Trade Policy Uncertainty and Optimal Monetary Policy*

Lu Han and Yuko Imura
Bank of Canada

23 October 2026

Preliminary

Latest version available here

Abstract

How does monetary policy shape the transmission of trade policy uncertainty shocks? In a canonical small open economy New Keynesian model, we derive closed-form impulse responses to a rise in export-tariff uncertainty. Under complete markets, uncertainty operates through two third-order expectation wedges: a New Keynesian Phillips curve wedge in firms' pricing and an uncovered interest parity wedge. The two wedges reinforce each other in lowering inflation but work against each other in output and the terms of trade. The monetary policy rule governs their relative strength, so uncertainty can be expansionary or contractionary. Solving the Ramsey problem in closed form, we show that the "divine coincidence" extends to tariff uncertainty: when the flexible-price allocation is efficient in every tariff state, strict producer-price-inflation targeting is optimal and uncertainty shocks have no real effects. Otherwise, optimal policy departs from strict targeting mainly to respond to realized tariffs, not to tariff uncertainty. Strict targeting captures almost all of the welfare gains from optimal policy, and across market structures, pricing assumptions, and export and import tariffs it stays close to Ramsey in nominal stabilization, while a standard Taylor rule departs substantially.

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, Jordi Galí, Erwan Gautier, Daniela Hauser, Oleksiy Kryvtsov, James MacGee, Ben Tomlin, Edouard Tsague-Djeutem, and Olivier Wang, as well as participants at the Bank of Canada Fellowship Learning Exchange, the Barcelona Summer Forum, the SED Annual Meeting, and the BIS Research Network Virtual Preparatory Workshop 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 re-emerged as a major source of macroeconomic risk. Alongside the tariff measures introduced by the US administration in 2025–2026, trade policy uncertainty rose sharply. In April 2025, the month of the "Liberation Day" tariff announcement, the US Trade Policy Uncertainty Index of Caldara, Iacoviello, Molligo, Prestipino and Raffo (2020) reached more than four times its pre-2024 peak, set in June 2019 during the US–China trade war.[1] For US trading partners, the concern is not only the tariffs currently in place but also the shifting distribution of future market access. Even if today's tariff rates are unchanged, a higher perceived likelihood of future tariff actions can affect current pricing, production, and exchange-rate expectations. This concern is especially salient for trade-dependent economies such as Canada. At the first joint review of the United States–Mexico–Canada Agreement (USMCA) on July 1, 2026, the United States did not confirm its extension. The USMCA remains in force but is now subject to annual reviews and is scheduled to expire in 2036 unless all three parties agree to extend it.[2]

This form of risk raises two questions central to monetary policy. First, how does trade policy uncertainty transmit to inflation, output, and the exchange rate, and how does monetary policy shape this transmission? Second, what is the optimal monetary policy response to it?

Answering these questions requires going beyond standard linearized models. A rise in tariff uncertainty leaves today's tariff unchanged but makes larger future tariff changes, in either direction, more likely. A first-order (log-linear) approximation misses such a shock entirely: volatility affects the equilibrium only through the nonlinear terms that it drops (Bloom, 2009; Fernández-Villaverde, Guerrón-Quintana, Rubio-Ramírez and Uribe, 2011; Benigno, Benigno and Nisticò, 2013). These terms are built from how inflation and the exchange rate respond to tariff level shocks, including how those responses change as the tariff rises. Understanding tariff uncertainty thus requires understanding the nonlinear responses to tariff levels, and monetary policy shapes those responses. Existing analyses of uncertainty shocks mostly compute these terms numerically. We characterize them, and the optimal monetary policy response, in closed form.

Our framework is a small open economy New Keynesian model. Home, the small open economy, exports to the rest of the world (Foreign), and Foreign's tariffs on Home's exports

[1] Online Appendix A plots the US Trade Policy Uncertainty Index for 2016–2026. See also White House (2025*a*,*b*, 2026); USTR (2025*a*) for recent US tariff actions covering automobiles, steel and aluminum, semiconductors, and China Section 301 exclusions.

[2] See USTR (2026) for the meeting and USTR (2025*b*); Kitamura and Trachtenberg (2026); Global Affairs Canada (2025) for the review process and preceding consultations.

follow a stochastic-volatility process. In the baseline, asset markets are complete, prices are subject to quadratic adjustment costs, and monetary policy follows a Taylor rule that responds to the inflation rate of Home-produced goods, measured by the producer price index (PPI), and to output. The key experiment is a pure uncertainty shock: the volatility of future export tariffs rises while the current tariff stays unchanged. This separates tariff uncertainty from shocks to the tariff level, which move relative prices directly. We solve the model with a third-order perturbation.[3]

Our analysis delivers three main results. First, under complete markets and this Taylor rule, tariff uncertainty pushes inflation in the opposite direction from a realized tariff. A realized export tariff lowers Home output and raises PPI inflation. Because the rule also leans against the fall in output, it raises the nominal rate by less than inflation, so the real interest rate falls, the Home currency depreciates in real terms, and firms' marginal cost rises. Tariff uncertainty instead lowers PPI inflation substantially, while output barely moves. Tariff increases typically arrive with heightened uncertainty about further tariff changes, and in our baseline the uncertainty component offsets much of the inflationary effect of the tariff itself. Under this rule, an analysis of the tariff level alone therefore overstates the inflationary effect of such episodes.

To explain this result, we show that under complete markets and this rule the equilibrium reduces exactly to two forward-looking equations in PPI inflation and the terms of trade (Lemma 1): the New Keynesian Phillips curve (NKPC), which summarizes firms' price setting, and an international asset-pricing condition that becomes uncovered interest parity (UIP) when log-linearized. Tariff uncertainty acts only through Jensen's-inequality terms in these two equations: because the future payoffs that drive price setting and asset prices depend nonlinearly on future tariffs, a wider spread of future tariffs changes their expected values. This two-equation system yields the responses in closed form (Proposition 1). The Jensen terms act as two wedges: a pricing wedge in the NKPC and an asset-market wedge in the UIP condition. Both are built from the slopes and curvatures of these level responses. In our baseline, both lower inflation.

The pricing wedge arises because firms weight future tariff states by their discounted sales, so they put less weight on high-tariff states, in which sales are low and inflation would be high. Moreover, inflation is concave in the tariff: each further increase raises inflation by less, because exports are already smaller. A third force works the other way: the adjustment-cost term in the Phillips curve is convex in inflation, so more dispersed future inflation on its own pushes today's inflation up. In our baseline the first two forces dominate, so higher tariff

[3]Our approach is related to Benigno, Benigno and Nisticò (2013), who develop approximation methods for models with time-varying exogenous risk.

uncertainty lowers the inflation that firms build into today's prices.

The asset-market wedge acts like an endogenous risk premium in the UIP condition. It arises for two reasons: the Home bond's return in Foreign currency is convex in depreciation, and depreciation, like inflation, rises less with each further tariff increase. Both make the Home bond more attractive when future tariffs become more dispersed, so the Home currency appreciates and firms' marginal cost falls. The two wedges move the currency in opposite directions: the pricing wedge weakens it through the rate cut that the Taylor rule makes in response to lower inflation, whereas the asset-market wedge strengthens it directly. Their effects on output nearly cancel, which is why output barely moves. Because the policy rule shapes the pricing wedge and how strongly that wedge moves the currency, the rule decides which wedge prevails, and hence whether tariff uncertainty is expansionary or contractionary (Section 3).

Second, optimal monetary policy neutralizes the effect of tariff uncertainty on producer prices: under Ramsey policy, the welfare-maximizing policy under commitment, PPI inflation barely moves. The reason is that tariff uncertainty, unlike a realized tariff, creates no trade-off of its own between price stability and efficiency. Under complete markets, both the flexible-price and the efficient allocation depend only on the current tariff, not on the distribution of future tariffs, so a pure rise in uncertainty moves neither. Consider first a benchmark in which a production tax or subsidy adjusts with the tariff, so that flexible prices always deliver the allocation a Home planner would choose. That planner uses Home's market power over its export prices, restraining supply to keep them high (the terms-of-trade motive). Strict PPI targeting, which keeps PPI inflation at zero, then reproduces the planner's allocation, so it is exactly optimal and uncertainty has no real effects. This is a version of the "divine coincidence," in which stabilizing inflation also delivers the efficient allocation (Proposition 2).[4]

In our baseline, by contrast, the tax or subsidy is fixed at the level that makes the no-tariff steady state efficient. A realized Foreign tariff then shrinks Home's exports and reduces the value of restraining supply, so the planner wants more output than flexible prices deliver. The tariff thus acts as a cost-push shock, one that opens a gap between the efficient and the flexible-price allocation. Ramsey policy responds by cutting the policy rate and accepting a small, brief rise in inflation that is later reversed, whereas the Taylor rule raises the nominal rate. Tariff uncertainty opens no such gap today, because it leaves the current tariff unchanged; it acts only through expectations of future tariffs and of the policy response

[4]Blanchard and Galí (2007) coined the term "divine coincidence" for a property of the standard New Keynesian model: stabilizing inflation also stabilizes the welfare-relevant output gap. In open economies the analogous result concerns producer prices; Galí and Monacelli (2005) and Corsetti, Dedola and Leduc (2010) characterize when stabilizing them is optimal. Our result extends this logic to tariff uncertainty.

to them. With no trade-off to manage, Ramsey policy almost fully offsets the downward pressure that tariff risk puts on firms' prices (Proposition 3).

Third, strict PPI targeting captures almost all of the welfare gains from optimal policy, whereas the Taylor rule falls well short. The Taylor rule lets producer prices move with every tariff realization and with tariff uncertainty; these movements waste output on adjustment costs and push production and relative prices away from the Ramsey allocation. Strict PPI targeting gives up mainly the small gains from responding to realized tariffs. The policy message is simple: a central bank should keep producer prices stable when tariff uncertainty rises. A realized tariff is a cost-push shock that justifies a small, temporary departure from price stability, but the uncertainty that accompanies it is not.

We extend these results in two directions. First, we limit international asset trade to a single noncontingent bond (Section 5). Home consumption is then no longer insured across tariff outcomes, so higher tariff uncertainty gives households a precautionary motive to save, building up net foreign assets. Under the Taylor rule, precautionary saving offsets most of the fall in PPI inflation. Ramsey policy, which we derive in closed form, still keeps PPI inflation nearly constant.

Second, we study uncertainty about Home *import* tariffs (Section 6). Import tariffs enter the consumer price index (CPI) directly, so they transmit differently from export tariffs, which work through foreign demand. Under complete markets and the PPI-based Taylor rule, realized import tariffs and import-tariff uncertainty both lower PPI inflation and raise output, so the two reinforce each other rather than offset as in the export-tariff case.

Our closed form also delivers a simple first-order representation: under complete markets and the Taylor rule, a fixed-weight combination of a productivity shock and a UIP (risk-premium) shock reproduces the effects of a tariff-uncertainty shock on PPI inflation and the terms of trade when these two shocks, the tariff level, and tariff volatility share a common persistence (Online Appendix Proposition E.1). Policy models solved at first order can therefore capture a tariff-uncertainty shock by combining the impulse responses to these two familiar shocks.

The optimal-policy prescription holds in every environment we compute, including alternative calibrations, incomplete markets, import tariffs, and extensions with Calvo pricing and local-currency pricing (Online Appendix L). In each, Ramsey policy keeps PPI inflation almost unchanged when tariff uncertainty rises, so it stays close to strict PPI targeting, whereas the Taylor-rule response varies in size and even sign across environments.

This paper contributes to three strands of the literature. First, it speaks to recent open-economy New Keynesian analyses of tariffs and trade wars (Lindé and Pescatori, 2019; Barattieri, Cacciatore and Ghironi, 2021; 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. Bergin and Corsetti (2023), Monacelli (2025), Werning, Lorenzoni and Guerrieri (2025) and Bianchi and Coulibaly (2025) characterize the optimal monetary policy response to tariffs. These studies mostly focus on realized or anticipated tariff *level* shocks. Closest to our focus, Merendino and Monacelli (2026) show empirically and in an incomplete-markets model that uncertainty about the persistence of realized tariffs shapes the exchange-rate response. We instead hold the current tariff fixed and shock the volatility of future tariffs, derive the resulting pricing and asset-market wedges in closed form, show how they interact with the policy rule, and solve for optimal policy. Our results show that tariff uncertainty, unlike tariff levels, calls for keeping producer prices nearly stable.

Second, it connects to the empirical and theoretical 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 Kalemli-Özcan, Soylu and Yildirim (2026) show that tariff uncertainty can depreciate the currency of the tariff-imposing country, through a risk premium demanded by financial intermediaries and precautionary saving by households. We study the other side: an economy whose exports face uncertain foreign tariffs. As in Kalemli-Özcan, Soylu and Yildirim (2026), tariff uncertainty acts like a risk premium in the UIP condition, which here strengthens the currency of the economy facing the tariffs. In our model it also works through firms' price setting, which weakens the currency through the policy rate, so the monetary policy rule determines whether tariff uncertainty appreciates or depreciates the currency.

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 saving and countercyclical markups. We show analytically that, under complete international asset markets and the PPI-based Taylor rule, the private equilibrium reduces exactly to a two-equation nonlinear system, which delivers closed-form third-order impulse responses and their first-order shock representation. Building on the open-economy optimal-policy literature (Benigno and Benigno, 2003; Galí and Monacelli, 2005; Corsetti, Dedola and Leduc, 2010), we also derive a closed-form third-order Ramsey response from the nonlinear commitment problem. We give an exact condition under which strict PPI targeting is optimal and show

that the optimal response to tariff uncertainty follows from the responses the central bank commits to make to each possible future tariff.

The remainder of the paper is organized as follows. Section 2 presents the model and the tariff-volatility experiment. Section 3 derives the complete-markets closed form under the Taylor rule, explains the two wedges, and contrasts tariff levels with tariff uncertainty. Section 4 solves the complete-markets Ramsey problem, contrasts its responses to tariff levels and to tariff uncertainty, and compares welfare across policies. Sections 5 and 6 study incomplete international asset markets and import-tariff uncertainty, and Section 7 concludes. The Online Appendix contains the proofs, derivations, numerical checks, and further extensions.

2 Model

Time is discrete. The model extends the canonical small-open-economy New Keynesian framework of Galí and Monacelli (2005) with tariffs and stochastic tariff volatility. Home is a small open economy and takes Foreign variables as exogenous. Home firms face Rotemberg price-adjustment costs (Rotemberg, 1982) and set prices in Home currency (producer-currency pricing, PCP), so exchange-rate movements pass through fully to the Foreign-currency price of Home exports. Home exporters pay a Foreign tariff T_t^*, and Home levies an import tariff $T_{F,t}$ on Foreign goods. Monetary policy follows a Taylor rule, strict PPI targeting, or the Ramsey plan (Section 2.6).

2.1 Key objects and notation

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

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

Thus $\Pi_{H,t}$ is PPI inflation and $\Pi_{C,t}$ is CPI inflation. The ratio G_t is the CPI-to-PPI relative price. It moves with the Home-currency price of imports, and hence with the exchange rate, Foreign prices, and the import tariff; it is what makes CPI inflation differ from PPI inflation. Let $\mathcal{E}_t$ be the nominal exchange rate (Home currency per unit of Foreign currency) and let P_t^* be the Foreign CPI. The terms of trade is

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

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

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

Foreign demand for Home goods is isoelastic with the same elasticity and is scaled by α:

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

For any level variable X_t with deterministic steady state $\bar{X}$, we use $\hat{x}_t \equiv \log(X_t/\bar{X})$ for log deviations. We normalize steady-state price levels and the terms of trade to one, so for those objects hats coincide with log levels. In Section 3 we work directly with log variables and suppress hats when no confusion arises.

2.2 Households

The representative Home household has constant-relative-risk-aversion (CRRA) utility

$$U(C_t, L_t) = u(C_t;\sigma) - \chi \frac{L_t^{1+\varphi}}{1+\varphi}, \qquad u(C;\sigma) = \begin{cases} C^{1-\sigma}/(1-\sigma), & \sigma \neq 1, \\ \log C, & \sigma = 1, \end{cases} \qquad \sigma > 0,\ \varphi \geq 0, \tag{4}$$

where σ is relative risk aversion (the inverse of the intertemporal elasticity of substitution), φ is the inverse Frisch elasticity of labor supply, 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, and T_t^{gov} is the government's net lump-sum transfer, which rebates the revenue from Home's import tariff $T_{F,t}$ and the net proceeds of the fixed production correction defined below; revenue from the Foreign tariff T_t^* on Home exports accrues abroad. The first-order conditions give labor supply and the Euler

equation

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

2.3 Firms, production, and price setting

Home firms are monopolistically competitive and produce output using labor:

$$Y_t = A_t L_t, \tag{7}$$

where A_t is aggregate productivity. Firms face Dixit–Stiglitz demand with elasticity $\epsilon > 1$ (Dixit and Stiglitz, 1977) and pay quadratic Rotemberg costs to adjust prices. Before any fiscal correction, real marginal cost in producer-price units is

$$m_t^r \equiv \frac{MC_t}{P_{H,t}} = \chi G_t C_t^{\sigma} Y_t^{\varphi} A_t^{-(1+\varphi)}. \tag{8}$$

At the deterministic no-tariff steady state, $\bar{\upsilon} \equiv m_{ss}^r = \chi$. The government applies a fixed production correction s^p, a constant tax or subsidy on production whose net proceeds are rebated lump sum, so the firm's effective marginal cost is $(1 - s^p)MC_t$. We set

$$1 - s^p = \frac{\epsilon - 1}{\epsilon \bar{\upsilon}}, \qquad mc_t \equiv \frac{m_t^r}{\bar{\upsilon}}, \tag{9}$$

so the flexible-price markup condition $mc_t = 1$ holds at the deterministic steady state. Together with the calibration of χ, this makes the flexible-price allocation efficient at the no-tariff steady state (Online Appendix F.2). The correction stays fixed when tariffs change, however, so flexible prices need not be efficient once a tariff is realized, which gives monetary policy a role beyond strict PPI targeting (Section 4). With the normalization in (9), the Rotemberg first-order condition is the nonlinear NKPC in PPI inflation:

$$\kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) + \beta\kappa \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \frac{\lambda_{t+1}}{\lambda_t} \frac{G_t}{G_{t+1}} \right], \tag{10}$$

where mc_t is normalized real marginal cost. Online Appendix B derives this equation from firm profits and records the lump-sum fiscal accounts.

2.4 Price indices, goods market clearing, and tariffs

With an import tariff $T_{F,t}$ levied on Foreign goods at the border, the CES aggregator (2) implies the CPI

$$P_t = \left[(1-\alpha)P_{H,t}^{1-\gamma} + \alpha\left(T_{F,t}\mathcal{E}_t P_{F,t}^*\right)^{1-\gamma}\right]^{1/(1-\gamma)}. \tag{11}$$

Foreign produces a single good, so $P_t^* = P_{F,t}^*$. In the small-open-economy limit we take Foreign prices as exogenous and constant, so the terms of trade (1) moves only with the exchange rate $\mathcal{E}_t$ and the Home producer price $P_{H,t}$.

Goods market clearing requires Home output to equal domestic demand, exports, and price-adjustment costs:

$$Y_t = C_{H,t} + C_{H,t}^* + \frac{\kappa}{2}(\Pi_{H,t}-1)^2 Y_t. \tag{12}$$

2.5 International asset markets: complete versus incomplete

The asset-market block has two parts: an asset-pricing relation and a closure that determines how much risk is shared internationally. Let R_t and R_t^* be the gross one-period nominal interest rates at Home and in Foreign. Under complete markets, we combine the Home nominal Euler equation with international risk sharing (14) below. With (C_t^*, P_t^*) normalized to constants and $R_t^* = \beta^{-1}$, this gives the exact nonlinear asset-pricing condition, augmented by an auxiliary disturbance ψ_t discussed next,

$$1 = \frac{R_t}{R_t^*} e^{\psi_t}\mathbb{E}_t\left[\frac{\mathcal{E}_t}{\mathcal{E}_{t+1}}\right], \tag{13}$$

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 term ψ_t is an auxiliary UIP (risk-premium) disturbance: at first order it shifts the log UIP condition $r_t - r_t^* + \psi_t = \mathbb{E}_t[\Delta e_{t+1}]$, with $r_t \equiv \log R_t$, $r_t^* \equiv \log R_t^*$, and $e_t \equiv \log \mathcal{E}_t$. It is zero in every baseline tariff experiment and serves only as a comparison shock for the first-order representation in Section 3.3 and Online Appendix E. Tariff uncertainty also shifts this relation, but *endogenously*: through the variance of future depreciation and the curvature of its response to the tariff, it adds a higher-order correction that acts like a risk premium on the Foreign bond. We call this correction the asset-market wedge.

