylor rule to a CPI-based variant. Under PCP, CPI inflation includes import-price inflation that moves with the exchange rate, so a CPI-based rule implicitly responds to external relative prices. Formally, the two rules differ only in the inflation concept inside the policy feedback:

$$\log \frac{R_t}{\bar{R}} = \phi_\pi \, \pi_t + \phi_y \, y_t \qquad \text{vs} \qquad \log \frac{R_t}{\bar{R}} = \phi_\pi \, \pi_{C,t} + \phi_y \, y_t,$$

where $\pi_t \equiv \log \Pi_{H,t}$ is PPI inflation and $\pi_{C,t} \equiv \log \Pi_{C,t}$ is CPI inflation. Under PCP and the CES CPI aggregator, CPI inflation differs from PPI inflation only through the CPI–PPI wedge G_t:

$$\pi_{C,t} = \pi_t + \Delta g_t, \qquad g_t \equiv \log G_t = \alpha q_t + O(2),$$

so (to first order) the CPI-based rule adds an *exchange-rate pass-through term* proportional to

changes in the terms of trade, $\phi_\pi \Delta g_t \approx \phi_\pi \alpha(q_t - q_{t-1})$. This single term drives all the differences between the responses under the two Taylor rules in Figure 9 (blue line versus red line), and its sign changes over time.

On impact ($t = 0$, with $q_{-1} = 0$), the CPI-based rule internalizes the import-price component of inflation and therefore changes the asset-pricing feedback through Δg_t. In the updated baseline simulation, it delivers milder deflation and a smaller policy-rate cut than the PPI-based rule (-0.91 vs. -1.05 pp), but the full equilibrium feedback implies a *larger* terms-of-trade response on impact ($+0.067$ vs. $+0.049\%$) and hence a larger output expansion ($+0.068$ vs. $+0.050\%$).

Because the CPI-based rule introduces the predetermined state q_{t-1}, its responses are no longer exact geometric decays in the volatility state. Quantitatively, this extra feedback makes the responses more persistent: cumulative output and export responses over 20 quarters are about two-thirds larger under the CPI-based rule in this calibration.

In short, the CPI-based rule inherits a terms-of-trade feedback through $\Delta g_t \approx \alpha(q_t - q_{t-1})$: it responds not only to domestic pricing distortions but also to changes in external relative prices. On impact this feedback tightens policy (import prices jump), but once the terms of trade starts to unwind it pushes CPI inflation below PPI inflation (import prices fall), slowing the return of q_t and real activity toward steady state. The same feedback can generate short-run overshooting, as in Figure 9.

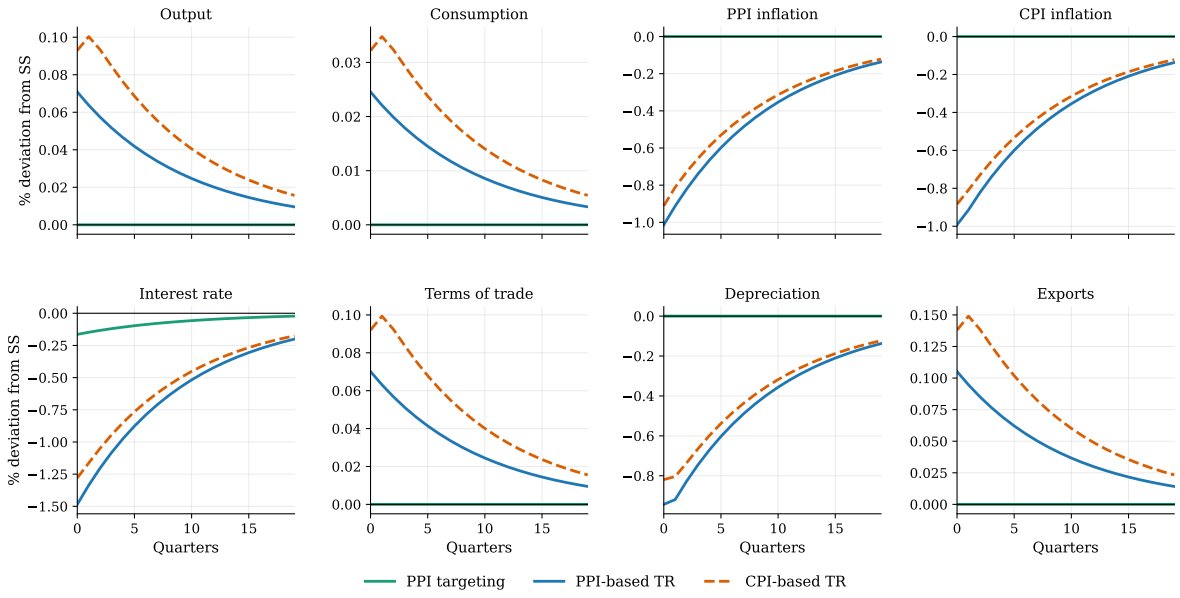

Figure 9: Transmission of tariff uncertainty shocks under different monetary policy rules.

Notes: The experiment is a one-standard-deviation volatility innovation at $t = 0$ (Definition 1) under the baseline calibration in Table A.1. Lines report third-order simulated impulse responses under strict PPI targeting (Ramsey benchmark in Proposition 3), a PPI-based Taylor rule, and a CPI-based Taylor rule. The interest-rate panel reports $100 \cdot \log(R_t/(1/\beta))$. Depreciation is computed as the first difference of the log nominal exchange rate. All variables are reported as percent deviations from steady state (log points multiplied by 100).

Online Appendix G extends the analysis to the exact bond-economy incomplete-markets baseline.

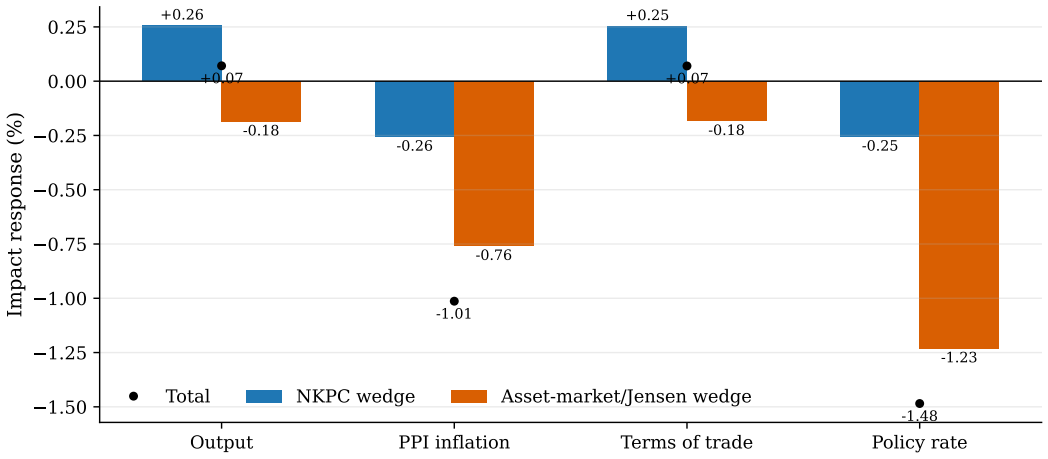

Figure 10: Closed-form wedge decomposition at impact under the PPI-based Taylor rule (baseline calibration).

Notes: The experiment is a one-standard-deviation volatility innovation at $t = 0$ (Definition 1) under the baseline calibration in Table A.1. Bars report the contribution of the NKPC uncertainty wedge (supply/marginal-cost channel) and the asset-market/Jensen wedge (external/asset-market channel) to *impact* (period-0) responses. The decomposition uses the linearity of the third-order system in wedge terms: we solve (29) twice—once shutting down the asset-market/Jensen forcing and once shutting down the NKPC forcing—and then translate each solution into impact responses via the static block and the Taylor rule. The marker reports the total response; up to $O(4)$ terms, the bars sum to the total. Units are percent deviations from steady state (log points multiplied by 100).

An additional Euler/Jensen (precautionary-savings) wedge becomes active at third order, enriching the mechanism from two channels to three (Corollary G.1). Under common persistence, the spanning result generalizes by adding a discount-factor shock (Proposition G.3); Online Appendix Figure G.1 provides an illustrative comparison. In the updated baseline comparison, complete markets with strict PPI targeting remain neutral in real allocations, whereas incomplete markets with strict PPI targeting already generate small but nonzero impact responses (about $+0.0084\%$ for output and -0.0068% for consumption), and a Taylor rule amplifies the consumption and terms-of-trade movements.

6 Conclusion

Tariff uncertainty is not a realized tariff. It is a change in how uncertain firms and households are about future trade barriers. In our model, such pure uncertainty shocks have real and nominal effects because key equilibrium conditions are forward-looking and *nonlinear*: higher dispersion changes the certainty equivalent of nonlinear expectations (Jensen's inequality). By construction, the commonly-used first-order linearizations thus fail to capture this mechanism at higher order.

Our main contribution is to derive a tractable closed-form characterization in a canonical small open economy New Keynesian model. Under complete markets and producer-currency pricing, the equilibrium collapses exactly to two forward-looking equations, and a third-order perturbation of this exact system yields analytical uncertainty responses that can be decomposed into two intuitive

wedges: a pricing wedge in the NKPC and an external-finance wedge in the asset-market block.

Three key results follow. Tariff uncertainty can generate surprising qualitative patterns—output and exports may rise even when inflation falls—not because tariffs are expected to fall, but because the two wedges push relative prices in opposing directions and the net effect depends on trade openness, price rigidity, and the policy rule. For practical policy applications, within our model and under common persistence (Assumption 2), the uncertainty impulse response is spanned by a constant-weight combination of two familiar first-order shocks up to the $O(4)$ remainder, offering a potential bridge to linear DSGE models that cannot account for uncertainty effects by construction. Furthermore, the Ramsey benchmark is clean: strict PPI targeting remains optimal and the divine coincidence extends to tariff uncertainty. This implies that any real effects of tariff uncertainty under implementable Taylor rules are policy-induced departures from the flexible-price benchmark allocation.

The online appendix shows that the analytical framework extends naturally to the exact bond-economy incomplete-markets baseline, where an additional precautionary-savings channel enriches the mechanism from two wedges to three and, under common persistence, the spanning result generalizes by adding a discount-factor shock (up to the $O(4)$ remainder). Strict PPI targeting no longer fully neutralizes uncertainty in this richer setting, pointing to a stabilization trade-off when international risk sharing is imperfect—but the core analytical structure carries over: the exact bond-economy Euler block generates a third wedge, while the remaining static recovery and constant-weight spanning logic remain intact. Whether the same toolkit can accommodate alternative invoicing currencies, endogenous trade-policy dynamics, and richer asset-market structures remains to be explored, and can be disciplined with data on exchange rates, inflation, and trade flows.

References

Auray, Stéphane, Michael B. Devereux, and Aurélien Eyquem. 2025. "Trade Wars, Nominal Rigidities, and Monetary Policy." *Review of Economic Studies*, 92(4): 2228–2270.

Basu, Susanto, and Brent Bundick. 2017. "Uncertainty shocks in a model of effective demand." *Econometrica*, 85(3): 937–958.

Benigno, Gianluca, and Pierpaolo Benigno. 2003. "Price Stability in Open Economies." *Review of Economic Studies*, 70(4): 743–764.

Benigno, Gianluca, Pierpaolo Benigno, and Salvatore Nisticò. 2012. "Risk, Monetary Policy, and the Exchange Rate." *NBER Macroeconomics Annual*, 26: 247–309.

Benigno, Gianluca, Pierpaolo Benigno, and Salvatore Nisticò. 2013. "Second-order approximation of dynamic models with time-varying risk." *Journal of Economic Dynamics and Control*, 37: 1231–1247.

Benigno, Pierpaolo, and Gianluca Benigno. 2008. "Exchange Rate Determination under Interest Rate Rules." *Journal of International Money and Finance*, 27(3): 271–284.

Bergin, Paul R., and Giancarlo Corsetti. 2023. "The macroeconomic stabilization of tariff shocks: What is the optimal monetary response?" *Journal of International Economics*, 143: 103758.

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

Caldara, Dario, Matteo Iacoviello, Patrick Molligo, Andrea Prestipino, and Andrea Raffo. 2020. "The Economic Effects of Trade Policy Uncertainty." *Journal of Monetary Economics*, 109: 38–59.

Cho, Daeha, Yoonshin Han, Joonseok Oh, and Anna Rogantini Picco. 2021. "Uncertainty shocks, precautionary pricing, and optimal monetary policy." *Journal of Macroeconomics*, 69: 103343.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc. 2010. "Optimal monetary policy in open economies." EUI Working Paper ECO 2010/35.

Dixit, Avinash K., and Joseph E. Stiglitz. 1977. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review*, 67(3): 297–308.

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martín Uribe. 2011. "Risk Matters: The Real Effects of Volatility Shocks." *American Economic Review*, 101(6): 2530–2561.

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.

Global Affairs Canada. 2025. "The Government of Canada launches public consultations ahead of the 2026 review of the Canada-United States-Mexico Agreement."

Gross, Isaac, and James Hansen. 2021. "Optimal policy design in nonlinear DSGE models: An n-order accurate approximation." *European Economic Review*, 140: 103918.

Handley, Kyle, and Nuno Limão. 2017. "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States." *American Economic Review*, 107(9): 2731–2783.

Handley, Kyle, and Nuno Limão. 2022. "Trade Policy Uncertainty." *Annual Review of Economics*, 14(1): 363–395.

KalemliÖzcan, Sebnem, Can Soylu, and Muhammed A. Yildirim. 2026. "Global Trade, Tariff Uncertainty and the U.S. Dollar." National Bureau of Economic Research Working Paper 34728.

Monacelli, Tommaso. 2025. "Tariffs and monetary policy."

Rotemberg, Julio J. 1982. "Sticky Prices in the United States." *Journal of Political Economy*, 90(6): 1187–1211.

USTR. 2025*a*. "Extends Certain Exclusions from China Section 301 Tariffs."

USTR. 2025*b*. "Seeks Public Comment on the Joint Review of USMCA."

Villarreal, M. Angeles, Ian F. Fergusson, and Joshua P. Meltzer. 2026. "United States-Mexico-Canada Agreement (USMCA)." Congressional Research Service Report R45944.

White House. 2025*a*. "Adjusting Imports of Aluminum and Steel Into the United States."

White House. 2025*b*. "Adjusting Imports of Automobiles and Automobile Parts Into the United States."

White House. 2026. "Adjusting Imports of Semiconductors and Semiconductor Manufacturing Equipment, and Their Derivative Products, Into the United States."

Online Appendix for "Trade Policy Uncertainty and Optimal Monetary Policy"

Lu Han and Yuko Imura
Bank of Canada

1 May 2026

Contents

A Full Nonlinear Equilibrium Conditions (Levels)

This appendix records a self-contained list of equilibrium conditions in levels implied by the primitives in Section 2 of the main text—"the model in one place." Most of the economic discussion works with the exact reduced system in Section B.

Roadmap. Section A.1 states the common equilibrium conditions in levels. Section A.2 then records the international asset-market block: the exact Euler-consistent complete-markets closure plus risk sharing, and the incomplete-markets bond-economy Euler system in which the domestic Euler equation becomes non-redundant. These conditions are the starting point for the exact reduction in Section B and the incomplete-markets extension in Section G.

A.1 Common block

Let $P^w_{F,t} \equiv \mathcal{E}_t P^*_{F,t}$ denote the Home-currency price of the Foreign good. In the SOE limit, Foreign variables $(C^*_t, P^*_t, P^*_{F,t}, R^*_t)$ are exogenous. In the one-Foreign-good case used throughout, $P^*_t = P^*_{F,t}$ and the terms of trade is $Q_t = P^w_{F,t}/P_{H,t} = \mathcal{E}_t P^*_t / P_{H,t}$.

The common equilibrium conditions are:

$$Y_t = A_t L_t, \tag{A.1}$$

$$\lambda_t = C_t^{-\sigma}, \tag{A.2}$$

$$\frac{W_t}{P_t} = \chi L_t^{\varphi} C_t^{\sigma}, \tag{A.3}$$

$$1 = \beta R_t \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \right], \tag{A.4}$$

$$P_t = \left[(1-\alpha) P_{H,t}^{1-\gamma} + \alpha \left(T_{F,t} P_{F,t}^{w} \right)^{1-\gamma} \right]^{1/(1-\gamma)}, \tag{A.5}$$

$$C_{H,t} = (1-\alpha) \left(\frac{P_{H,t}}{P_t} \right)^{-\gamma} C_t, \tag{A.6}$$

$$C_{F,t} = \alpha \left(\frac{T_{F,t} P_{F,t}^{w}}{P_t} \right)^{-\gamma} C_t, \tag{A.7}$$

$$G_t^{1-\gamma} = (1-\alpha) + \alpha \left(T_{F,t} Q_t \right)^{1-\gamma}, \tag{A.8}$$

$$Q_t \equiv \frac{\mathcal{E}_t P_t^*}{P_{H,t}}, \tag{A.9}$$

$$\frac{Q_t}{Q_{t-1}} = \frac{\mathcal{E}_t}{\mathcal{E}_{t-1}} \cdot \frac{P_t^*}{P_{t-1}^*} \cdot \Pi_{H,t}^{-1}, \tag{A.10}$$

$$\Pi_{C,t} = \Pi_{H,t} \cdot \frac{G_t}{G_{t-1}}, \tag{A.11}$$

$$C_{H,t}^* = \alpha \left(\frac{T_t^*}{Q_t} \right)^{-\gamma} C_t^*, \tag{A.12}$$

$$Y_t = C_{H,t} + C_{H,t}^* + \frac{\kappa}{2} (\Pi_{H,t} - 1)^2 Y_t, \tag{A.13}$$

$$mc_t \equiv \frac{MC_t}{P_{H,t}} = G_t \frac{W_t / P_t}{A_t}, \tag{A.14}$$

$$\kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) + \beta \kappa \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \frac{\lambda_{t+1}}{\lambda_t} \frac{G_t}{G_{t+1}} \right], \tag{A.15}$$

$$\frac{R_t}{\bar{R}} = \Pi_{H,t}^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y}. \tag{A.16}$$

In (A.15), the stochastic discount factor term λ_{t+1}/λ_t can be written as $(C_{t+1}/C_t)^{-\sigma}$ using (A.2).

A.2 International asset markets: complete vs incomplete

Under **complete markets**, the exact nonlinear asset-pricing relation is

$$1 = \frac{R_t}{R_t^*} e^{\psi_t} \mathbb{E}_t \left[\frac{\mathcal{E}_t}{\mathcal{E}_{t+1}} \right]. \tag{A.17}$$

At first order, (A.17) implies the familiar log UIP restriction

$$\hat{r}_t - \hat{r}_t^* + \psi_t = \mathbb{E}_t[\Delta e_{t+1}].$$

Add risk sharing:

$$\frac{C_t^{-\sigma}}{(C_t^*)^{-\sigma}} = \vartheta \cdot \frac{P_t}{\mathcal{E}_t P_t^*}. \tag{A.18}$$

Under the normalization $(C_t^*, P_t^*) = \text{constant}$, (A.18) implies the static consumption mapping (13) and yields the two-equation minimal representation in Lemma B.1. Under **incomplete markets**, (A.18) is dropped; consumption is determined by the Euler equation (A.4). In quantitative work we use the bond-economy closure in Section F.2.1: the domestic Euler equation (F.21), the foreign-bond Euler equation (F.22), and the external-budget law for bond holdings. Dividing the two Euler equations yields first-order UIP with an endogenous premium, but there is no separate exact nonlinear UIP equation in the bond economy.

A.3 Baseline Calibration

Table A.1 reports the baseline calibration used throughout the quantitative illustrations in the main text and this appendix.

Table A.1: Baseline calibration

Parameter	Description	Value
β	Discount factor	0.99
σ	Risk aversion	2
φ	Inverse Frisch elasticity	1
α	Import share	0.30
γ	Trade elasticity	1.5
ϵ	Demand elasticity	10
κ	Rotemberg cost	40
ρ_τ	Export-tariff persistence	0.90
$\bar{\sigma}_\tau$	Export-tariff innovation scale	0.10
ρ_σ	Volatility persistence	0.90
η_σ	Volatility innovation size	$\log 2 \approx 0.693$
ϕ_π	Taylor rule: PPI inflation	1.50
ϕ_y	Taylor rule: output	0.50

Notes: This quarterly calibration is used for the illustrative figures and tables in the main text (Sections 3.5 and 5.1) and throughout this appendix. The "tariff uncertainty shock" is a one-standard-deviation innovation to the volatility state in (17) (Definition 1), holding tariff-level innovations at zero. With $\eta_\sigma = \log 2$, that innovation doubles the conditional standard deviation of next-period tariff innovations from $\bar{\sigma}_\tau = 0.10$ to 0.20; under complete markets, leading uncertainty effects scale with $\bar{\sigma}_\tau^2 \eta_\sigma$ (up to $O(4)$ terms).

B Minimal Nonlinear Representation (Complete Markets)

This appendix derives the minimal nonlinear representation used throughout the paper. Under complete markets and PCP, the full SOE equilibrium collapses to a 2×2 forward-looking system in producer-price inflation and the terms of trade, with all other variables recovered from a static mapping.

Roadmap. Section B.1 states the starting point (the full system plus risk sharing) and the key ingredients behind the reduction. Section B.2 derives the static block that expresses (Y_t, mc_t) as functions of $(\Pi_{H,t}, Q_t)$ and exogenous states. Section B.3 then proves Lemma B.1, giving the exact two-equation nonlinear system. Section B.4 collects identities that recover all remaining variables from $(\Pi_{H,t}, Q_t)$. Section B.5 provides first-order tariff level and news benchmarks used for interpretation.

Endogenous variables, exogenous states, and shocks. In the complete-markets reduction, the two forward-looking endogenous variables are $(\Pi_{H,t}, Q_t)$. The system is driven by exogenous states $(T_{F,t}, T_t^*, A_t, \psi_t)$, where $T_{F,t}$ is the Home import tariff, T_t^* is the Foreign tariff on Home exports, A_t is productivity, and ψ_t is an exogenous UIP wedge (risk-premium shifter). For the uncertainty analysis, we additionally work with the log export-tariff state $\tau_t \equiv \log T_t^*$ and its volatility state $\sigma_{T,t}$ as in (16)–(17); shocks are the innovations $\varepsilon_{\tau,t}$ (tariff), $\varepsilon_{\sigma,t}$ (volatility), $\varepsilon_{A,t}$ (productivity), and $\varepsilon_{\psi,t}$ (UIP wedge).

B.1 Starting point: the full nonlinear system under complete markets

The complete-markets equilibrium is characterized by the common block of conditions in Section A together with risk sharing (A.18). The reduction relies on three ingredients: (i) static demand and price-index relations (which pin down G_t and $C_{H,t}$ as functions of Q_t and tariffs), (ii) complete-markets risk sharing (which pins down C_t as a function of G_t/Q_t), and (iii) the fact that the Rotemberg NKPC depends on intertemporal terms only through the stochastic discount factor, which simplifies under risk sharing.