Under complete international asset markets, international risk sharing holds:

$$\frac{C_t^{-\sigma}}{(C_t^*)^{-\sigma}} = \vartheta \cdot \frac{P_t}{\mathcal{E}_t P_t^*}, \tag{14}$$

for a constant $\vartheta > 0$. With (C_t^*, P_t^*) constant and $\vartheta(C^*)^{-\sigma} = 1$, risk sharing pins Home consumption to the real exchange rate:

$$C_t = \left(\frac{G_t}{Q_t}\right)^{-1/\sigma}. \tag{15}$$

2.6 Monetary policy

We study three monetary-policy regimes. Under the Taylor rule, the nominal interest rate responds to PPI inflation and output relative to its deterministic steady state:

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

where $\bar{R} = 1/\beta$ and $\bar{Y}$ are deterministic steady-state values and (ϕ_π, ϕ_y) satisfy the Taylor principle. In the baseline, $(\phi_\pi, \phi_y) = (1.5, 0.5)$, with the policy rate and inflation measured in quarterly rates; Online Appendix Table B.1 reports the full calibration. Under *strict PPI targeting*, the central bank sets $\Pi_{H,t} = 1$ after every history and uses the nominal rate to implement the resulting flexible-price allocation. Under Ramsey policy, a planner with commitment chooses a state-contingent allocation to maximize expected welfare subject to the private equilibrium conditions. Section 4 shows when the Ramsey plan coincides exactly with strict PPI targeting and derives the Ramsey response when the fixed production correction does not adjust with tariffs.

2.7 Equilibrium

Given the exogenous processes $\{T_{F,t}, T_t^*, A_t, \psi_t\}$ and the exogenous Foreign variables, an equilibrium is a set of stochastic processes for allocations and prices that satisfy (i) household optimality (labor supply and the Euler equation), (ii) the nonlinear NKPC (10), (iii) goods market clearing, (iv) the asset-market block, and (v) the monetary policy rule. Under complete markets, the asset-market block consists of (13) and risk sharing (14). Under incomplete markets, it consists of the Euler equations for the domestic bond and the Foreign-currency bond and the law of motion for net foreign assets, with no risk-sharing condition. Online Appendix B records the full nonlinear system in levels.

Under complete markets, risk sharing pins consumption to the real exchange rate, so, with the PPI-based Taylor rule, the equilibrium reduces to a 2×2 nonlinear system in $(\Pi_{H,t}, Q_t)$ (Lemma 1). Under incomplete markets, consumption and net foreign assets become dynamic variables: the Euler equation for the Foreign-currency bond and the law of motion for net foreign assets join the pricing block, and the domestic-bond Euler equation determines the nominal rate that implements the allocation (Online Appendix H).

Lemma 1 (Exact minimal nonlinear system under complete markets and the PPI Taylor rule)**.** *Assume producer-currency pricing and complete international asset markets, hold the Foreign objects (C_t^*, P_t^*, R_t^*) constant with $R_t^* = \bar{R}$, and let monetary policy follow the PPI-based Taylor rule* (16)*. Then the private equilibrium under this rule is pinned down by the following* exact *nonlinear 2×2 system in $(\Pi_{H,t}, Q_t)$:*

$$\begin{aligned} \kappa(\Pi_{H,t} - 1)\Pi_{H,t} &= (\epsilon - 1)(mc_t - 1) + \beta\kappa\mathbb{E}_t\left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1}\frac{Y_{t+1}}{Y_t}\frac{Q_t}{Q_{t+1}}\right], \\ 1 &= \Pi_{H,t}^{\phi_\pi}\left(\frac{Y_t}{\bar{Y}}\right)^{\phi_y} e^{\psi_t}\mathbb{E}_t\left[\frac{Q_t}{Q_{t+1}\Pi_{H,t+1}}\right], \end{aligned} \tag{17}$$

where all remaining variables are static functions of $(\Pi_{H,t}, Q_t)$ and exogenous states:

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

Two features deliver this reduction. First, risk sharing removes any independent intertemporal margin for Home consumption: the stochastic discount factor is tied to the terms of trade, so the stochastic discount factor in the NKPC collapses to $\beta Q_t/Q_{t+1}$. Second, combining the asset-pricing condition (13) with the Taylor rule and the definition of Q_t gives (17). Once $(\Pi_{H,t}, Q_t)$ are known, output, marginal cost, consumption, exports, and the policy rate follow from static relations. The only approximation in our uncertainty analysis is the third-order perturbation of this system; Online Appendix C gives the derivation.

2.8 Tariff uncertainty shocks

Our focus is *export-tariff uncertainty*: uncertainty about the future Foreign tariff on Home exports. The experiment is deliberately clean. On impact, agents learn that future tariff realizations will be more dispersed, but the tariff level does not change. We model this as stochastic volatility in the export tariff. Let $\tau_t \equiv \log T_t^*$ denote the log export tariff and let $\sigma_{T,t}$ denote its volatility state. We assume

$$\tau_t = \rho_\tau \tau_{t-1} + \bar{\sigma}_\tau e^{\sigma_{T,t-1}} \varepsilon_{\tau,t}, \tag{18}$$

$$\sigma_{T,t} = \rho_\sigma \sigma_{T,t-1} + \eta_\sigma \varepsilon_{\sigma,t}, \tag{19}$$

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

$$\begin{aligned} a_t \equiv \log A_t &= \rho_A a_{t-1} + \sigma_A \varepsilon_{A,t}, \\ \psi_t &= \rho_\psi \psi_{t-1} + \sigma_\psi \varepsilon_{\psi,t}. \end{aligned} \tag{20}$$

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:

$$\mathrm{Var}_t(\tau_{t+1}) = \bar{\sigma}_\tau^2 e^{2\sigma_{T,t}}. \tag{21}$$

Because the NKPC and the asset-pricing condition are forward-looking, a change in (21) moves date-t allocations even when the tariff level does not change.

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 $\tau_t \equiv 0$.

Unless stated otherwise, every reported uncertainty response is computed along this *no-realization history*: no tariff innovation is ever realized, while agents continue to price the risk of future tariff changes. Each response is the difference between this path and the same policy's *no-innovation path*, which has neither a tariff nor a volatility innovation but the same baseline tariff risk.

In the baseline calibration, $\bar{\sigma}_\tau = 0.10$ and $\eta_\sigma = \log 2$, so in the exact shock process the innovation doubles the conditional standard deviation of next period's log tariff, from 0.10 to 0.20.

Even though the log tariff stays at zero, greater dispersion raises the expected gross tariff: with Gaussian innovations, $\mathbb{E}_t T^*_{t+1} = \exp\{\mathrm{Var}_t(\tau_{t+1})/2\}$. This is a Jensen effect of greater dispersion, not a realized change in the tariff.

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 enters only through the scale $\bar{\sigma}_\tau e^{\sigma_{T,t}}$ of the next tariff innovation in (18). The forward-looking blocks (the NKPC and the asset-pricing condition) take expectations of nonlinear functions of next-period variables. A third-order expansion replaces each such expectation by the expectation of the linear term plus a Jensen correction proportional to the conditional variance. Under symmetric innovations (Assumption 1), the third-order terms that involve conditional skewness vanish, so only conditional variances remain. □

We compute the responses to uncertainty shocks with a third-order perturbation around the deterministic steady state, the lowest order at which a volatility shock has any effect. Throughout, $O(k)$ denotes terms of order δ^k or higher, where the parameter δ scales the tariff innovation and the volatility innovation jointly. Baseline tariff variance $\bar{\sigma}_\tau^2$ is then second order and the volatility state $\sigma_{T,t}$ first order. By Lemma 2, up to third order the volatility state moves equilibrium only through terms proportional to $\bar{\sigma}_\tau^2 \sigma_{T,t}$, which are third order (Online Appendix D.4.1). This structure underlies the closed-form responses in Section 3 and the first-order shock representation in Online Appendix E.

3 Tariff Uncertainty under a Taylor Rule

This section shows why, under a Taylor rule, tariff uncertainty causes sizable producer-price deflation but a much more muted output response. A rise in uncertainty leaves today's tariff unchanged; it only widens the distribution of future tariffs, making larger tariff changes in either direction more likely. Firms setting prices today and investors pricing Home and Foreign assets must average over these possible tariffs, and because their payoffs depend nonlinearly on future tariffs, a wider distribution shifts these averages. We collect these Jensen corrections into two wedges: a *pricing wedge* in the Phillips curve and an *asset-market wedge*

in the asset-pricing condition whose log-linear counterpart is UIP. Under complete markets, Lemma 1 reduces the equilibrium to these two forward-looking equations in producer-price inflation and the terms of trade; output, exports, and the policy rate then follow from static relationships and the policy rule. Uncertainty therefore moves the economy only through these two wedges. Section 3.4 then contrasts uncertainty with a tariff-level shock.

For exposition, we hold the import tariff $T_{F,t}$ constant and study uncertainty in the export tariff T_t^*. A time-varying import tariff changes only the static block, and hence the closed-form coefficients, not the two-equation reduction. Online Appendix D gives the coefficient-matching details behind the closed forms below.

3.1 The minimal two-equation system

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

$$\kappa_{mc} \equiv \frac{\epsilon - 1}{\kappa}. \tag{22}$$

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

$$(e^{\pi_t} - 1)e^{\pi_t} = \kappa_{mc}\big(e^{q_t + \varphi y_t - (1+\varphi)a_t} - 1\big) + \beta \mathbb{E}_t[(e^{\pi_{t+1}} - 1)e^{\pi_{t+1}}\, \Omega_{t+1}], \tag{23}$$

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

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 $a_t \equiv \log A_t$. Equations (23)–(24) are an exact two-equation system in the forward-looking variables (π_t, q_t): they hold in the nonlinear equilibrium, before any approximation.

Uncertainty enters only through Jensen corrections. The volatility experiment in Definition 1 changes only the conditional distribution of τ_{t+1}, where $\tau_t \equiv \log T_t^*$. Because the only expectations in the system are those in (23)–(24), tariff uncertainty enters only through Jensen corrections to these two expectations. Expanding the NKPC gives 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} is the change in the conditional second-moment terms relative to the same path without the volatility innovation. Taking logs of the asset-pricing condition (24) and expanding the log expectation gives the familiar log-linear UIP condition plus an additive

asset-market wedge:

$$(\phi_\pi \pi_t + \phi_y y_t + \psi_t) = \mathbb{E}_t[(q_{t+1} - q_t) + \pi_{t+1}] + \Delta_t^{UIP} + O(4), \tag{26}$$

where Δ_t^{UIP} is the corresponding change generated by the log expectation in (24). It has two parts: a variance term, which is nonpositive, and a curvature term (Section 3.2). Step 1 of the proof in Online Appendix E derives Δ_t^{UIP} by a cumulant expansion of this log expectation.

The economic source of both wedges is Jensen's inequality. With a known future tariff, each forward-looking equation is evaluated at a single tariff state. A volatility innovation widens the distribution of future tariffs without changing today's tariff. If the relevant pricing or payoff object, written schematically as $f(\tau)$, is nonlinear, then $\mathbb{E}_t f(\tau_{t+1})$ differs from $f(\mathbb{E}_t \tau_{t+1})$. The gap is negative for a concave object, as Figure 1 illustrates, positive for a convex object, and zero for a linear one. In the model, its sign depends on the curvature of the responses to a realized tariff and, in the pricing condition, on how those responses covary with the weight firms place on each future tariff state (Section 3.2).

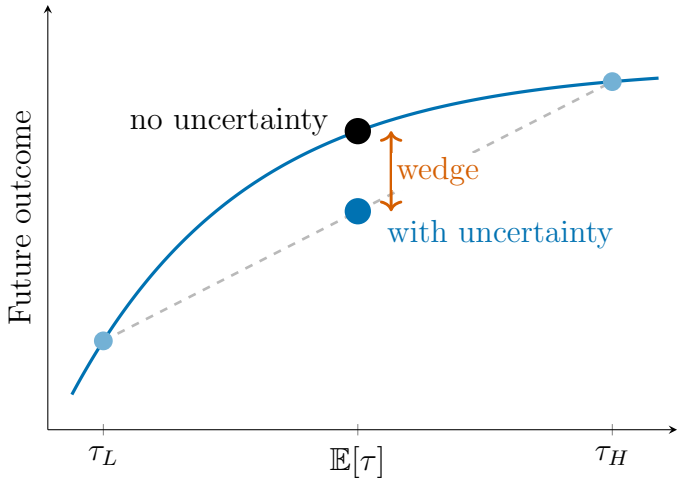


Figure 1: Tariff risk and the expectation wedge

Notes: Next period's log tariff equals τ_L or τ_H with equal probability, so its mean $\mathbb{E}[\tau]$ is the same as without uncertainty. For a concave outcome, the expected outcome with uncertainty lies below the outcome without it; the gap is the expectation wedge. For a convex outcome the wedge has the opposite sign. The curve is illustrative.

Order of the effect and the uncertainty coefficients. Both wedges are third-order objects: they move only through terms proportional to $\bar{\sigma}_\tau^2 \sigma_{T,t}$ (Section 2.8). At that order, each wedge is proportional to the volatility state,

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

so the pair of loadings (u^{NK}, u^{UIP}) is a sufficient statistic for the uncertainty impulse response.

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

$$\pi_t = \bar{\pi} + \chi_\pi \, \sigma_{T,t} + O(4), \qquad q_t = \bar{q} + \chi_q \, \sigma_{T,t} + O(4), \tag{27}$$

where $(\bar{\pi}, \bar{q}) = O(\bar{\sigma}_\tau^2)$ are the constant shifts in the stochastic means caused by baseline tariff risk. These shifts are common to the shocked and unshocked paths, so the impulse response depends only on the coefficients on $\sigma_{T,t}$. A volatility innovation raises the conditional variance of next period's tariff by

$$\mathrm{Var}_t(\tau_{t+1}) - \bar{\sigma}_\tau^2 = \bar{\sigma}_\tau^2 \left(e^{2\sigma_{T,t}} - 1\right) = 2\bar{\sigma}_\tau^2 \sigma_{T,t} + O(4).$$

Under the Taylor rule the reduced system has no endogenous state variable, so inflation and the terms of trade inherit the AR(1) path of the volatility state. Under the volatility experiment in Definition 1,

$$\Delta\pi_t = \chi_\pi \, \rho_\sigma^t \eta_\sigma, \qquad \Delta q_t = \chi_q \, \rho_\sigma^t \eta_\sigma, \qquad t \geq 0, \tag{28}$$

with impact effects $\Delta\pi_0 = \eta_\sigma \chi_\pi$ and $\Delta q_0 = \eta_\sigma \chi_q$.[5] The next subsection characterizes the loadings (u^{NK}, u^{UIP}) and maps them into (χ_π, χ_q).

3.2 The two uncertainty wedges

Tariff risk shifts the two forward-looking expectations in the system: firms' expected future pricing term in the NKPC and the expected return on the Home bond in the asset-pricing condition. These shifts are the pricing wedge and the asset-market wedge. Both depend on how inflation and the terms of trade differ between high- and low-tariff states. We summarize these differences by the slopes and curvatures of the responses to a realized tariff, derive how the two wedges map into inflation and the terms of trade, and then interpret each wedge.

Coefficient matching via a common 2×2 template. Under complete markets, the static block provides a smooth map $y_t = y(\pi_t, q_t, \tau_t)$, 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 D.1),

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

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

The second-order terms of this map, which come from goods-market clearing and the Rotemberg resource cost, are reported in equation (D.1) of the same appendix. Along deterministic tariff paths with no innovations, equilibrium decision rules admit the expansion

$$\pi_t = A_\pi \tau_t + \tfrac{1}{2} B_\pi \tau_t^2 + O(\tau_t^3), \qquad q_t = A_q \tau_t + \tfrac{1}{2} B_q \tau_t^2 + O(\tau_t^3), \tag{29}$$

where the slopes (A_π, A_q) measure first-order exposure to the tariff level and the curvatures (B_π, B_q) measure how that exposure bends. A key simplification is that the slopes, the curvatures, and the uncertainty coefficients all solve a 2×2 linear system with the *same* matrix, evaluated at different persistence parameters. Define

$$\mathbf{M}(\rho) \equiv \begin{pmatrix} 1 - \beta\rho & -\tilde{\kappa} \\ \phi_\pi - \rho & \phi_y \theta_q + 1 - \rho \end{pmatrix}, \qquad \tilde{\kappa} \equiv \kappa_{mc}(1 + \varphi\theta_q), \tag{30}$$

where κ_{mc} is given in (22). At first order, substituting $y_t = \theta_q q_t - \theta_\tau \tau_t + O(2)$ into (23)–(24) gives

$$\mathbf{M}(\rho_\tau) \begin{pmatrix} A_\pi \\ A_q \end{pmatrix} = \begin{pmatrix} -\kappa_{mc}\varphi\theta_\tau \\ \phi_y \theta_\tau \end{pmatrix}. \tag{31}$$

The curvatures (B_π, B_q) solve the same system at persistence ρ_τ^2, with a right-hand side built from the slopes and the second-order static terms (Online Appendix D.3). They matter because tariff uncertainty depends on how the response to a realized tariff bends, not only on its slope. At the volatility order, (χ_π, χ_q) solve

$$\mathbf{M}(\rho_\sigma) \begin{pmatrix} \chi_\pi \\ \chi_q \end{pmatrix} = \begin{pmatrix} u^{NK} \\ u^{UIP} \end{pmatrix} \equiv \bar{\sigma}_\tau^2 \begin{pmatrix} \beta\left(B_\pi + 3A_\pi^2 + 2A_\pi s_1\right) \\ (B_\pi + B_q) - (A_\pi + A_q)^2 \end{pmatrix}, \tag{32}$$

where $s_1 \equiv (\theta_q - 1)A_q - \theta_\tau$ is the loading of the NKPC discount term Ω_{t+1} on a realized tariff τ_{t+1}. In the baseline, $s_1 < 0$, because $\theta_q A_q - \theta_\tau < 0$ and $A_q > 0$ (Section 3.4). When $\det \mathbf{M}(\rho_\sigma) \neq 0$, the system has a unique solution. Its right-hand side depends on the lower-order coefficients (A_π, A_q, B_π, B_q); Online Appendix D records explicit expressions for all layers.

Anatomy of the wedges. Equation (32) separates tariff risk into one shift in firms' pricing condition and one shift in international asset pricing. Both are built from the response to a *realized* tariff: the slopes (A_π, A_q) determine how high- and low-tariff states differ, and the curvatures (B_π, B_q) determine whether averaging across those states raises or lowers the relevant expectation.

The pricing wedge,

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

is negative in the baseline: tariff risk puts downward pressure on Home PPI inflation. It comes from second-order expansions of the forward-looking Rotemberg term, $(e^{\pi_{t+1}} - 1)e^{\pi_{t+1}} = \pi_{t+1} + \frac{3}{2}\pi_{t+1}^2 + O(3)$, and of the discount term, $\Omega_{t+1} = 1 + s_1\tau_{t+1} + O(2)$. First, the Rotemberg term is convex in inflation, so more dispersed future inflation raises its expected value and therefore raises today's inflation; this adjustment-cost term, $3\beta\bar{\sigma}_\tau^2 A_\pi^2$, is nonnegative. Second, firms do not weight all future tariff states equally when choosing today's price. A high Foreign tariff reduces future sales of Home goods and their discounted value ($s_1 < 0$), so pricing in that state matters less for profits than pricing in a low-tariff, high-sales state. In the baseline, high-tariff states are also the states with higher inflation ($A_\pi > 0$), so giving them less weight lowers the expected forward-looking term: the covariance term $2\beta\bar{\sigma}_\tau^2 A_\pi s_1$ is negative. Third, the curvature term $\beta\bar{\sigma}_\tau^2 B_\pi$ records whether inflation rises with the tariff at an increasing or a decreasing rate. In the baseline $B_\pi < 0$: a higher Foreign tariff shrinks Home exports, so further tariff increases bite on a smaller export base and raise inflation by less. Because inflation is concave in the tariff, a wider spread of future tariffs lowers expected future inflation. These two negative terms outweigh the adjustment-cost term in the baseline.

The asset-market wedge,

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

has two pieces. Condition (13) prices the Home bond's return in Foreign currency, $R_t\mathcal{E}_t/\mathcal{E}_{t+1}$, which is convex in nominal depreciation: a depreciation lowers this return by less than an equal-sized appreciation raises it. More dispersed future depreciation therefore raises the bond's expected Foreign-currency return at given interest rates, which is the nonpositive variance term $-\bar{\sigma}_\tau^2(A_\pi + A_q)^2$. The curvature term $\bar{\sigma}_\tau^2(B_\pi + B_q)$ records whether nominal depreciation, $(q_{t+1} - q_t) + \pi_{t+1}$, rises with the tariff at an increasing or a decreasing rate. In the baseline the rate is decreasing, for the same shrinking-export-base reason as in the pricing wedge, so averaging over high- and low-tariff states lowers expected future nominal depreciation, which also raises the bond's expected Foreign-currency return. Both pieces are therefore negative, $u^{UIP} < 0$: tariff risk makes the Home bond more attractive at given interest rates. Restoring the asset-pricing condition requires more expected depreciation of the Home currency, which comes from an appreciation today. The asset-market wedge therefore puts direct appreciation pressure on the Home currency.[6]

[6]Set $u^{NK} = 0$ in (32). Then $\chi_q = (1 - \beta\rho_\sigma)u^{UIP}/\det\mathbf{M}(\rho_\sigma)$, so the single-channel terms-of-trade response has the sign of u^{UIP} in the determinacy region used below: $u^{UIP} < 0$ lowers q_t, a Home real appreciation.

These signs are not imposed: they vary with openness and the two monetary-policy coefficients, as we show below.

3.3 Responses to a tariff uncertainty shock

The central result is a nominal–real asymmetry. In the baseline, both wedges push PPI inflation down, but the pricing wedge causes a real depreciation while the asset-market wedge causes a real appreciation. The two wedges therefore reinforce each other in inflation and offset each other in output. Proposition 1 characterizes this pattern in closed form.

Proposition 1 (Closed-form uncertainty responses under a Taylor rule)**.** *Under Assumption 1 and the volatility experiment in Definition 1, maintain PCP and complete international asset markets as in Lemma 1, impose the PPI-based Taylor rule* (16)*, assume that the resulting equilibrium is locally determinate, and suppose* $\det \mathbf{M}(\rho) \neq 0$ *for* $\rho \in \{\rho_\tau, \rho_\tau^2, \rho_\sigma\}$*. Let* (u^{NK}, u^{UIP}) *denote the loadings of the pricing wedge and the asset-market wedge on the right-hand side of* (32)*. Then the third-order uncertainty coefficients are*

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

and the impulse responses of (π_t, q_t) *are geometric,* $\Delta\pi_t = \chi_\pi \rho_\sigma^t \eta_\sigma$ *and* $\Delta q_t = \chi_q \rho_\sigma^t \eta_\sigma$*, as in* (28)*. The remaining real allocations and relative prices follow from the static block, and the policy rate follows from the Taylor rule:* $\Delta y_t = \theta_q\, \Delta q_t + O(4)$*,* $\Delta \log X_t = \gamma\, \Delta q_t + O(4)$*, and* $\Delta\hat{r}_t = \phi_\pi\, \Delta\pi_t + \phi_y\, \Delta y_t + O(4)$*, where* $\hat{r}_t \equiv \log(R_t/\bar{R})$*.*

Proposition 1 compresses the dynamic response into two scalars. PPI inflation, the terms of trade, and the static real allocations recovered from them all share the geometric profile ρ_σ^t. Variables that involve lagged relative prices or cumulative nominal objects, such as nominal depreciation or nominal wages, follow from accounting identities. The value of the closed form is that it shows how the two wedges combine.

Two features of the solution organize everything that follows. In the determinacy region used in the calibration, $D_\sigma \equiv \det \mathbf{M}(\rho_\sigma) > 0$. First, the inflation coefficient χ_π loads the two wedges with the *same* sign: the weight on u^{NK} is $\phi_y\theta_q + 1 - \rho_\sigma > 0$ and the weight on u^{UIP} is $\tilde{\kappa} > 0$. When both wedges are deflationary, as in the baseline, they *reinforce* one another, so the inflation response is large and its sign is unambiguous. Second, the terms-of-trade coefficient χ_q loads the two wedges with *opposite* signs: u^{UIP} enters with the positive weight $1 - \beta\rho_\sigma$, and u^{NK} enters with the weight $-(\phi_\pi - \rho_\sigma)$, which is negative whenever the Taylor

rule obeys $\phi_\pi > \rho_\sigma$. On the real side the two channels therefore *offset*. Because output and exports move in the same direction as the terms of trade ($\theta_q, \gamma > 0$ in Proposition 1), the sign of the real response is the sign of χ_q, which is not restricted: it depends on which wedge dominates.

Monetary policy enters this mapping twice. First, it shapes the wedges themselves, because it governs how the economy responds to a realized tariff. Second, for given wedges, it determines how they are transmitted into inflation and the terms of trade. A deflationary pricing wedge leads to Taylor-rule easing and a cut in the real interest rate; a stronger response of the policy rate to inflation (a larger ϕ_π) means a larger cut and therefore a larger real depreciation. Formally, the pricing wedge enters χ_q with weight $-(\phi_\pi - \rho_\sigma)/D_\sigma$, whose magnitude rises with ϕ_π even though D_σ also rises. Online Appendix D.6 writes (χ_π, χ_q) in terms of primitive parameters.

The clearest way to see the mechanism is to ask what each wedge would do on its own. Figure 2 sets one wedge in (32) to zero at a time; Online Appendix Corollary D.2 gives the exact single-wedge coefficients and their flexible-price limits.

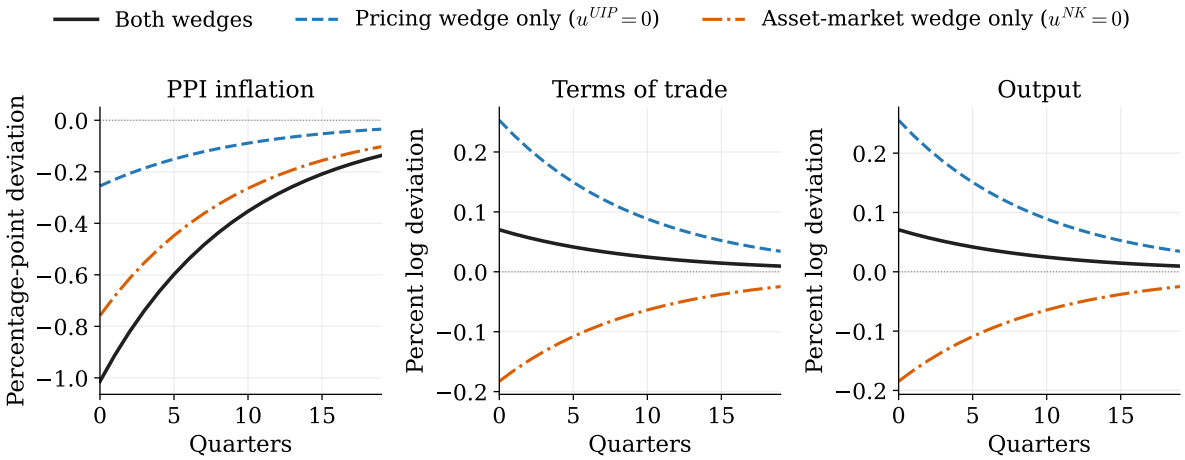


Figure 2: Responses to a tariff uncertainty shock: single-wedge counterfactuals

Notes: Third-order impulse responses to the tariff uncertainty shock of Definition 1 under the baseline calibration (Online Appendix Table B.1) and the PPI-based Taylor rule (16). All other innovations, including realized tariff innovations, are zero along the displayed path; future tariff risk remains in agents' expectations. Each panel shows the response with both wedges and two single-wedge counterfactuals, each of which sets one wedge in (32) to zero. PPI inflation is in quarterly percentage points (not annualized); the terms of trade and output are in percent log deviations. A rise in the terms of trade is a Home real depreciation.

Figure 2 delivers the paper's first main result: tariff uncertainty can have a pronounced nominal effect and a muted real effect, because the same two channels add in inflation but offset in output. The "pricing wedge only" curve isolates firms' response to the distribution of future sales. Low-tariff states have high export sales and so receive greater weight in current

price setting; in the baseline, those states call for a lower Home producer price. The resulting deflation induces Taylor-rule easing, a real depreciation, stronger export demand, and higher output.

The "asset-market wedge only" curve isolates the asset-pricing channel. Because this wedge makes the Home bond more attractive, restoring the asset-market condition (26) calls for a real appreciation today. The appreciation lowers marginal cost, so PPI inflation also falls and the policy rate is cut, but the rate cut only partly offsets the appreciation. Home goods become more expensive abroad, so exports and output fall. With both wedges present, their deflationary effects reinforce each other, whereas the pricing wedge's real depreciation and the asset-market wedge's real appreciation largely cancel. The economy therefore experiences sizable PPI deflation with only a small real depreciation and a small output expansion. Because the two terms-of-trade effects nearly offset, the policy rule can tip the net terms-of-trade response either way (see below). Online Appendix Figure G.4 reports further equilibrium responses.

What determines the size of the wedges? The baseline response is not a generic consequence of greater risk; it depends on trade exposure and on policy. In the symmetric steady state, greater Home import openness also means greater exposure of Home production to Foreign demand, which magnifies the differences between high- and low-tariff states that generate both wedges. Policy feedback matters because it changes inflation and depreciation after a realized tariff, and thereby the slopes, curvatures, and covariances that firms and households average over when risk rises.

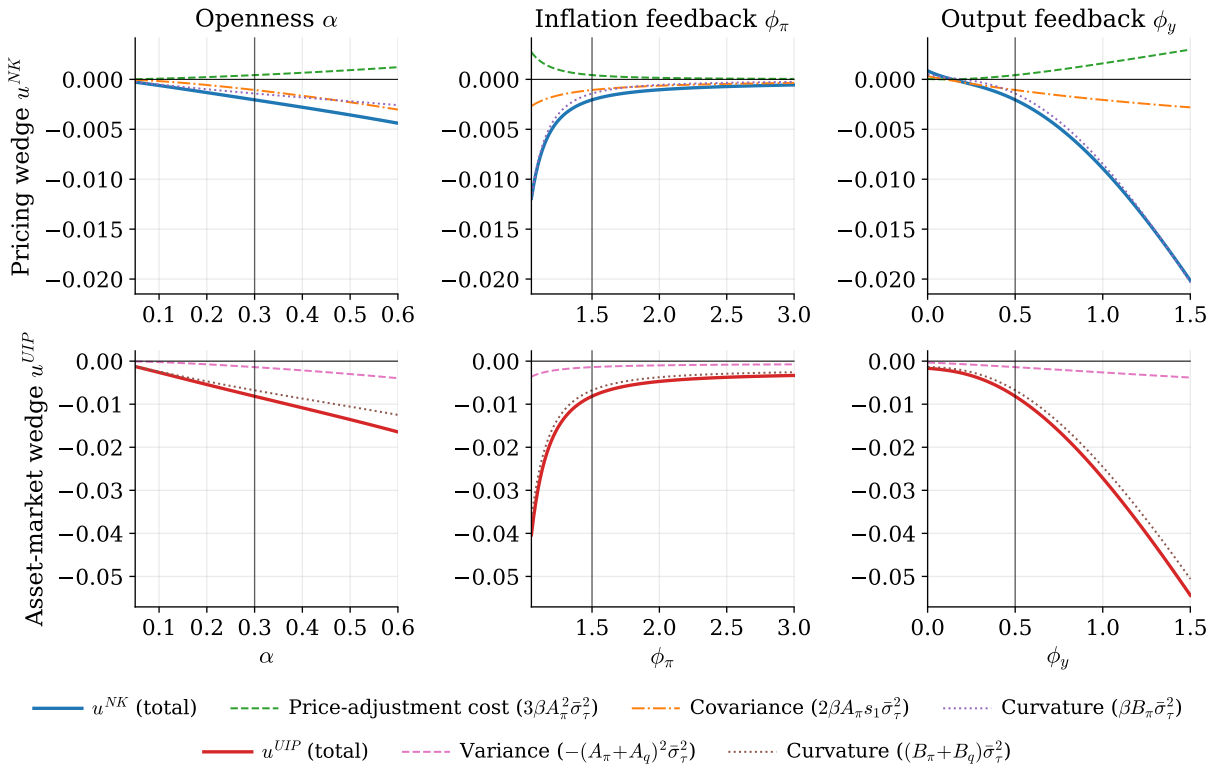


Figure 3: How trade exposure and monetary policy shape the uncertainty wedges

Notes: Each column varies one parameter while holding the others at the baseline calibration; vertical lines mark the baseline value. The top row decomposes the pricing wedge into price-adjustment-cost, covariance, and curvature terms. The bottom row decomposes the asset-market wedge into variance and curvature terms. The wedges are in model units, so their signs and component contributions are more informative than their levels. Openness is the Home expenditure share on Foreign goods. Inflation and output feedback are the coefficients in the Taylor rule (16).

Figure 3 shows how trade exposure and policy shape the two wedges. Greater openness enlarges both net wedges. The asset-market wedge keeps its sign: both of its components are negative throughout the displayed range, with the curvature component the larger. The pricing wedge, by contrast, can change sign as policy varies. The Taylor-rule coefficients change how inflation and export sales covary across tariff states, and hence the relative size of the pricing terms; when output feedback is very weak, the pricing wedge turns positive. Since the pricing wedge enters χ_q with the negative weight $-(\phi_\pi - \rho_\sigma)/D_\sigma$, a positive value reinforces the asset-market wedge's push toward real appreciation and lower exports and output. The output response then changes sign (Online Appendix Figure G.3): with no output feedback, the uncertainty shock is contractionary (Online Appendix Figure G.2).

A first-order representation. The two wedges also let standard linear models mimic the effects of tariff uncertainty. In the linearized equilibrium, the pricing wedge enters where a

productivity disturbance would, and the asset-market wedge where a UIP (risk-premium) disturbance would. When the tariff level, the volatility state and these two shocks share a common persistence, a fixed-weight combination of the two familiar shocks reproduces the third-order uncertainty responses of PPI inflation, the terms of trade, and the variables recovered from them, except those whose first-order equation contains productivity directly, such as labor and marginal cost. In the baseline, neither shock alone matches both the nominal and the real responses. The representation is interpretive: it does not claim that productivity or risk-premium shocks are the primitive disturbance. Online Appendix Proposition E.1 and Figure E.1 derive the weights and verify the fit.

3.4 Tariff level versus uncertainty shocks

So far we have studied a pure tariff-*volatility* innovation with the export-tariff level held fixed. Actual trade-policy episodes rarely deliver such a clean experiment: a realized tariff change usually arrives together with a shift in the perceived distribution of *future* tariffs. This subsection compares the uncertainty shock of Definition 1 with the more familiar tariff-*level* shock, a first-order surprise change in the current export tariff. The two shocks load on different features of the same equilibrium and need not move inflation or output in the same direction.

The contrast is cleanest when both shocks are read off the same decision rules. Along a deterministic export-tariff path with no innovations, the complete-markets equilibrium admits the expansion in (29),

$$\pi(\tau) = A_\pi\,\tau + \tfrac{1}{2}B_\pi\,\tau^2 + O(\tau^3), \qquad q(\tau) = A_q\,\tau + \tfrac{1}{2}B_q\,\tau^2 + O(\tau^3), \tag{33}$$

with tariff-exposure slopes (A_π, A_q) and curvature coefficients (B_π, B_q). A tariff *level* shock moves the economy *along* this schedule. Consider a one-standard-deviation surprise, $\varepsilon_{\tau,0} = 1$, with the volatility state held at zero. The log export tariff then follows $\tau_t = \rho_\tau^t \bar{\sigma}_\tau$, rising by $\bar{\sigma}_\tau = 0.10$ on impact. To first order, $\pi_t = A_\pi \tau_t$, $q_t = A_q \tau_t$, and, from the static block, $y_t = (\theta_q A_q - \theta_\tau)\tau_t$, so the impact signs are the signs of these coefficients. In the baseline, $A_\pi > 0$, $A_q > 0$, and $\theta_q A_q - \theta_\tau < 0$: a surprise export-tariff increase raises inflation, depreciates the Home currency in real terms, and lowers output, because the real depreciation only partly offsets the direct loss of export demand.

The shock raises inflation despite lower demand for Home goods. Under complete risk sharing, $C_t^\sigma = Q_t/G_t$, where $G_t \equiv P_t/P_{H,t}$ is the consumer price level relative to the producer price level, so real marginal cost (8) is proportional to $Q_t Y_t^\varphi$: it rises one for one with the terms of trade and with elasticity φ in output. The tariff lowers output, which on its own

lowers marginal cost. But because output falls, the rule's output term keeps the rise in the policy rate below the rise in expected inflation, so the real interest rate falls and the Home currency depreciates in real terms, which raises marginal cost. A more persistent tariff keeps the real rate low for longer and so depreciates the currency more. Solving (31) gives $\text{sign}(A_\pi) = \text{sign}\{\phi_y - \varphi(1-\rho_\tau)\}$ (Online Appendix D.6). In the baseline, $\phi_y = 0.5$ exceeds $\varphi(1-\rho_\tau) = 0.1$, so the depreciation dominates and the realized tariff is inflationary; under a rule with no response to output, it would be deflationary.

A level shock thus has a first-order effect, which a log-linear approximation captures; in the baseline, higher-order terms change its size (1.39 rather than 1.20 percentage points on impact) but not its sign.

A tariff *uncertainty* shock is qualitatively different. By Definition 1, the realized log-tariff path is unchanged ($\tau_t \equiv 0$), so the slopes (A_π, A_q) no longer move the economy by themselves. What moves is the conditional dispersion $\text{Var}_t(\tau_{t+1}) = \bar{\sigma}_\tau^2 e^{2\sigma_{T,t}}$. Because the pricing and asset-market blocks are forward-looking, this dispersion moves the equilibrium only through Jensen corrections in their nonlinear expectations. These corrections are the two wedges u^{NK} and u^{UIP} of Section 3.2. Both are built from the slopes and curvatures in (33), which determine how far apart high- and low-tariff states lie and whether averaging across them raises or lowers each expectation, and (32) propagates them to (χ_π, χ_q). At second order, the volatility state enters only multiplied by future tariff innovations, whose expectation is zero, so the uncertainty response vanishes at first and second order and first appears at third order. Its signs are set by the signs and relative strength of the two wedges under the policy rule (Proposition 1), and there is no reason for them to match the slope conditions of a level shock.

In the baseline the two shocks move both inflation and output in opposite directions. The positive slope A_π makes a realized export-tariff increase inflationary. A volatility innovation instead reweights future tariff states inside the NKPC: high-tariff states are low-sales states, and, as explained in Section 3.2, the covariance term and the local concavity $B_\pi < 0$ dominate the Rotemberg adjustment-cost term, so the pricing wedge u^{NK} is deflationary. The asset-market wedge u^{UIP} also lowers inflation, by appreciating the Home currency and so lowering marginal cost, and together they make the uncertainty shock deflationary (Section 3.3). The level shock contracts output, whereas the uncertainty shock expands it only slightly, because its two wedges largely offset in output.

Figure 4 shows the sign difference. Because tariff actions typically also raise uncertainty about further tariff changes, an observed episode combines the two shocks; under the Taylor rule, the deflationary uncertainty component offsets much of the inflationary effect of the realized tariff.