B.2 Static block

Normalize (C_t^*, P_t^*, R_t^*) to constants, and consider the steady state where $\Pi_H = Q = T_F = T^* = A = 1$ and $\psi = 0$. Define the CES price-index object and CPI–PPI wedge

$$\Phi_t \equiv (1-\alpha) + \alpha(T_{F,t}Q_t)^{1-\gamma}, \qquad G_t = \Phi_t^{1/(1-\gamma)}. \tag{B.1}$$

This follows by dividing the CPI price index (A.5) by $P_{H,t}$ and using $Q_t \equiv \mathcal{E}_t P_t^*/P_{H,t}$ with $P_t^* = P_{F,t}^*$ in the one-Foreign-good case.

Under risk sharing (13), consumption satisfies $C_t = (G_t/Q_t)^{-1/\sigma}$ so

$$C_t = \Phi_t^{-1/((1-\gamma)\sigma)} Q_t^{1/\sigma}. \tag{B.2}$$

Let $\xi \equiv (\gamma - 1/\sigma)/(1-\gamma)$. Domestic absorption of Home goods can be written as

$$C_{H,t} = (1-\alpha)\Phi_t^{\xi} Q_t^{1/\sigma}, \tag{B.3}$$

obtained from the demand system implied by the CES aggregator (2) and the CPI (A.5): $C_{H,t} = (1-\alpha)(P_{H,t}/P_t)^{-\gamma} C_t = (1-\alpha) G_t^{\gamma} C_t$, and substituting G_t and C_t from (B.1)–(B.2).

Export demand (using the normalization $C_t^* \equiv 1$ for notational simplicity) is

$$X_t \equiv C_{H,t}^* = \alpha \left(\frac{Q_t}{T_t^*} \right)^{\gamma}. \tag{B.4}$$

Goods market clearing (A.13) implies

$$Y_t = \frac{C_{H,t} + X_t}{1 - \frac{\kappa}{2}(\Pi_{H,t} - 1)^2}. \tag{B.5}$$

Finally, combining labor supply (A.3), production (A.1), and risk sharing gives the PPI-based real marginal cost mapping

$$mc_t = \chi \, Q_t \, Y_t^{\varphi} A_t^{-(1+\varphi)}, \tag{B.6}$$

To see (B.6), start from the definition $mc_t \equiv MC_t/P_{H,t} = G_t(W_t/P_t)/A_t$ and use labor supply (A.3) with $L_t = Y_t/A_t$ to obtain $mc_t = \chi \, G_t \, C_t^{\sigma} \, Y_t^{\varphi} A_t^{-(1+\varphi)}$. Under risk sharing (13), $C_t^{\sigma} = Q_t/G_t$, implying (B.6).

B.3 Two structural equations

Lemma B.1 (Exact minimal nonlinear system under complete markets)**.** *Under complete markets and PCP, the equilibrium dynamics are pinned down by the following two equations in* $(\Pi_{H,t}, Q_t)$*:*

$$\kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) + \beta\kappa \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \frac{Q_t}{Q_{t+1}} \right], \tag{B.7}$$

$$1 = \Pi_{H,t}^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y} e^{\psi_t} \mathbb{E}_t \left[\frac{Q_t}{Q_{t+1}\Pi_{H,t+1}} \right], \tag{B.8}$$

where (Y_t, mc_t) *are given by* (B.5)–(B.6) *as functions of* $(\Pi_{H,t}, Q_t, T_{F,t}, T_t^*, A_t)$*.*

Intuition. Complete markets eliminate an independent intertemporal margin for Home consumption: risk sharing ties the stochastic discount factor to the terms of trade. As a result, the only dynamic equilibrium restrictions are the price-setting block (the Rotemberg NKPC) and the asset-market block (Taylor rule combined with the Euler-consistent asset-pricing relation in terms of the

terms of trade). Everything else—output, marginal cost, absorption and exports—is recovered from a static block. This two-equation reduction is what makes closed-form uncertainty characterizations possible in the main text.

Derivation of (B.7). Start from the nonlinear NKPC in levels (A.15). Under complete markets, risk sharing implies

$$\left(\frac{C_t}{C_{t+1}}\right)^{\sigma} \frac{G_t}{G_{t+1}} = \frac{Q_t}{Q_{t+1}},$$

so the stochastic discount factor term in (A.15) collapses to Q_t/Q_{t+1}, delivering (B.7).

Derivation of (B.8). Combine the Taylor rule (A.16) with the Euler-consistent asset-pricing relation (A.17) and the definition of the terms of trade $Q_t \equiv \mathcal{E}_t P_t^*/P_{H,t}$. Under constant Foreign prices ($P_{t+1}^* = P_t^*$), the nominal inverse depreciation satisfies $\mathcal{E}_t/\mathcal{E}_{t+1} = Q_t/(Q_{t+1}\Pi_{H,t+1})$, so (A.17) implies

$$1 = \Pi_{H,t}^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}}\right)^{\phi_y} e^{\psi_t} \mathbb{E}_t\left[\frac{Q_t}{Q_{t+1}\Pi_{H,t+1}}\right],$$

which is (B.8).

Remark B.1 (Fully substituted form). For analytical work it is convenient to substitute out the static block. Define

$$Z_t \equiv (1-\alpha)\Phi_t^{\xi} Q_t^{1/\sigma} + \alpha \left(\frac{Q_t}{T_t^*}\right)^{\gamma}, \tag{B.9}$$

$$\mathrm{ADJ}_t \equiv 1 - \frac{\kappa}{2}(\Pi_{H,t} - 1)^2, \tag{B.10}$$

so that $Y_t = Z_t/\mathrm{ADJ}_t$ and $mc_t = \chi Q_t A_t^{-(1+\varphi)} (Z_t/\mathrm{ADJ}_t)^{\varphi}$. Substituting into (B.7)–(B.8) yields a closed system in $(\Pi_{H,t}, Q_t)$ and exogenous states.

$$\begin{aligned}\kappa(\Pi_{H,t}-1)\Pi_{H,t} = (\epsilon-1)\left[\chi Q_t A_t^{-(1+\varphi)}\left(\frac{Z_t}{\mathrm{ADJ}_t}\right)^{\varphi} - 1\right] \\ + \beta\kappa\,\mathbb{E}_t\left[(\Pi_{H,t+1}-1)\Pi_{H,t+1}\cdot\frac{Z_{t+1}}{Z_t}\cdot\frac{\mathrm{ADJ}_t}{\mathrm{ADJ}_{t+1}}\cdot\frac{Q_t}{Q_{t+1}}\right],\end{aligned} \tag{B.11}$$

$$1 = \Pi_{H,t}^{\phi_\pi}\left(\frac{Z_t}{\bar{Y}\,\mathrm{ADJ}_t}\right)^{\phi_y} e^{\psi_t}\mathbb{E}_t\left[\frac{Q_t}{Q_{t+1}\Pi_{H,t+1}}\right]. \tag{B.12}$$

B.4 Recovering Other Variables from the Minimal System

This subsection records identities and static mappings that recover the remaining variables from $(\Pi_{H,t}, Q_t)$.

B.4.1 Prices, inflation, terms of trade, and the exchange rate

Let $p_{H,t} \equiv \log P_{H,t}$, $e_t \equiv \log \mathcal{E}_t$, and $p_t^* \equiv \log P_t^*$. Then (A.9) implies

$$q_t \equiv \log Q_t = e_t + p_t^* - p_{H,t}. \tag{B.13}$$

Producer-price inflation satisfies

$$\pi_t \equiv \log \Pi_{H,t} = p_{H,t} - p_{H,t-1}. \tag{B.14}$$

Combining (B.13)–(B.14) yields the exact law of motion in (A.10) written in logs:

$$\Delta e_t = (q_t - q_{t-1}) + \pi_t + (p_t^* - p_{t-1}^*), \qquad \Delta e_t \equiv e_t - e_{t-1}. \tag{B.15}$$

Under constant Foreign prices ($p_t^* \equiv p^*$), (B.15) reduces to $\Delta e_t = (q_t - q_{t-1}) + \pi_t$.

B.4.2 CPI inflation and the CPI–PPI wedge

Define $G_t \equiv P_t / P_{H,t}$ and $g_t \equiv \log G_t$. Then (A.11) implies the identity

$$\Pi_{C,t} = \Pi_{H,t} \cdot \frac{G_t}{G_{t-1}} \qquad \Longleftrightarrow \qquad \pi_{C,t} = \pi_t + (g_t - g_{t-1}), \tag{B.16}$$

where $\pi_{C,t} \equiv \log \Pi_{C,t}$. In the one-Foreign-good case, the CPI–PPI wedge in (A.8) can be written as

$$G_t^{1-\gamma} = (1-\alpha) + \alpha (T_{F,t} Q_t)^{1-\gamma} \equiv \Phi_t, \qquad g_t = \frac{1}{1-\gamma} \log \Phi_t. \tag{B.17}$$

B.4.3 Consumption (risk sharing) and demand components

Under complete markets, risk sharing (A.18) pins down consumption as $C_t = (G_t / Q_t)^{-1/\sigma}$. In log deviations,

$$c_t \equiv \log(C_t / \bar{C}) = \frac{1}{\sigma}(q_t - g_t). \tag{B.18}$$

Home absorption of Home goods and export demand follow from the CES demand system (A.6) and (A.12):

$$C_{H,t} = (1-\alpha) \left(\frac{P_{H,t}}{P_t} \right)^{-\gamma} C_t = (1-\alpha) G_t^{\gamma} C_t, \tag{B.19}$$

$$X_t = C_{H,t}^* = \alpha \left(\frac{T_t^*}{Q_t} \right)^{-\gamma} C_t^* = \alpha \left(\frac{Q_t}{T_t^*} \right)^{\gamma} \qquad (\text{with } C_t^* \equiv 1). \tag{B.20}$$

Goods market clearing (A.13) implies $Y_t = (C_{H,t} + X_t)/\mathrm{ADJ}_t$ with ADJ_t in (B.10). Equivalently, defining Z_t as in (B.9) gives $Y_t = Z_t/\mathrm{ADJ}_t$. Marginal cost (in PPI units) can be written as

$$mc_t = \chi\, Q_t A_t^{-(1+\varphi)} \left(\frac{Z_t}{\mathrm{ADJ}_t}\right)^{\varphi}, \tag{B.21}$$

as in Remark B.1.

B.4.4 First-order log-linear mappings

For first-order impulse responses, it is convenient to work in logs around the deterministic steady state with $T_{F,t} \equiv 1$. From (B.17), a first-order expansion yields $g_t = \alpha q_t + O(2)$, so (B.18) gives $c_t = \frac{1-\alpha}{\sigma} q_t + O(2)$. Moreover, the static block implies

$$y_t = \theta_q q_t - \theta_\tau \tau_t + O(2), \qquad \tau_t \equiv \log T_t^*, \tag{B.22}$$

and marginal cost in logs satisfies

$$\log mc_t = q_t + \varphi y_t - (1+\varphi) a_t, \qquad a_t \equiv \log A_t. \tag{B.23}$$

Export volume (in log deviations) is $x_t \equiv \log(X_t/\bar{X}) = \gamma(q_t - \tau_t)$ from (B.20). Finally, under the PPI-based Taylor rule (A.16), the policy rate in logs is

$$\hat{r}_t \equiv \log\left(\frac{R_t}{\bar{R}}\right) = \phi_\pi \pi_t + \phi_y y_t. \tag{B.24}$$

B.4.5 Flexible-price benchmark output and the output gap

Define the benchmark flexible-price allocation as the one obtained when price adjustment costs vanish (Rotemberg $\kappa \to 0$), so that $\Pi_{H,t} \equiv 1$ and hence $\pi_t^{\mathrm{eff}} \equiv 0$. Under the normalized marginal-cost definition above, the flexible-price pricing condition implies $mc_t^{\mathrm{eff}} \equiv 1$, so at first order

$$0 = \widehat{mc}_t^{\mathrm{eff}} = q_t^{\mathrm{eff}} + \varphi y_t^{\mathrm{eff}} - (1+\varphi) a_t, \tag{B.25}$$

where $\widehat{mc}_t \equiv \log mc_t$. Using the first-order static mapping (B.22) gives

$$(1 + \varphi\theta_q)\, q_t^{\mathrm{eff}} = \varphi\theta_\tau \tau_t + (1+\varphi) a_t, \qquad q_t^{\mathrm{eff}} = \frac{\varphi\theta_\tau}{1+\varphi\theta_q}\tau_t + \frac{1+\varphi}{1+\varphi\theta_q} a_t, \tag{B.26}$$

and benchmark output is $y_t^{\mathrm{eff}} = \theta_q q_t^{\mathrm{eff}} - \theta_\tau \tau_t$. The output gap is $\tilde{y}_t \equiv y_t - y_t^{\mathrm{eff}}$ (as in Section F.1 below). Under complete markets with PCP, strict PPI targeting implies $\pi_t \equiv 0$ and therefore reproduces the benchmark flexible-price allocation at first order (divine coincidence), so $\tilde{y}_t \equiv 0$ state-by-state.

B.5 Supplementary: First-Order Tariff Level and News Shocks

This subsection collects closed-form first-order responses to *tariff-level* shocks and deterministic tariff *news* paths, which serve as benchmarks for the spanning arguments in Proposition 2.

B.5.1 Surprise tariff-level shock under a PPI-based Taylor rule

Consider a tariff-only experiment $(a_t, \psi_t) \equiv 0$ with an AR(1) export tariff state

$$\tau_t = \rho_\tau \tau_{t-1} + \bar{\sigma}_\tau \varepsilon_{\tau,t}, \qquad \varepsilon_{\tau,t} \sim \mathcal{N}(0,1). \tag{B.27}$$

For a unit surprise shock at $t = 0$ with $\tau_{-1} = 0$, one has $\tau_t = \rho_\tau^t \bar{\sigma}_\tau$. The first-order solution in Section C.2 implies

$$\pi_t = A_\pi \tau_t, \qquad q_t = A_q \tau_t, \qquad (A_\pi, A_q) \equiv (A_\pi^{(\tau)}, A_q^{(\tau)}), \tag{B.28}$$

where $(A_\pi^{(\tau)}, A_q^{(\tau)})$ are given by (C.11)–(C.13) evaluated at $\rho_s = \rho_\tau$. Using (B.22) and $g_t = \alpha q_t + O(2)$ gives the first-order responses of output, CPI wedge, consumption, and exports:

$$y_t = (\theta_q A_q - \theta_\tau)\tau_t, \tag{B.29}$$

$$g_t = \alpha A_q \tau_t, \qquad c_t = \frac{1-\alpha}{\sigma} A_q \tau_t, \qquad x_t = \gamma(A_q - 1)\tau_t. \tag{B.30}$$

The policy rate is

$$\hat{r}_t = \phi_\pi A_\pi \tau_t + \phi_y(\theta_q A_q - \theta_\tau)\tau_t, \tag{B.31}$$

and CPI inflation follows from (B.16):

$$\pi_{C,t} = \pi_t + (g_t - g_{t-1}) = A_\pi \tau_t + \alpha A_q(\tau_t - \tau_{t-1}) \quad \text{with } \tau_{-1} = 0. \tag{B.32}$$

Under constant Foreign prices, the exchange rate satisfies $\Delta e_t = (q_t - q_{t-1}) + \pi_t$ from (B.15).

B.5.2 Strict PPI targeting benchmark

Under strict PPI targeting, $\Pi_{H,t} \equiv 1$ so $\pi_t \equiv 0$. Then the NKPC implies $mc_t \equiv 1$, which together with (B.23) and (B.22) yields $q_t = -\varphi y_t$ for $a_t \equiv 0$. Combining with (B.22) gives the closed-form slope of output on the tariff state:

$$\xi_{T^*} \equiv \frac{\partial y_t}{\partial \tau_t}\Big|_{ss,\ \pi \equiv 0} = -\frac{\theta_\tau}{1 + \varphi\theta_q} = -\frac{\alpha\gamma}{1 + \varphi\theta_q}. \tag{B.33}$$

Hence, for $\tau_t = \rho_\tau^t \bar{\sigma}_\tau$,

$$y_t = \xi_{T^*}\tau_t, \qquad q_t = -\varphi\xi_{T^*}\tau_t, \qquad g_t = -\alpha\varphi\xi_{T^*}\tau_t, \qquad c_t = -\frac{(1-\alpha)\varphi}{\sigma}\xi_{T^*}\tau_t, \tag{B.34}$$

$$x_t = \gamma(q_t - \tau_t) = -\gamma(1+\varphi\xi_{T^*})\tau_t, \qquad e_t = q_t \ \text{ (up to } p_{H,t} \text{ with } \pi \equiv 0), \qquad \pi_{C,t} = g_t - g_{t-1}. \tag{B.35}$$

B.5.3 Deterministic tariff news path

Fix an announcement horizon $k \geq 1$ and consider the deterministic "news" path

$$\tau_t = 0 \ \ (0 \leq t < k), \qquad \tau_t = \rho_\tau^{t-k}\bar{\sigma}_\tau \ \ (t \geq k), \tag{B.36}$$

with all other shocks set to zero. Under strict PPI targeting, the equilibrium is static in the sense that (q_t, y_t) depend only on current τ_t, so $q_t = y_t = c_t = g_t = x_t = 0$ for $0 \leq t < k$.

Under a PPI-based Taylor rule, the forward-looking system generates anticipation effects even though $\tau_t = 0$ for $t < k$. Let $z_t \equiv (\pi_t, q_t)^\top$ and suppress (a_t, ψ_t). Under perfect foresight, (C.9)–(C.10) with $\tau_t \equiv 0$ imply

$$\underbrace{\begin{pmatrix} 1 & -\tilde{\kappa} \\ \phi_\pi & \phi_y\theta_q + 1 \end{pmatrix}}_{\mathbf{A}_0} z_t = \underbrace{\begin{pmatrix} \beta & 0 \\ 1 & 1 \end{pmatrix}}_{\mathbf{A}_1} z_{t+1}, \qquad 0 \leq t < k. \tag{B.37}$$

If $\mathbf{A}_0$ is invertible, define $\mathbf{F} \equiv \mathbf{A}_0^{-1}\mathbf{A}_1$ so that $z_t = \mathbf{F} z_{t+1}$ during the anticipation window. At implementation, the policy functions give $z_k = (A_\pi, A_q)^\top \bar{\sigma}_\tau$, so the anticipation path is

$$z_t = \mathbf{F}^{k-t}(A_\pi, A_q)^\top \bar{\sigma}_\tau, \qquad 0 \leq t < k, \tag{B.38}$$

while for $t \geq k$ one has $z_t = (A_\pi, A_q)^\top \tau_t$ with τ_t from (B.36). All remaining variables follow from (B.16)–(B.24) and the static mappings above.

C Proof of Proposition 1

Roadmap. This appendix records the coefficient-matching steps that deliver the closed-form system (29) and hence proves Proposition 1. We proceed in four steps. Section C.1 derives the second-order expansion of the static block and collects the needed derivatives. Section C.2 solves the first-order (π_t, q_t) system and expresses the solution in closed form. Section C.3 characterizes the second-order curvature induced by the tariff state under the volatility specification. Section C.4 then combines these ingredients to obtain the third-order uncertainty coefficients (χ_π, χ_q), derives the geometric IRFs, and maps the objects to Dynare's perturbation outputs.

C.1 Log representation and static derivatives

Let $\pi_t \equiv \log \Pi_{H,t}$, $q_t \equiv \log Q_t$, and $y_t \equiv \log(Y_t/\bar{Y})$. Under complete markets, the static block in Section B implies a smooth mapping $y_t = y(\pi_t, q_t, \tau_t)$, where $\tau_t \equiv \log T_t^*$ is the export tariff state. A second-order expansion around the deterministic steady state takes the form

$$y_t = \theta_q q_t - \theta_\tau \tau_t + \frac{1}{2}\Big(y_{qq} q_t^2 + 2y_{q\tau} q_t \tau_t + y_{\tau\tau} \tau_t^2 + y_{\pi\pi} \pi_t^2\Big) + O(3). \tag{C.1}$$

The coefficients in (C.1) are functions of primitive parameters $(\alpha, \gamma, \sigma, \kappa)$:

C.1.1 Deriving the static derivatives

For the closed-form calculations in Proposition 1 we treat the import tariff $T_{F,t}$ as constant and focus on the export tariff T_t^*, so that $\tau_t \equiv \log T_t^*$ is the relevant tariff state in the static block. Let $Z_t \equiv C_{H,t} + X_t$ denote the goods-market numerator and recall $Y_t = Z_t/\mathrm{ADJ}_t$ with $\mathrm{ADJ}_t = 1 - \frac{\kappa}{2}(\Pi_{H,t} - 1)^2$ from (B.10). Then

$$y_t = \log(Y_t/\bar{Y}) = \log Z_t - \log(\mathrm{ADJ}_t), \qquad \Pi_{H,t} = e^{\pi_t}, \quad Q_t = e^{q_t}, \quad T_t^* = e^{\tau_t}.$$

Since Z_t depends on (q_t, τ_t) but not on π_t, the cross-derivatives satisfy $y_{\pi q} = y_{\pi\tau} = 0$ at the deterministic steady state.

Step 1 (express Z_t in (q_t, τ_t)). Under the complete-markets static block in Section B.2, Home absorption of Home goods can be written as

$$C_{H,t} = (1-\alpha)\Phi_t^{\xi} Q_t^{1/\sigma}, \qquad \Phi_t = (1-\alpha) + \alpha Q_t^{1-\gamma}, \qquad \xi \equiv \frac{\gamma - 1/\sigma}{1-\gamma},$$

and export demand is

$$X_t = \alpha \left(\frac{Q_t}{T_t^*}\right)^{\gamma} = \alpha e^{\gamma(q_t - \tau_t)}.$$

Hence $Z_t = C_{H,t} + X_t$ is a smooth function of (q_t, τ_t). At the deterministic steady state $(q, \tau, \pi) = (0, 0, 0)$ one has $(C_H, X, Z) = (1-\alpha, \alpha, 1)$ and $\mathrm{ADJ} = 1$.

Step 2 (first derivatives). Because ADJ does not depend on (q_t, τ_t) and equals one at the steady state,

$$\left.\frac{\partial y}{\partial q}\right|_{ss} = \left.\frac{\partial \log Z}{\partial q}\right|_{ss}, \qquad \left.\frac{\partial y}{\partial \tau}\right|_{ss} = \left.\frac{\partial \log Z}{\partial \tau}\right|_{ss}.$$

Moreover,

$$\left.\frac{\partial \log C_H}{\partial q}\right|_{ss} = \frac{1}{\sigma} + \xi \left.\frac{\Phi_q}{\Phi}\right|_{ss}, \qquad \left.\frac{\Phi_q}{\Phi}\right|_{ss} = \alpha(1-\gamma), \qquad \frac{\partial \log X}{\partial q} = \gamma, \quad \frac{\partial \log X}{\partial \tau} = -\gamma.$$

Using $\xi(1-\gamma) = \gamma - 1/\sigma$ and the steady-state weights $C_H/Z = 1-\alpha$ and $X/Z = \alpha$ yields

$$\theta_q \equiv \left.\frac{\partial y}{\partial q}\right|_{ss} = (1-\alpha)\left[\frac{1}{\sigma} + \alpha\left(\gamma - \frac{1}{\sigma}\right)\right] + \alpha\gamma, \qquad \theta_\tau \equiv -\left.\frac{\partial y}{\partial \tau}\right|_{ss} = \alpha\gamma.$$

Step 3 (second derivatives). For $y = \log Z$ one has

$$y_{ij} = \left.\frac{Z_{ij}}{Z}\right|_{ss} - \left.\frac{Z_i Z_j}{Z^2}\right|_{ss}, \qquad i,j \in \{q,\tau\},$$

and since $Z = 1$ at steady state this reduces to $y_{ij} = Z_{ij} - Z_i Z_j$. Combining the derivatives of $X_t = \alpha e^{\gamma(q_t - \tau_t)}$ with those of $C_{H,t} = (1-\alpha)\Phi_t^{\xi} e^{q_t/\sigma}$ yields the expressions for $(y_{qq}, y_{q\tau}, y_{\tau\tau})$ reported in (C.4)–(C.6). Finally, because ADJ_t enters y_t as $-\log \mathrm{ADJ}_t$ and $\mathrm{ADJ}_t = 1 - \frac{\kappa}{2}(e^{\pi_t} - 1)^2$, one has $y_{\pi\pi} = \kappa$ at the steady state, delivering (C.7).

$$\theta_q = (1-\alpha)\left[\frac{1}{\sigma} + \alpha\left(\gamma - \frac{1}{\sigma}\right)\right] + \alpha\gamma, \tag{C.2}$$

$$\theta_\tau = \alpha\gamma, \tag{C.3}$$

$$y_{qq} = -\alpha(1-\alpha)^2(\gamma\sigma - 1)\,\frac{\alpha\gamma\sigma - \alpha - \sigma + 1}{\sigma^2}, \tag{C.4}$$

$$y_{q\tau} = -\alpha\gamma(1-\alpha)^2\,\frac{\gamma\sigma - 1}{\sigma}, \tag{C.5}$$

$$y_{\tau\tau} = \alpha\gamma^2(1-\alpha), \tag{C.6}$$

$$y_{\pi\pi} = \kappa. \tag{C.7}$$

C.2 First-order system and closed-form coefficients

Using $\log mc_t = q_t + \varphi y_t - (1+\varphi)a_t$ and the normalization $mc = 1$ at the deterministic steady state, define

$$\Gamma_q \equiv 1 + \varphi\theta_q, \qquad \Gamma_\tau \equiv \varphi\theta_\tau, \qquad \Gamma_a \equiv 1 + \varphi, \tag{C.8}$$

and recall $\kappa_{mc} = (\epsilon - 1)/\kappa$ from (20). Linearizing (21)–(22) around the deterministic steady state and using the first-order static mapping $y_t = \theta_q q_t - \theta_\tau \tau_t + O(2)$ yields the 2×2 forward-looking system

$$\pi_t = \tilde{\kappa} q_t - c_\tau \tau_t - c_a a_t + \beta \mathbb{E}_t[\pi_{t+1}], \tag{C.9}$$

$$(\phi_\pi)\pi_t + (\phi_y\theta_q)q_t - (\phi_y\theta_\tau)\tau_t + \psi_t = \mathbb{E}_t[(q_{t+1} - q_t) + \pi_{t+1}], \tag{C.10}$$

where $\tilde{\kappa} = \kappa_{mc}\Gamma_q$, $c_\tau = \kappa_{mc}\Gamma_\tau$, and $c_a = \kappa_{mc}\Gamma_a$.

Consider a generic AR(1) state s_t with persistence ρ_s and conjecture $\pi_t = A_\pi^{(s)} s_t$ and $q_t = A_q^{(s)} s_t$.

Then (C.9)–(C.10) imply

$$\mathbf{M}(\rho_s)\begin{pmatrix} A_\pi^{(s)} \\ A_q^{(s)} \end{pmatrix} = \begin{pmatrix} b_1^{(s)} \\ b_2^{(s)} \end{pmatrix}, \tag{C.11}$$

with $\mathbf{M}(\cdot)$ defined in (28) and forcing terms

$$(b_1^{(\tau)}, b_2^{(\tau)}) = (-c_\tau, \phi_y\theta_\tau), \qquad (b_1^{(a)}, b_2^{(a)}) = (-c_a, 0), \qquad (b_1^{(\psi)}, b_2^{(\psi)}) = (0, -1).$$

Whenever $\det \mathbf{M}(\rho_s) \neq 0$, Cramer's rule yields

$$A_\pi^{(s)} = \frac{b_1^{(s)}(\phi_y\theta_q + 1 - \rho_s) + \tilde{\kappa}\, b_2^{(s)}}{\det \mathbf{M}(\rho_s)}, \tag{C.12}$$

$$A_q^{(s)} = \frac{(1 - \beta\rho_s)b_2^{(s)} - (\phi_\pi - \rho_s)b_1^{(s)}}{\det \mathbf{M}(\rho_s)}, \tag{C.13}$$

where $\det \mathbf{M}(\rho_s) = (1 - \beta\rho_s)(\phi_y\theta_q + 1 - \rho_s) + \tilde{\kappa}(\phi_\pi - \rho_s)$.

C.3 Second-order curvature in the tariff state

Consider a deterministic path with no innovations and tariff dynamics $\tau_{t+1} = \rho_\tau\tau_t$. Expand the equilibrium decision rules in powers of the tariff level state:

$$\pi_t = A_\pi\tau_t + \frac{1}{2}B_\pi\tau_t^2 + O(\tau_t^3), \qquad q_t = A_q\tau_t + \frac{1}{2}B_q\tau_t^2 + O(\tau_t^3), \tag{C.14}$$

where $(A_\pi, A_q) \equiv (A_\pi^{(\tau)}, A_q^{(\tau)})$ are the first-order coefficients given by (C.11)–(C.13). Define the composite curvature term

$$\mathcal{C}_y \equiv y_{qq}A_q^2 + 2y_{q\tau}A_q + y_{\tau\tau} + y_{\pi\pi}A_\pi^2, \tag{C.15}$$

and the two repeated scalar combinations

$$m_1 \equiv \Gamma_q A_q - \Gamma_\tau, \qquad s_1 \equiv (\theta_q - 1)A_q - \theta_\tau. \tag{C.16}$$

To make the coefficient matching explicit, first substitute (C.14) into the static expansion (C.1). Since

$$q_t^2 = A_q^2\tau_t^2 + O(\tau_t^3), \qquad q_t\tau_t = A_q\tau_t^2 + O(\tau_t^3), \qquad \pi_t^2 = A_\pi^2\tau_t^2 + O(\tau_t^3),$$

the output expansion becomes

$$y_t = (\theta_q A_q - \theta_\tau)\tau_t + \frac{1}{2}\big(\theta_q B_q + \mathcal{C}_y\big)\tau_t^2 + O(\tau_t^3).$$

Along the deterministic tariff path $\tau_{t+1} = \rho_\tau\tau_t$, this implies

$$\pi_{t+1} = \rho_\tau A_\pi\tau_t + \frac{1}{2}\rho_\tau^2 B_\pi\tau_t^2 + O(\tau_t^3), \qquad q_{t+1} = \rho_\tau A_q\tau_t + \frac{1}{2}\rho_\tau^2 B_q\tau_t^2 + O(\tau_t^3),$$

$$y_{t+1} = \rho_\tau(\theta_q A_q - \theta_\tau)\tau_t + \frac{1}{2}\rho_\tau^2(\theta_q B_q + \mathcal{C}_y)\tau_t^2 + O(\tau_t^3).$$

Hence the two nonlinear objects that recur in (21) are

$$q_t + \varphi y_t = m_1\tau_t + \frac{1}{2}\big(\Gamma_q B_q + \varphi\mathcal{C}_y\big)\tau_t^2 + O(\tau_t^3),$$

and

$$(y_{t+1} - y_t) + (q_t - q_{t+1}) = (\rho_\tau - 1)s_1\tau_t + O(\tau_t^2),$$

so

$$\Omega_{t+1} = 1 + (\rho_\tau - 1)s_1\tau_t + O(\tau_t^2).$$

Now expand the two equilibrium equations to order $O(\tau_t^2)$. For the NKPC, with $a_t = 0$ along this deterministic tariff path,

$$(e^{\pi_t} - 1)e^{\pi_t} = \pi_t + \frac{3}{2}\pi_t^2 + O(\tau_t^3), \qquad e^{q_t + \varphi y_t} - 1 = m_1\tau_t + \frac{1}{2}\big(\Gamma_q B_q + \varphi\mathcal{C}_y + m_1^2\big)\tau_t^2 + O(\tau_t^3).$$

Using the expansion for Ω_{t+1} above and matching the τ_t^2 coefficients in (21) gives

$$(1 - \beta\rho_\tau^2)B_\pi - \tilde{\kappa}B_q = \kappa_{mc}\big(\varphi\mathcal{C}_y + m_1^2\big) + 2\beta\rho_\tau(\rho_\tau - 1)A_\pi s_1 - 3(1 - \beta\rho_\tau^2)A_\pi^2.$$

For the asset-pricing block, because the path is deterministic and $\psi_t = 0$, (22) reduces exactly to

$$\phi_\pi\pi_t + \phi_y y_t = (q_{t+1} - q_t) + \pi_{t+1}.$$

Substituting the same expansions and matching the τ_t^2 coefficients yields

$$(\phi_\pi - \rho_\tau^2)B_\pi + (\phi_y\theta_q + 1 - \rho_\tau^2)B_q = -\phi_y\mathcal{C}_y.$$

Stacking these two coefficient conditions gives the 2×2 system

$$\mathbf{M}(\rho_\tau^2)\begin{pmatrix} B_\pi \\ B_q \end{pmatrix} = \begin{pmatrix} \kappa_{mc}\big(\varphi\mathcal{C}_y + m_1^2\big) + 2\beta\rho_\tau(\rho_\tau - 1)A_\pi s_1 - 3(1 - \beta\rho_\tau^2)A_\pi^2 \\ -\phi_y\mathcal{C}_y \end{pmatrix}. \tag{C.17}$$

Since (C.17) is a 2×2 system, the curvature coefficients (B_π, B_q) are available in closed form by Cramer's rule whenever $\det \mathbf{M}(\rho_\tau^2) \neq 0$.

Figure C.1 reports the same one-dimensional parameter sweeps used in the main-text comparative statics, but now directly for the coefficient objects (A_π, A_q, B_π, B_q). Panel A isolates the first-order tariff-exposure coefficients, while Panel B isolates the second-order curvature coefficients. At the baseline calibration, $(A_\pi, A_q) \approx (0.12, 0.26)$ and $(B_\pi, B_q) \approx (-0.13, -0.52)$, so the tariff state loads more strongly on the terms of trade than on inflation at the first order, while both second-order mappings are locally concave. Across the displayed range, the first-order coefficients remain comparatively stable and mostly positive, whereas the curvature terms are substantially more

sensitive to nominal rigidity and Taylor-rule feedback. This is useful for reading (C.24): the variance terms are governed by the relatively smooth first-order exposure coefficients, but changes in $(\kappa, \phi_\pi, \phi_y)$ reshape the uncertainty wedges mainly through the curvature block (B_π, B_q). Section D returns to these objects and gives the economic interpretation of the signs of B_π, B_q, u_1, and u_2.

Remark C.1 (Curvature objects under log vs level states and variables). The curvature coefficients in (C.14) are defined with respect to the *log* tariff state $\tau_t = \log T_t^*$. If instead one uses a *level* state such as $x_{t-1} \equiv T_{t-1}^* - 1$, then $\tau_t = \rho_\tau \log(1 + x_{t-1}) = \rho_\tau x_{t-1} - \frac{1}{2}\rho_\tau x_{t-1}^2 + O(x_{t-1}^3)$ along a deterministic tariff path. For a scalar decision rule written in a log variable y_t,

$$y_t = A\,\tau_t + \frac{1}{2}B\,\tau_t^2 + O(\tau_t^3), \qquad y \in \{\pi, q\},$$

this implies

$$y_t = \rho_\tau A\,x_{t-1} + \frac{1}{2}\big(\rho_\tau^2 B - \rho_\tau A\big)x_{t-1}^2 + O(x_{t-1}^3).$$

Hence, if a perturbation package reports derivatives of y_t with respect to the level state x_{t-1} at the steady state, the correct mapping is

$$B = \frac{y_{xx} + y_x}{\rho_\tau^2}, \tag{C.18}$$

where y_x and y_{xx} denote first and second derivatives with respect to x_{t-1}. If the reported object is a *level* variable $Y_t = e^{y_t}$ (e.g., $Y = \Pi_H$ or $Y = Q$), then $Y_{xx} = y_{xx} + y_x^2$ at the steady state, so

$$B = \frac{Y_{xx} - Y_x^2 + Y_x}{\rho_\tau^2}. \tag{C.19}$$

Therefore an apparent mismatch in curvature coefficients can reflect comparing different curvature objects (log vs level variables, log vs level states) rather than an economic difference. The economically invariant objects are the implied second- and third-order policy functions and the resulting impulse responses.

C.4 Second-order mean shifts and third-order uncertainty coefficients

C.4.1 Why uncertainty IRFs are third order (and the link to Benigno, Benigno and Nisticò (2013))

Write the export-tariff process under stochastic volatility as

$$\tau_{t+1} = \rho_\tau \tau_t + u_t\,\varepsilon_{\tau,t+1}, \qquad u_t = \bar{\sigma}_\tau e^{\sigma_{T,t}}, \qquad \sigma_{T,t} = \rho_\sigma \sigma_{T,t-1} + \eta_\sigma \varepsilon_{\sigma,t}, \tag{C.20}$$

where $\mathrm{Var}(\varepsilon_{\tau,t+1}) = \mathrm{Var}(\varepsilon_{\sigma,t}) = 1$. Under the volatility experiment in Definition 1, we set $\varepsilon_{\tau,t} = 0$ for all t, so $\tau_t \equiv 0$ along the baseline and along the uncertainty-shock path; the only object that changes at time t is the *conditional distribution* of τ_{t+1} through u_t. As emphasized by Benigno, Benigno and Nisticò (2013), when the volatility/risk state u_t is not an argument of the equilibrium conditions (it only scales primitive innovations), its effect on time-t equilibrium can arise only through conditional

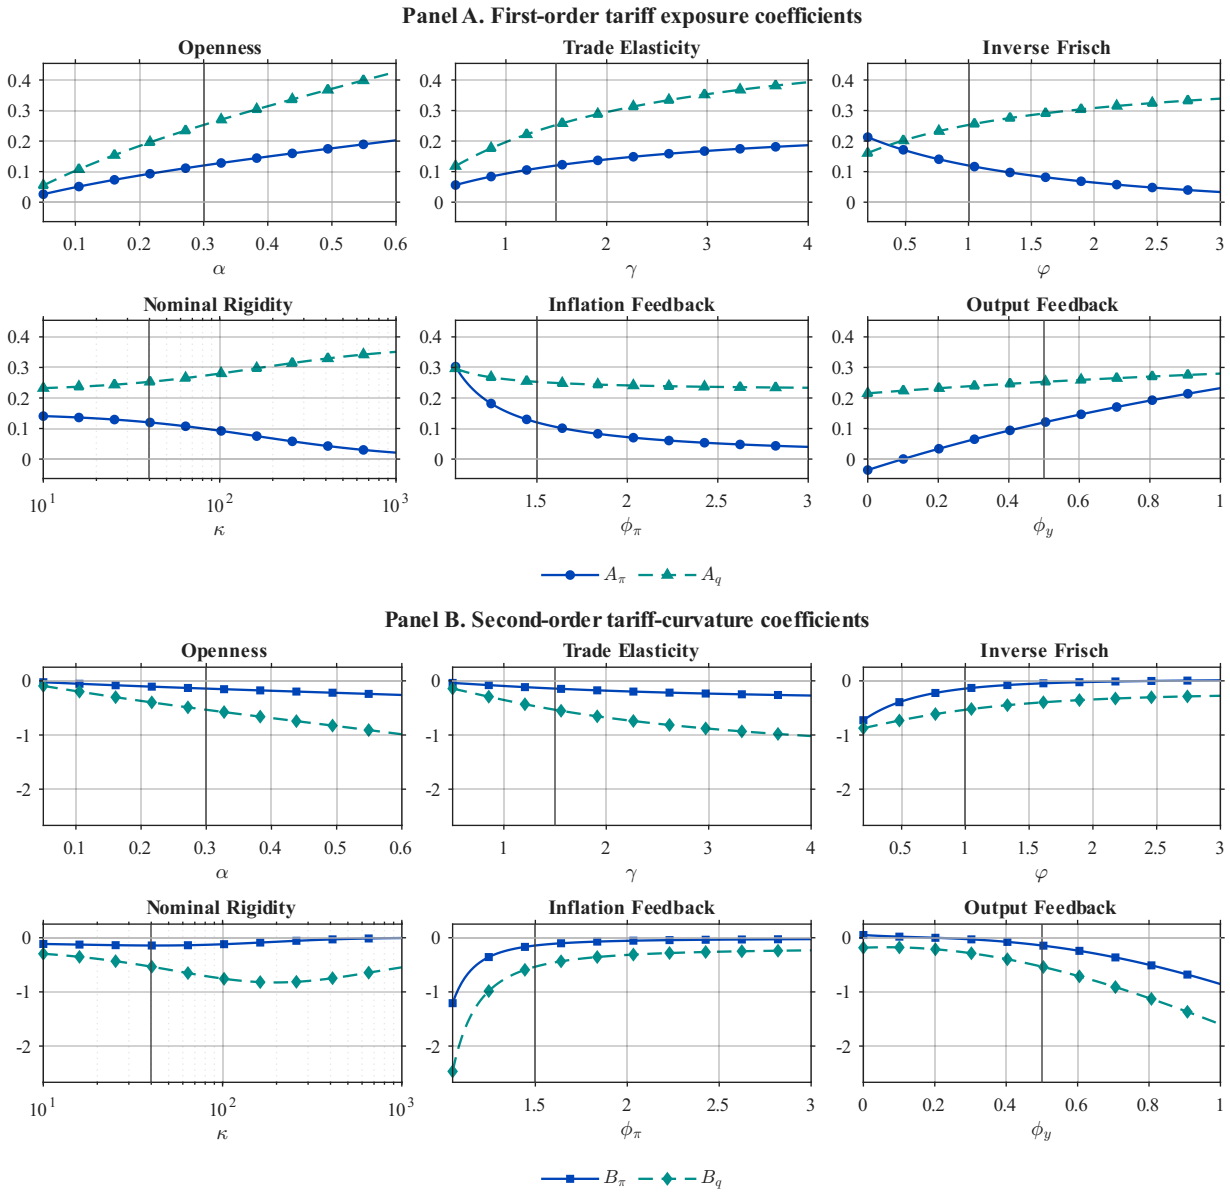

Figure C.1: Sensitivity of the first- and second-order tariff coefficients to structural parameters.

Notes: Each subplot varies one parameter at a time around the baseline calibration in Table A.1; the vertical line marks the baseline value. Panel A plots the first-order tariff-exposure coefficients (A_π, A_q) from (C.11)–(C.13), and Panel B plots the second-order tariff-curvature coefficients (B_π, B_q) from (C.17). Colors identify inflation vs. terms-of-trade loadings, while marker shapes distinguish first-order from second-order objects.

moments such as $\mathbb{E}_t[\tau_{t+1}^2] = u_t^2$ that enter expectations in Euler equations and pricing conditions.

In our log-volatility parameterization, $\mathbb{E}_t[\tau_{t+1}^2] = \bar{\sigma}_\tau^2 e^{2\sigma_{T,t}} = \bar{\sigma}_\tau^2(1 + 2\sigma_{T,t}) + O(4)$ (cf. (C.22) below), so the deviation of conditional variance from the $\sigma_{T,t} \equiv 0$ benchmark is $O(\bar{\sigma}_\tau^2 \sigma_{T,t})$. Under standard perturbation order counting, $\bar{\sigma}_\tau$ is first order and $\sigma_{T,t} = O(\eta_\sigma)$, so uncertainty affects equilibrium at $O(\bar{\sigma}_\tau^2 \eta_\sigma) = O(3)$—the first nonzero responses appear at third order (Figure I.2).

The distinction between "second order" and "third order" is primarily one of state parameterization. Benigno, Benigno and Nisticò (2013) parameterize the conditional standard deviation as a *linear* exogenous process, so a second-order approximation suffices. In our log-volatility specification, the time-varying component of u_t^2 is $\bar{\sigma}_\tau^2(e^{2\sigma_{T,t}} - 1)$, which is inherently higher-order under SGU/Dynare scaling and requires third-order perturbation (cf. Section C.4.6).

At the evaluation point $\tau_t = 0$, a second-order approximation implies that tariff risk shifts the stochastic mean of (π_t, q_t) by $(\bar{\pi}, \bar{q}) = O(\bar{\sigma}_\tau^2)$. These mean shifts solve

$$\underbrace{\begin{pmatrix} 1-\beta & -\tilde{\kappa} \\ \phi_\pi - 1 & \phi_y \theta_q \end{pmatrix}}_{\mathbf{M}_{ss}} \begin{pmatrix} \bar{\pi} \\ \bar{q} \end{pmatrix} = \bar{\sigma}_\tau^2 \begin{pmatrix} \beta\left(\frac{1}{2}B_\pi + \frac{3}{2}A_\pi^2 + A_\pi s_1\right) \\ \frac{1}{2}\left((B_\pi + B_q) - (A_\pi + A_q)^2\right) \end{pmatrix}. \tag{C.21}$$

C.4.2 From mean shifts to uncertainty coefficients

Equation (C.21) is derived under constant conditional tariff variance $\mathbb{E}_t[\tau_{t+1}^2] = \bar{\sigma}_\tau^2$ (i.e. $\sigma_{T,t} \equiv 0$). Under stochastic volatility, the conditional variance is time-varying:

$$\mathbb{E}_t[\tau_{t+1}^2] = \bar{\sigma}_\tau^2 e^{2\sigma_{T,t}} = \bar{\sigma}_\tau^2\big(1 + 2\sigma_{T,t}\big) + O(4), \tag{C.22}$$

where the last equality uses that the volatility state $\sigma_{T,t}$ affects equilibrium only through terms scaled by $\bar{\sigma}_\tau^2$, so the quadratic term $\bar{\sigma}_\tau^2 \sigma_{T,t}^2$ is $O(4)$ and can be dropped at third order.

Define the uncertainty coefficients as in the main text (Definition 1 and (26)):

$$\pi_t = \bar{\pi} + \chi_\pi\, \sigma_{T,t} + O(4), \qquad q_t = \bar{q} + \chi_q\, \sigma_{T,t} + O(4). \tag{C.23}$$

Because $\sigma_{T,t}$ follows (17), $\mathbb{E}_t[\sigma_{T,t+1}] = \rho_\sigma \sigma_{T,t}$. Hence, collecting the $O(\bar{\sigma}_\tau^2 \sigma_{T,t})$ terms in the same coefficient-matching steps that lead to (C.21) yields a linear system for (χ_π, χ_q) with the same 2×2 matrix $\mathbf{M}(\rho_\sigma)$ that governs first-order dynamics (cf. (28)):

$$\mathbf{M}(\rho_\sigma) \begin{pmatrix} \chi_\pi \\ \chi_q \end{pmatrix} = \bar{\sigma}_\tau^2 \begin{pmatrix} \beta\left(B_\pi + 3A_\pi^2 + 2A_\pi s_1\right) \\ (B_\pi + B_q) - (A_\pi + A_q)^2 \end{pmatrix}, \tag{C.24}$$

which is exactly the closed-form uncertainty-coefficient system (29) stated in the main text. Relative to (C.21), the right-hand side doubles because $\partial(\mathbb{E}_t[\tau_{t+1}^2])/\partial\sigma_{T,t} = 2\bar{\sigma}_\tau^2$ at $\sigma_{T,t} = 0$ by (C.22).