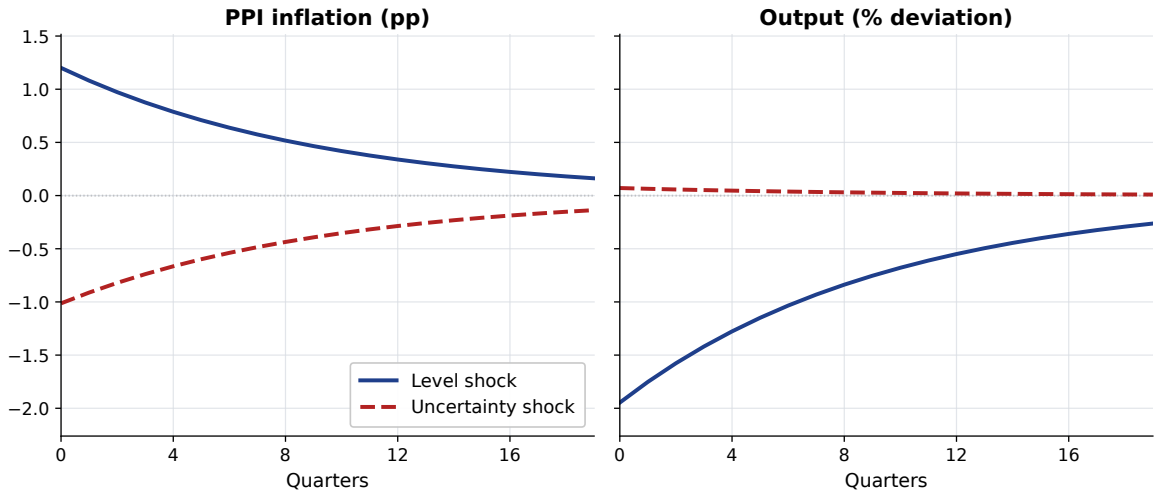


Figure 4: Export-tariff level versus uncertainty shocks

Notes: Impulse responses under the baseline calibration (Online Appendix Table B.1) and the PPI-based Taylor rule (16). The *level* shock (solid blue) is a first-order one-standard-deviation surprise increase in the export tariff T_t^* ($\tau_0 = \bar{\sigma}_\tau = 0.10$), computed from $\pi_t = A_\pi \tau_t$, $q_t = A_q \tau_t$, and $y_t = (\theta_q A_q - \theta_\tau)\tau_t$; at third order the impact PPI response is 1.39 percentage points (Table 1). The *uncertainty* shock (dashed red) is the one-standard-deviation tariff-volatility innovation of Definition 1, computed from the third-order closed form in Proposition 1 with the realized log-tariff path held fixed. The panels share a vertical scale, but the two shocks have different normalizations, so the panels are not a unit-for-unit magnitude comparison. PPI inflation is in quarterly percentage points (not annualized); output is in percent log deviations.

4 Optimal Monetary Policy

The preceding sections hold the interest-rate rule fixed and ask how tariff uncertainty is transmitted. We now ask how a central bank with commitment should respond, both today and after each possible future tariff realization. The central bank chooses one state-contingent plan at date 0, subject to private-sector equilibrium after every history, and later policy honors that plan rather than reoptimizing once a tariff is realized.

To define the problem, let h^t be a history through date t and write $q_t \equiv \log Q_t$, $\pi_t \equiv \log \Pi_{H,t}$, and $\tau_t \equiv \log T_t^*$, holding the Home import tariff fixed (the same argument applies when the import tariff varies instead). Under complete markets, the static equilibrium conditions recover consumption and output from (q_t, π_t, τ_t). Holding productivity at its normalized value $A_t = 1$, as in the tariff experiments, labor equals output by (7). The resulting exact period utility is

$$\mathcal{V}(q_t, \pi_t; \tau_t) \equiv U(C(e^{q_t}), Y(q_t, \pi_t, \tau_t)) \,. \tag{34}$$

Let $\mathcal{F}_t(h^t) = 0$ denote firms' exact nonlinear price-setting condition after substituting these

static allocations. This restriction is what prevents the central bank from choosing inflation and the real allocation independently. No approximation has been made so far. Because the nominal rate is unconstrained, the household Euler equation then recovers the interest rate that implements any allocation satisfying this restriction.

The Ramsey commitment policy solves

$$\max_{\{q_t(h^t),\pi_t(h^t)\}_{t\geq 0,h^t}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \mathcal{V}(q_t(h^t),\pi_t(h^t);\tau_t(h^t)) \quad \text{subject to} \quad \mathcal{F}_t(h^t)=0 \quad \text{for every } (t,h^t). \tag{35}$$

The plan starts a new commitment at date 0. The Lagrange multiplier on the price-setting constraint records policy promises; its inherited value at date 0 is zero because no promise was made before that date, and it becomes an endogenous history-dependent state afterwards. Online Appendix F derives the exact first-order conditions from the primitive resource and price-setting equations, states the initial and terminal conditions it imposes, and solves the local policy response through third order.

4.1 Exact benchmark: divine coincidence under tariff uncertainty

When the production correction can adjust with the tariff, stabilizing domestic producer prices also delivers the allocation that maximizes welfare subject to the economy's trade constraints. This is the New Keynesian *divine coincidence*: absent an independent cost-push distortion, price stability eliminates adjustment costs and closes the welfare-relevant output gap (Clarida, Galí and Gertler, 2001; Blanchard and Galí, 2007). In a small open economy where exporters set prices in their own currency, the corresponding target is domestic PPI inflation, provided the flexible-price allocation is efficient (Galí and Monacelli, 2005). Under tariff risk, stabilizing PPI implements the allocation that private agents would choose under flexible prices; it is optimal if that allocation also maximizes welfare in every tariff state, including states that agents consider possible but that may not be realized.

Let $Q^f(\tau)$ be the terms of trade in the flexible-price equilibrium when the current tariff is τ, and let $Q^e(\tau)$ be the constrained-efficient terms of trade. In this benchmark the production correction s^p of (9) may vary with the tariff; the fixed correction is a special case. Both terms of trade depend on the current tariff but not on tariff volatility. Strict PPI targeting is exactly optimal if the two coincide at every tariff the process can reach:

$$Q^f(\tau) = Q^e(\tau) \quad \text{for every reachable tariff } \tau. \tag{36}$$

We call this the state-by-state efficiency condition. Equivalently, the production correction

must make firms' flexible-price marginal-cost condition coincide with the planner's condition at $Q^e(\tau)$ in every tariff state. Online Appendix F.2 derives this equivalence from preferences, demand, technology, and the resource constraint.

Proposition 2 (Divine coincidence under tariff uncertainty)**.** *Suppose international asset markets are complete, prices are set in producer currency, there is no primitive cost-push shock, the nominal rate is unconstrained, and productivity and the auxiliary disturbance are constant ($A_t = 1$, $\psi_t = 0$). Suppose also that, in every tariff state, the static planner has a unique global optimum, the flexible-price equilibrium condition has a unique interior solution, and* (36) *holds. Then strict PPI inflation targeting is exactly Ramsey optimal:*

$$\Pi^R_{H,t} = 1, \qquad Q^R_t = Q^e_t, \qquad R^R_t = \left\{ \beta \mathbb{E}_t \left[\frac{Q^e_t}{Q^e_{t+1}} \right] \right\}^{-1}. \tag{37}$$

For a pure tariff-uncertainty innovation, along the path on which no tariff innovation is realized, PPI inflation and the real allocation do not respond at any order; in particular, their third-order impulse responses are zero. The nominal rate may nevertheless move to implement the allocation through the household Euler equation.

Strict PPI eliminates price-adjustment costs in every tariff state, and condition (36) makes the resulting flexible-price allocation welfare maximizing in each of those states, including those in which a tariff is realized. Strict PPI also removes the pricing wedge itself: with PPI inflation at zero after every future tariff realization, tariff risk has no future inflation to act on. Under complete markets, risk sharing makes this allocation the solution of a static problem in each tariff state: it depends on the current tariff state but not on tariff volatility, so a pure rise in tariff uncertainty leaves it unchanged. Because this allocation maximizes welfare in every state, no policy can raise welfare by reallocating resources across dates or states. Tariff risk still enters the asset-pricing condition (13), but under strict PPI it can move only the nominal rate, which implements this allocation. The result is exact: it does not rely on a quadratic welfare approximation or a third-order policy solution. Online Appendix F.3 gives the proof.

Condition (36) can hold if fiscal policy adjusts with the tariff state: a state-contingent production correction can decentralize the efficient flexible-price allocation after every tariff history. It can also hold for special parameter combinations that make the required correction constant.[7] The monetary authority does not choose this correction; the proposition treats

[7]For example, when the experiment varies the Home import tariff while the export tariff is fixed, $\gamma\sigma = 1$ makes the planner's marginal-cost target equal to $1 - \alpha$ in every import-tariff state, so a constant production correction delivers the divine coincidence despite nonzero tariff risk (Online Appendix F.2). The same restriction does not remove the export-tariff distortion in our baseline: with export-tariff variation, the

the fiscal regime as part of the economic environment.

4.2 General case: small deviation from PPI inflation targeting

We now fix the production correction, the production tax or subsidy of Section 2, at the level that makes the flexible-price allocation efficient at the no-tariff steady state (Online Appendix F.2). A Foreign tariff lowers export demand and so changes the allocation that the planner would choose. Because the correction does not adjust with the tariff, flexible prices no longer implement the efficient allocation once a tariff is realized, and strict PPI targeting is not exactly optimal.

The fixed correction matters even when no tariff change is realized, because Ramsey policy would depart from price stability if a tariff change were realized. Firms set prices before the next tariff innovation, so they take into account how monetary policy has committed to respond to each possible tariff. Because firms' forward-looking pricing term depends nonlinearly on the future tariff, through these committed responses, a wider tariff distribution changes its expected value, and with it firms' pricing today. Ramsey policy therefore moves the terms of trade and output even though today's tariff and today's efficient allocation are unchanged. In the baseline, higher tariff risk shifts firms' pricing toward deflation, and Ramsey leans against this shift with a small, persistent real depreciation (a rise in the terms of trade), which raises real marginal cost (under complete risk sharing, real marginal cost rises one for one with the terms of trade and, with elasticity φ, with the accompanying rise in output) and keeps PPI inflation close to zero (Figure 5).

To solve the Ramsey problem, we first characterize how the committed allocation responds when a tariff innovation is actually realized. Tariff risk enters through the slope and the curvature of that response with respect to the realized tariff. Averaging these state-contingent responses over the wider tariff distribution, with commitment starting at date 0, gives the response to an uncertainty shock. With $\sigma_{T,t} = \eta_\sigma \rho_\sigma^t$ denoting the response of log tariff volatility (Definition 1), define

$$v_t \equiv \bar{\sigma}_\tau^2 \sigma_{T,t} = v_0 \rho_\sigma^t, \qquad v_0 \equiv \bar{\sigma}_\tau^2 \eta_\sigma. \tag{38}$$

Because next period's tariff innovation has conditional variance $\bar{\sigma}_\tau^2 e^{2\sigma_{T,t}}$, v_t equals, to third order, one-half of the change in that variance. It is the only exogenous risk term in the Ramsey equations, so a pure volatility innovation affects the allocation only from third order on (Online Appendix D.4.1).

Let Δ^U denote the response to the volatility innovation on the *no-realization path*, along

planner's target still moves with the Foreign tariff (same Online Appendix section).

which no tariff change is realized although agents continue to expect that one might. As in Section 2.8, it is measured relative to the same path without the innovation (the *no-innovation path*). As in Proposition 2, asset markets are complete, prices are set in producer currency, the nominal rate is unconstrained, and productivity and the UIP disturbance stay at their steady-state values ($A_t = 1$, $\psi_t = 0$). The export tariff and its volatility follow (18)–(19), with the symmetric innovations of Assumption 1 and the joint scaling of Section 2.8. The production correction is fixed, and the central bank solves the commitment problem (35). Online Appendix F.4 states the local regularity conditions.

Proposition 3 (Ramsey response on the no-realization path)**.** *Under the preceding joint perturbation scaling, the local regularity conditions, and the terminal condition that rules out explosive policy paths, the first- and second-order coefficients of the incremental volatility response are zero. Conditional on no subsequent tariff innovations being realized, the leading response is*

$$\Delta^U q_t = v_0 \Gamma_v \sum_{j=0}^{t} \rho_R^{t-j} \rho_\sigma^j + O(\delta^4), \tag{39}$$

$$\Delta^U \pi_t = -\omega_R \big(\Delta^U q_t - \Delta^U q_{t-1}\big) + O(\delta^4), \qquad \Delta^U q_{-1} = 0, \tag{40}$$

for $t \geq 0$. Here $\rho_R \in (0,1)$ is the persistence with which a policy promise made under the commitment affects later allocations, Γ_v maps the variance shift into the terms of trade, and the targeting coefficient $\omega_R > 0$ measures the welfare trade-off between PPI inflation and changes in the terms of trade. Equations (42) *and* (43) *give these coefficients in closed form; Γ_v is obtained analytically from the slope and curvature of the Ramsey allocation after a realized tariff. If flexible prices implement the efficient allocation after every tariff history, then $\Gamma_v = 0$ and Proposition 2 applies: strict PPI targeting is exactly optimal.*

Proposition 3 gives the entire local path in closed form, as a combination of geometric terms in the volatility persistence ρ_σ and the commitment root ρ_R. To define the coefficients, let $\mathcal{D}(q, \pi, \tau) \equiv Y(q, \pi, \tau) e^{-q}$ and let bars denote steady-state values. Define

$$\begin{aligned} a_0 \equiv \kappa \bar{\mathcal{D}}, \qquad b_0 \equiv (\epsilon - 1)\bar{\mathcal{D}}(1 + \varphi\theta_q), \qquad h_q \equiv -\mathcal{V}_{qq}|_{(0,0,0)}, \qquad h_\pi \equiv -\mathcal{V}_{\pi\pi}|_{(0,0,0)}, \\ \lambda_R \equiv \frac{b_0^2 h_\pi}{a_0^2 h_q}, \qquad P_R(r) \equiv \beta r^2 - (1 + \beta + \lambda_R) r + 1. \end{aligned} \tag{41}$$

Here $a_0 = \mathcal{F}_\pi$ is the inflation slope of firms' pricing restriction $\mathcal{F}$ and $b_0 = -\mathcal{F}_q$ is the magnitude of its terms-of-trade slope; h_q and h_π are the corresponding welfare curvatures. The model's primitives make a_0, b_0, and h_π positive, and h_q is positive at the baseline

($h_q \approx 0.79$); we maintain $h_q > 0$ throughout, so $\lambda_R > 0$. The commitment persistence ρ_R (the stable root of $P_R(r) = 0$) and the targeting coefficient are

$$\rho_R = \frac{1 + \beta + \lambda_R - \sqrt{(1 + \beta + \lambda_R)^2 - 4\beta}}{2\beta}, \qquad \omega_R = \frac{a_0 h_q}{b_0 h_\pi}. \tag{42}$$

The risk loading Γ_v comes directly from the Ramsey response to a realized tariff. Firms setting prices at date 0 anticipate how inflation and the terms of trade will respond to next period's tariff, so a wider tariff distribution shifts their pricing condition even though the expected log tariff is unchanged. Let g_v denote the resulting shift in firms' date-0 pricing condition. It depends on the impact slopes $q_{1,0}$ and $\pi_{1,0}$ of the terms of trade and inflation with respect to a realized tariff and on the curvature $\pi_{2,0}$ of impact inflation (Online Appendix equation (F.75)). Then

$$\Gamma_v = \frac{g_v b_0 h_\pi \rho_R}{a_0^2 h_q (1 - \beta \rho_R \rho_\sigma)}. \tag{43}$$

Through the slope of the Ramsey response to a realized tariff, a wider tariff distribution changes two things: how future inflation covaries with firms' discounted sales, and, because the adjustment-cost term in the Phillips curve is convex in inflation, the expected value of that term. Together these form the conditional-covariance contribution to g_v. Through the curvature, it changes expected future inflation itself (the conditional-mean contribution). At the baseline both are positive, so $g_v > 0$ and $\Gamma_v > 0$ (Online Appendix Table F.2), which gives the deflationary shift in firms' pricing described above.

The first equation in Proposition 3 is a distributed lag with two sources of persistence. Tariff risk carries the exogenous persistence ρ_σ. Commitment adds endogenous history dependence: once policy has promised a response in possible tariff states, the effect of that promise on later allocations decays geometrically with persistence ρ_R, which is about 0.18 at the baseline, so a promise fades within a quarter or two. The second equation is the Ramsey targeting relation: PPI inflation moves with the change in the terms of trade, not with their level. When the terms of trade first rise, PPI inflation dips slightly, by less than 0.01 percentage point at the baseline; as the depreciation persists, inflation stays close to zero. The initialization $\Delta^U q_{-1} = 0$ states that commitment begins at date 0, with no promise inherited from an earlier plan.

Online Appendix F derives ρ_R, ω_R, and Γ_v from the nonlinear Ramsey first-order conditions at the deterministic steady state and shows how the finite-sum representation follows; Online Appendix F.11 checks it against a separately computed third-order perturbation solution.

In Figure 5, strict PPI keeps PPI inflation at zero and, because it implements the flexible-price allocation, leaves the terms of trade and output unchanged. Ramsey and the Taylor rule

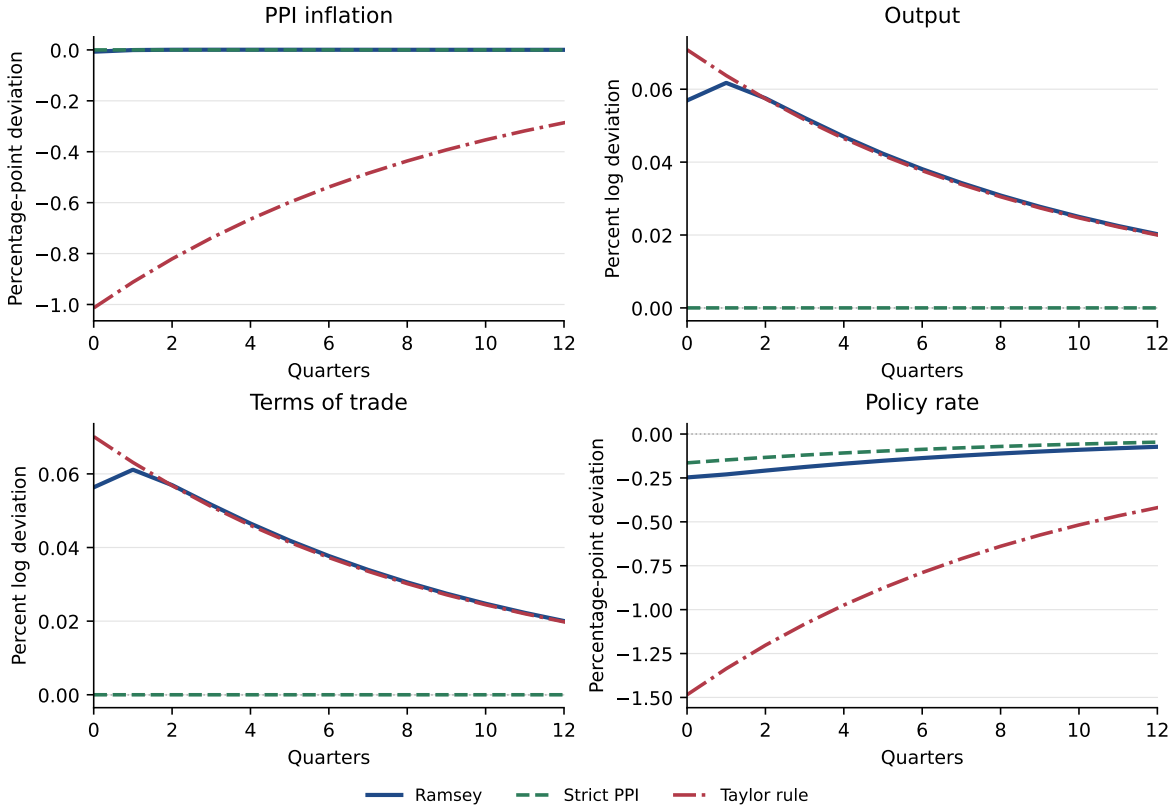


Figure 5: Pure tariff-uncertainty responses under Ramsey, strict PPI, and the Taylor rule

Notes: Complete-markets, third-order responses to uncertainty about the Foreign tariff on Home exports. The shock is the volatility innovation of Definition 1, which raises the conditional standard deviation of next quarter's tariff innovation from 0.10 to 0.20. Responses are on the no-realization path and measured relative to each policy's own no-innovation path (Section 4.2). Ramsey solves (35) with a new commitment at date 0; strict PPI imposes $\Pi_{H,t} = 1$ after every history; the PPI-based Taylor rule (16) responds to inflation and to output relative to deterministic steady-state output. A rise in the terms of trade is a Home real depreciation. Output and the terms of trade are percent log deviations; PPI inflation and the quarterly policy rate are percentage-point deviations, not annualized. Each panel has its own vertical scale.

produce very different inflation and interest-rate responses but similar real responses: from the second quarter on, both generate similarly small real depreciations and output expansions. Under complete markets, output moves with the terms of trade: to third order, $\Delta y_t = \theta_q \Delta q_t$, so the two real panels show the same adjustment in production and international relative prices.