C.4.3 Economically transparent special cases for (χ_π, χ_q)

For compact notation in this subsection, define

$$u_1 \equiv \bar{\sigma}_\tau^2 \beta \left(B_\pi + 3A_\pi^2 + 2A_\pi s_1\right), \qquad u_2 \equiv \bar{\sigma}_\tau^2 \left((B_\pi + B_q) - (A_\pi + A_q)^2\right),$$

with $s_1 \equiv (\theta_q - 1)A_q - \theta_\tau$, and let

$$D_\sigma \equiv \det \mathbf{M}(\rho_\sigma), \qquad a_\sigma \equiv 1 - \rho_\sigma + \phi_y \theta_q, \qquad c_\sigma \equiv 1 - \beta\rho_\sigma, \qquad d_\sigma \equiv \phi_\pi - \rho_\sigma.$$

Then (C.24) is equivalent to

$$\chi_\pi = \frac{a_\sigma u_1 + \tilde{\kappa} u_2}{D_\sigma}, \qquad \chi_q = \frac{c_\sigma u_2 - d_\sigma u_1}{D_\sigma}, \tag{C.25}$$

with $D_\sigma = c_\sigma a_\sigma + \tilde{\kappa} d_\sigma$.

To obtain Corollary 2, set $\rho_\sigma = 0$, so $a_\sigma = 1 + \phi_y\theta_q$, $c_\sigma = 1$, $d_\sigma = \phi_\pi$, and $D_\sigma = (1 + \phi_y\theta_q) + \tilde{\kappa}\phi_\pi$. Substituting into (C.25) yields the stated formulas. Since (27) implies $\Delta\pi_t = \chi_\pi \rho_\sigma^t \eta_\sigma$ and $\Delta q_t = \chi_q \rho_\sigma^t \eta_\sigma$, responses are impact-only when $\rho_\sigma = 0$.

To obtain Corollary 3, set $\phi_y = 0$, so $a_\sigma = 1 - \rho_\sigma$ and $D_\sigma = (1 - \beta\rho_\sigma)(1 - \rho_\sigma) + \tilde{\kappa}(\phi_\pi - \rho_\sigma)$. Substituting these into (C.25) gives the closed forms in the corollary.

To obtain Corollary 1, impose either $u_1 = 0$ (asset-only forcing) or $u_2 = 0$ (NKPC-only forcing) in (C.25). The ratio formulas follow by direct division of the resulting expressions. For the low-rigidity quantitative illustrations in Figure 4, we evaluate both $u_2 = 0$ (NKPC-only) and $u_1 = 0$ (asset-only) mappings at $\kappa = 4$ (baseline $\kappa = 40$), holding all other parameters fixed.

To obtain the flexible-price statements in Corollary 1, write the two single-channel solutions as

$$\chi_\pi^{NK} = \frac{a_\sigma u_1}{D_\sigma}, \qquad \chi_q^{NK} = -\frac{d_\sigma u_1}{D_\sigma}, \tag{C.26}$$

$$\chi_\pi^{UIP} = \frac{\tilde{\kappa} u_2}{D_\sigma}, \qquad \chi_q^{UIP} = \frac{c_\sigma u_2}{D_\sigma}, \qquad D_\sigma = c_\sigma a_\sigma + \tilde{\kappa} d_\sigma, \tag{C.27}$$

where $\tilde{\kappa} = \kappa_{mc}(1 + \varphi\theta_q)$ and $d_\sigma = \phi_\pi - \rho_\sigma$. Assume $\phi_\pi > \rho_\sigma$, so $d_\sigma > 0$. Define $K \equiv (\epsilon - 1)(1 + \varphi\theta_q)$, so $\tilde{\kappa} = K/\kappa$.

For the NKPC-only channel,

$$\chi_q^{NK} = -\frac{u_1}{\tilde{\kappa} + \frac{c_\sigma a_\sigma}{d_\sigma}} = -\frac{\kappa\, u_1}{K + \kappa \frac{c_\sigma a_\sigma}{d_\sigma}}, \tag{C.28}$$

$$\chi_\pi^{NK} = \frac{a_\sigma}{d_\sigma} \frac{u_1}{\tilde{\kappa} + \frac{c_\sigma a_\sigma}{d_\sigma}} = \frac{a_\sigma}{d_\sigma} \frac{\kappa\, u_1}{K + \kappa \frac{c_\sigma a_\sigma}{d_\sigma}}, \tag{C.29}$$

For the asset-only channel,

$$\chi_q^{UIP} = \frac{c_\sigma u_2}{\tilde{\kappa} d_\sigma + c_\sigma a_\sigma} = \frac{c_\sigma \kappa\, u_2}{K d_\sigma + \kappa c_\sigma a_\sigma}, \tag{C.30}$$

$$\chi_\pi^{UIP} = \frac{\tilde{\kappa} u_2}{\tilde{\kappa} d_\sigma + c_\sigma a_\sigma} = \frac{K\, u_2}{K d_\sigma + \kappa c_\sigma a_\sigma}. \tag{C.31}$$

Hence, if $u_1(\kappa) = O(1)$, then (C.28)–(C.29) imply $\chi_q^{NK} \to 0$ and $\chi_\pi^{NK} \to 0$. If $u_2(\kappa) = O(1)$, then (C.30) implies $\chi_q^{UIP} \to 0$. If additionally $u_2(\kappa) \to u_2^0$, then (C.31) implies $\chi_\pi^{UIP} \to u_2^0/d_\sigma$.

It remains to verify $u_1(\kappa), u_2(\kappa) = O(1)$. Recall $u_1 = \bar{\sigma}_\tau^2 \beta(B_\pi + 3A_\pi^2 + 2A_\pi s_1)$ with $s_1 = (\theta_q - 1)A_q - \theta_\tau$, and $u_2 = \bar{\sigma}_\tau^2\big((B_\pi + B_q) - (A_\pi + A_q)^2\big)$. From the first-order system $\mathbf{M}(\rho_\tau)(A_\pi, A_q)' = b_\tau$, $b_\tau = (-\kappa_{mc}\Gamma_\tau, \phi_y\theta_\tau)'$ and $\kappa_{mc} = (\epsilon - 1)/\kappa$. As $\kappa \to 0^+$, both $\det \mathbf{M}(\rho_\tau)$ and the leading term in b_τ are $O(1/\kappa)$, so $(A_\pi, A_q) = O(1)$. For (B_π, B_q), $\mathbf{M}(\rho_\tau^2)(B_\pi, B_q)' = \mathrm{rhs}_B$ with $\mathrm{rhs}_B = O(1/\kappa)$ and $\det \mathbf{M}(\rho_\tau^2) = O(1/\kappa)$, implying $(B_\pi, B_q) = O(1)$. Therefore both u_1 and u_2 are $O(1)$ as $\kappa \to 0^+$.

C.4.4 Fully substituted primitive closed form for (χ_π, χ_q)

To remove matrix shorthand and all (A_π, A_q, B_π, B_q) objects, define

$$\Delta(\rho) \equiv (1 - \beta\rho)\big(1 - \rho + \phi_y\theta_q\big) + \kappa_{mc}\big(1 + \varphi\theta_q\big)(\phi_\pi - \rho), \tag{C.32}$$

and set $\Delta_\tau \equiv \Delta(\rho_\tau)$, $\Delta_{\tau^2} \equiv \Delta(\rho_\tau^2)$, and $\Delta_\sigma \equiv \Delta(\rho_\sigma)$. **Primitive numerator blocks.** Let $\kappa_{mc} \equiv (\epsilon - 1)/\kappa$ and define

$$N_\pi \equiv \theta_\tau \kappa_{mc}\left[-\varphi(1 - \rho_\tau + \phi_y\theta_q) + (1 + \varphi\theta_q)\phi_y\right], \tag{C.33}$$

$$N_q \equiv \theta_\tau\left[(1 - \beta\rho_\tau)\phi_y + \kappa_{mc}\varphi(\phi_\pi - \rho_\tau)\right]. \tag{C.34}$$

Define

$$N_C \equiv y_{qq}N_q^2 + 2y_{q\tau}N_q\Delta_\tau + y_{\tau\tau}\Delta_\tau^2 + y_{\pi\pi}N_\pi^2, \tag{C.35}$$

$$N_m \equiv (1 + \varphi\theta_q)N_q - \varphi\theta_\tau\Delta_\tau, \qquad N_s \equiv (\theta_q - 1)N_q - \theta_\tau\Delta_\tau, \tag{C.36}$$

and

$$N_{D_1} \equiv \kappa_{mc}\big(\varphi N_C + N_m^2\big) + 2\beta\rho_\tau(\rho_\tau - 1)N_\pi N_s - 3(1 - \beta\rho_\tau^2)N_\pi^2, \tag{C.37}$$

$$N_{D_2} \equiv -\phi_y N_C. \tag{C.38}$$

Then

$$N_{B_\pi} \equiv N_{D_1}(1 - \rho_\tau^2 + \phi_y\theta_q) + \kappa_{mc}(1 + \varphi\theta_q)N_{D_2}, \tag{C.39}$$

$$N_{B_q} \equiv (1 - \beta\rho_\tau^2)N_{D_2} - (\phi_\pi - \rho_\tau^2)N_{D_1}. \tag{C.40}$$

Define two final composite blocks

$$N_1 \equiv N_{B_\pi} + \Delta_{\tau^2}\big(3N_\pi^2 + 2N_\pi N_s\big), \tag{C.41}$$

$$N_2 \equiv \Delta_{\tau^2}(N_\pi + N_q)^2 + N_{B_\pi} + N_{B_q}. \tag{C.42}$$

Then the uncertainty coefficients from (C.24) are

$$\chi_\pi = \frac{\bar{\sigma}_\tau^2}{\Delta_\sigma \, \Delta_\tau^2 \, \Delta_{\tau^2}} \left[\beta(1-\rho_\sigma+\phi_y\theta_q)N_1 + \kappa_{mc}(1+\varphi\theta_q)N_2\right], \tag{C.43}$$

$$\chi_q = \frac{\bar{\sigma}_\tau^2}{\Delta_\sigma \, \Delta_\tau^2 \, \Delta_{\tau^2}} \left[(1-\beta\rho_\sigma)N_2 - \beta(\phi_\pi-\rho_\sigma)N_1\right]. \tag{C.44}$$

Equations (C.2)–(C.44) provide a fully explicit recursive closed form in primitives

$$\{\alpha, \gamma, \sigma, \varphi, \beta, \kappa, \epsilon, \phi_\pi, \phi_y, \rho_\tau, \rho_\sigma, \bar{\sigma}_\tau\}.$$

C.4.5 Geometric IRFs under the volatility experiment

Under the volatility experiment in Definition 1, $\sigma_{T,t} = \rho_\sigma^t \eta_\sigma$ for $t \geq 0$ (starting from $\sigma_{T,-1} = 0$). Substituting this path into (C.23) gives the geometric impulse responses in (27).

C.4.6 Mapping to Dynare perturbation objects (`ghu`, `ghs2`, `ghxss`, `ghuss`)

This subsection records how our analytic coefficient objects map to Dynare's perturbation outputs (SGU convention). The key point is that Dynare expresses decision rules as functions of lagged predetermined states x_{t-1} and current innovations u_t.

First-order impact responses (`ghu`). Let a generic state s_t satisfy $s_t = \rho_s s_{t-1} + \sigma_s \varepsilon_{s,t}$ with $\mathrm{Var}(\varepsilon_{s,t}) = 1$ and suppose the first-order solution satisfies

$$\pi_t = A_\pi^{(s)} s_t, \qquad q_t = A_q^{(s)} s_t.$$

Then the impact response to a unit innovation $\varepsilon_{s,t}$ is

$$\texttt{ghu}[\pi, \varepsilon_s] = \sigma_s A_\pi^{(s)}, \qquad \texttt{ghu}[q, \varepsilon_s] = \sigma_s A_q^{(s)}. \tag{C.45}$$

Second-order curvature in the tariff state (`ghx` and `ghxx`). In coefficient matching we define curvature coefficients with respect to the *log* tariff state $\tau_t \equiv \log T_t^*$ and work with log endogenous variables $y_t \in \{\pi_t, q_t\}$. In contrast, Dynare implementations often use the *level* state T_{t-1}^* (or its level deviation $x_{t-1} \equiv T_{t-1}^* - 1$) as the predetermined state in `ghx` and `ghxx`. This difference matters for mapping B coefficients.

To see this transparently, consider the scalar policy rule (suppressing other states)

$$y_t = A\,\tau_t + \frac{1}{2} B\,\tau_t^2 + O(\tau_t^3), \qquad y \in \{\pi, q\}, \qquad \tau_t = \rho_\tau \tau_{t-1},$$

and suppose Dynare's predetermined state is the level deviation $x_{t-1} \equiv T_{t-1}^* - 1$. Since $\tau_{t-1} =$

$\log(1+x_{t-1}) = x_{t-1} - \frac{1}{2}x_{t-1}^2 + O(x_{t-1}^3)$, one obtains

$$\tau_t = \rho_\tau x_{t-1} - \frac{1}{2}\rho_\tau x_{t-1}^2 + O(x_{t-1}^3), \qquad \Rightarrow \qquad y_t = \rho_\tau A\, x_{t-1} + \frac{1}{2}\big(\rho_\tau^2 B - \rho_\tau A\big)x_{t-1}^2 + O(x_{t-1}^3).$$

Hence, if Dynare reports `ghx` and `ghxx` for the same *log* variable y and the level state T^*, the correct mapping is

$$\texttt{ghx}[y, T^*] = \rho_\tau A, \qquad \texttt{ghxx}[y, T^*, T^*] = \rho_\tau^2 B - \rho_\tau A \quad \Rightarrow \quad B = \frac{\texttt{ghxx}[y, T^*, T^*] + \texttt{ghx}[y, T^*]}{\rho_\tau^2}. \tag{C.46}$$

If instead Dynare reports `ghx`/`ghxx` for the corresponding *level* variable $Y_t = e^{y_t}$ (e.g. $Y \in \{\Pi_H, Q\}$), then $Y_x = y_x$ but $Y_{xx} = y_{xx} + y_x^2$ at the steady state, so

$$B = \frac{\texttt{ghxx}[Y, T^*, T^*] - \texttt{ghx}[Y, T^*]^2 + \texttt{ghx}[Y, T^*]}{\rho_\tau^2}, \qquad Y \in \{\Pi_H, Q\}. \tag{C.47}$$

Equations (C.46)–(C.47) explain why a naive "divide by ρ_τ^2" comparison can yield apparent B mismatches: it typically reflects comparing different curvature objects (log-vs-level variables and log-vs-level states) rather than an economic discrepancy. For verification, the robust objects are the higher-order perturbation terms `ghs2`, `ghxss`, `ghuss` and the implied IRFs, which are invariant to these parameterizations.

Practical note (Dynare state ordering). In Dynare, the columns of `ghx` and `ghxx` are ordered by the internal lead–lag incidence mapping, not by the declaration order in the `.mod` file. A robust way to identify the correct state column is to use Dynare's state index list (e.g. `oo_.dr.order_var(1:M_.nspred)` in recent versions) or, for simple AR(1) states, to locate the column whose coefficient equals the persistence parameter (e.g. ρ_τ for the tariff state).

Second-order mean shifts (`ghs2`). Dynare reports the second-order constant (stochastic-mean) corrections as

$$\texttt{ghs2}[\pi] = 2\bar{\pi}, \qquad \texttt{ghs2}[q] = 2\bar{q}, \tag{C.48}$$

where $(\bar{\pi}, \bar{q})$ are characterized by the linear system (C.21).

Deterministic vs stochastic steady state (SSS-based IRFs). In quantitative work it is common to report uncertainty-shock impulse responses as deviations from the *stochastic* steady state (SSS), i.e. to subtract the mean corrections generated by nonzero level-shock variances. Our analytic derivations, by contrast, use a perturbation expansion around the *deterministic* steady state (DSS). In our volatility experiment, this difference is immaterial at the third order at which uncertainty affects equilibrium. Moreover, our IRFs are defined as deviations from the baseline path with $\varepsilon_{\sigma,0} = 0$ (equivalently, $\sigma_{T,t} \equiv 0$), so the second-order mean shifts captured by $(\bar{\pi}, \bar{q})$ cancel mechanically.

Centering IRFs at the SSS rather than the DSS can matter only through the way stochastic volatility shifts the unconditional second moments that enter equilibrium via Jensen/second-moment terms. But since $\sigma_{T,t}$ has stationary variance $\mathrm{Var}(\sigma_{T,t}) = \eta_\sigma^2/(1-\rho_\sigma^2)$ under (17), a Taylor expansion implies

$$\mathbb{E}\big[e^{2\sigma_{T,t}}\big] = 1 + 2\,\mathrm{Var}(\sigma_{T,t}) + O(\eta_\sigma^4), \tag{C.49}$$

so the unconditional tariff variance differs from its $\sigma_{T,t} \equiv 0$ benchmark only by $\bar{\sigma}_\tau^2\big(\mathbb{E}[e^{2\sigma_{T,t}}] - 1\big) = O(\bar{\sigma}_\tau^2\eta_\sigma^2) = O(4)$. Therefore, the difference between SSS-centered and DSS-centered uncertainty IRFs is itself $O(4)$, beyond the accuracy of our third-order characterization.

Third-order uncertainty terms (`ghxss` and `ghuss`). Our main-text uncertainty coefficients are defined with respect to the *time-t* volatility state (cf. (C.23)):

$$\pi_t = \bar{\pi} + \chi_\pi \sigma_{T,t} + O(4), \qquad q_t = \bar{q} + \chi_q \sigma_{T,t} + O(4).$$

Using the volatility law of motion (17), $\sigma_{T,t} = \rho_\sigma \sigma_{T,t-1} + \eta_\sigma \varepsilon_{\sigma,t}$, gives the equivalent representation

$$\pi_t = \bar{\pi} + (\rho_\sigma\chi_\pi)\sigma_{T,t-1} + (\eta_\sigma\chi_\pi)\varepsilon_{\sigma,t} + O(4), \qquad q_t = \bar{q} + (\rho_\sigma\chi_q)\sigma_{T,t-1} + (\eta_\sigma\chi_q)\varepsilon_{\sigma,t} + O(4).$$

Hence Dynare's third-order policy-function objects satisfy

$$\texttt{ghxss}[\pi, \sigma_T] = 2\rho_\sigma\chi_\pi, \qquad \texttt{ghxss}[q, \sigma_T] = 2\rho_\sigma\chi_q, \tag{C.50}$$

and

$$\texttt{ghuss}[\pi, \varepsilon_\sigma] = 2\eta_\sigma\chi_\pi, \qquad \texttt{ghuss}[q, \varepsilon_\sigma] = 2\eta_\sigma\chi_q. \tag{C.51}$$

In particular, when $\rho_\sigma \neq 0$, (C.50)–(C.51) imply the convenient identity

$$\texttt{ghuss}[y, \varepsilon_\sigma] = \frac{\eta_\sigma}{\rho_\sigma}\,\texttt{ghxss}[y, \sigma_T], \qquad y \in \{\pi, q\}. \tag{C.52}$$

Finally, the period-0 uncertainty-shock impact reported by Dynare is $\Delta\pi_0 = \frac{1}{2}\,\texttt{ghuss}[\pi, \varepsilon_\sigma] = \eta_\sigma\chi_\pi$ and $\Delta q_0 = \frac{1}{2}\,\texttt{ghuss}[q, \varepsilon_\sigma] = \eta_\sigma\chi_q$, which matches the main-text IRF formula (27).

D Economic Interpretation of Curvature and Wedge Signs

Figure C.1 shows that the first-order tariff-exposure coefficients (A_π, A_q) are comparatively stable across the displayed parameter sweeps, while the curvature coefficients (B_π, B_q) move more sharply and are negative near the baseline. This matters because, as shown in (C.24), the uncertainty wedges are

$$u_1 = \beta\bar{\sigma}_\tau^2\left(B_\pi + 3A_\pi^2 + 2A_\pi s_1\right), \qquad u_2 = \bar{\sigma}_\tau^2\left((B_\pi + B_q) - (A_\pi + A_q)^2\right).$$

The variance term in the pricing block, $3A_\pi^2$, is mechanically nonnegative, while the variance term in the asset-market block enters with the opposite sign, $-(A_\pi + A_q)^2$, because the exact closure prices the inverse gross return. The signs of the wedges are therefore governed by two economically distinct objects: the covariance term $A_\pi s_1$ in the pricing block and the curvature terms B_π and $B_\pi + B_q$. This section explains why these objects are negative in the baseline calibration and why Figure C.1 is informative for the sign of uncertainty transmission.

D.1 Pricing Wedge: Why Can $B_\pi < 0$ and $u_1 < 0$?

Start from the sharpest partial-equilibrium benchmark. Suppose an exporter faces CES demand

$$y^*(p, \tau) = D^*(e^\tau p)^{-\gamma},$$

where p is the producer price, e^τ is the export tariff, and real marginal cost is fixed at mc. If the firm can reset its price after observing τ, it chooses p to maximize

$$(p - mc)\, y^*(p, \tau).$$

Because the tariff enters only as the demand shifter $e^{-\gamma\tau}$, the first-order condition delivers the standard constant markup

$$p = \mu\, mc, \qquad \mu \equiv \frac{\gamma}{\gamma - 1},$$

which is independent of τ. Thus a realized tariff shift does not affect the desired price when marginal cost is fixed and demand is isoelastic.

The same conclusion survives in the simplest preset-price environment. Suppose the firm must choose p before τ is realized and solves

$$\max_p \mathbb{E}_t\left[(p - mc)D^*(e^\tau p)^{-\gamma}\right] = \max_p \left\{\mathbb{E}_t[e^{-\gamma\tau}]\,(p - mc)D^* p^{-\gamma}\right\}.$$

The expectation over tariffs factors out of the pricing problem, so the same markup condition obtains. Therefore, in the static CES knife-edge with fixed marginal cost, neither the tariff level nor tariff uncertainty affects the optimal preset price. This is a useful benchmark because it shows that the pricing wedge is not a generic consequence of tariff risk.

The same logic extends to a Calvo reset-price problem if marginal cost is held fixed. The reset price solves

$$P_t^{\#} = \mu\, \frac{\mathbb{E}_t \sum_{j\geq 0} (\beta\theta)^j \Lambda_{t,t+j} Y_{t+j} MC_{t+j} P_{H,t+j}^{\epsilon}}{\mathbb{E}_t \sum_{j\geq 0} (\beta\theta)^j \Lambda_{t,t+j} Y_{t+j} P_{H,t+j}^{\epsilon-1}}.$$

If $MC_{t+j} \equiv mc$ is constant and tariffs only rescale demand Y_{t+j}, the common demand shifters enter numerator and denominator symmetrically and cancel, leaving $P_t^{\#} = \mu mc$. Hence even Calvo price setting delivers no tariff-risk effect in the fixed-cost knife-edge.

These knife-edge examples also suggest the right simplified problem for obtaining a negative

pricing effect. Let $z^*(\tau)$ denote the state-dependent desired price, or desired reset price, and suppose a firm chooses a preset price p to minimize a weighted quadratic loss

$$\min_p \mathbb{E}\Big[w(\tau)\big(p - z^*(\tau)\big)^2\Big].$$

The solution is

$$p^* = \frac{\mathbb{E}[w(\tau)z^*(\tau)]}{\mathbb{E}[w(\tau)]}.$$

This immediately separates the two forces that matter in the full model. First, if $w(\tau)$ is constant and $z^*(\tau)$ is concave, then greater tariff dispersion lowers p^* by Jensen's inequality. This is the simplified preset-price counterpart to a negative B_π. Second, if high-tariff states receive relatively low weights, then $\mathrm{Cov}(w(\tau), z^*(\tau)) < 0$, which lowers p^* even further. This is the simplified counterpart to the covariance term in u_1. A Calvo reset-price problem has exactly this structure: the reset price is a weighted average of future desired prices, so concavity of the desired-price schedule and low weights on high-price states both push the reset price downward.

This benchmark also clarifies what is *not* the right intuition. It is not that firms mechanically "cut prices to avoid demand losses" under uncertainty. Under CES demand with fixed marginal cost, that logic is false: uncertainty does not change the desired markup. The negative pricing wedge appears only once tariffs change desired future prices or the payoff-relevant weights attached to those future states.

Rotemberg can be analyzed just as explicitly. In the reduced complete-markets system, the nonlinear pricing condition is

$$\kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) + \beta\kappa\,\mathbb{E}_t[(\Pi_{H,t+1} - 1)\Pi_{H,t+1}\Omega_{t+1}],$$

with $\Omega_{t+1} \equiv (Y_{t+1}/Y_t)(Q_t/Q_{t+1})$ from (B.8). In the fixed-cost knife-edge, $mc_t \equiv 1$, and if tariffs only rescale current demand without changing future marginal cost or continuation weights, then $\Pi_{H,t} = 1$ solves the equation regardless of tariff uncertainty. So even Rotemberg by itself does not generate a pricing wedge from a pure multiplicative demand shifter.

These no-effect results identify what must break the knife-edge in the full model. To obtain a nonzero pricing wedge, tariffs must make future inflation and the NKPC weighting kernel state-dependent. In the model, tariffs change output composition, marginal cost, and continuation values, so the price-setting problem is no longer just a static markup problem under a multiplicative demand shifter. With preset prices, the firm is solving an insurance problem across future tariff realizations; with Calvo, the reset price becomes a weighted average of state-dependent desired future prices. In Rotemberg, the object is inflation rather than an explicit reset price, but the same forward-looking force is present.

This is seen directly from the second-order Rotemberg NKPC,

$$\hat{\pi}_{H,t} + \frac{3}{2}\hat{\pi}_{H,t}^2 = \kappa_{mc}\left(\widehat{mc}_t + \frac{1}{2}\widehat{mc}_t^2\right) + \beta\mathbb{E}_t\left[\hat{\pi}_{H,t+1} + \frac{3}{2}\hat{\pi}_{H,t+1}^2 + \hat{\pi}_{H,t+1}\log\Omega_{t+1}\right] + O(3),$$

which is equation (F.4) below. Now write

$$\pi_{t+1} = A_\pi \tau_{t+1} + \frac{1}{2} B_\pi \tau_{t+1}^2 + O(\tau_{t+1}^3), \qquad \log \Omega_{t+1} = s_1 \tau_{t+1} + O(\tau_{t+1}^2),$$

and consider mean-zero tariff risk. Then the expectation term contributes

$$\beta \mathbb{E}_t \left[\pi_{t+1} + \frac{3}{2} \pi_{t+1}^2 + \pi_{t+1} \log \Omega_{t+1} \right] \approx \frac{\beta}{2} \left(B_\pi + 3A_\pi^2 + 2A_\pi s_1 \right) \mathrm{Var}_t(\tau_{t+1}).$$

After baseline subtraction and rewriting the conditional variance in terms of the volatility state, this is exactly the source of

$$u_1 = \beta \bar{\sigma}_\tau^2 \left(B_\pi + 3A_\pi^2 + 2A_\pi s_1 \right).$$

The term $3\beta\bar{\sigma}_\tau^2 A_\pi^2$ is therefore the pure Rotemberg-convexity effect. It is absent in the static knife-edge and would not appear in the same form in a Calvo decomposition. Its sign is always positive: if tariffs make inflation more state-dependent, inflation risk raises expected adjustment costs.

The term $2\beta\bar{\sigma}_\tau^2 A_\pi s_1$ is a weighting effect. The Rotemberg NKPC does not average future inflation with fixed weights, but with the kernel Ω_{t+1} in (21). In the baseline calibration, tariffs raise inflation at first order, $A_\pi > 0$, but the same high-tariff states are weak-demand states in which the kernel is low, $s_1 < 0$, so the inflationary states receive relatively little weight in firms' pricing decisions. Formally, the covariance between inflation and the NKPC weighting kernel is negative. Economically, firms care less about the states in which future pricing pressure is strongest.

"Low weight" here means low payoff relevance. In the baseline calibration, the tariff-induced depreciation does not fully offset the export-demand loss, so high-tariff states are states in which firms sell fewer units and the profit cost of having the wrong price is smaller. The kernel Ω_{t+1} captures exactly that idea: future pricing distortions matter less when future sales are low and the continuation value of that state is weak. Hence a high-tariff state can be both an inflationary state and a low-weight state. This is why the covariance term is negative: the states in which firms would like to raise prices more are also the states in which getting the price exactly right matters less for discounted profits.

The remaining term, $\beta\bar{\sigma}_\tau^2 B_\pi$, is a Jensen term coming from the curvature of the deterministic inflation schedule

$$\pi(\tau) = A_\pi \tau + \frac{1}{2} B_\pi \tau^2 + O(\tau^3).$$

If τ is mean-preserving but more dispersed, then

$$\mathbb{E}[\pi(\tau)] \approx \pi(\mathbb{E}[\tau]) + \frac{1}{2} B_\pi \, \mathrm{Var}(\tau),$$

so the sign of the mean-shift contribution is the sign of B_π. In the baseline calibration, $B_\pi < 0$: once export tariffs are already high, the export base is already compressed, so another tariff increase has a smaller marginal effect on demand, marginal cost, and pricing pressure. Monetary policy

and nominal rigidities reinforce this flattening by damping incremental price adjustment as the distortion grows. Hence moving from a moderate tariff to a high tariff raises inflation by less than moving from a moderate tariff to a low tariff reduces it.

It is important to distinguish this curvature force from the Ω_{t+1} force. The kernel Ω_{t+1} explains the covariance term in u_1; it does not by itself explain why $B_\pi < 0$. The sign of B_π is a statement about the deterministic mapping $\pi(\tau)$ and therefore about the full equilibrium response of inflation to tariff levels. In our calibration, that curvature is shaped importantly by the terms-of-trade block: a high tariff already depresses exports and induces a large relative-price adjustment, so the additional inflation required from an even higher tariff becomes smaller. Through the joint (π, q) system, this flattening of the terms-of-trade response feeds back into inflation curvature.

Quantitatively, the overall negative NKPC wedge is not driven by Rotemberg convexity. In the baseline calibration, the decomposition is approximately

$$u_1^{\text{Rot}} = +0.00041, \qquad u_1^{\text{Cov}} = -0.00104, \qquad u_1^{\text{Curv}} = -0.00133,$$

so both negative terms matter and the curvature term is somewhat larger in absolute value. The most intuitive summary is therefore: tariff uncertainty lowers the certainty-equivalent pricing pressure because high-inflation states are low-value states *and* because the inflation schedule is concave. A Calvo version would keep exactly these two ingredients, even though it would not generate the same standalone $3A_\pi^2$ term.

D.2 UIP Wedge: Why Can $B_\pi + B_q < 0$ and $u_2 < 0$?

Start from a one-period bond-arbitrage condition. A Home investor compares a Home bond that pays one unit of Home currency tomorrow with a Foreign bond whose Home-currency payoff is $e^{\Delta_{t+1}}$, where

$$\Delta_{t+1} \equiv (q_{t+1} - q_t) + \pi_{t+1}$$

is next period's nominal depreciation. Under the Euler-consistent complete-markets closure, the relevant pricing object is not $\log \mathbb{E}_t[e^{\Delta_{t+1}}]$ but

$$-\log \mathbb{E}_t[e^{-\Delta_{t+1}}],$$

because the exact asset-pricing condition prices the inverse gross return. The key point is that this wedge is still not an ad hoc risk-premium shock and not a symptom of incomplete markets. It is the Jensen gap between the log of the expected inverse return and the expected log return inside a no-arbitrage condition. Economically, a Foreign bond is valuable because it pays more in Home currency when the Home currency depreciates. To see how tariff risk moves this object, look at the state-by-state tariff-level mapping rather than the uncertainty IRF. At first order,

$$\Delta(\tau) \approx (A_\pi + A_q)\tau.$$

In the baseline calibration, $A_\pi + A_q > 0$, so a higher export tariff is locally a nominal-depreciation state. Since a high export tariff is also a bad external-demand state for Home, the Foreign bond pays more in Home currency exactly in those bad states. That is why foreign currency exposure acts as a hedge. The asset-market wedge measures how tariff uncertainty changes the certainty-equivalent value of that hedge. If Δ_{t+1} is risky, then

$$-\log \mathbb{E}_t[e^{-\Delta_{t+1}}] \approx \mathbb{E}_t[\Delta_{t+1}] - \frac{1}{2}\operatorname{Var}_t(\Delta_{t+1}),$$

which immediately shows why a volatility shock creates an asset-market wedge even when the mean tariff path is unchanged. In a symmetric two-state example with $\Delta_{t+1} \in \{-d, +d\}$, mean log depreciation is zero but $\mathbb{E}_t[e^{-\Delta_{t+1}}] = \cosh(d) > 1$, so the Euler-consistent log pricing object is negative. Under inverse-return pricing, gross inverse returns are convex in depreciation, and uncertainty therefore lowers the certainty-equivalent depreciation term through Jensen's inequality alone.

The same point also clarifies the current exchange-rate adjustment. Rearranging the reduced asset-market condition gives

$$\phi_\pi \pi_t + \phi_y y_t + \psi_t - q_t = -\log \mathbb{E}_t[e^{-q_{t+1}-\pi_{t+1}}].$$

Holding the expected future distribution fixed, a more negative Jensen term on the right-hand side must be matched by a lower current q_t. Since a fall in q_t is a real appreciation in our convention, the pure variance component now pushes toward current *real appreciation*. This distinction matters for interpretation: the underlying object is still nominal depreciation, Δ_{t+1}, but the contemporaneous adjustment margin in the reduced system is the terms of trade q_t. Hence the wedge is nominal in origin, while its contemporaneous implementation works through real exchange-rate adjustment.

This is the negative variance component in u_2. In our notation, the first-order tariff sensitivity of nominal depreciation is $A_\pi + A_q$, so the Jensen component is the mechanically negative term $-\bar{\sigma}_\tau^2(A_\pi + A_q)^2$. If that were the whole story, tariff uncertainty would always act like a negative reduced-form asset-market wedge. The intuition is straightforward: under inverse-return pricing, more tariff uncertainty raises the expected inverse payoff relevant for the Euler-consistent bond-pricing condition, so the Home currency must be stronger today to offset that extra hedge value. In the reduced system this appears as downward pressure on q_t, that is, on current real appreciation.

But UIP also inherits the curvature of the deterministic depreciation schedule. Write

$$\Delta(\tau) = (A_\pi + A_q)\tau + \frac{1}{2}(B_\pi + B_q)\tau^2.$$

If tariff innovations are mean zero, then a second-order expansion gives

$$\mathbb{E}[\Delta(\tau)] \approx \frac{1}{2}(B_\pi + B_q)\operatorname{Var}(\tau), \qquad \operatorname{Var}(\Delta(\tau)) \approx (A_\pi + A_q)^2 \operatorname{Var}(\tau).$$

Hence

$$-\log \mathbb{E}[e^{-\Delta(\tau)}] \approx \frac{1}{2}\Big((B_\pi + B_q) - (A_\pi + A_q)^2\Big)\,\mathrm{Var}(\tau),$$

which is exactly the intuition behind the decomposition in (E.5). The term $B_\pi + B_q$ is therefore the curvature of nominal depreciation with respect to the tariff level. If that mapping is concave, more tariff dispersion lowers average depreciation and reinforces the negative variance term.

A simple way to see why $B_q < 0$ is to abstract from the nominal block and ask how much real depreciation is needed to stabilize output as tariffs rise. Using the static reduced-form map,

$$y(\pi, q, \tau) \approx \theta_q q - \theta_\tau \tau + \frac{1}{2}\mathcal{C}_y \tau^2,$$

and holding inflation fixed for the moment, a constant-output schedule satisfies

$$0 \approx \theta_q q(\tau) - \theta_\tau \tau + \frac{1}{2}\mathcal{C}_y \tau^2,$$

so

$$q(\tau) \approx \frac{\theta_\tau}{\theta_q}\tau - \frac{\mathcal{C}_y}{2\theta_q}\tau^2.$$

Hence if $\mathcal{C}_y > 0$, the real-depreciation schedule is concave and its curvature is negative. This stripped-down calculation makes the sign transparent: once export tariffs are already high, the export base is already small, so an additional tariff increase destroys fewer marginal exports and calls for less extra real depreciation than at low tariff levels.

The full coefficient system says the same thing. In (C.17), the second row that governs B_q is driven by $-\phi_y \mathcal{C}_y$. In the baseline calibration, $\mathcal{C}_y > 0$ and $\phi_y > 0$, so the Taylor rule's output-stabilization motive pushes toward $B_q < 0$: when tariffs are already high, the policy problem requires less incremental terms-of-trade adjustment because the marginal output damage has flattened out. The full model then adds feedback from the nominal block through the matrix inversion in (C.17), but the sign intuition is the same.

Why is $B_\pi + B_q$ negative near the baseline? The same diminishing-exposure force that makes $q(\tau)$ concave also makes $\pi(\tau)$ concave. Once export tariffs are already high, the export base is small, so another tariff increase has a weaker marginal effect on competitiveness, output, and relative prices. The required extra depreciation to stabilize demand is therefore smaller when tariffs are already high. Monetary policy reinforces this flattening because the Taylor rule leans against large nominal adjustments as the distortion grows. Since nominal depreciation is the sum $\pi(\tau) + q(\tau)$, concavity in these two schedules makes the composite depreciation mapping locally concave as well. This is the economic content of $B_\pi + B_q < 0$: the upper tail of the tariff distribution does not buy much extra depreciation because the exchange-rate schedule has already flattened out there.

This also clarifies the sign reversal. The variance term says that tariff uncertainty increases the value of the foreign bond as a hedge. The curvature term says that this hedge becomes less effective exactly in the very high-tariff states where it is most needed, because depreciation responds less there. So the negative result does not come from "more risk" by itself. It comes from a deterioration

in the state-contingent payoff profile of the hedge: a mean-preserving increase in tariff dispersion puts more weight in the bad tail, but the exchange rate no longer provides proportionally more insurance in those states.

When the curvature term is also negative, it reinforces the negative Jensen component and makes

$$u_2 = \bar{\sigma}_\tau^2\Big((B_\pi + B_q) - (A_\pi + A_q)^2\Big) < 0.$$

Then the implied asset-market adjustment pushes toward current real appreciation through a lower q_t. In the baseline calibration, both components are negative, so the contractionary sign does not rely on a curvature reversal of an otherwise positive Jensen benchmark. Rather, the curvature term amplifies the negative variance term inherited from the Euler-consistent inverse-return pricing object. This is why the asset-market channel now exerts stronger appreciation pressure on the terms of trade and output than in the old reverse-order UIP formulation.

D.3 Level Shocks versus Uncertainty Shocks

It is useful to end this section by separating two distinct comparative statics that are easy to conflate. A tariff *level* shock asks what happens when the expected tariff itself rises. A tariff *uncertainty* shock asks what happens when the mean tariff is unchanged, but the conditional distribution of future tariffs becomes more dispersed.

For a tariff-level shock, the relevant objects are the first-order coefficients in the deterministic decision rules:

$$\pi(\tau) = A_\pi \tau + \frac{1}{2} B_\pi \tau^2 + O(\tau^3), \qquad q(\tau) = A_q \tau + \frac{1}{2} B_q \tau^2 + O(\tau^3).$$

Hence the local effect of a small increase in the tariff level is governed by the slopes A_π and A_q:

$$\frac{\partial \pi}{\partial \tau}\Big|_{\tau=0} = A_\pi, \qquad \frac{\partial q}{\partial \tau}\Big|_{\tau=0} = A_q.$$

In the baseline calibration, $A_\pi > 0$ and $A_q > 0$, so a surprise export-tariff increase is locally inflationary and depreciating at first order; see (B.28). This is the comparative static along the deterministic schedule itself.

For an uncertainty shock, by contrast, the tariff level remains at its baseline path and only conditional second moments move. Along the volatility experiment, $\tau_t \equiv 0$ on the equilibrium path and

$$\mathbb{E}_t[\tau_{t+1}^2] = \bar{\sigma}_\tau^2 e^{2\sigma_{T,t}}$$

changes with the volatility state. The relevant question is therefore not the slope of $\pi(\tau)$ at zero, but how the certainty-equivalent value of the pricing and asset-market blocks responds to more dispersion. In the pricing block, the uncertainty wedge is

$$u_1 = \beta \bar{\sigma}_\tau^2 \left(B_\pi + 3A_\pi^2 + 2A_\pi s_1\right),$$

and in the UIP block it is

$$u_2 = \bar{\sigma}_\tau^2 \left((B_\pi + B_q) - (A_\pi + A_q)^2 \right).$$

Thus uncertainty responses depend on curvature and state weighting, not just on the first-order slopes.

This is why a model can easily have a tariff level that is locally inflationary but tariff uncertainty that is deflationary. A simple scalar analogy is

$$f(\tau) = a\tau - \frac{b}{2}\tau^2, \qquad a, b > 0.$$

Then $f'(0) = a > 0$, so a small increase in the level raises f. But for a mean-preserving spread in τ around the same mean,

$$\mathbb{E}[f(\tau)] \approx f(\mathbb{E}[\tau]) - \frac{b}{2}\,\mathrm{Var}(\tau),$$

so more dispersion lowers the certainty-equivalent value because the function is concave. The same logic applies here. In our baseline calibration, the level effect is locally inflationary because $A_\pi > 0$, while the uncertainty effect is deflationary because the curvature term B_π and the covariance term $A_\pi s_1$ are negative and dominate the positive Rotemberg-convexity term.

The most useful interpretation is therefore the following. A tariff-level shock moves the economy *along* the deterministic schedule $(\pi(\tau), q(\tau))$. A tariff-uncertainty shock does not move the mean tariff path; instead it changes the *dispersion* of future states, so the equilibrium response is governed by Jensen effects and by how the pricing equations reweight those states. This is why the signs of the level responses and the uncertainty responses need not coincide.

E Proof of Proposition 2

Roadmap. The proof constructs two higher-order wedge processes induced by the volatility innovation: an asset-market uncertainty wedge that enters the log linearized UIP restriction through conditional second moments (a variance component plus a curvature component), and an NKPC uncertainty wedge that enters the Rotemberg NKPC through conditional second moments. It then shows that, in the linearized system, these wedges are observationally equivalent to a pair of first-order shocks (productivity and a UIP wedge). Under the common-persistence restriction (Assumption 2), the resulting spanning weights are time-invariant.

Proof. Work under Assumption 2 and consider the third-order perturbation around the deterministic steady state.

Step 1 (Asset-market uncertainty wedge). Define the log "depreciation-plus-inflation" term

$$\Delta_{t+1} \equiv (q_{t+1} - q_t) + \pi_{t+1}.$$

Taking logs of (22) yields

$$\phi_\pi \pi_t + \phi_y y_t + \psi_t = -\log \mathbb{E}_t \left[e^{-\Delta_{t+1}} \right]. \tag{E.1}$$

By a cumulant expansion around the deterministic steady state,

$$-\log \mathbb{E}_t\left[e^{-\Delta_{t+1}}\right] = \mathbb{E}_t[\Delta_{t+1}] - \frac{1}{2}\operatorname{Var}_t(\Delta_{t+1}) + \frac{1}{6}\kappa_{3,t}(\Delta_{t+1}) + O(4), \tag{E.2}$$

where $\kappa_{3,t}(X) \equiv \mathbb{E}_t[(X - \mathbb{E}_t[X])^3]$ is the conditional third cumulant (third central moment). Under Assumption 1 and the volatility experiment, $\kappa_{3,t}(\Delta_{t+1}) = O(4)$, so third-order corrections in (E.2) are governed by conditional second moments.

Under the volatility experiment (Definition 1), the time-t conditional dispersion of Δ_{t+1} is driven by the next tariff innovation $\varepsilon_{\tau,t+1}$ scaled by $\bar{\sigma}_\tau e^{\sigma_{T,t}}$. Let ξ_Δ denote the first-order loading of Δ_{t+1} on $\varepsilon_{\tau,t+1}$ and let ζ_Δ denote the second-order loading on $\varepsilon^2_{\tau,t+1}$, both evaluated at the deterministic steady state:

$$\Delta_{t+1} = \xi_\Delta\, \bar{\sigma}_\tau e^{\sigma_{T,t}} \varepsilon_{\tau,t+1} + \frac{1}{2}\,\zeta_\Delta\, \bar{\sigma}^2_\tau e^{2\sigma_{T,t}} \varepsilon^2_{\tau,t+1} + O(3).$$

Then $\operatorname{Var}_t(\Delta_{t+1}) = \xi^2_\Delta \bar{\sigma}^2_\tau e^{2\sigma_{T,t}} + O(3)$, so the baseline-subtracted variance (Jensen) component is

$$\widetilde{\Delta}^{UIP}_t \equiv -\frac{1}{2}\operatorname{Var}_t(\Delta_{t+1}) + \frac{1}{2}\operatorname{Var}_t(\Delta_{t+1})\Big|_{\sigma_{T,t}=0} = -\frac{1}{2}\xi^2_\Delta \bar{\sigma}^2_\tau\left(e^{2\sigma_{T,t}} - 1\right) + O(4). \tag{E.3}$$

The curvature term in (E.2) arises from $\mathbb{E}_t[\Delta_{t+1}]$ through the second-order component in the expansion of Δ_{t+1} above. Since $\mathbb{E}[\varepsilon^2_{\tau,t+1}] = 1$, the baseline-subtracted curvature component is

$$\Delta^{UIP,curv}_t \equiv \mathbb{E}_t[\Delta_{t+1}] - \mathbb{E}_t[\Delta_{t+1}]\Big|_{\sigma_{T,t}=0} = \frac{1}{2}\,\zeta_\Delta\,\bar{\sigma}^2_\tau\left(e^{2\sigma_{T,t}} - 1\right) + O(4). \tag{E.4}$$

Combining (E.3)–(E.4), the asset-market uncertainty wedge that shifts the *linearized* log UIP restriction is

$$\Delta^{UIP}_t \equiv \Delta^{UIP,curv}_t + \widetilde{\Delta}^{UIP}_t = \frac{1}{2}\left(\zeta_\Delta - \xi^2_\Delta\right)\bar{\sigma}^2_\tau\left(e^{2\sigma_{T,t}} - 1\right) + O(4). \tag{E.5}$$

In the notation of (23), $\xi_\Delta = A_\pi + A_q$ and $\zeta_\Delta = B_\pi + B_q$, so (E.5) implies $\Delta^{UIP}_t = u_2\,\sigma_{T,t} + O(4)$ with u_2 given by the asset-market component of (29). Since $e^{2\sigma_{T,t}} - 1 = 2\sigma_{T,t} + O(4)$ and $\sigma_{T,t}$ follows (17), Δ^{UIP}_t inherits AR(1) dynamics with persistence ρ at third order under Assumption 2.