The two real paths are similar for different reasons. Under the Taylor rule, the pricing and asset-market effects of tariff risk largely cancel in the terms of trade and output but reinforce each other in PPI inflation, as Figure 2 shows. The rule reacts to the resulting deflation with a large interest-rate cut; the small rise in output raises the rate only slightly. Under Ramsey, the deflationary shift comes instead from the risk loading $g_v > 0$, which reflects the responses the planner has promised to make if a tariff change is realized later. The small real depreciation that leans against this shift, raising marginal cost and keeping PPI inflation near zero, lasts about as long as the rise in risk. With almost no deflation to offset, Ramsey cuts the policy rate far less than the Taylor rule. Its smaller cut, like the cut under strict PPI, comes almost entirely from the direct effect of tariff risk on the Euler equation, which lowers the interest rate that implements a given allocation (Online Appendix F.6).

Under the Taylor rule the terms of trade follow $\eta_\sigma \chi_q \rho_\sigma^t$ (Proposition 1). Under Ramsey the commitment term fades quickly (ρ_R is far below $\rho_\sigma = 0.9$), and the terms of trade then follow approximately $v_0 \Gamma_v \rho_\sigma (\rho_\sigma - \rho_R)^{-1} \rho_\sigma^t$ (Proposition 3). At the baseline the two coefficients are almost equal (0.070 and 0.071 percent), so the paths nearly coincide after the impact quarter, although Ramsey's impact response is smaller (0.056 against 0.070 percent). Online Appendix Figure F.2 shows how volatility persistence and price rigidity shape the two paths.

Strict PPI is therefore close to Ramsey in the inflation panel but forgoes Ramsey's small real adjustment, because the flexible-price allocation it implements depends on the realized tariff but not on tariff volatility. Section 4.5 compares the policies by welfare.

4.3 Two restricted Ramsey benchmarks

Two restricted Ramsey benchmarks show where the small but persistent real depreciation under Ramsey comes from. A Ramsey plan has two parts: the response on impact, at date 0, and the policy it promises after every future tariff history, which we call the continuation. Each benchmark restricts one part. The *Level-Ramsey benchmark* restricts the continuation: from date 1 onward, policy follows the linear-quadratic (LQ) rule that is optimal for small realized changes in the tariff level (hence "Level"), and the date-0 response minimizes the corresponding quadratic welfare loss given that rule. The *zero-impact-PPI benchmark* restricts impact inflation: the rise in tariff risk may not move date-0 PPI inflation, and the central

bank chooses its promises about future policy optimally under that restriction.

Level-Ramsey benchmark. This benchmark asks how the central bank should respond on impact to higher tariff risk if, from date 1 onward, it must follow the policy that is optimal for small realized tariff changes. The rule is fixed, but the continuation allocation remains state contingent because it responds to every realized tariff history. Full Ramsey differs because it also chooses, at date 0, how its promised responses bend across tariff states, that is, how much more it would respond to a large tariff change than to a small one; the linear Level-Ramsey rule fixes that curvature.

The LQ problem uses the same utility and price-setting conditions that define Ramsey in (34)–(35). Let y_t^a denote log output net of Rotemberg resource costs, that is, the output used for Home absorption and exports, let $y_t^e(\tau_t)$ denote its constrained-efficient level at the current tariff, the level a Home planner would choose given the economy's trade constraints, and define the welfare-relevant gap $\widetilde{y}_t^a \equiv y_t^a - y_t^e(\tau_t)$. A second-order expansion of period utility around the constrained-efficient allocation, including the Rotemberg resource cost, and a first-order expansion of firms' exact pricing condition give the loss and constraint

$$\mathcal{L}_0^{LQ} = \frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\lambda_\pi \pi_t^2 + \lambda_y (\widetilde{y}_t^a)^2\right], \tag{44}$$

$$\pi_t = \kappa_y \widetilde{y}_t^a + \beta \mathbb{E}_t \pi_{t+1} + z_t. \tag{45}$$

The loss is normalized by the preference weight χ in (4). Here $\lambda_\pi = \kappa$ is the welfare weight generated by price-adjustment costs (κ is the Rotemberg adjustment-cost parameter), λ_y is the curvature of utility along the open-economy feasible allocation (positive at the baseline calibration, as the second-order condition requires), $\kappa_y > 0$ is the Phillips-curve slope, and z_t is a tariff-induced cost-push term: the pricing distortion that a realized tariff creates because the fixed production correction does not adjust. Online Appendix F.8 derives κ_y, z_t, and λ_y from the primitive utility, goods-market, and pricing equations.

Minimizing (44) subject to (45), with commitment beginning at date 0, gives the Level-Ramsey targeting rule

$$\lambda_\pi \pi_t + \frac{\lambda_y}{\kappa_y}(\widetilde{y}_t^a - \widetilde{y}_{t-1}^a) = 0, \qquad \widetilde{y}_{-1}^a = 0. \tag{46}$$

As in the standard commitment solution, the rule trades inflation off against the change in the welfare-relevant gap.

A pure volatility innovation leaves today's tariff unchanged but still moves firms' date-0

pricing condition through two channels. It changes expected date-1 inflation under the Level-Ramsey continuation, and it shifts the condition directly, through the same conditional-mean and covariance terms that create the pricing wedge of Section 3.2, now evaluated under the Level-Ramsey continuation. We denote this direct third-order shift by $u_0^{NK,L}$; like the pricing wedge, it is negative when tariff risk is deflationary. Taking the continuation from date 1 onward as given, the central bank minimizes the loss by splitting the total date-0 pricing shift between PPI inflation and the output gap, subject to the date-0 pricing condition. Online Appendix F.8 gives the closed form (equation (F.121)).

At the baseline calibration, $u_0^{NK,L} < 0$: higher tariff risk shifts firms' pricing condition toward deflation, and the Level-Ramsey optimum accepts a small PPI deflation (about 0.006 percentage point on impact) and a positive output gap. The positive gap raises marginal cost and offsets part of the deflationary pricing shift.[8]

Zero-impact-PPI benchmark. The second benchmark asks whether Ramsey's real response depends on letting uncertainty move inflation on impact. It sets to zero only the *additional* date-0 PPI response caused by the uncertainty innovation. Let $\Delta^U \pi_0$ denote the difference between impact inflation with and without the date-0 volatility innovation, holding baseline tariff risk $\bar{\sigma}_\tau$ fixed. We impose

$$\Delta^U \pi_0 = 0, \tag{47}$$

and leave all other date-0 choices free; total inflation need not be zero, because baseline tariff risk can itself move inflation. From date 1 onward, policy satisfies the full Ramsey first-order conditions, starting from the promise selected by the restricted date-0 problem, and keeps that promise after every tariff realization rather than re-optimizing. Online Appendix F.9 gives its closed form.

Figure 6 shows which part of the Ramsey plan drives its response to uncertainty. The Level-Ramsey benchmark delivers about four-fifths of full Ramsey's terms-of-trade and output response at every date. A central bank that follows the optimal rule for realized tariff changes, and chooses its impact response optimally, therefore already delivers most of Ramsey's real response to tariff risk. The reason is that firms and households average over the future allocations promised by the continuation rule, so policy toward *realized* tariffs feeds back into date-0 decisions under tariff *risk*. The remaining fifth comes from the curvature that full Ramsey chooses for its promised responses.

The zero-impact-PPI benchmark isolates the role of Ramsey's impact response. It removes

[8]Because tariff risk operates through curvature, we evaluate the efficient target $y_t^e(\tau_t)$ through third order before applying the linear targeting rule.

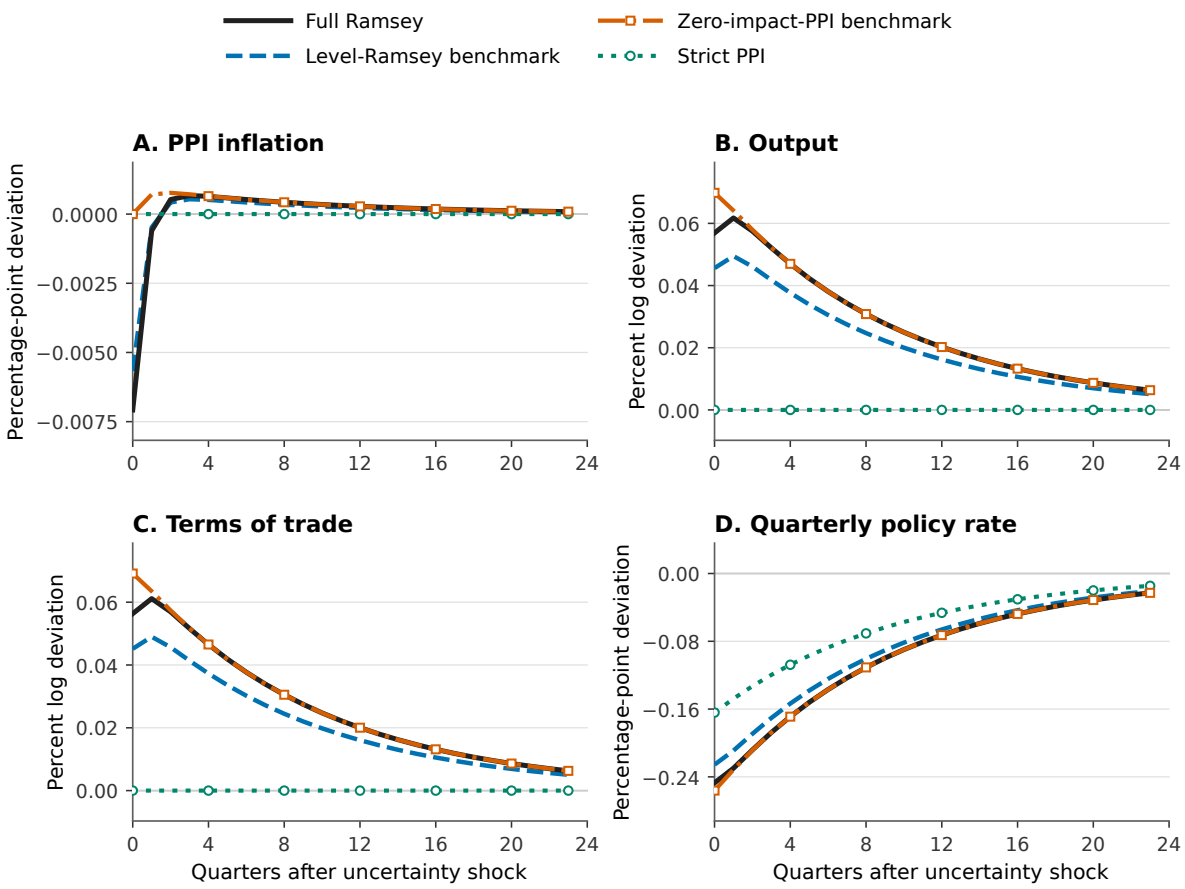


Figure 6: Ramsey, strict PPI, and two restricted Ramsey benchmarks

Notes: Complete-markets, third-order responses to the date-0 tariff-volatility innovation of Figure 5 (Foreign tariff on Home exports; current tariff unchanged; later realized tariff innovations set to zero). Each line is the shocked path minus the same policy's no-innovation path; units are as in Figure 5. Solid black: full Ramsey. Dashed blue: Level-Ramsey benchmark, which chooses the LQ-optimal date-0 inflation–output-gap pair and then applies the fixed Level-Ramsey rule from date 1 onward. Orange dash-dotted with squares: zero-impact-PPI benchmark, which sets the additional date-0 PPI response to zero and then continues the state-contingent plan selected under that restriction. Green dotted with circles: strict PPI after every history. Full Ramsey and both benchmarks begin commitment at date 0 with no inherited planner promise. A rise in the terms of trade is a Home real depreciation. Each panel has its own vertical scale.

the small additional movement in date-0 inflation but keeps the state-contingent continuation selected under that restriction. Its terms of trade and output stay within 0.013 percent of full Ramsey's in every quarter shown. The largest gap is on impact, where its terms of trade rise by 0.069 percent against Ramsey's 0.056: to hold impact PPI inflation at zero rather than accept Ramsey's slight deflation, it needs a somewhat larger depreciation to raise marginal cost. Ramsey's real response thus comes not from its tiny impact inflation movement but from the promises about future policy made at date 0, and most of their effect is already

contained in the optimal rule for realized tariffs that the Level-Ramsey benchmark follows. Optimal policy therefore departs from strict PPI targeting mainly to respond to realized tariffs, not to tariff uncertainty: its response to tariff risk is a by-product of the responses it promises to realized tariffs, which firms and households anticipate. Strict PPI removes both margins, the impact response and these promises, by requiring price stability after every history. On the displayed no-realization history, its output and terms of trade therefore stay on the no-innovation path, and the interest rate merely implements that allocation.

4.4 Tariff levels versus tariff uncertainty under optimal policy

Section 3.4 showed that, under the Taylor rule, a realized export-tariff increase is inflationary whereas a rise in tariff uncertainty is deflationary, so the two partly offset in a combined episode. We now ask how optimal policy treats the two shocks. One targeting rule covers both, written in terms of the gap between the terms of trade and their constrained-efficient value defined in Section 4.1.

Proposition 4 (Ramsey targeting rule with tariff cost-push)**.** *Maintain the assumptions of Proposition 3, in particular complete markets, producer-currency pricing, the export-tariff experiment, the fixed production correction, and a new commitment at date 0. Let* $q^f(\tau) \equiv \log Q^f(\tau)$ *and* $q^e(\tau) \equiv \log Q^e(\tau)$ *be the flexible-price and constrained-efficient terms of trade of Section 4.1 when the current log export tariff is* τ*. Let* q_1^f *and* q_1^e *be their slopes at* $\tau = 0$*, and define the terms-of-trade gap* $\widetilde{q}_t \equiv q_t - q^e(\tau_t)$*. Consider either* (L) *the first-order response to a realized date-0 tariff innovation* ζ*, so that* $\tau_t = \zeta\rho_\tau^t$*, or* (U) *the leading third-order response* Δ^U *to the volatility innovation of Proposition 3, whose remainder is* $O(\delta^4)$*; in case (U),* π_t *and* q_t *denote* $\Delta^U\pi_t$ *and* $\Delta^U q_t$*. In both cases, along the path on which no later innovation is realized, the Ramsey responses satisfy, for* $t \geq 0$*,*

$$\pi_t = \beta\pi_{t+1} + \frac{b_0}{a_0}\,\widetilde{q}_t + \mu_t, \qquad \pi_t = -\omega_R\big(\widetilde{q}_t - \widetilde{q}_{t-1}\big), \qquad \widetilde{q}_{-1} = 0. \tag{48}$$

Apart from the efficient benchmark $q^e(\tau_t)$ *in the gap, tariffs enter only through the cost-push term* μ_t *(in case (U), an endogenous term that depends on the Ramsey continuation rather than a primitive shock), which is*

$$\mu_t^L = \frac{b_0}{a_0}\big(q_1^e - q_1^f\big)\,\zeta\rho_\tau^t \quad \textit{in case (L)}, \qquad \mu_t^U = -\frac{g_v}{a_0}\,v_0\rho_\sigma^t \quad \textit{in case (U)}. \tag{49}$$

For a cost-push term $\mu_t = \mu_0 \rho^t$ with $0 \leq \rho < 1$, the unique bounded solution is

$$\widetilde{q}_t = -\frac{\rho_R \, \mu_0}{\omega_R (1 - \beta \rho_R \rho)} \sum_{j=0}^{t} \rho_R^{t-j} \rho^j, \qquad \pi_0 = \frac{\rho_R \, \mu_0}{1 - \beta \rho_R \rho}, \qquad \sum_{s=0}^{t} \pi_s = -\omega_R \, \widetilde{q}_t \xrightarrow[t \to \infty]{} 0. \tag{50}$$

Consequently, (i) strict PPI targeting is the first-order Ramsey response to a realized tariff if and only if $q_1^f = q_1^e$; otherwise, on impact, the Ramsey gap is the fraction $(1 - \rho_R)(1 - \beta \rho_R)/(1 - \beta \rho_R \rho_\tau) \in (0, 1)$ of the gap under strict PPI. (ii) The uncertainty cost-push term vanishes whenever Ramsey's impact PPI response to a realized tariff has zero slope and curvature, $\pi_{1,0} = \pi_{2,0} = 0$.

Online Appendix F.7 derives the proposition from the Ramsey first-order conditions used for Proposition 3. Case (U) is that proposition, rewritten as a cost-push problem. The last equality in (50) gives the price-level form of commitment (Woodford, 2003): the cumulative change in the log PPI is proportional to the current gap. As a result, the PPI level returns to its initial value as the gap closes.

The proposition places the two shocks in one trade-off. Under the state-by-state efficiency condition (36) of Proposition 2, both cost-push terms vanish and strict PPI is optimal for both shocks. With the fixed correction, a realized export tariff instead raises the constrained-efficient terms of trade by more than the flexible-price terms of trade: at the baseline, $q_1^e = 0.50$ against $q_1^f = 0.22$. The efficient allocation restrains Home supply to exploit Home's market power over its exports; a Foreign tariff shrinks export demand and weakens this motive, so the planner wants more output and a weaker terms of trade than the flexible-price allocation delivers when the correction is held fixed. Strict PPI implements the flexible-price allocation and so leaves a terms-of-trade gap; narrowing that gap costs PPI inflation, a standard cost-push trade-off. Werning, Lorenzoni and Guerrieri (2025) obtain a trade-off of the same form for a tariff on imported inputs; here the tariff is levied abroad on Home's exports. Tariff uncertainty, by contrast, moves neither today's flexible-price allocation nor today's efficient allocation. It matters for policy only through the risk loading g_v: higher risk makes larger tariff changes, in either direction, more likely, and firms set prices today anticipating the response the planner has promised to those changes. The uncertainty cost-push term is thus a by-product of the trade-off created by realized tariffs: each term in g_v contains $\pi_{1,0}$ or $\pi_{2,0}$, the slope or curvature of Ramsey's impact PPI response to a realized tariff, so the term is small whenever that PPI response is small, even if the terms of trade move substantially.

Table 1 quantifies the two responses at the baseline. After a 0.10 log-point export tariff, Ramsey is more expansionary than strict PPI: it cuts the policy rate by about 0.6 percentage point, against 0.2 under strict PPI, its impact terms-of-trade gap is about four fifths of strict

Table 1: Optimal and rule-based responses to tariff levels and tariff uncertainty

Shock	Impact response	Ramsey	Strict PPI	Taylor rule
Tariff level	PPI inflation (pp)	0.262	0	1.393
	Output (%)	−1.616	−2.146	−1.982
	Policy rate (pp)	−0.607	−0.201	1.099
Tariff uncertainty	PPI inflation (pp)	−0.007	0	−1.013
	Output (%)	0.057	0	0.071
	Policy rate (pp)	−0.247	−0.164	−1.485
Both	PPI inflation (pp)	0.255	0	0.380
	Output (%)	−1.559	−2.146	−1.911
	Policy rate (pp)	−0.854	−0.365	−0.386
Uncertainty offset of the level PPI response		2.7%	–	73%

Notes: Complete markets, export tariff, baseline calibration. Entries are third-order impact responses, with pruning, measured relative to each policy's own path without the innovation and conditional on no later tariff realizations. Ramsey begins a new commitment at date 0. The level shock is a 0.10 log-point innovation in the export tariff; the uncertainty shock is the volatility innovation of Definition 1; "Both" is the same-quarter joint impulse, which equals the sum of the two. PPI inflation and the policy rate are quarterly percentage points; output is a percent log deviation. The offset row is minus the ratio of the uncertainty to the level PPI response. The two shocks are normalized differently, so the rows are not a unit-for-unit comparison. Figure 4 plots the first-order level response, 1.20 percentage points under the Taylor rule; the third-order entry in the table also includes curvature in the tariff level and the interaction of the realized tariff with background tariff risk ($\bar{\sigma}_\tau^2$).