Step 2 (NKPC uncertainty wedge). Rewrite the nonlinear NKPC (21) in the form

$$\kappa\,\Psi(\Pi_{H,t}) = (\epsilon - 1)\left(mc_t - 1\right) + \beta\kappa\,\mathbb{E}_t[\Psi(\Pi_{H,t+1})\,\Omega_{t+1}]\,, \qquad \Psi(\Pi) \equiv (\Pi - 1)\Pi,$$

where $\Omega_{t+1} \equiv (Y_{t+1}/Y_t)\exp(q_t - q_{t+1})$ in logs coincides with the intertemporal term in (21). Expanding around the deterministic steady state and collecting third-order terms yields the linear NKPC augmented by an additive wedge:

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + \tilde{\kappa} q_t - c_\tau \tau_t - c_a a_t + \Delta^{NK}_t + O(4), \tag{E.6}$$

where $(\tilde{\kappa}, c_\tau, c_a)$ are defined in (C.9) and Δ^{NK}_t collects the baseline-subtracted conditional second-moment terms generated by the expectation in the Rotemberg NKPC. Under the volatility ex-

periment, the only time-varying component of these conditional second moments is again driven by $\varepsilon_{\tau,t+1}$ scaled by $\bar{\sigma}_\tau e^{\sigma_{T,t}}$, so Δ_t^{NK} is proportional to $\bar{\sigma}_\tau^2(e^{2\sigma_{T,t}}-1)+O(4)$ and therefore inherits persistence ρ at third order under Assumption 2.

Step 3 (Equivalence to two first-order shocks). Consider the first-order linearized system for (π_t, q_t), given by (C.9)–(C.10). In that system, productivity a_t enters the NKPC through $-c_a a_t$, while the UIP wedge ψ_t enters the UIP condition additively. Hence the additive wedges $(\Delta_t^{NK}, \Delta_t^{UIP})$ constructed in Steps 1–2 are observationally equivalent (at first order) to the pair of fictitious shocks

$$a_t^{eq} \equiv -\frac{\Delta_t^{NK}}{c_a}, \qquad \psi_t^{eq} \equiv -\Delta_t^{UIP},$$

in the sense that substituting $(a_t, \psi_t) = (a_t^{eq}, \psi_t^{eq})$ into the linearized system reproduces the same paths for (π_t, q_t) (and thus, through the static block, for any equilibrium variable) up to $O(4)$ terms.

Step 4 (Spanning and time-invariant weights). Linearity of the first-order equilibrium mapping implies that the impulse response of any endogenous variable x_t to the combined disturbance (a_t^{eq}, ψ_t^{eq}) is the sum of the impulse responses to each component. Under Assumption 2, the equivalent processes (a_t^{eq}, ψ_t^{eq}) and the basis shocks (a_t, ψ_t) share the same persistence ρ, so all impulse responses share a common geometric time profile. It follows that there exist constants $(\lambda_A, \lambda_\psi)$ such that (30) holds for all horizons.

To obtain closed-form weights, match the impact uncertainty responses of (π_t, q_t)—namely $(\Delta\pi_0, \Delta q_0) = \eta_\sigma(\chi_\pi, \chi_q)$ from (26)–with the span of the *first-order* impact responses to the two basis shocks, as in (31). See Proposition 2 (and (32)) for the explicit closed forms. When the two basis directions are linearly independent, the weights are uniquely determined. □

F Proof of Proposition 3

Section F.1 develops the QC welfare approximation and the associated targeting criterion under complete markets, and uses it to prove Proposition 3. Section F.2 then discusses why the same strict-PPI-targeting benchmark does not generally extend to incomplete markets and records a bond-economy closure useful for implementing Ramsey policy under incomplete markets.

F.1 Quadratic–cubic welfare loss and targeting criterion (complete markets)

This appendix expands on the Quadratic–Cubic (QC) approach used to characterize Ramsey policy in our environment (Gross and Hansen, 2021). The QC approach combines a third-order approximation of welfare with a second-order approximation of equilibrium constraints.

Roadmap. We first provide a compact QC representation of welfare and the Rotemberg NKPC in gap form. We then derive the associated (history-dependent) targeting criterion and use it to prove Proposition 3. The key takeaway is that, in the complete-markets PCP economy around the

flexible-price benchmark allocation, strict stabilization of PPI inflation eliminates all policy-relevant distortions. Uncertainty affects welfare only through terms independent of monetary policy ("t.i.p.").

Gaps and marginal-cost wedge. Let $\hat{x}_t \equiv \log(X_t/\bar{X})$ denote log deviations from the deterministic steady state, and define (log) PPI inflation $\hat{\pi}_{H,t} \equiv \log \Pi_{H,t}$. Let $\hat{y}_t^{\text{eff}}$ denote the benchmark flexible-price output allocation under complete markets for the same exogenous states, and define the output gap $\tilde{y}_t \equiv \hat{y}_t - \hat{y}_t^{\text{eff}}$. Around that flexible-price benchmark, log real marginal cost in PPI units can be written as

$$\widehat{mc}_t = \omega_y \, \tilde{y}_t + u_t + O(2), \qquad \omega_y > 0, \tag{F.1}$$

where u_t is a (potential) cost-push wedge capturing components of marginal cost that are not proportional to the output gap. In our complete-markets PCP economy, $u_t = 0$.

Welfare loss. Define the (period) welfare loss $\mathcal{L}_t^{(3)}$ as minus the third-order approximation to the welfare gain relative to the benchmark flexible-price allocation, ignoring terms independent of monetary policy. A QC approximation yields

$$\mathcal{L}_t^{(3)} = \frac{1}{2}\Big(\lambda_\pi \hat{\pi}_{H,t}^2 + \lambda_y \tilde{y}_t^2\Big) + \frac{1}{6}\Big(\lambda_{\pi 3} \hat{\pi}_{H,t}^3 + \lambda_{y3} \tilde{y}_t^3 + \lambda_{y\pi} \tilde{y}_t \hat{\pi}_{H,t}^2\Big) + \text{t.i.p.} + O(4), \tag{F.2}$$

where "t.i.p." denotes terms independent of monetary policy (including efficient-allocation risk terms). Under the normalization $\bar{C} = \bar{Y} = \bar{L} = \bar{G} = 1$, a convenient calibration-consistent choice is

$$\lambda_\pi = \kappa, \qquad \lambda_y = \sigma + \varphi, \qquad \lambda_{\pi 3} = 3\lambda_\pi, \qquad \lambda_{y3} = \varphi(\varphi - 1), \tag{F.3}$$

and $\lambda_{y\pi}$ is typically zero (or negligible) in our benchmark.[1]

Second-order NKPC. The Rotemberg NKPC admits the second-order approximation

$$\hat{\pi}_{H,t} + \frac{3}{2}\hat{\pi}_{H,t}^2 = \kappa_{mc}\left(\widehat{mc}_t + \frac{1}{2}\widehat{mc}_t^2\right) + \beta \mathbb{E}_t\left[\hat{\pi}_{H,t+1} + \frac{3}{2}\hat{\pi}_{H,t+1}^2 + \hat{\pi}_{H,t+1}\log \Omega_{t+1}\right] + O(3), \tag{F.4}$$

where $\kappa_{mc} = (\epsilon - 1)/\kappa$ is defined in (20) and $\Omega_{t+1} \equiv (Y_{t+1}/Y_t)\exp(q_t - q_{t+1})$ is the intertemporal term inside the discounted Rotemberg NKPC. Using (F.1) with $u_t = 0$ expresses the NKPC as a second-order constraint in $(\hat{\pi}_{H,t}, \tilde{y}_t)$.

Targeting criterion. We now derive the targeting rule implied by the QC problem. Throughout this derivation we work in the complete-markets PCP economy around the flexible-price benchmark allocation, so $u_t = 0$ in (F.1). The Ramsey planner chooses an allocation (equivalently, a sequence $\{\hat{\pi}_{H,t}, \tilde{y}_t\}_{t\geq 0}$) to minimize expected discounted welfare loss (F.2) subject to the (approximated) NKPC (F.4).

[1] Equation (F.2) follows from (i) the Rotemberg resource cost, which generates a quadratic and cubic loss in inflation, and (ii) the second-order Taylor expansion of household utility around the benchmark flexible-price allocation, which generates a quadratic (and possibly cubic) loss in the output gap.

Step 1: LQ (first-order) targeting rule. Linearizing (F.4) and using (F.1) delivers the standard first-order NKPC in gap form,

$$\hat{\pi}_{H,t} = \kappa_{mc}\omega_y\,\tilde{y}_t + \beta\mathbb{E}_t[\hat{\pi}_{H,t+1}] + O(2). \tag{F.5}$$

Consider the associated LQ Ramsey problem that minimizes the quadratic loss $\frac{1}{2}(\lambda_\pi\hat{\pi}^2_{H,t} + \lambda_y\tilde{y}_t^2)$ subject to (F.5). Form the Lagrangian

$$\mathcal{J} = \mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\left\{\frac{1}{2}\Big(\lambda_\pi\hat{\pi}^2_{H,t} + \lambda_y\tilde{y}_t^2\Big) + \phi_t\Big(\hat{\pi}_{H,t} - \kappa_{mc}\omega_y\,\tilde{y}_t - \beta\mathbb{E}_t[\hat{\pi}_{H,t+1}]\Big)\right\}, \tag{F.6}$$

where $\{\phi_t\}$ are Lagrange multipliers. Using the law of iterated expectations and the fact that ϕ_t is $\mathcal{F}_t$-measurable, the term involving $\mathbb{E}_t[\hat{\pi}_{H,t+1}]$ can be rewritten as

$$\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\Big(-\beta\phi_t\mathbb{E}_t[\hat{\pi}_{H,t+1}]\Big) = \mathbb{E}_0\sum_{t=1}^{\infty}\beta^t\Big(-\phi_{t-1}\hat{\pi}_{H,t}\Big).$$

Differentiating (F.6) with respect to $(\hat{\pi}_{H,t}, \tilde{y}_t)$ therefore yields the LQ first-order conditions (with the convention $\phi_{-1} \equiv 0$):

$$\lambda_\pi\hat{\pi}_{H,t} + \phi_t - \phi_{t-1} = 0, \tag{F.7}$$

$$\lambda_y\tilde{y}_t - \kappa_{mc}\omega_y\,\phi_t = 0. \tag{F.8}$$

Eliminating the multiplier using (F.8) gives $\phi_t = (\lambda_y/(\kappa_{mc}\omega_y))\tilde{y}_t$. Substituting into (F.7) yields the familiar history-dependent targeting criterion:

$$\hat{\pi}_{H,t} = -\frac{\lambda_y}{\lambda_\pi\kappa_{mc}\omega_y}\big(\tilde{y}_t - \tilde{y}_{t-1}\big). \tag{F.9}$$

Step 2: QC (second-order) targeting rule. Next, return to the QC objective (F.2) and the second-order NKPC (F.4). Let $\Gamma_{t+1} \equiv \log\Omega_{t+1}$ and use (F.1) (with $u_t = 0$) to express marginal cost in terms of $\tilde{y}_t$. Up to second order, it is convenient to write the NKPC constraint as

$$\mathcal{F}_t \equiv \hat{\pi}_{H,t} + \frac{3}{2}\hat{\pi}^2_{H,t} - \kappa_{mc}\left(\omega_y\tilde{y}_t + \frac{1}{2}\omega_y^2\tilde{y}_t^2\right) - \beta\mathbb{E}_t\left[\hat{\pi}_{H,t+1} + \frac{3}{2}\hat{\pi}^2_{H,t+1} + \hat{\pi}_{H,t+1}\Gamma_{t+1}\right] = O(3). \tag{F.10}$$

The corresponding QC Lagrangian is

$$\mathcal{J}^{QC} = \mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\Big\{\mathcal{L}_t^{(3)} + \phi_t\,\mathcal{F}_t\Big\}. \tag{F.11}$$

Differentiating (F.11) and using the same time-shifting logic as in Step 1 delivers the QC Euler–Lagrange conditions. In the *general* Rotemberg/QC problem, these conditions contain additional terms involving Γ_t and expectations of future multipliers. For completeness, keeping these terms

yields the (benchmark-consistent) QC FOCs

$$\lambda_\pi \hat{\pi}_{H,t} + \frac{\lambda_{\pi 3}}{2}\hat{\pi}_{H,t}^2 + \frac{\lambda_{y\pi}}{3}\tilde{y}_t\hat{\pi}_{H,t} + (1+3\hat{\pi}_{H,t})\phi_t - (1+3\hat{\pi}_{H,t}+\Gamma_t)\phi_{t-1} = O(3), \tag{F.12}$$

$$\lambda_y \tilde{y}_t + \frac{\lambda_{y3}}{2}\tilde{y}_t^2 + \frac{\lambda_{y\pi}}{6}\hat{\pi}_{H,t}^2 - \kappa_{mc}\omega_y(1+\omega_y\tilde{y}_t)\phi_t + \beta\mathbb{E}_t[\kappa_{mc}\omega_y\,\phi_{t+1}\Gamma_{t+1}] = O(3), \tag{F.13}$$

where $\Gamma_t \equiv \log\Omega_t$ and we adopt the convention $\phi_{-1} \equiv 0$. In our benchmark complete-markets PCP economy, $u_t = 0$ and the Ramsey allocation features $\hat{\pi}_{H,t} = \tilde{y}_t = 0$ at first order (divine coincidence). Hence the Ramsey multipliers are themselves second-order objects, so any terms that multiply the multipliers by first-order objects such as Γ_t are of order $O(3)$ along the Ramsey path and can be absorbed into the remainder.[2] Under this benchmark-relevant simplification, the QC FOCs reduce to

$$\lambda_\pi \hat{\pi}_{H,t} + \frac{\lambda_{\pi 3}}{2}\hat{\pi}_{H,t}^2 + \frac{\lambda_{y\pi}}{3}\tilde{y}_t\hat{\pi}_{H,t} + (1+3\hat{\pi}_{H,t})(\phi_t - \phi_{t-1}) = O(3), \tag{F.14}$$

$$\lambda_y \tilde{y}_t + \frac{\lambda_{y3}}{2}\tilde{y}_t^2 + \frac{\lambda_{y\pi}}{6}\hat{\pi}_{H,t}^2 - \kappa_{mc}\omega_y(1+\omega_y\tilde{y}_t)\phi_t = O(3). \tag{F.15}$$

Equation (F.15) implies $\phi_t = (\lambda_y/(\kappa_{mc}\omega_y))\tilde{y}_t + O(2)$ and therefore

$$\phi_t - \phi_{t-1} = \frac{\lambda_y}{\kappa_{mc}\omega_y}\Delta\tilde{y}_t + O(2). \tag{F.16}$$

Substituting (F.16) into (F.14) and dividing by λ_π yields the following convenient representation of the QC targeting criterion:

$$\hat{\pi}_{H,t} + \frac{\lambda_{\pi 3}}{2\lambda_\pi}\hat{\pi}_{H,t}^2 + \frac{\lambda_{y\pi}}{3\lambda_\pi}\tilde{y}_t\hat{\pi}_{H,t} = -\frac{\lambda_y}{\lambda_\pi\kappa_{mc}\omega_y}\Delta\tilde{y}_t + O(3), \tag{F.17}$$

where $\Delta\tilde{y}_t \equiv \tilde{y}_t - \tilde{y}_{t-1}$.

Proof. Work in the complete-markets PCP economy around the benchmark flexible-price allocation.

Step 1 (Divine coincidence under PPI targeting). Suppose policy implements strict PPI targeting, $\Pi_{H,t} = 1$ for all t (equivalently, $\hat{\pi}_{H,t} = 0$). Then Rotemberg adjustment costs are zero and the NKPC (A.15) reduces to

$$0 = (\epsilon - 1)(mc_t - 1),$$

so $mc_t = 1$ state-by-state. Around the benchmark flexible-price allocation, real marginal cost in PPI units is proportional to the output gap. To see this, use $mc_t = \exp(q_t + \varphi y_t - (1+\varphi)a_t)$ together with the first-order static mapping $y_t = \theta_q q_t - \theta_\tau \tau_t + O(2)$ to obtain

$$\widehat{mc}_t = \Gamma_q q_t - \Gamma_\tau \tau_t - \Gamma_a a_t + O(2), \qquad (\Gamma_q, \Gamma_\tau, \Gamma_a) \text{ as in (C.8)}.$$

[2] Formally, with $u_t = 0$ the first-order Ramsey solution sets $(\hat{\pi}_{H,t}, \tilde{y}_t) = (0,0)$ for all t, implying $\phi_t = O(2)$ in the perturbation around the benchmark flexible-price steady state. Since $\Gamma_t = O(1)$, products such as $\phi_{t-1}\Gamma_t$ are $O(3)$ and therefore do not enter the second-order targeting criterion.

Under flexible prices, the benchmark allocation satisfies $\widehat{mc}_t^{\text{eff}} = 0$, implying $q_t^{\text{eff}} = (\Gamma_\tau \tau_t + \Gamma_a a_t)/\Gamma_q + O(2)$. Hence

$$\tilde{y}_t = \theta_q\big(q_t - q_t^{\text{eff}}\big) + O(2), \qquad \widehat{mc}_t = \Gamma_q\big(q_t - q_t^{\text{eff}}\big) + O(2) = \underbrace{\frac{\Gamma_q}{\theta_q}}_{\omega_y > 0} \tilde{y}_t + O(2).$$

Therefore $mc_t = 1$ implies $\tilde{y}_t = 0$ locally: strict PPI targeting delivers both zero inflation and a zero output gap (divine coincidence).

Step 2 (Ramsey optimality of strict PPI targeting). The Ramsey planner maximizes expected lifetime utility subject to the equilibrium constraints in Section 2, taking foreign processes as exogenous. By the QC approximation, the policy-relevant component of the welfare loss is summarized by (F.2). Since the allocation with $\hat{\pi}_{H,t} = 0$ is feasible and, by Step 1, implies $\tilde{y}_t = 0$, it sets the distortion terms in (F.2) to zero, leaving only terms independent of monetary policy. Because $\lambda_\pi > 0$ and $\lambda_y > 0$, any feasible allocation with $(\hat{\pi}_{H,t}, \tilde{y}_t) \neq (0, 0)$ yields a strictly higher loss locally. Hence Ramsey policy coincides with strict PPI targeting under complete markets and PCP.

Step 3 (Uncertainty shocks). Under the volatility experiment (Definition 1) the tariff level τ_t is unchanged and only the conditional variance of future tariff innovations varies. Under Ramsey/PPI targeting, $\hat{\pi}_{H,t} = 0$ and $\tilde{y}_t = 0$ state-by-state, so the allocation coincides with the benchmark flexible-price allocation, which depends on exogenous *levels* but not on the volatility state. Hence tariff uncertainty shocks do not move allocations (up to third order), though they can affect welfare through risk terms evaluated at the benchmark allocation. □

F.2 Incomplete markets: Ramsey policy and the breakdown of strict PPI targeting

The main text focuses on complete markets, where strict PPI targeting is Ramsey optimal. Under incomplete markets, risk sharing is absent and consumption is no longer pinned by the real exchange rate. Consequently, the planner faces an additional intertemporal stabilization motive. A convenient way to summarize this distinction is that the QC welfare loss generically includes a consumption-gap term:

Proposition F.1 (QC welfare loss under incomplete markets (gap form))**.**

$$\mathcal{L}_t^{IM,(3)} = \frac{1}{2}\Big(\lambda_\pi \hat{\pi}_{H,t}^2 + \lambda_y \tilde{y}_t^2 + \lambda_c \tilde{c}_t^2\Big) + \textit{higher-order terms} + \textit{t.i.p.} + O(4), \tag{F.18}$$

where $\tilde{c}_t \equiv \hat{c}_t - \hat{c}_t^{\text{eff}}$ *is the consumption gap relative to the benchmark flexible-price allocation under incomplete markets.*[3]

[3]The presence of $\tilde{c}_t$ reflects that, without international risk sharing, consumption dynamics are pinned down by intertemporal optimality (Euler) conditions rather than by a static risk-sharing relationship.

Under the normalization $\bar{C} = \bar{G} = 1$, a convenient benchmark coefficient is

$$\lambda_c = \sigma - 1, \tag{F.19}$$

so the quadratic consumption-gap term is absent under log utility ($\sigma = 1$) and strictly positive for $\sigma > 1$.

Remark F.1 (Log-utility benchmark). When $\sigma = 1$, the quadratic consumption-gap term drops out of (F.18). Intuitively, with log utility the marginal utility of consumption is proportional to $1/C_t$, so (locally) the welfare cost of intertemporal consumption fluctuations is weaker. In this case strict PPI stabilization can remain a particularly accurate approximation even under incomplete markets, though higher-order terms (and external-balance considerations in a bond economy) can still generate small Ramsey deviations.

At third order, the consumption block also generates a cubic term $\frac{1}{6}\lambda_{c3}\tilde{c}_t^3$ with $\lambda_{c3} = -(\sigma - 1)\lambda_c$ under the same normalization. Even when strict PPI targeting forces $mc_t = 1$ and eliminates price-adjustment costs, it does not in general eliminate the consumption gap; hence $\hat{\pi}_{H,t} = 0$ need not be Ramsey optimal. Section G provides the corresponding positive-equilibrium restrictions under incomplete markets and shows that stochastic volatility introduces an additional Euler/Jensen wedge.

Proposition F.2 (Strict PPI targeting is generally not Ramsey optimal under incomplete markets)**.** *Consider the incomplete-markets economy and the QC welfare representation in Proposition F.1. Suppose $\lambda_c > 0$ and that the consumption gap $\tilde{c}_t$ is locally policy-relevant (i.e., feasible policy perturbations can change $\tilde{c}_t$ around the benchmark flexible-price allocation). If strict PPI targeting, $\hat{\pi}_{H,t} = 0$ for all t, implies a nonzero consumption gap for some histories, then strict PPI targeting cannot be Ramsey optimal under incomplete markets.*

Proof. Under the QC approximation, Ramsey policy minimizes the expected discounted sum of the policy-relevant loss terms in (F.18) subject to the incomplete-markets equilibrium constraints. Under strict PPI targeting, $\hat{\pi}_{H,t} = 0$ for all t, so the inflation-loss component is identically zero. If strict PPI targeting implies $\tilde{c}_t \neq 0$ for some histories, then the consumption-gap term $\frac{1}{2}\lambda_c\tilde{c}_t^2$ is strictly positive at those histories. Since $\tilde{c}_t$ is locally policy-relevant by assumption, there exists a feasible local policy perturbation that changes $\tilde{c}_t$ at second order. With $\lambda_c > 0$, such a perturbation can (locally) reduce the objective relative to strict PPI targeting unless strict PPI already minimizes the consumption-gap component. Therefore strict PPI targeting cannot, in general, coincide with the Ramsey optimum under incomplete markets. □

Remark F.2. Proposition F.2 formalizes the sense in which "divine coincidence" breaks under incomplete markets: even if policy neutralizes the marginal-cost/inflation distortion, an intertemporal distortion remains because consumption is pinned down by an Euler equation (rather than by static international risk sharing).