PPI's, and output falls by 1.6 rather than 2.1 percent. The cost is about 0.26 percentage point of impact PPI inflation, and it is temporary: at first order, Ramsey PPI inflation turns slightly negative two quarters after impact, and the producer-price level returns to its pre-shock value, as (50) requires. The Taylor rule instead raises the policy rate and lets inflation decay only with the tariff, so the producer-price level rises permanently. The uncertainty shock, by contrast, barely moves Ramsey's producer prices: the impact PPI response is −0.007 percentage point under Ramsey, against −1.01 under the Taylor rule. At third order, uncertainty therefore offsets 2.7 percent of Ramsey's level-shock PPI response, against 73 percent under the Taylor rule.

Through third order, a combined episode is the sum of the two responses. The optimal response to a tariff increase accompanied by higher uncertainty is therefore essentially the level response, plus a further third-order rate cut of about 0.25 percentage point. Under the Taylor rule, by contrast, observed inflation can look moderate even though the rule's responses to the level and uncertainty shocks are each far from Ramsey's.

The direction in which Ramsey departs from strict PPI after a realized tariff depends on whether it is an export or an import tariff and on whether asset markets are complete.

On impact, Ramsey is more expansionary than strict PPI after a realized export tariff in all 27 complete-markets calibrations of Online Appendix L and after an import tariff under complete markets. After an export tariff under incomplete markets, it also cuts the policy rate by more than strict PPI (Online Appendix L). After an import tariff under incomplete markets, it is more contractionary: PPI inflation falls, and the policy rate rises by more than under strict PPI (Online Appendix Table I.1). By contrast, Ramsey's impact PPI response to uncertainty remains small in every environment we compute, at most about 0.03 percentage point (Online Appendix Table L.1).

The policy message is to keep producer prices stable when tariff uncertainty rises, which calls for a modest rate cut, and to depart from price stability in response to realized tariffs, in a direction that depends on the environment. Strict PPI targeting gets the first part right and forgoes mainly Ramsey's modest departures after realized tariffs. The welfare comparison below shows how little that costs.

4.5 Welfare comparison

The welfare ranking cannot be read from the no-realization paths alone. Strict PPI visibly removes Ramsey's small real adjustment in Figures 5–6, yet remains close to Ramsey in welfare. After the impact quarter, the Taylor rule produces a real path similar to Ramsey's (Figure 5), but a much larger welfare loss. The figures condition on no tariff innovation being realized, whereas welfare integrates over all future tariff histories and includes the resource cost of price adjustment. The welfare criterion is the household utility function (4), not a separate loss function. Substituting the production identity $L_t = Y_t/A_t$ from (7), the welfare of policy j and its gap relative to Ramsey commitment (R) are

$$W^j \equiv \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(C_t^j; \sigma) - \chi \frac{(Y_t^j/A_t)^{1+\varphi}}{1+\varphi} \right], \qquad D_j \equiv W^R - W^j. \tag{51}$$

All policies start from the same initial state at date -1 and face the same tariff process. At date 0 the tariff innovation is zero and the volatility innovation raises the conditional standard deviation of next quarter's tariff innovation from 0.10 to 0.20 in the exact shock process (Definition 1). Unlike in the plotted responses, future tariff and volatility innovations remain stochastic and are integrated out. The reported welfare numbers are the corresponding fourth-order local approximations. Because fourth-order welfare depends on fourth moments of the shocks, the calculation uses Gaussian innovations.

We express D_j as a permanent consumption equivalent. With steady-state consumption normalized to $\bar{C} = 1$ and steady-state labor held fixed, λ_j is the proportional increase in

steady-state consumption whose utility value equals the gap:

$$u(1+\lambda_j;\sigma) - u(1;\sigma) = (1-\beta)D_j. \tag{52}$$

We report $10^4\lambda_j$ in basis points of permanent consumption.

Table 2 reports each policy's loss relative to Ramsey after the date-0 volatility innovation. Online Appendix F.10 splits this gap into a no-innovation baseline and the increment created by the volatility innovation (Online Appendix Table F.1). The increment measures how much the innovation widens a policy's gap to Ramsey.

Table 2: Fourth-order complete-markets welfare losses relative to Ramsey

Policy	Consumption-equivalent loss (basis points)
Ramsey commitment	0.000
Zero-impact-PPI benchmark	0.000†
Strict PPI	4.072 (0.056)
Taylor rule	331.168 (5.753)

Notes: Entries are fourth-order approximations to $10^4\lambda_j$, with λ_j defined in (52); parentheses report quasi-Monte Carlo standard errors. Ramsey is zero by construction ($D_R = 0$). † Restricting incremental date-0 PPI inflation changes the Ramsey allocation only at third order, so the envelope condition makes the loss $O(\delta^6)$, which is zero in a fourth-order approximation. Online Appendix F.10 gives the order derivation, the baseline decomposition (Online Appendix Table F.1), the estimator, and numerical checks.

The contrast is sharp. Strict PPI costs about four basis points of permanent consumption relative to Ramsey, whereas the Taylor rule costs about 331 basis points. The difference reflects how much each policy lets producer prices move. Strict PPI keeps them fixed, so it pays no price-adjustment costs, but it gives up Ramsey's modest gain from adjusting production and international relative prices across tariff states. The Taylor rule lets producer prices move a lot, both when tariffs change and when tariff volatility rises. These price movements use up output in adjustment costs and push production and relative prices away from Ramsey's. The gaps come entirely from tariff risk and, to leading order, grow with its variance; in our calibration the standard deviation of quarterly innovations to the log tariff is 0.10. Most of each gap is present even without the date-0 volatility innovation: the no-innovation gaps are 3.95 of 4.07 basis points for strict PPI and 322.8 of 331.2 basis points for the Taylor rule (Online Appendix Table F.1). The volatility innovation thus widens the Taylor rule's gap by 8.4 basis points, about seventy times the 0.12 basis points it adds to strict PPI's. With the production tax or subsidy fixed, strict PPI comes within about four basis points of Ramsey both before and after the volatility innovation.

5 Incomplete Markets

Complete markets insure Home households against tariff outcomes, so the terms of trade determine consumption through a static risk-sharing condition. We now remove that insurance: Home households can trade only a single noncontingent bond internationally. A domestic nominal bond remains, and monetary policy still sets its return. Consumption becomes a forward-looking saving choice, and net foreign assets become a state variable that links current policy to future external adjustment. We ask two separate questions. First, holding the Taylor rule fixed, how does the loss of insurance change the transmission of tariff uncertainty? Second, within the bond economy, does strict PPI stabilization still come close to Ramsey nominal stabilization?

Two conclusions emerge. Incomplete markets turn tariff risk into a self-insurance problem: precautionary saving sharply dampens the impact deflation and output expansion under the Taylor rule. Strict PPI remains close to Ramsey in nominal stabilization, and the two policies distribute the real adjustment differently.

5.1 What incomplete markets add

Demand, production, price setting, and the fixed production correction are the same as under complete markets; only the international asset-market block changes. Let $\mathcal{E}_t$ be the nominal exchange rate, and let b_t be net foreign assets: the CPI-deflated Home-currency value of the end-of-period position chosen at date t (hence predetermined at $t+1$), measured relative to steady-state output. Home households trade a domestic nominal bond and a Foreign-currency bond whose premium depends on b_t. The two Euler equations and the external-asset law are

$$1 = \beta R_t \, \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \right], \tag{53}$$

$$1 = \beta R_t^* e^{-\psi_b b_t} \, \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} \right], \tag{54}$$

$$b_t = R_{t-1}^* e^{-\psi_b b_{t-1}} \frac{\mathcal{E}_t}{\mathcal{E}_{t-1}} \frac{b_{t-1}}{\Pi_{C,t}} + tb_t, \tag{55}$$

where tb_t is the CPI-deflated trade balance. The small debt-elastic premium $e^{-\psi_b b_t}$ (Schmitt-Grohé and Uribe, 2003), with $\psi_b > 0$, lowers the return on foreign assets as the position grows, so net foreign assets return gradually to their balanced-trade steady state.

The two Euler equations play different roles. Given an allocation, the domestic Euler equation pins down the nominal rate that implements it. The Foreign Euler equation and

the asset law instead restrict the allocation itself: consumption, the terms of trade, and the accumulation of external wealth. Higher tariff risk can therefore induce precautionary saving: households reduce current consumption, run a trade surplus, and carry more wealth into future tariff states. This margin is absent when consumption is insured state by state. Online Appendix Proposition H.1 derives (53)–(55) from the household budget constraint and the two bond first-order conditions.

5.2 Transmission under the Taylor rule

Without state-contingent claims, the same uncertainty shock has much weaker impact effects on prices and output under the Taylor rule (Figure 7). The impact PPI deflation falls from about 1 quarterly percentage point to about 0.1, the output expansion is much smaller, and consumption, which rises on impact under complete markets, falls in the bond economy. The two economies differ only in their asset markets: households trade either state-contingent claims or a single international bond.

Under complete markets, households are insured against the loss of export income in high-tariff states, so the pricing and asset-market wedges operate without a saving motive: they generate a large PPI deflation, and the Taylor-rule rate cut supports a modest real depreciation and output expansion. In the bond economy, high-tariff states have high marginal utility because export income and consumption are low. The Foreign bond pays off in every future state rather than in the adverse states alone, but it is still valuable for self-insurance: an extra unit of future wealth is worth most when marginal utility is high. Greater tariff risk widens the spread of future consumption, and because marginal utility is convex in consumption, this raises expected future marginal utility and strengthens the precautionary saving motive. Households cut current consumption and imports, run a trade surplus, and accumulate net foreign assets.

This saving response weakens the initial effects relative to complete markets. The shift toward saving changes how firms value future profits, which weakens the downward pressure that tariff risk puts on their prices. PPI deflation is therefore smaller, the Taylor rule cuts the policy rate by much less, and the real depreciation and output expansion are smaller. Later, as the accumulated foreign assets finance more consumption and imports, the trade surplus unwinds, and the terms of trade and output fall below the paths they would follow without the volatility innovation.

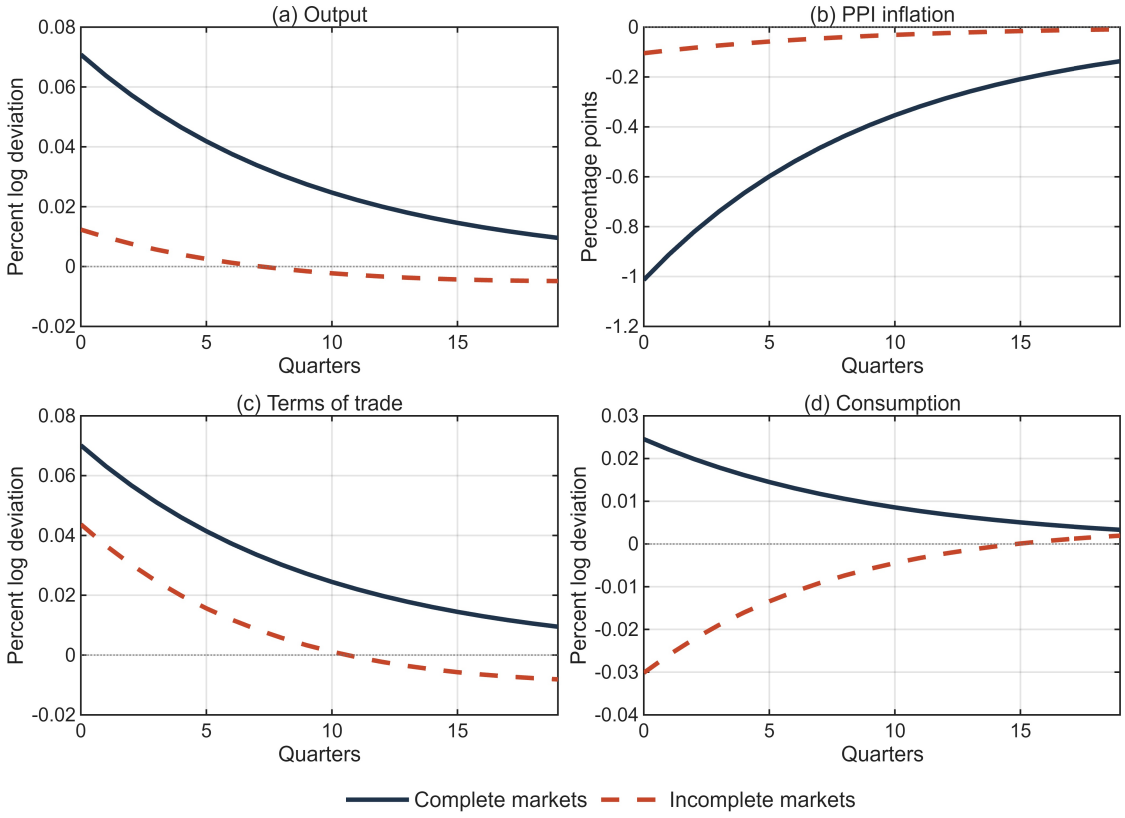


Figure 7: Tariff uncertainty and market completeness under the Taylor rule

Notes: Output, PPI inflation, the terms of trade, and consumption under the PPI-based Taylor rule (16), whose output term uses deterministic steady-state output. The shock is the date-0 volatility innovation in the Foreign tariff on Home exports (Definition 1), which raises the exact conditional standard deviation of next quarter's tariff innovation from 0.10 to 0.20; responses are local third-order approximations. The current tariff is unchanged, and later realized tariff innovations are zero while future risk remains in expectations. The solid line is the complete-markets economy and the dashed line is the bond (incomplete-markets) economy. Both start from the same deterministic physical state, and each is measured relative to its own no-innovation path. Inflation is in quarterly percentage points; the other variables are percent log deviations. A rise in the terms of trade is a Home real depreciation.

5.3 Ramsey policy under incomplete markets

The case for stabilizing producer prices carries over to the bond economy: strict PPI targeting stays close to Ramsey policy. We compare three policies in the bond economy, all starting from the same deterministic steady state: Ramsey policy, strict PPI targeting, and the Taylor rule. The question is how well each policy stabilizes PPI inflation and how it distributes the adjustment across consumption, production, and external saving.

For Ramsey policy, we study a central bank that continues the plan it committed to before tariff uncertainty rose, the timeless-perspective form of commitment: it honors promises about future policy made before date 0, which its inherited planner multipliers record. Strict PPI instead imposes $\Pi_{H,t} = 1$ after every history, and the Taylor rule sets the interest rate as a fixed function of PPI inflation and output.

Figure 8 shows that strict PPI targeting closely matches Ramsey policy in inflation. Strict PPI holds PPI inflation at zero, Ramsey policy keeps it close to zero, and the Taylor rule produces a much larger deflation. Both price-stabilizing policies still generate precautionary saving: stabilizing PPI inflation removes price-adjustment costs, not the motive to self-insure. Relative to strict PPI, Ramsey policy raises output more on impact (about 0.03 percent, against about 0.01 under both strict PPI and the Taylor rule), with a smaller initial fall in consumption and greater foreign-asset accumulation. In consumption and in peak foreign-asset accumulation, strict PPI is closer to Ramsey than the Taylor rule, which produces the largest impact fall in consumption and the largest foreign-asset accumulation over the displayed horizon. As under complete markets, Ramsey policy departs from strict PPI mainly in response to realized tariffs: a realized 0.10 log-point export tariff moves its impact PPI inflation by 0.13 percentage point, whereas the uncertainty shock moves it by only 0.002 (Online Appendix L).

Online Appendix H.5 derives this Ramsey response in closed form at third order and verifies it numerically, and Online Appendix Proposition H.3 states the conditions under which strict PPI coincides with Ramsey policy in the bond economy. Section 6 ranks the policies by welfare for import tariffs.

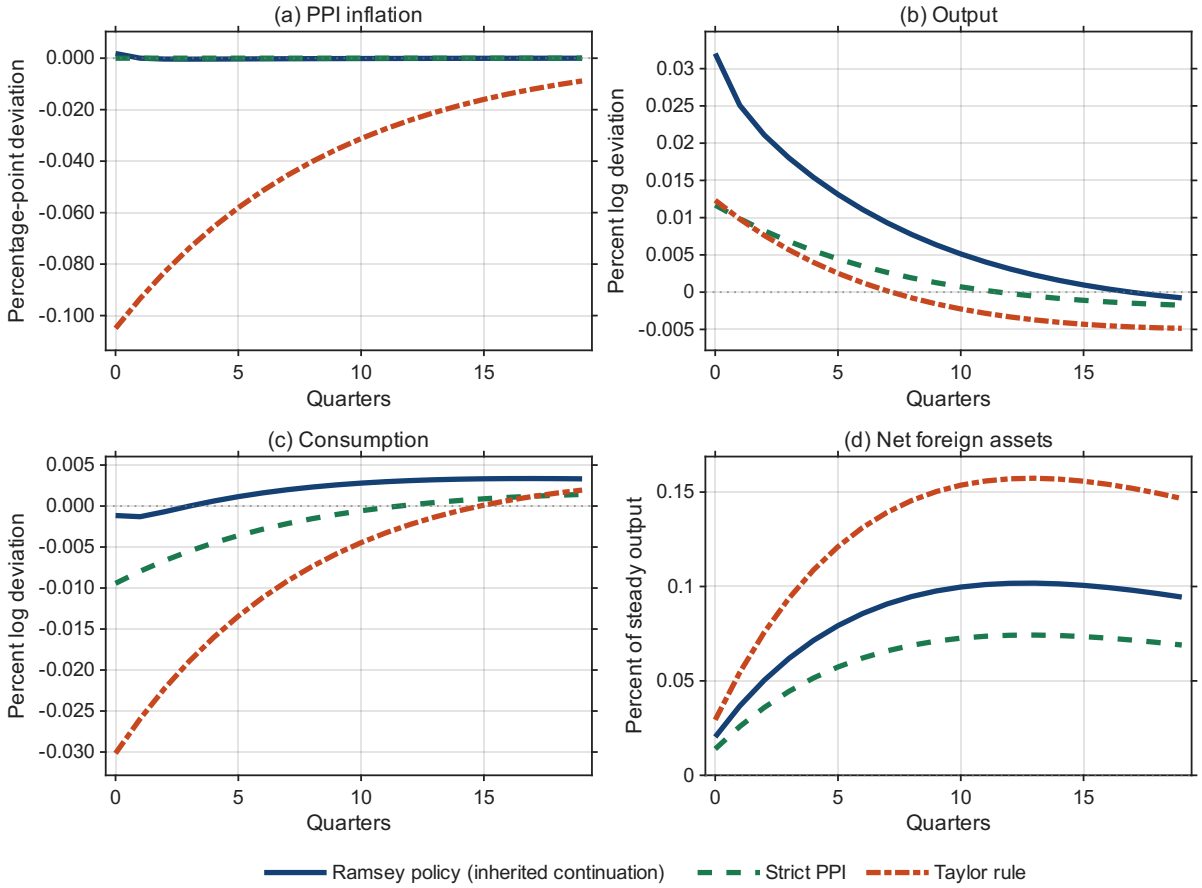


Figure 8: Tariff uncertainty under incomplete markets: Ramsey policy, strict PPI targeting, and the Taylor rule

Notes: Bond economy. The shock is the date-0 export-tariff volatility innovation of Definition 1, which raises the exact conditional standard deviation of next quarter's tariff innovation from 0.10 to 0.20; the curves are local third-order responses. Realized tariff innovations are zero along the displayed path, while agents continue to price future tariff risk. Each curve is the shocked path minus the same policy's path without the volatility innovation; all policies start from the same deterministic steady state. PPI inflation is in quarterly percentage points; output and consumption are percent log responses, and net foreign assets are a percent of steady-state output. The solid line is Ramsey policy continuing a plan chosen before date 0 (Section 5.3).

6 Import Tariffs

This section shows how the results change when Home taxes imports rather than Foreign taxing Home exports. Three findings stand out. First, the tariff-inclusive import price enters the Home consumer price index directly, so CPI inflation contains a tariff-price component that PPI inflation does not. Second, under incomplete markets, precautionary saving and current-account adjustment add to the expenditure-switching channel. Under the PPI-based Taylor rule and complete markets, a realized import tariff lowers PPI inflation and raises output. With incomplete markets, consumption no longer moves to offset the tariff's effect on marginal cost, which helps turn the PPI response positive, and the output gain essentially disappears. The response to import-tariff uncertainty, by contrast, keeps its sign and becomes larger. Third, strict PPI targeting comes much closer to Ramsey in welfare than either the PPI-based or the CPI-based Taylor rule. Online Appendix I contains the derivations, and Online Appendix I.1 defines the level and uncertainty experiments used throughout this section.

The economic difference begins with who pays and receives the tariff. A Foreign tariff on Home exports reduces foreign demand for Home goods, and its revenue accrues abroad. A Home import tariff instead raises the price of foreign goods paid by Home households. The Home government rebates the tariff revenue, together with the net proceeds of the fixed production correction, to households as a lump-sum transfer. Because the revenue returns to households, tariff receipts are not a resource loss. The tariff still distorts the allocation, however, because households face a higher relative price for imports, which changes import demand and expenditure switching. Online Appendix I.2 derives the consumer price index and records the government budget and the net household transfer.

Under complete markets, state-contingent insurance absorbs household-income risk, and the main new channel is that the tariff drives a wedge between consumer and producer prices. The PPI-based Taylor rule responds to domestic producer-price inflation, whereas the CPI-based rule also responds to the tariff-driven movement in import prices, so the two rules can prescribe different interest rates after the same tariff disturbance. Under the PPI-based rule, a realized increase in the import tariff raises Home output and lowers PPI inflation on impact. A rise in tariff risk with no realized tariff change has the same effect, so, unlike in the export-tariff case, the level and uncertainty shocks reinforce each other in output and PPI inflation. They move the terms of trade in opposite directions. Online Appendix I.3 derives the coefficient conditions that determine these signs and reports the full complete-markets responses.

Incomplete markets change the adjustment because households can no longer trade state-

contingent claims and instead insure themselves through external saving. Greater import-tariff risk can then reduce current consumption and imports, improve the trade balance, and raise net foreign assets, even when no tariff change is realized. Online Appendix Table I.1 compares the impact responses under the two financial structures. This saving channel also explains why strict PPI targeting need not remove the real effects of import-tariff uncertainty: with PPI inflation at zero, households still adjust consumption and external saving to insure against tariff risk. Online Appendix I.4 derives the foreign-bond Euler equation, the law of motion for net foreign assets, and the trade-balance dynamics, and reports the incomplete-markets responses to the uncertainty shock.

Where the tariff is levied and whether financial markets are complete determine which real margins a tariff shock moves. An export tariff works mainly through foreign demand for Home goods; an import tariff also moves the consumer-price wedge; and under incomplete markets, external saving adds a further margin of adjustment.

Strict PPI again comes close to Ramsey in welfare. Its fourth-order lifetime-utility gap to Ramsey is about 1 percent of the better Taylor rule's gap under complete markets and about 0.3 percent under incomplete markets, and these shares barely change whether a level shock, an uncertainty shock, both, or neither occurs at date 0 (Online Appendix L). The comparison holds the initial state, the tariff process, and the shock normalization fixed across policies within each financial structure, and Ramsey begins a new commitment at date 0.

Ramsey's impact PPI response to an import-tariff uncertainty shock is again small: 0.0001 percentage point under complete markets and -0.03 under incomplete markets, against -0.13 to -0.16 and -0.48 to -0.71, respectively, under the two Taylor rules.[9] Online Appendix I.5 reports the welfare calculation and numerical checks.

7 Conclusion

This paper characterizes analytically how tariff uncertainty transmits through a small open economy and how monetary policy should respond. We show that in a canonical New Keynesian model, under complete markets, producer-currency pricing, and the PPI-based Taylor rule, the full nonlinear private equilibrium reduces exactly to two equations in producer-price inflation and the terms of trade. This reduction yields closed-form third-order responses to tariff uncertainty and decomposes them into two expectation wedges: a pricing (NKPC) wedge, which captures the effect of tariff risk on firms' price setting, and an asset-market

[9] Unlike the welfare comparison, the incomplete-markets Ramsey response, here and in Online Appendix I.4, continues a plan chosen before date 0; under complete markets, a new and an inherited commitment give the same responses.

(UIP) wedge, which captures its effect on the exchange rate.

Under the Taylor rule, both wedges lower PPI inflation after an export-tariff uncertainty shock, but they move the currency in opposite directions. The pricing wedge weakens the currency through the rate cut it triggers; the asset-market wedge strengthens the currency directly, more than offsetting the rate cut it induces. Their effects on exports and output therefore nearly cancel: PPI inflation falls sharply while output barely moves, and the sign of the output response depends on the rule's output coefficient.

Optimal policy, by contrast, is simple, because tariff uncertainty, unlike a realized tariff, creates no trade-off of its own. Strict PPI targeting is exactly optimal when a production tax or subsidy aligns the flexible-price allocation with the allocation a Home planner would choose in every tariff state. When the tax or subsidy is instead fixed at the level that makes the no-tariff steady state efficient, as in our baseline, a realized export tariff weakens Home's terms-of-trade motive to restrain export supply, so the planner wants more output than flexible prices deliver. The tariff then acts as a cost-push shock that creates a trade-off between price stability and efficiency, and Ramsey policy accepts a temporary rise in PPI inflation. Tariff uncertainty leaves the current tariff unchanged and so opens no gap today between the efficient and flexible-price allocations. With no trade-off to manage, Ramsey policy offsets the expectation-driven pressure on producer prices almost fully (Proposition 3), whereas the Taylor rule lets PPI inflation fall by about one percentage point. Strict PPI targeting comes close to Ramsey in welfare, because responding to realized tariffs gains little, whereas the Taylor rule loses much more by letting producer prices move with realized tariffs and tariff risk. When tariff uncertainty rises, central banks should keep domestic producer-price inflation stable.

We then show how the mechanism changes across several policy-relevant environments. First, under the Taylor rule, tariff *level* and tariff *uncertainty* shocks move inflation in opposite directions in our baseline, so the uncertainty that accompanies a tariff increase dampens its inflationary effect. Second, incomplete markets add a precautionary-saving channel: households self-insure against tariff risk by running a trade surplus and carrying more wealth into future tariff states. Third, import tariffs enter the consumer price index directly, so PPI- and CPI-based Taylor rules prescribe different interest rates after the same tariff disturbance; under complete markets and the PPI-based rule, import-tariff level and uncertainty shocks reinforce each other in PPI inflation and output.

In the Calvo- and local-currency-pricing extensions (Online Appendices J and K), the sign of the Taylor-rule response to tariff uncertainty can change; under complete markets and Calvo pricing, for example, PPI inflation rises. Across all the environments we compute, strict PPI targeting stays close to Ramsey in nominal stabilization, and in each of our welfare

comparisons its gap to Ramsey is at most about 1.2 percent of the Taylor rule's (Online Appendix L).

Taken together, these results show that trade policy uncertainty is not mechanically inflationary or contractionary. Its effects depend on how tariff risk interacts with firms' pricing decisions, asset-market risk premia, trade exposure, market completeness, and the monetary policy regime. Monetary policy plays a central role in this transmission: by shaping the relative strength of the pricing and asset-market wedges, it can alter both the magnitude and the sign of the macroeconomic response.

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.

Barattieri, Alessandro, Matteo Cacciatore, and Fabio Ghironi. 2021. "Protectionism and the Business Cycle." *Journal of International Economics*, 129: 103417.

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.

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.

Bianchi, Javier, and Louphou Coulibaly. 2025. "The Optimal Monetary Policy Response to Tariffs." National Bureau of Economic Research Working Paper 33560.

Blanchard, Olivier, and Jordi Galí. 2007. "Real Wage Rigidities and the New Keynesian Model." *Journal of Money, Credit and Banking*, 39(s1): 35–65.

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.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2001. "Optimal Monetary Policy in Open versus Closed Economies: An Integrated Approach." *American Economic Review*, 91(2): 248–252.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc. 2010. "Optimal Monetary Policy in Open Economies." In *Handbook of Monetary Economics*. Vol. 3, ed. Benjamin M. Friedman and Michael Woodford, 861–933. Elsevier.

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. "Canada Launches Public Consultations on the Canada-United States-Mexico Agreement." Government of Canada News release September 19.

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.

Kalemli-Özcan, Şebnem, Can Soylu, and Muhammed A. Yildirim. 2026. "Global Trade, Tariff Uncertainty, and the US Dollar." *AEA Papers and Proceedings*, 116: 47–52.

Kitamura, Kyla H., and Danielle M. Trachtenberg. 2026. "USMCA Joint Review: Process and Role of Congress." Congressional Research Service Report R48787.

Lindé, Jesper, and Andrea Pescatori. 2019. "The Macroeconomic Effects of Trade Tariffs: Revisiting the Lerner Symmetry Result." *Journal of International Money and Finance*, 95: 52–69.

Merendino, Alfonso, and Tommaso Monacelli. 2026. "Tariffs, Uncertainty, and the Exchange Rate." Bocconi University Working Paper.

Monacelli, Tommaso. 2025. "Tariffs and Monetary Policy." Centre for Economic Policy Research Discussion Paper 20142.

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

Schmitt-Grohé, Stephanie, and Martín Uribe. 2003. "Closing Small Open Economy Models." *Journal of International Economics*, 61(1): 163–185.

USTR. 2025*a*. "Extends Certain Exclusions from China Section 301 Tariffs." Office of the United States Trade Representative Press release August 28.

USTR. 2025*b*. "Seeks Public Comment on the Joint Review of USMCA." Office of the United States Trade Representative Press release September 16.

USTR. 2026. "Ambassador Greer Issues Statement on the USMCA Joint Review." Office of the United States Trade Representative Press release July 1.

Werning, Iván, Guido Lorenzoni, and Veronica Guerrieri. 2025. "Tariffs as Cost-Push Shocks: Implications for Optimal Monetary Policy." National Bureau of Economic Research Working Paper 33772.

White House. 2025*a*. "Adjusting Imports of Aluminum and Steel Into the United States." The White House Presidential proclamation June 3.

White House. 2025*b*. "Adjusting Imports of Automobiles and Automobile Parts Into the United States." The White House Presidential proclamation March 26.

White House. 2026. "Adjusting Imports of Semiconductors and Semiconductor Manufacturing Equipment, and Their Derivative Products, Into the United States." The White House Presidential proclamation January 14.

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

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

Lu Han and Yuko Imura
Bank of Canada

23 October 2026
Preliminary

Disclaimer: The views expressed herein are those of the authors and not necessarily those of the Bank of Canada.

Contents

I Motivation and Foundations **1**

A Trade Policy Uncertainty Index **1**

B Full Nonlinear Equilibrium Conditions (Levels) **2**
B.1 Common block 2
B.2 International asset markets: complete versus incomplete 4
B.3 Baseline calibration 5

C Minimal Nonlinear Representation (Complete Markets) **6**
C.1 Starting point: the full nonlinear system under complete markets 7
C.2 Static block 7
C.3 Two structural equations 8
C.4 Recovering other variables from the minimal system 9

II Analytics **12**

D Proof of Proposition 1 and Its Corollaries **12**
D.1 Log representation and static derivatives 12
D.2 First-order system and closed-form coefficients 14
D.3 Second-order curvature in the tariff state 15
D.4 Second-order mean shifts and third-order uncertainty coefficients 17
D.5 Corollaries and special cases 19
D.6 Fully substituted primitive closed form for (χ_π, χ_q) 21

E A Two-Channel First-Order Representation **22**

F Complete-Markets Optimal Policy: Exact Conditions and Closed-Form Tariff-Uncertainty Responses **28**
F.1 Policy experiment and contingent histories 29
F.2 Primitive static maps and the constrained-efficient target 29
F.3 Strict PPI: implementation versus optimality 32
F.4 Exact nonlinear commitment problem 33
F.5 Closed-form tariff-level continuation 35
F.6 Third-order risk shift and closed-form response 40
F.7 Proof of the Ramsey targeting rule with tariff cost-push 44
F.8 The LQ Level-Ramsey approximation and the optimal date-0 choice 48
F.9 Incremental date-0 PPI stabilization with Ramsey conditions from date 1 onward 53
F.10 Welfare accounting and approximation orders 55
F.11 Numerical verification and economic interpretation 60
F.12 Why the baseline Taylor and Ramsey real paths nearly overlap 62

G Wedge Decomposition, Special Cases, and Comparative Statics **65**
G.1 Why uncertainty responses are third order 65
G.2 Decomposition of the uncertainty wedges 65
G.3 Special cases of the closed-form perturbation solution 65
G.4 Comparative statics of the uncertainty responses 67
G.5 Selected equilibrium responses and numerical check 70

III Extensions and Robustness 71

H Incomplete Markets: Exact Closure and Optimal Policy **71**
H.1 Experiment, approximation order, and normalization 71
H.2 Exact private equilibrium and bond closure 73
H.3 Third-order risk terms in private and planner conditions 75
H.4 Ramsey conditions and the strict-PPI test 79
H.5 Closed-form response and numerical verification 87
H.6 Requirements for a fourth-order welfare comparison 93

I Import Tariffs: Derivations and Policy Comparisons **98**
I.1 Experiments, units, and commitment boundary 98
I.2 Primitives, fiscal ledger, and price setting 99
I.3 Complete markets 100
I.4 Incomplete markets: rebated bond economy 113
I.5 Approximation order and welfare 126

J Calvo Pricing **128**
J.1 Setup 129
J.2 Higher-order coefficient system under fixed rules 131
J.3 Ramsey continuation and quantitative comparison 135

K Local-Currency Export Pricing **138**
K.1 The export-pricing block from primitives 138
K.2 A three-equation local representation 139
K.3 Ramsey continuation and quantitative comparison 141

L Strict PPI Targeting across Environments **144**

Part I

Motivation and Foundations

Numbers without a letter prefix refer to the main text; lettered numbers refer to this Online Appendix.

A Trade Policy Uncertainty Index

This appendix documents the empirical motivation for the analysis: large, persistent swings in U.S. trade-policy uncertainty over 2016–2026. Figure A.1 plots the monthly Trade Policy Uncertainty Index of Caldara, Iacoviello, Molligo, Prestipino and Raffo (2020) and marks U.S. trade-policy events relevant for Canada.

Two features motivate the focus on tariff uncertainty as a distinct disturbance.[1] First, in April 2025, the month of the "Liberation Day" tariff announcement, the index reached more than four times its pre-2024 peak. Second, the increase was persistent rather than confined to that month: the index remained elevated for several quarters and rose again in early 2026. At the first joint review of the USMCA on July 1, 2026, the United States did not confirm the agreement's extension; the USMCA remains in force, subject to annual reviews, and is scheduled to expire in 2036 unless all three parties agree to extend it (USTR, 2026; Kitamura and Trachtenberg, 2026).

[1]The index measures the frequency of joint references to trade policy and uncertainty in major U.S. newspapers. Caldara, Iacoviello, Molligo, Prestipino and Raffo (2020) also construct a separate firm-level measure from earnings-call transcripts. The newspaper index motivates time variation in perceived tariff risk, but it does not directly measure the model's conditional tariff variance, $\bar{\sigma}_\tau^2 \exp(2\sigma_{T,t})$, where $\sigma_{T,t}$ is the log volatility state. We therefore do not map the index numerically into that state.

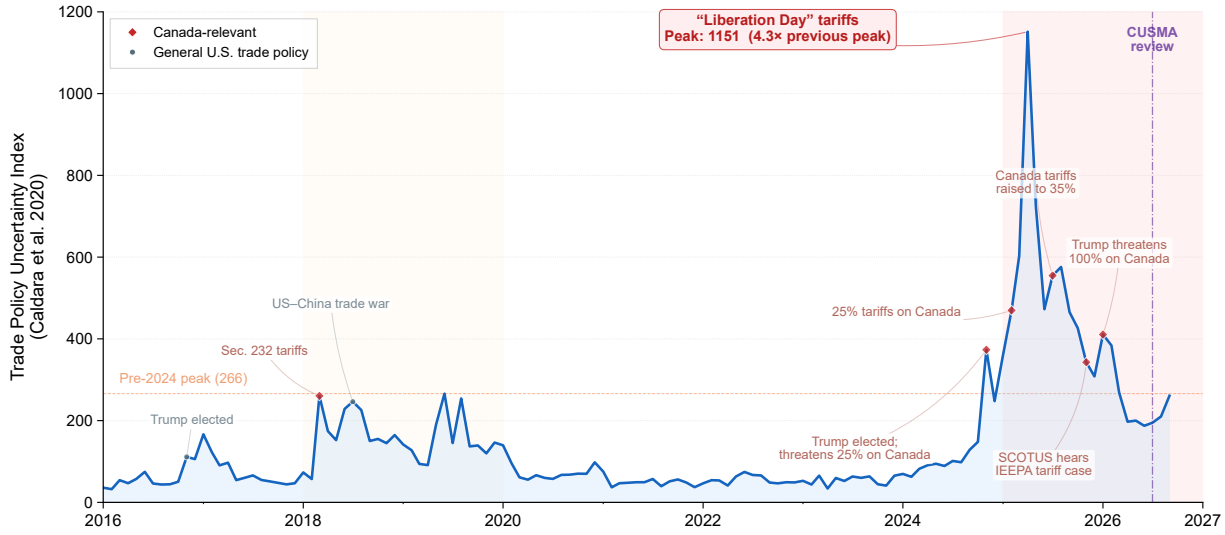


Figure A.1: U.S. Trade Policy Uncertainty Index, 2016–2026

Notes: Monthly newspaper-based Trade Policy Uncertainty Index of Caldara, Iacoviello, Molligo, Prestipino and Raffo (2020); the sample ends in September 2026. Diamond markers denote events with direct implications for Canada, and circles denote broader U.S. trade-policy events. The dashed horizontal line marks the pre-2024 peak (June 2019). The vertical dash-dot line marks the July 1, 2026 USMCA (CUSMA) joint review. Shaded regions indicate the first U.S.–China trade-war period (2018–2019) and the renewed tariff period beginning in 2025.

B Full Nonlinear Equilibrium Conditions (Levels)

This appendix lists the equilibrium conditions in levels implied by the primitives in Section 2 of the main text. Appendix C reduces them to the exact system used in most of the analysis.

Roadmap. Section B.1 states the conditions common to all monetary-policy closures. Section B.2 records the international asset-market block: under complete markets, the asset-pricing relation plus risk sharing; under incomplete markets, the bond-economy Euler system, in which the domestic Euler equation is no longer redundant. Section B.3 reports the baseline calibration. These conditions are the starting point for the exact reduction in Appendix C and the incomplete-markets extension in Appendix H.

B.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 small-open-economy (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 index is $Q_t = P^w_{F,t}/P_{H,t} = \mathcal{E}_t P^*_t/P_{H,t}$.

The private equilibrium conditions common to all monetary-policy closures are:

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

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

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

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

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

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

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

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

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

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

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

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

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

$$m_t^r \equiv \frac{MC_t}{P_{H,t}} = G_t \frac{W_t / P_t}{A_t}, \tag{B.14}$$

$$\bar{\upsilon} \equiv m_{ss}^r = \chi, \qquad mc_t \equiv \frac{m_t^r}{\bar{\upsilon}}, \tag{B.15}$$

$$1 - s^p = \frac{\epsilon - 1}{\epsilon \bar{\upsilon}}, \tag{B.16}$$

$$\kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) + \beta \kappa \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \frac{\lambda_{t+1}}{\lambda_t} \frac{G_t}{G_{t+1}} \right]. \tag{B.17}$$

For the PPI Taylor-rule experiments, the additional closure is

$$\frac{R_t}{\bar{R}} = \Pi_{H,t}^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y}. \tag{B.18}$$

Strict PPI targeting replaces this equation with $\Pi_{H,t} = 1$; under Ramsey commitment, R_t

implements the planner's state-contingent allocation. The constant s^p is the fixed production correction: firm i's effective nominal marginal cost is $(1-s^p)MC_t$. Its sign is not imposed. The Home fiscal authority rebates import-tariff revenue, finances the production correction, and returns the residual lump sum:

$$T_t^{gov} = (T_{F,t}-1)P_{F,t}^w C_{F,t} - s^p MC_t Y_t. \tag{B.19}$$

Neither item uses real resources, and Foreign export-tariff revenue accrues abroad. To derive the pricing equation, firm i faces $Y_t(i) = (P_{H,t}(i)/P_{H,t})^{-\epsilon} Y_t$ and the current nominal profit flow

$$[P_{H,t}(i) - (1-s^p)MC_t]Y_t(i) - P_{H,t}\frac{\kappa}{2}\left(\frac{P_{H,t}(i)}{P_{H,t-1}(i)} - 1\right)^2 Y_t. \tag{B.20}$$

Differentiating with respect to $P_{H,t}(i)$, evaluating at the symmetric allocation and dividing by $Y_t > 0$ gives the current-period derivative $1 - \epsilon + \epsilon(1-s^p)m_t^r - \kappa(\Pi_{H,t}-1)\Pi_{H,t}$. The same price appears in the denominator of next period's adjustment rate, generating the expectation term in (B.17). Finally,

$$1 - \epsilon + \epsilon(1-s^p)m_t^r = (\epsilon - 1)\left(\frac{m_t^r}{\bar{v}} - 1\right) = (\epsilon-1)(mc_t - 1), \tag{B.21}$$

where the first equality uses (B.16). This derives (B.17) from the firm's problem; its stochastic discount factor term λ_{t+1}/λ_t equals $(C_{t+1}/C_t)^{-\sigma}$ by (B.2).