F.2.1 A bond-economy incomplete-markets closure (for Ramsey)

To make the Ramsey problem under incomplete markets well-defined under perturbation (i.e., to avoid non-stationary net foreign assets around a balanced-trade steady state), we adopt a standard *bond economy* closure with a debt-elastic premium. Let B_t denote Home-currency nominal bond holdings (in zero net supply) paying gross return R_t (the policy instrument), and let B_t^* denote holdings of a Foreign-currency nominal bond paying the exogenous gross return R_t^* times a stationarity premium. Let the CPI-deflated Home-currency value of net foreign assets be

$$b_t \equiv \frac{\mathcal{E}_t B_t^*}{P_t}, \tag{F.20}$$

where $b_t > 0$ means net foreign assets and $b_t < 0$ means net foreign debt. We model the stationarity device by assuming the effective foreign gross return is $R_t^* \exp(-\psi_b b_t)$ with $\psi_b > 0$, so borrowing becomes more expensive as b_t falls (debt rises).

The two Euler equations are

$$1 = \beta R_t \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \right], \tag{F.21}$$

$$1 = \beta R_t^* e^{-\psi_b b_t} \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} \right]. \tag{F.22}$$

Dividing (F.22) by (F.21) yields UIP with an endogenous premium. At first order (around $\bar{b} = 0$), this takes the familiar form

$$\hat{r}_t - \hat{r}_t^* = \mathbb{E}_t[\Delta e_{t+1}] - \psi_b b_t + O(2), \tag{F.23}$$

where $\hat{r}_t \equiv \log(R_t/\bar{R})$, $\Delta e_{t+1} \equiv \log(\mathcal{E}_{t+1}/\mathcal{E}_t)$, and hats denote log deviations from the deterministic steady state. The law of motion for b_t follows from the Home household budget constraint and goods-market clearing; the key point is that $\psi_b > 0$ induces mean reversion in b_t and therefore a well-defined stochastic steady state.

F.2.2 Why strict PPI targeting is not Ramsey optimal under incomplete markets

Even though strict PPI targeting enforces $mc_t = 1$ and eliminates Rotemberg price-adjustment costs, it does not in general eliminate the *consumption gap* $\tilde{c}_t$ in (F.18), because consumption is determined intertemporally by (F.21)–(F.22). The planner therefore faces a genuine stabilization trade-off among inflation, the output gap, and intertemporal allocations.

Implementation note. The bond-economy Ramsey problem is most conveniently implemented by treating $(\hat{r}_t, \hat{\pi}_{H,t})$ as endogenous and imposing the Euler blocks (F.21)–(F.22) together with the external-budget law for b_t and the remaining equilibrium conditions. This avoids relying on an ad hoc UIP truncation and makes the role of the stationarity device transparent.

G Incomplete Markets Extension

Under incomplete markets, risk sharing (12) is dropped. The static demand/price-index block is unchanged, but consumption is no longer pinned by Q_t via (13). The baseline nonlinear closure is therefore the *exact bond-economy Euler block* from Section F.2.1: the domestic Euler equation, the foreign-bond Euler equation, and the external-budget law for bond holdings. At first order, dividing the two Euler equations yields UIP with an endogenous premium, but that UIP object is only a linearized implication of the bond economy, not a separate exact nonlinear restriction.

Roadmap. We first state the exact bond-economy incomplete-markets representation (Proposition G.1) and highlight how it differs from the complete-markets reduction. We then isolate the additional Euler/precautionary-savings Jensen channel that becomes active under tariff uncertainty (Proposition G.2) and show how it modifies the spanning logic and the Ramsey benchmark. The remainder of the section reports the resulting quantitative comparison, including Figure G.1.

$$\Phi_t \equiv (1-\alpha) + \alpha (T_{F,t} Q_t)^{1-\gamma}, \qquad G_t \equiv \Phi_t^{1/(1-\gamma)}, \tag{G.1}$$

$$C_{H,t} = (1-\alpha) G_t^{\gamma} C_t, \qquad X_t \equiv C^*_{H,t} = \alpha \left(\frac{Q_t}{T_t^*} \right)^{\gamma} \quad (\text{with } C_t^* \equiv 1), \tag{G.2}$$

$$Y_t = \frac{C_{H,t} + X_t}{1 - \frac{\kappa}{2}(\Pi_{H,t} - 1)^2}, \qquad mc_t = \chi \, G_t \, C_t^{\sigma} \, Y_t^{\varphi} \, A_t^{-(1+\varphi)}, \tag{G.3}$$

$$\Pi_{C,t} = \Pi_{H,t} \cdot \frac{G_t}{G_{t-1}}. \tag{G.4}$$

Proposition G.1 (Exact bond-economy incomplete-markets system)**.** *Under incomplete markets and PCP, equilibrium is summarized by the static block (G.1)–(G.4), the nonlinear NKPC*

$$\kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) + \beta \kappa \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \left(\frac{C_t}{C_{t+1}} \right)^{\sigma} \frac{G_t}{G_{t+1}} \right], \tag{G.5}$$

and the exact bond-economy asset-market block

$$1 = \beta R_t \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \right],$$

$$1 = \beta R_t^* e^{-\psi_b b_t} \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} \right],$$

together with the external-budget law for b_t from Section F.2.1. Here $(G_t, Y_t, mc_t, \Pi_{C,t})$ are defined by (G.1)–(G.4) as functions of $(\Pi_{H,t}, Q_t, C_t, T_{F,t}, T_t^*, A_t)$, R_t is determined by the policy rule in the main text, and the displayed Euler equations are precisely (F.21)–(F.22). Dividing the foreign-bond Euler equation by the domestic Euler equation yields the first-order UIP-with-premium relation (F.23), but there is no separate exact nonlinear UIP equation in the baseline bond economy.

Relative to complete markets, four differences matter. First, risk sharing no longer pins

consumption, so the domestic Euler equation becomes a non-redundant equilibrium condition. Second, the stochastic discount factor in the NKPC no longer collapses to Q_t/Q_{t+1}. Third, bond holdings become an additional state variable that matters for determinacy and the stochastic steady state. Fourth, at third order, stochastic volatility enters the domestic Euler equation through an additional Jensen/precautionary-savings wedge. As a result, the exact two-channel decomposition in Proposition 2 becomes approximate (though often quantitatively tight), and strict PPI targeting no longer neutralizes tariff uncertainty in general (Corollary G.1).

Remark G.1 (First-order UIP shorthand for comparison). For linear comparison exercises it is convenient to summarize the foreign-asset block by the implied first-order UIP-with-premium relation (F.23) and to keep the reduced-form ψ_t shock as a basis shock in spanning exercises below. This is only a first-order comparison device. The nonlinear incomplete-markets baseline is always the bond-economy block (F.21)–(F.22) together with the external-budget law for b_t.

Proposition G.2 (Incomplete markets add an Euler Jensen channel)**.** *Consider the exact bond-economy incomplete-markets economy in Proposition G.1 under the volatility experiment in Definition 1. Let $\hat{c}_t \equiv \log(C_t/\bar{C})$, $\hat{r}_t \equiv \log(R_t/\bar{R})$, and $\hat{\pi}_{C,t} \equiv \log(\Pi_{C,t})$. Then the third-order perturbation of the domestic Euler equation (F.21) admits the representation*

$$\hat{c}_t = \mathbb{E}_t[\hat{c}_{t+1}] - \frac{1}{\sigma}\big(\hat{r}_t - \mathbb{E}_t[\hat{\pi}_{C,t+1}]\big) - \frac{1}{\sigma}\,\widetilde{\Delta}^E_t + O(4), \tag{G.6}$$

where the Euler Jensen wedge is the baseline-subtracted conditional-variance correction induced by the log-expectation in the Euler equation,

$$\widetilde{\Delta}^E_t \equiv \frac{1}{2}\operatorname{Var}_t\big(X_{t+1}\big) - \frac{1}{2}\operatorname{Var}_t\big(X_{t+1}\big)\Big|_{\sigma_{T,t}=0} + O(4), \tag{G.7}$$

with

$$X_{t+1} \equiv -\sigma(\hat{c}_{t+1} - \hat{c}_t) + \hat{r}_t - \hat{\pi}_{C,t+1}. \tag{G.8}$$

Under stochastic volatility, the Euler Jensen wedge $\widetilde{\Delta}^E_t$ is generically nonzero and provides an additional precautionary-savings transmission channel for tariff uncertainty. Consequently, the exact two-shock spanning result of Proposition 2 need not hold under incomplete markets.

Proof. Rewrite the domestic Euler equation (F.21) as a log-expectation. Using $\bar{R} = 1/\beta$ and $\bar{\Pi}_C = 1$, define X_{t+1} as in (G.8) so that

$$1 = \mathbb{E}_t\big[e^{X_{t+1}}\big] \qquad \Longleftrightarrow \qquad 0 = \log \mathbb{E}_t\big[e^{X_{t+1}}\big]\,.$$

Applying the same third-order cumulant expansion as in (E.2) gives

$$0 = \mathbb{E}_t[X_{t+1}] + \frac{1}{2}\operatorname{Var}_t(X_{t+1}) + \frac{1}{6}\kappa_{3,t}(X_{t+1}) + O(4).$$

Under Assumption 1 and the volatility experiment, $\kappa_{3,t}(X_{t+1}) = O(4)$, so it does not affect the third-

order representation. Since $\hat{r}_t$ and $\hat{c}_t$ are $\mathcal{F}_t$-measurable, $\mathbb{E}_t[X_{t+1}] = -\sigma(\mathbb{E}_t[\hat{c}_{t+1}] - \hat{c}_t) + \hat{r}_t - \mathbb{E}_t[\hat{\pi}_{C,t+1}]$ and $\mathrm{Var}_t(X_{t+1}) = \mathrm{Var}_t(-\sigma\hat{c}_{t+1} - \hat{\pi}_{C,t+1})$. Rearranging yields

$$\hat{c}_t = \mathbb{E}_t[\hat{c}_{t+1}] - \frac{1}{\sigma}\big(\hat{r}_t - \mathbb{E}_t[\hat{\pi}_{C,t+1}]\big) - \frac{1}{2\sigma}\,\mathrm{Var}_t(X_{t+1}) + O(4).$$

Baseline subtraction (relative to the $\sigma_{T,t} \equiv 0$ path) yields (G.6)–(G.7).

To emphasize the economic content, note the variance decomposition

$$\mathrm{Var}_t(X_{t+1}) = \sigma^2\,\mathrm{Var}_t(\hat{c}_{t+1}) + \mathrm{Var}_t(\hat{\pi}_{C,t+1}) + 2\sigma\;\mathrm{Cov}_t(\hat{c}_{t+1}, \hat{\pi}_{C,t+1}). \tag{G.9}$$

In addition, CPI inflation satisfies $\hat{\pi}_{C,t} = \hat{\pi}_{H,t} + (\hat{g}_t - \hat{g}_{t-1})$ from (A.11). Under constant import tariffs and around the deterministic steady state, $\hat{g}_t = \alpha\hat{q}_t + O(2)$, so the Euler wedge inherits the terms-of-trade component through CPI inflation.

Under the volatility experiment, the only time-varying component of (G.9) is driven by the next tariff innovation scaled by $\bar{\sigma}_\tau e^{\sigma_{T,t}}$. Hence $\widetilde{\Delta}^E_t$ is proportional to $\bar{\sigma}^2_\tau(e^{2\sigma_{T,t}} - 1) + O(4)$ and is generically nonzero. □

Remark G.2 (Closed-form wedge under the volatility experiment). Under the volatility experiment, let (ξ_c, ξ_π, ξ_q) denote the first-order loadings (evaluated at the deterministic steady state) such that

$$\begin{aligned}
\hat{c}_{t+1} &= \xi_c\,\bar{\sigma}_\tau e^{\sigma_{T,t}}\varepsilon_{\tau,t+1} + O(2),\\
\hat{\pi}_{H,t+1} &= \xi_\pi\,\bar{\sigma}_\tau e^{\sigma_{T,t}}\varepsilon_{\tau,t+1} + O(2),\\
\hat{q}_{t+1} &= \xi_q\,\bar{\sigma}_\tau e^{\sigma_{T,t}}\varepsilon_{\tau,t+1} + O(2).
\end{aligned}$$

Using the identities

$$\hat{\pi}_{C,t+1} = \hat{\pi}_{H,t+1} + (\hat{g}_{t+1} - \hat{g}_t), \qquad \hat{g}_{t+1} = \alpha\hat{q}_{t+1} + O(2)$$

under constant import tariffs, the CPI-inflation loading is $\xi_{\pi_C} = \xi_\pi + \alpha\xi_q$. Then the variance term in (G.7) satisfies

$$\widetilde{\Delta}^E_t = \frac{1}{2}(\sigma\xi_c + \xi_{\pi_C})^2\bar{\sigma}^2_\tau\big(e^{2\sigma_{T,t}} - 1\big) + O(4). \tag{G.10}$$

Since $(\sigma\xi_c + \xi_{\pi_C})^2 \geq 0$ and $e^{2\sigma_{T,t}} - 1 \geq 0$ after a positive volatility innovation, $\widetilde{\Delta}^E_t$ is (locally) nonnegative and therefore acts as a precautionary-savings wedge that lowers current consumption in (G.6), all else equal.

Corollary G.1 (Incomplete markets: strict PPI targeting does not eliminate uncertainty effects)**.** *Consider the exact bond-economy incomplete-markets economy in Proposition G.1 and the tariff-volatility experiment in Definition 1. Suppose monetary policy implements strict PPI targeting, $\Pi_{H,t} = 1$ for all t. Then the Rotemberg NKPC implies $mc_t = 1$ state-by-state, eliminating price-adjustment distortions. Nevertheless, if the Euler Jensen wedge $\widetilde{\Delta}^E_t$ in (G.6) is nonzero for some histories (generically under stochastic volatility), consumption and therefore output and*

external variables depend on the volatility state at third order. Hence, unlike under complete markets (Proposition 3), strict PPI targeting does not in general make allocations invariant to tariff uncertainty under incomplete markets.

Proposition G.3 (Three-channel spanning under incomplete markets)**.** *Impose Assumption 2 and consider the exact bond-economy incomplete-markets economy in Proposition G.1. Introduce a (classical) discount-factor shock by letting the effective discount factor be $\beta_t \equiv \beta e^{d_t}$, where $d_t = \rho d_{t-1} + \sigma_d \varepsilon_{d,t}$ is independent of the other innovations. Let $\mathrm{IRF}_t^{\mathrm{Dem}}(x)$ denote the first-order impulse response to a one-standard-deviation innovation in d_t.*

Then there exist constants $(\lambda_A, \lambda_\psi, \lambda_D)$ such that for every horizon $t \geq 0$,

$$\mathrm{IRF}_t^{\mathrm{unc}}(x) = \lambda_A \, \mathrm{IRF}_t^{\mathrm{Prod}}(x) + \lambda_\psi \, \mathrm{IRF}_t^{\psi}(x) + \lambda_D \, \mathrm{IRF}_t^{\mathrm{Dem}}(x) + O(4). \tag{G.11}$$

In general $\lambda_D \neq 0$, so a two-shock decomposition need not be exact under incomplete markets.

Proof. The proof follows the same wedge-equivalence logic as Proposition 2, with one additional wedge coming from the Euler equation.

Step 1 (Euler Jensen wedge ⇒ discount-factor shock). Proposition G.2 shows that, at third order, the Euler equation can be written as the usual log-linear Euler equation plus the additive term $-(1/\sigma)\widetilde{\Delta}_t^E$. At first order, a discount-factor shock d_t enters the log Euler equation additively in the same way (it shifts the log SDF). Hence the wedge path is observationally equivalent, to first order, to a fictitious discount-factor shock process $d_t^{eq} \equiv \widetilde{\Delta}_t^E$.

Step 2 (Uncertainty generates three wedges). Under the volatility experiment, the exact bond-economy equilibrium conditions generate three distinct higher-order wedges: (i) an NKPC wedge Δ_t^{NK} from conditional second moments inside the discounted Rotemberg NKPC term, (ii) an asset-market/UIP wedge Δ_t^{UIP} from conditional second moments in the foreign-bond Euler equation, equivalently in the implied first-order UIP relation (E.5), and (iii) the Euler Jensen wedge $\widetilde{\Delta}_t^E$ from the log-expectation term in the domestic Euler equation. Therefore the third-order equilibrium under tariff uncertainty is observationally equivalent, up to $O(4)$ terms, to the equilibrium of the *linearized* incomplete-markets system under the three fictitious first-order shocks $(a_t^{eq}, \psi_t^{eq}, d_t^{eq})$ that replicate these wedge processes.

Step 3 (Common persistence ⇒ constant weights). Under Assumption 2, each wedge process is proportional to $\bar{\sigma}_\tau^2(e^{2\sigma_{T,t}} - 1) = 2\bar{\sigma}_\tau^2\sigma_{T,t} + O(4)$ and hence inherits AR(1) dynamics with persistence ρ at third order. The basis shocks (a_t, ψ_t, d_t) also have persistence ρ, so all first-order impulse responses share a common geometric time profile. Linearity then implies the constant-weight spanning representation (G.11). □

Remark G.3 (A closed-form expression for the demand weight). In the normalization of Proposition G.3, the discount-factor shock has impact σ_d. Since the equivalent shock path is $d_t^{eq} = \widetilde{\Delta}_t^E$, a

convenient choice of the demand weight is

$$\lambda_D \equiv \frac{\widetilde{\Delta}_0^E}{\sigma_d}. \tag{G.12}$$

Using (G.10) gives an economically transparent expression for λ_D in terms of risk aversion, openness, and the first-order sensitivity of $(\hat{c}_{t+1}, \hat{\pi}_{H,t+1}, \hat{q}_{t+1})$ to tariff innovations.

G.1 Robustness: incomplete markets (quantitative illustration)

Figure G.1 provides a quantitative comparison of tariff-uncertainty impulse responses under complete and incomplete markets, illustrating the mechanisms formalized above. The updated baseline comparison makes the distinction sharp. Under complete markets and strict PPI targeting, output, consumption, exports, and the terms of trade are zero on impact up to numerical error, though the policy rate still adjusts to satisfy the asset-pricing condition. Under incomplete markets and strict PPI targeting, output rises by about 0.0084%, consumption falls by about 0.0068%, exports rise by about 0.0259%, and the terms of trade rise by about 0.0172% on impact (equivalently, Home experiences a small real depreciation). Under incomplete markets with a PPI-based Taylor rule, output rises by about 0.0079%, consumption falls by about 0.0219%, exports rise by about 0.0455%, and the terms of trade rise by about 0.0304%, reflecting the additional interaction between the Euler/Jensen wedge and the policy rule.

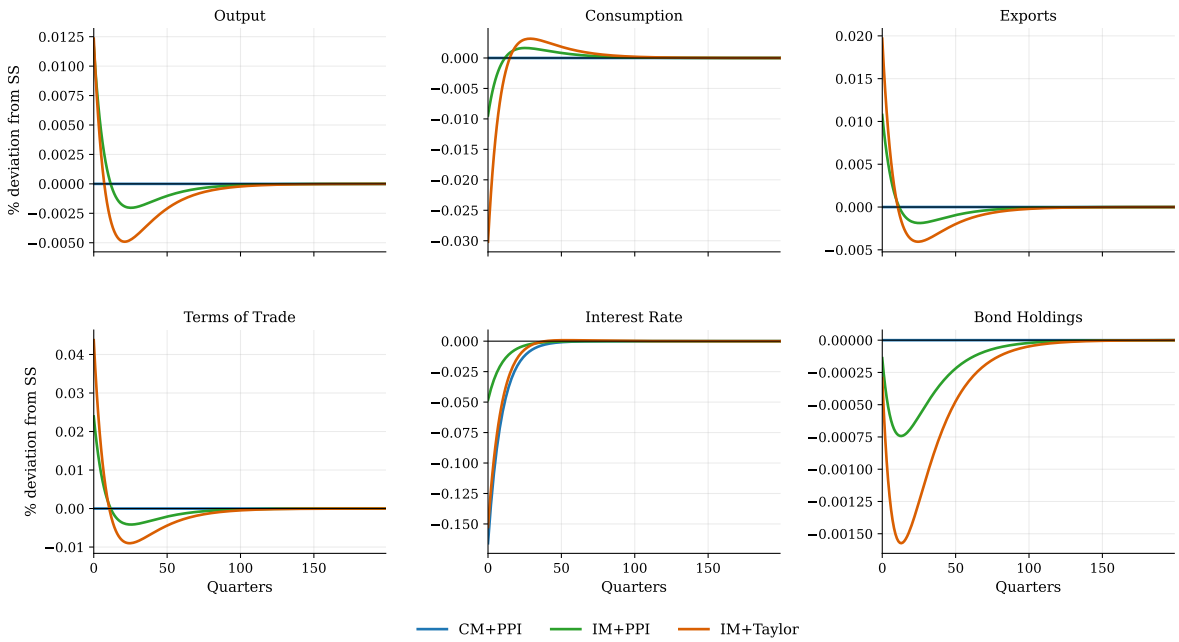

Figure G.1: Tariff uncertainty shock under complete vs incomplete markets.

Notes: The figure plots third-order impulse responses to a one-standard-deviation innovation in tariff volatility at $t = 0$ (Definition 1), using the baseline calibration in Table A.1. Blue (CM+PPI) imposes complete markets and strict PPI targeting; green (IM+PPI) imposes incomplete markets under strict PPI targeting; orange (IM+Taylor) imposes incomplete markets under a PPI-based Taylor rule. All variables are reported as percent deviations from steady state (log points multiplied by 100), except bond holdings which are shown in levels. Under complete markets, strict PPI targeting neutralizes uncertainty shocks in real allocations (Proposition 3). Under incomplete markets, uncertainty effects are generically nonzero even under strict PPI targeting because an Euler/precautionary-savings (Jensen) channel becomes active at third order (Proposition G.2).

H Policy-Rule Comparison: PPI vs CPI Taylor Rules

This appendix records the algebra behind the comparison between a PPI-based Taylor rule and a CPI-based Taylor rule in the complete-markets economy. Because CPI inflation embeds exchange-rate movements through the CPI–PPI wedge, a CPI-based rule responds to terms-of-trade movements in addition to domestic price-setting distortions, even at first order.

Roadmap. We proceed in three steps. First, we derive the Euler-consistent asset-market closure implied by a CPI-based Taylor rule. Second, we relate CPI to PPI inflation via the CPI–PPI wedge and show how CPI targeting introduces additional state dependence (through q_{t-1} at first order) in the minimal log system. Third, we solve the first-order dynamics under CPI targeting and report the quantitative comparison used in the paper.

H.1 CPI-based Taylor rule and the asset-market closure

The main text imposes the PPI-based Taylor rule (14). A CPI-based alternative replaces PPI inflation with CPI inflation:

$$\frac{R_t}{\bar{R}} = \Pi_{C,t}^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y}. \tag{H.1}$$

Combining (H.1) with the Euler-consistent asset-pricing relation (11) and the terms-of-trade identity $\mathcal{E}_t/\mathcal{E}_{t+1} = Q_t/(Q_{t+1}\Pi_{H,t+1})$ (under constant Foreign prices) yields the closure

$$1 = \Pi_{C,t}^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y} e^{\psi_t} \mathbb{E}_t \left[\frac{Q_t}{Q_{t+1}\Pi_{H,t+1}} \right], \tag{H.2}$$

which differs from (B.8) only by the inflation measure inside the Taylor rule.

H.2 CPI inflation vs PPI inflation

The CPI–PPI wedge $G_t \equiv P_t/P_{H,t}$ implies the identity (A.11):

$$\Pi_{C,t} = \Pi_{H,t} \cdot \frac{G_t}{G_{t-1}} \qquad \Longleftrightarrow \qquad \pi_{C,t} = \pi_t + (g_t - g_{t-1}), \tag{H.3}$$

where $\pi_{C,t} \equiv \log \Pi_{C,t}$, $\pi_t \equiv \log \Pi_{H,t}$, and $g_t \equiv \log G_t$. Under complete markets, G_t is a static function of $(Q_t, T_{F,t})$ through (A.8). Around the deterministic steady state with $T_{F,t} \equiv 1$, a first-order expansion gives

$$g_t = \alpha q_t + O(2), \tag{H.4}$$

so CPI inflation inherits an exchange-rate/terms-of-trade component:

$$\pi_{C,t} = \pi_t + \alpha(q_t - q_{t-1}) + O(2). \tag{H.5}$$

H.3 Log form in the minimal two-equation system (complete markets)

Under complete markets, the log form of the asset-market closure under a CPI-based Taylor rule is obtained by replacing π_t with $\pi_{C,t}$ in (22):

$$1 = \exp\big(\phi_\pi \pi_{C,t} + \phi_y y_t + \psi_t\big) \mathbb{E}_t \big[\exp\big(-(q_{t+1} - q_t) - \pi_{t+1} \big) \big], \tag{H.6}$$

where $\pi_{C,t}$ is given by (H.3) and $y_t = y(\pi_t, q_t, \tau_t)$ is pinned down by the static block (Section B). Relative to the PPI-based system (21)–(22), the CPI-based rule introduces the predetermined term g_{t-1} (equivalently q_{t-1} at first order) through $\pi_{C,t}$.

H.4 First-order solution under a CPI-based Taylor rule

At first order, the NKPC (C.9) is unchanged. Using (H.5) in the first-order log asset-market restriction associated with (H.6) yields

$$\phi_\pi \pi_t + (\phi_y \theta_q + \alpha\phi_\pi) q_t - (\phi_y \theta_\tau)\tau_t - \alpha\phi_\pi q_{t-1} + \psi_t = \mathbb{E}_t[(q_{t+1} - q_t) + \pi_{t+1}], \tag{H.7}$$

so the CPI-based rule introduces the endogenous predetermined component q_{t-1} into the linearized system.

H.4.1 Homogeneous dynamics

Suppress all exogenous shocks and set $(\tau_t, a_t, \psi_t) \equiv 0$. Guess a homogeneous solution of the form $\pi_t = B_\pi q_{t-1}$ and $q_t = B_q q_{t-1}$. Substituting into (C.9) and (H.7) yields

$$B_\pi = \frac{\tilde{\kappa} B_q}{1 - \beta B_q}, \tag{H.8}$$

and B_q must satisfy the cubic equation

$$\beta B_q^3 - \Big[(\tilde{\kappa} + 1) + \beta\big(1 + \phi_y \theta_q + \alpha\phi_\pi\big)\Big] B_q^2 + \Big[\tilde{\kappa}\phi_\pi + \big(1 + \phi_y \theta_q + \alpha\phi_\pi\big) + \beta\alpha\phi_\pi\Big] B_q - \alpha\phi_\pi = 0. \tag{H.9}$$

The relevant root is the stable one with $|B_q| < 1$; then B_π follows from (H.8).

H.4.2 Shock loadings and impulse responses

For an AR(1) tariff state with persistence ρ_τ, seek a solution of the form

$$\pi_t = A_\pi \tau_t + B_\pi q_{t-1}, \qquad q_t = A_q \tau_t + B_q q_{t-1}.$$

Matching coefficients on τ_t in (C.9) and (H.7) yields the 2×2 linear system

$$\begin{pmatrix} 1 - \beta\rho_\tau & -(\tilde{\kappa} + \beta B_\pi) \\ \phi_\pi - \rho_\tau & 1 + \phi_y \theta_q + \alpha\phi_\pi - \rho_\tau - B_q - B_\pi \end{pmatrix} \begin{pmatrix} A_\pi \\ A_q \end{pmatrix} = \begin{pmatrix} -c_\tau \\ \phi_y \theta_\tau \end{pmatrix}, \tag{H.10}$$

so given (B_π, B_q), Cramer's rule delivers closed forms for (A_π, A_q).

For a one-time surprise tariff shock at $t = 0$ with $\tau_t = \rho_\tau^t \bar{\sigma}_\tau$ and initial condition $q_{-1} = 0$, the terms of trade follows the recursion

$$q_t = A_q \tau_t + B_q q_{t-1}, \qquad q_{-1} = 0. \tag{H.11}$$

If $B_q \neq \rho_\tau$, this recursion has the closed-form solution

$$q_t = A_q \bar{\sigma}_\tau \frac{\rho_\tau^{t+1} - B_q^{t+1}}{\rho_\tau - B_q}. \tag{H.12}$$

Inflation is then $\pi_t = A_\pi \tau_t + B_\pi q_{t-1}$, and all remaining variables follow from the static mappings in Section B.4 together with the exchange-rate identity (B.15).

H.4.3 Deterministic tariff news under a CPI-based Taylor rule

Consider the deterministic news path in (B.36) with initial condition $q_{-1} = 0$. For $t \geq k$ (at and after implementation), the law of motion is the same as in the surprise-shock case:

$$q_t = A_q \tau_t + B_q q_{t-1}, \qquad \pi_t = A_\pi \tau_t + B_\pi q_{t-1}, \qquad \tau_t = \rho_\tau^{t-k} \bar{\sigma}_\tau.$$

Defining $n \equiv t - k \geq 0$ and assuming $B_q \neq \rho_\tau$, this yields the closed form

$$q_{k+n} = B_q^{n+1} \, q_{k-1} + A_q \bar{\sigma}_\tau \, \frac{\rho_\tau^{n+1} - B_q^{n+1}}{\rho_\tau - B_q}, \qquad n \geq 0, \tag{H.13}$$

with $\pi_{k+n} = A_\pi \rho_\tau^n \bar{\sigma}_\tau + B_\pi q_{k+n-1}$. Therefore, to characterize the entire news-shock path it suffices to determine the pre-implementation state q_{k-1}.

During the anticipation window $0 \leq t < k$, one has $\tau_t \equiv 0$. Under perfect foresight, (C.9) and (H.7) reduce to

$$\pi_t = \tilde{\kappa} q_t + \beta \pi_{t+1}, \tag{H.14}$$

$$\phi_\pi \pi_t + (1 + \phi_y \theta_q + \alpha \phi_\pi) q_t - \alpha \phi_\pi q_{t-1} = \pi_{t+1} + q_{t+1}. \tag{H.15}$$

Define the augmented state vector $s_t \equiv (\pi_t, q_t, q_{t-1})^\top$. Since q_t is the predetermined component in s_{t+1}, (H.14)–(H.15) imply a backward recursion of the form

$$s_t = \mathbf{H} \, s_{t+1}, \qquad 0 \leq t < k, \tag{H.16}$$

where

$$\mathbf{H} \equiv \begin{pmatrix} \beta & 0 & \tilde{\kappa} \\ 0 & 0 & 1 \\ \dfrac{\phi_\pi \beta - 1}{\alpha \phi_\pi} & -\dfrac{1}{\alpha \phi_\pi} & \dfrac{\phi_\pi \tilde{\kappa} + 1 + \phi_y \theta_q + \alpha \phi_\pi}{\alpha \phi_\pi} \end{pmatrix}. \tag{H.17}$$

At implementation, the policy functions imply

$$s_k = \bar{\sigma}_\tau \underbrace{\begin{pmatrix} A_\pi \\ A_q \\ 0 \end{pmatrix}}_{v_0} + q_{k-1} \underbrace{\begin{pmatrix} B_\pi \\ B_q \\ 1 \end{pmatrix}}_{v_1}. \tag{H.18}$$

Iterating (H.16) gives $s_0 = \mathbf{H}^k s_k$, whose third component equals the initial condition q_{-1}. Imposing

$q_{-1} = 0$ therefore pins down q_{k-1}:

$$q_{k-1} = -\bar{\sigma}_\tau \frac{e_3^\top \mathbf{H}^k v_0}{e_3^\top \mathbf{H}^k v_1}, \qquad e_3 \equiv (0,0,1)^\top. \tag{H.19}$$

Finally, the anticipation path follows from $s_t = \mathbf{H}^{k-t} s_k$ for $0 \leq t < k$ and the post-implementation path from (H.13). All remaining variables are recovered using the static mappings in Section B.4 and the exchange-rate identity (B.15).

I Additional Figures

This appendix collects additional figures that complement the main-text quantitative illustrations.

I.1 All model variables under tariff uncertainty

Figure I.1 reports the impulse responses of all endogenous variables under the PPI-based Taylor rule and complete markets; the close agreement between the closed-form predictions and Dynare order-3 simulations confirms the accuracy of the analytical characterization.

The top row reports real quantities. Output, consumption, and exports all rise on impact: the NKPC uncertainty wedge is deflationary, inducing monetary easing that depreciates the real exchange rate and boosts competitiveness. Imports decline because the real depreciation raises the relative price of Foreign goods. The second row shows price and exchange-rate variables. Both PPI and CPI inflation fall, reflecting the deflationary NKPC wedge; the terms of trade rise (real depreciation); and the nominal depreciation rate is negative (appreciation) because the asset-market/Jensen wedge dominates the nominal side even as PPI deflation dominates the real side. The bottom row confirms that the policy rate drops sharply (Taylor-rule easing in response to deflation), marginal cost rises (driven by the terms-of-trade movement), and both the real and nominal wage respond accordingly—the real wage rises with output and the terms of trade, while the nominal wage falls as the cumulating deflation dominates.

A notable feature is that marginal cost rises even though PPI inflation falls. Decomposing the impact NKPC, $\pi_{H,0} = \beta E_0[\pi_{H,1}] + \kappa_{mc}\widehat{mc}_0 + \Delta_0^J$, the contemporaneous cost-push term ($\kappa_{mc}\widehat{mc}_0 \approx +0.02$ pp) is dwarfed by strongly negative expected future inflation ($\beta E_0[\pi_{H,1}] \approx -0.64$ pp) and the deflationary Jensen correction ($\Delta_0^J \approx -0.10$ pp): uncertainty works mainly through forward-looking pricing wedges.

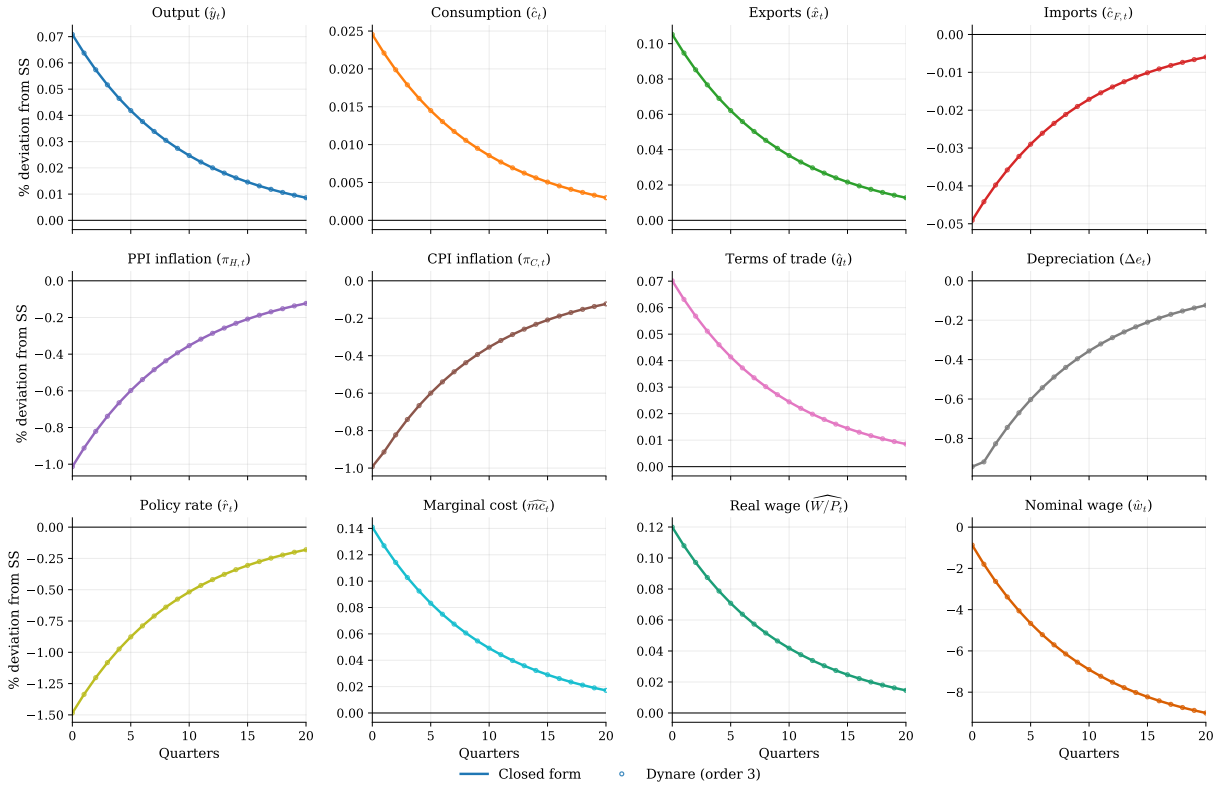

Figure I.1: Tariff uncertainty shock: all model variables (Taylor rule).

Notes: Each panel plots the third-order impulse response to a one-standard-deviation innovation in tariff volatility at $t = 0$ (Definition 1), with all other shocks shut down, using the baseline calibration in Table A.1 and the PPI-based Taylor rule. Solid lines: closed-form predictions from (χ_π, χ_q) via the static mappings in Section B.4. Circle markers: Dynare order-3 simulation. For the eight core variables $(\hat{y}_t, \hat{c}_t, \pi_{H,t}, \pi_{C,t}, \hat{q}_t, \hat{x}_t, \hat{r}_t, \Delta e_t)$ the markers come directly from the saved model output; for derived variables $(\widehat{mc}_t, \hat{c}_{F,t}, \widehat{W/P}_t, \hat{w}_t)$ they are computed from the simulated paths using exact (marginal cost, real wage) or first-order (imports, nominal wage) static mappings. Variables are reported as percent deviations from the deterministic steady state. The depreciation rate $\Delta e_t \equiv e_t - e_{t-1} = (q_t - q_{t-1}) + \pi_t$ is plotted rather than the exchange-rate level to avoid cumulative drift; the nominal wage $\hat{w}_t$ is a cumulative (level) deviation and is therefore not purely geometric in the volatility state. Labor is omitted because $L_t = Y_t/A_t$ and A_t is at steady state under the volatility experiment, so $\hat{l}_t = \hat{y}_t$ exactly.

I.2 Higher-order nature and the policy benchmark

Under the volatility experiment in Definition 1, the tariff *level* remains at its steady state, so the shock affects equilibrium only through time-varying conditional variances and the associated Jensen corrections. In a perturbation around the deterministic steady state, this implies that impulse responses to the volatility innovation are zero at first order *and* at second order: second-order terms capture stochastic-mean shifts from baseline risk, not responses to a volatility innovation. The first nonzero impulse responses therefore appear at third order through terms proportional to $\bar{\sigma}_\tau^2 \sigma_{T,t}$.

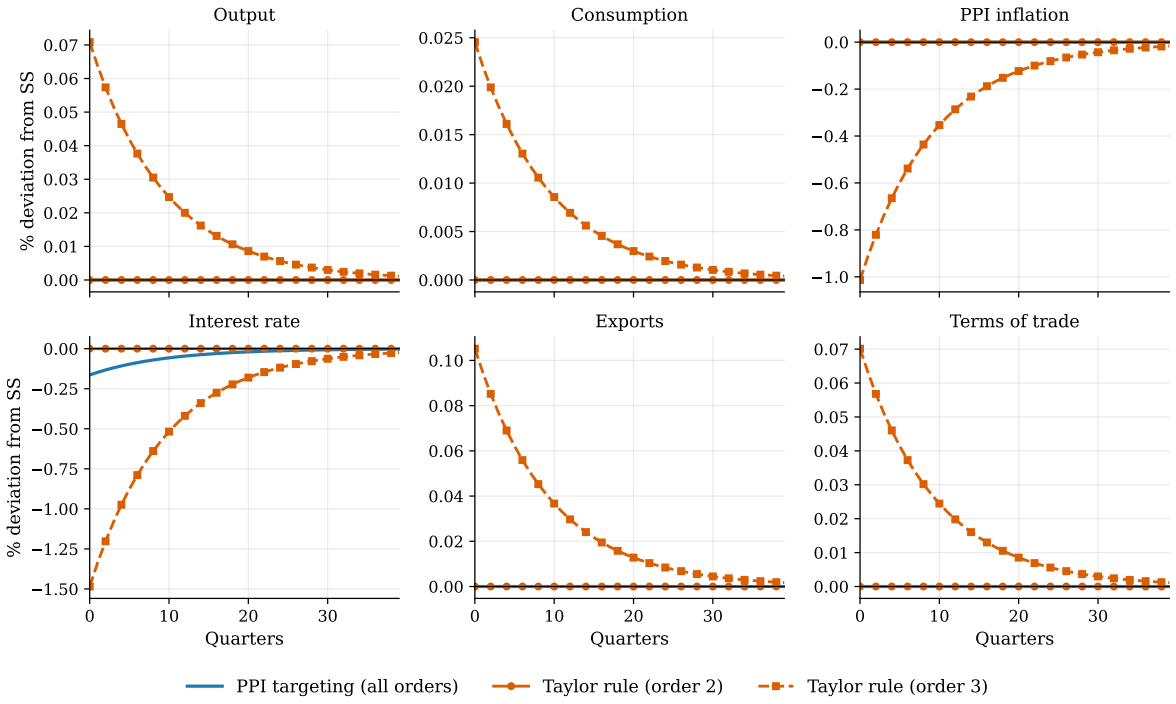

Figure I.2: Uncertainty shocks are higher-order objects.

Notes: In this model, a volatility innovation affects equilibrium through conditional-variance (Jensen) terms and therefore has no first- or second-order impulse responses. The figure plots impulse responses to the volatility innovation in Definition 1 across perturbation orders, contrasting the nonzero third-order responses under a Taylor rule with the (approximately) zero responses under strict PPI targeting in the complete-markets economy.

I.3 Decomposition and market structure

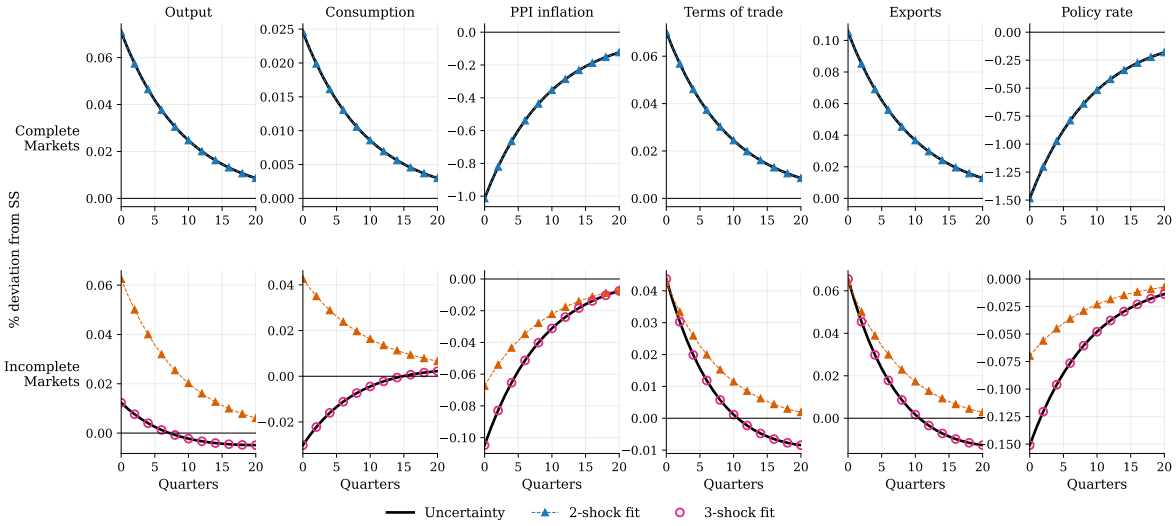

Figure I.3: Decomposition across market structures.

Notes: Dots report third-order uncertainty impulse responses to the one-standard-deviation volatility innovation in Definition 1 under the baseline calibration in Table A.1 and the PPI-based Taylor rule. Solid lines plot the best constant-weight two-shock fit using first-order productivity and UIP-wedge shocks; dashed lines add a third first-order discount-factor (Euler) shock. Under complete markets, two shocks span the uncertainty responses essentially exactly (baseline $R^2 = 1.000$; Proposition 2). Under incomplete markets, the same two shocks leave a material gap (baseline $R^2 = 0.545$); adding the discount-factor/Euler shock restores an exact fit ($R^2 = 1.000$), consistent with the additional Euler/Jensen wedge at third order (Proposition G.3).

I.4 Trade policy uncertainty index

Figure I.4: U.S. Trade Policy Uncertainty Index, 2014–2026.

Notes: Monthly Trade Policy Uncertainty (TPU) index of Caldara, Iacoviello, Molligo, Prestipino and Raffo (2020). Diamond markers denote events with direct implications for Canada; circles denote general U.S. trade-policy events. The dashed horizontal line marks the pre-2024 peak. The vertical dash-dot line marks the scheduled CUSMA/USMCA joint review (July 2026). Shaded regions indicate the first U.S.–China trade-war era (2018–2019) and the second trade-war era beginning in 2025.

References

Benigno, Gianluca, Pierpaolo Benigno, and Salvatore Nisticò. 2013. "Second-order approximation of dynamic models with time-varying risk." *Journal of Economic Dynamics and Control*, 37: 1231–1247.

Caldara, Dario, Matteo Iacoviello, Patrick Molligo, Andrea Prestipino, and Andrea Raffo. 2020. "The Economic Effects of Trade Policy Uncertainty." *Journal of Monetary Economics*, 109: 38–59.

Gross, Isaac, and James Hansen. 2021. "Optimal policy design in nonlinear DSGE models: An n-order accurate approximation." *European Economic Review*, 140: 103918.