B.2 International asset markets: complete versus incomplete

Under **complete markets**, the Home nominal Euler equation (B.4) and international risk sharing (B.23), with (C_t^*, P_t^*) constant and $R_t^* = \beta^{-1}$, imply the exact nonlinear asset-pricing relation (B.22) with $\psi_t = 0$. For later comparison, we write the relation with an auxiliary multiplicative disturbance e^{ψ_t}:

$$1 = \frac{R_t}{R_t^*} e^{\psi_t} \mathbb{E}_t\left[\frac{\mathcal{E}_t}{\mathcal{E}_{t+1}}\right]. \tag{B.22}$$

At first order, (B.22) implies the familiar log UIP restriction

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

The auxiliary disturbance ψ_t is used only in the two-channel comparison in Appendix E and is zero in every baseline tariff experiment. It is distinct from the endogenous asset-market uncertainty wedge, built from conditional-mean and Jensen terms, that tariff uncertainty

generates in the same relation. Risk sharing is

$$\frac{C_t^{-\sigma}}{(C_t^*)^{-\sigma}} = \vartheta \cdot \frac{P_t}{\mathcal{E}_t P_t^*}. \tag{B.23}$$

Here $\vartheta > 0$ is a constant. Holding (C_t^*, P_t^*) fixed, (B.23) implies the static consumption mapping (15) and yields the two-equation minimal representation in Lemma C.1. Under **incomplete markets**, (B.23) is dropped; consumption is determined by the Euler equation (B.4). We use the bond-economy closure in Appendix H.2: the domestic Euler equation (H.17), the Foreign-bond Euler equation (H.18), and the net-foreign-asset (NFA) law (H.19). 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.

B.3 Baseline calibration

Table B.1 reports the quarterly baseline calibration used in the quantitative illustrations in the main text and this appendix unless stated otherwise.

Table B.1: Baseline calibration

Parameter	Description	Value
β	Discount factor	0.99
σ	Risk aversion	2
φ	Inverse Frisch elasticity	1
χ	Labor-disutility weight	$35/101 \approx 0.3465$
α	Import share	0.30
γ	Trade elasticity	1.5
ϵ	Demand elasticity	10
κ	Rotemberg adjustment 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	$\log 2 \approx 0.693$
ψ_b	Debt-elastic risk-premium slope (incomplete markets)	0.007
ϕ_π	Taylor rule: PPI inflation	1.50
ϕ_y	Taylor rule: output	0.50

Notes: The labor-disutility weight χ makes the deterministic no-tariff allocation satisfy the constrained Home planner's terms-of-trade condition; it also equals raw steady-state PPI marginal cost $\bar{v}$, while the pricing equation uses normalized $mc_t = m_t^r/\bar{v}$. The "tariff uncertainty shock" is a one-standard-deviation innovation to the volatility state in (19) (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).

C Minimal Nonlinear Representation (Complete Markets)

This appendix shows that, under complete markets and producer-currency pricing (PCP), the full nonlinear SOE equilibrium reduces exactly to a 2×2 forward-looking system in producer-price inflation and the terms of trade. All other variables follow from a static mapping.

Roadmap. Section C.1 states the starting point and the three ingredients of the reduction. Section C.2 derives the static block, which expresses (Y_t, mc_t) as functions of $(\Pi_{H,t}, Q_t)$ and the exogenous states. Section C.3 proves Lemma C.1 (main-text Lemma 1), the exact two-equation system, and Section C.4 collects the identities that recover the remaining variables.

Endogenous variables, exogenous states, and shocks. The two forward-looking endogenous variables are $(\Pi_{H,t}, Q_t)$. The exogenous states are the Home import tariff $T_{F,t}$, the Foreign tariff on Home exports T_t^*, and productivity A_t. An auxiliary UIP disturbance ψ_t enters only in the comparison exercise of Appendix E and is zero in the baseline tariff experiments. For the uncertainty analysis, we also work with the log export-tariff state $\tau_t \equiv \log T_t^*$ and its volatility state $\sigma_{T,t}$ as in (18)–(19); the corresponding innovations are $\varepsilon_{\tau,t}$ (tariff), $\varepsilon_{\sigma,t}$ (volatility), $\varepsilon_{A,t}$ (productivity), and, only in the comparison exercise, $\varepsilon_{\psi,t}$ (UIP disturbance).

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

The complete-markets equilibrium is characterized by the common block of conditions in Section B together with risk sharing (B.23). 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 property that the Rotemberg NKPC depends on intertemporal terms only through the stochastic discount factor, which risk sharing ties to the terms of trade.

C.2 Static block

Hold (C_t^*, P_t^*, R_t^*) constant and evaluate at the deterministic steady state $\Pi_H = Q = T_F = T^* = A = 1$, $\psi = 0$. Define the CES price-index term Φ_t and the CPI–PPI wedge $G_t \equiv P_t/P_{H,t}$,

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

This follows by dividing the CPI price index (B.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 (15), consumption satisfies $C_t = (G_t/Q_t)^{-1/\sigma}$ so

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

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

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

obtained from the demand system implied by the CES aggregator (2) and the CPI (B.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 (C.1)–(C.2).

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

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

Goods market clearing (B.13) implies

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

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

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

To see (C.6), start from raw marginal cost (B.14) and use labor supply (B.3) with $L_t = Y_t/A_t$ to obtain $m_t^r = \chi G_t C_t^{\sigma} Y_t^{\varphi} A_t^{-(1+\varphi)}$. Dividing by $\bar{v} = \chi$ gives $mc_t = G_t C_t^{\sigma} Y_t^{\varphi} A_t^{-(1+\varphi)}$. Under risk sharing (15), $C_t^{\sigma} = Q_t/G_t$, implying (C.6).

C.3 Two structural equations

Lemma C.1 (Exact minimal nonlinear system under complete markets and the PPI Taylor rule)**.** *Under complete markets and PCP, impose the PPI Taylor rule* (B.18)*, hold* (C_t^*, P_t^*) *constant, set* $R_t^* = \beta^{-1}$*, and allow the auxiliary comparison disturbance* ψ_t *in* (B.22)*. The equilibrium dynamics are then characterized by the following two equations in* $(\Pi_{H,t}, Q_t)$*:*

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

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

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

Intuition. Complete markets remove an independent intertemporal margin for Home consumption: risk sharing ties the stochastic discount factor to the terms of trade. The only dynamic restrictions are therefore the Rotemberg NKPC and the asset-market block (the Taylor rule combined with the Euler-consistent UIP relation, written in the terms of trade). Output, marginal cost, absorption and exports follow from the static block. This exact reduction underlies the closed-form uncertainty results in the main text.

Proof begins: Lemma C.1.

Derivation of (C.7)**.** Start from the nonlinear NKPC in levels (B.17). 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 (B.17) collapses to Q_t/Q_{t+1}, delivering (C.7).

Derivation of (C.8)**.** Combine the Taylor rule (B.18) with the Euler-consistent asset-pricing relation (B.22) 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 (B.22) 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 (C.8).

End of proof: Lemma C.1.

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

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

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

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

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

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

C.4 Recovering other variables from the minimal system

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

C.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 (B.9) implies

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

Producer-price inflation satisfies

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

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

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

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

C.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 (B.11) implies the identity

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

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

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

C.4.3 Consumption (risk sharing) and demand components

Under complete markets, risk sharing (B.23) pins down consumption as $C_t = (G_t / Q_t)^{-1/\sigma}$. In log deviations,

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

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

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

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

Goods market clearing (B.13) gives $Y_t = (C_{H,t} + X_t)/\mathrm{ADJ}_t = Z_t/\mathrm{ADJ}_t$, with Z_t and ADJ_t defined in (C.9)–(C.10), and PPI-based real marginal cost is

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

as in Remark C.1.

C.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 (C.17), a first-order expansion yields $g_t = \alpha q_t + O(2)$, so (C.18) gives $c_t = \frac{1-\alpha}{\sigma} q_t + O(2)$. Moreover, with θ_q and θ_τ defined in (D.2)–(D.3), the static block implies

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

and marginal cost in logs satisfies

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

Export volume (in log deviations) is $x_t \equiv \log(X_t/\bar{X}) = \gamma(q_t - \tau_t)$ from (C.20). Finally, under the PPI-based Taylor rule (B.18), the policy rate in logs is

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

C.4.5 Decentralized flexible-price output and the output gap

The decentralized flexible-price allocation removes price-adjustment costs and imposes the static markup condition, which in our normalization is $mc_t^f = 1$. This condition pins down the real allocation but not inflation: with flexible prices the price level can adjust between dates. At first order,

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

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

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

and decentralized flexible-price output is $y_t^f = \theta_q q_t^f - \theta_\tau \tau_t$. Under strict PPI targeting, $\Pi_{H,t} = 1$ and the exact Rotemberg NKPC imply $mc_t = 1$, so under complete markets and PCP strict PPI implements the decentralized flexible-price allocation. Its output gap relative

to the constrained-efficient allocation y_t^e is zero only if, in addition, $y_t^f = y_t^e$ in every state, that is, only under the state-by-state efficiency condition (F.15) (Appendices F.2 and F.3).

Part II

Analytics

D Proof of Proposition 1 and Its Corollaries

Roadmap. This appendix proves Proposition 1 by coefficient matching, which delivers the closed-form system (32). Section D.1 expands the static block to second order. Section D.2 solves the first-order (π_t, q_t) system in closed form. Section D.3 derives the second-order curvature in the tariff state along deterministic tariff paths. Section D.4 combines these layers into the third-order uncertainty coefficients (χ_π, χ_q) and the geometrically decaying impulse responses, which completes the proof. Section D.5 then collects the corollaries and special cases, and Section D.6 records the closed form in primitives only.

Proof begins: Proposition 1.

D.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, set the Home import tariff to its baseline value, $T_{F,t} \equiv 1$. The static block in Section C then 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

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

The coefficients in (D.1) depend only on the primitives $(\alpha, \gamma, \sigma, \kappa)$; they are derived below and collected in (D.2)–(D.7).

D.1.1 Static derivatives

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 (C.10). Then

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

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 C.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 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 (D.4)–(D.6). Finally, because ADJ_t enters y_t as $-\log \text{ADJ}_t$ and $\text{ADJ}_t = 1 - \frac{\kappa}{2}(e^{\pi_t} - 1)^2$, one has $y_{\pi\pi} = \kappa$ at the steady state, delivering (D.7).

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

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

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

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

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

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

D.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

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

and recall $\kappa_{mc} = (\epsilon - 1)/\kappa$ from (22). Linearizing (23)–(24) 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

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

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

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 (D.9)–(D.10) imply

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

with $\mathbf{M}(\cdot)$ defined in (30) and right-hand-side loadings

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

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

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

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

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

D.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 state:

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

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

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

and two scalar combinations that recur below,

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

Substitute (D.14) into the static expansion (D.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 exponents of the two nonlinear objects in (23), marginal cost and the discounting term $\Omega_{t+1} \equiv \exp\big((y_{t+1} - y_t) + (q_t - q_{t+1})\big)$, satisfy

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

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 (23) 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 $\psi_t = 0$, (24) 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

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

Whenever $\det \mathbf{M}(\rho_\tau^2) \neq 0$, Cramer's rule gives the curvature coefficients (B_π, B_q) in closed form.

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

D.4.1 Why uncertainty IRFs are third order

Write the export-tariff process under stochastic volatility as

$$\tau_{t+1} = \rho_\tau \tau_t + \varsigma_t \, \varepsilon_{\tau,t+1}, \qquad \varsigma_t \equiv \bar{\sigma}_\tau e^{\sigma_{T,t}}, \qquad \sigma_{T,t} = \rho_\sigma \sigma_{T,t-1} + \eta_\sigma \varepsilon_{\sigma,t}, \tag{D.18}$$

where $\mathrm{Var}(\varepsilon_{\tau,t+1}) = \mathrm{Var}(\varepsilon_{\sigma,t}) = 1$. The volatility experiment in Definition 1 sets $\varepsilon_{\tau,t} = 0$ for all t, so $\tau_t \equiv 0$ on both the baseline and the shocked path; the shock changes only the *conditional distribution* of τ_{t+1}, through ς_t. As emphasized by Benigno, Benigno and Nisticò (2013), because the volatility state is not itself an argument of the equilibrium conditions and only scales primitive innovations, it affects date-t equilibrium only through conditional moments, such as $\mathbb{E}_t[\tau_{t+1}^2] = \varsigma_t^2$, that enter the expectations in the Euler and pricing conditions.

By (D.20) below, the conditional variance deviates from its $\sigma_{T,t} \equiv 0$ value by $2\bar{\sigma}_\tau^2 \sigma_{T,t} + O(4)$. Under standard perturbation order counting, $\bar{\sigma}_\tau$ is first order and $\sigma_{T,t} = O(\eta_\sigma)$, so this deviation is $O(\bar{\sigma}_\tau^2 \eta_\sigma) = O(3)$: the first nonzero uncertainty responses appear at third order (Figure G.1, Section G.1).

Whether uncertainty effects appear at second or third order is not primarily a matter of log versus linear volatility. Under the same scaling, a linear conditional standard deviation $\varsigma_t^{\mathrm{lin}} = \bar{\sigma}_\tau(1 + \sigma_{T,t})$ gives $(\varsigma_t^{\mathrm{lin}})^2 - \bar{\sigma}_\tau^2 = 2\bar{\sigma}_\tau^2 \sigma_{T,t} + O(4)$, again an $O(3)$ variance effect. Whether this effect appears at second or third order depends on how the volatility state and the scale parameter are ordered in the perturbation, and on whether heteroskedasticity is retained. Comparisons with second-order formulations such as Benigno, Benigno and Nisticò (2013) must therefore use the same order assignment; the difference cannot be attributed to a linear rather than log specification of the standard deviation.

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

$$\underbrace{\begin{pmatrix} 1-\beta & -\tilde{\kappa} \\ \phi_\pi - 1 & \phi_y \theta_q \end{pmatrix}}_{\mathbf{M}_{ss}} \begin{pmatrix} \bar{\pi} \\ \bar{q} \end{pmatrix} = \bar{\sigma}_\tau^2 \begin{pmatrix} \beta \left(\frac{1}{2} B_\pi + \frac{3}{2} A_\pi^2 + A_\pi s_1 \right) \\ \frac{1}{2} \left((B_\pi + B_q) - (A_\pi + A_q)^2 \right) \end{pmatrix}. \tag{D.19}$$

D.4.2 From mean shifts to uncertainty coefficients

Equation (D.19) 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:

$$\mathbb{E}_t[\tau_{t+1}^2] = \bar{\sigma}_\tau^2 e^{2\sigma_{T,t}} = \bar{\sigma}_\tau^2\big(1 + 2\sigma_{T,t}\big) + O(4), \tag{D.20}$$

where the last equality uses that $\sigma_{T,t}$ affects equilibrium only through terms scaled by $\bar{\sigma}_\tau^2$, so the omitted terms, starting with $2\bar{\sigma}_\tau^2\sigma_{T,t}^2$, are $O(4)$ and drop out at third order.

Define the uncertainty coefficients as in the main text (Definition 1 and (27)):

$$\pi_t = \bar{\pi} + \chi_\pi\, \sigma_{T,t} + O(4), \qquad q_t = \bar{q} + \chi_q\, \sigma_{T,t} + O(4). \tag{D.21}$$

Because $\sigma_{T,t}$ follows (19), $\mathbb{E}_t[\sigma_{T,t+1}] = \rho_\sigma \sigma_{T,t}$. Collecting the $O(\bar{\sigma}_\tau^2\sigma_{T,t})$ terms in the coefficient-matching steps behind (D.19) therefore yields a linear system for (χ_π, χ_q) with the first-order matrix $\mathbf{M}(\rho_\sigma)$ of (30):

$$\mathbf{M}(\rho_\sigma)\begin{pmatrix}\chi_\pi\\ \chi_q\end{pmatrix} = \bar{\sigma}_\tau^2\begin{pmatrix}\beta\,(B_\pi + 3A_\pi^2 + 2A_\pi s_1)\\ (B_\pi + B_q) - (A_\pi + A_q)^2\end{pmatrix}, \tag{D.22}$$

which is the main-text system (32). Relative to (D.19), 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 (D.20).

Define the wedge loadings, as in the main text,

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

with s_1 from (D.16), and let

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

Solving (D.22) by Cramer's rule gives

$$\chi_\pi = \frac{a_\sigma u^{NK} + \tilde{\kappa} u^{UIP}}{D_\sigma}, \qquad \chi_q = \frac{c_\sigma u^{UIP} - d_\sigma u^{NK}}{D_\sigma}, \tag{D.23}$$

with $D_\sigma = c_\sigma a_\sigma + \tilde{\kappa} d_\sigma$. Under the volatility experiment in Definition 1, $\sigma_{T,t} = \rho_\sigma^t \eta_\sigma$ for $t \geq 0$ (from $\sigma_{T,-1} = 0$), so (D.21) gives the geometric impulse responses (28).

End of proof: Proposition 1.

D.5 Corollaries and special cases

This subsection uses the compact formula (D.23) to derive three special cases: a transitory volatility state, a pure inflation rule, and the single-wedge limits.

Corollary D.1 (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 $t \geq 1$.

Proof begins: Corollary D.1.

Setting $\rho_\sigma = 0$ gives $a_\sigma = 1 + \phi_y\theta_q$, $c_\sigma = 1$, $d_\sigma = \phi_\pi$, and $D_\sigma = (1+\phi_y\theta_q) + \tilde{\kappa}\phi_\pi$; substituting into (D.23) yields the stated coefficients. Because (28) gives $\Delta\pi_t = \chi_\pi\rho_\sigma^t\eta_\sigma$ and $\Delta q_t = \chi_q\rho_\sigma^t\eta_\sigma$, both responses vanish for $t \geq 1$.

End of proof: Corollary D.1.

Under a pure inflation rule ($\phi_y = 0$), $a_\sigma = 1-\rho_\sigma$ and $D_\sigma = (1-\beta\rho_\sigma)(1-\rho_\sigma) + \tilde{\kappa}(\phi_\pi - \rho_\sigma)$; substituting into (D.23) gives the closed-form coefficients behind the $\phi_y = 0$ response in Figure G.2.

Corollary D.2 (Single-channel limits)**.** *Let $(\chi_\pi^{UIP}, \chi_q^{UIP})$ denote the coefficients when only the UIP wedge is active ($u^{NK} = 0$), and $(\chi_\pi^{NK}, \chi_q^{NK})$ those when only the NKPC wedge is active ($u^{UIP} = 0$). With only the UIP wedge,*

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

With only the NKPC wedge,

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

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

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

The low-rigidity variants in Figure G.2 evaluate the NKPC-wedge-only ($u^{UIP} = 0$) and UIP-wedge-only ($u^{NK} = 0$) coefficients at $\kappa = 4$ (baseline $\kappa = 40$), holding all other parameters fixed.

Proof begins: Corollary D.2.

Setting $u^{NK} = 0$ or $u^{UIP} = 0$ in (D.23) gives the single-wedge coefficients; the ratios follow by division.

For the flexible-price limits, write the single-wedge coefficients as

$$\chi_\pi^{NK} = \frac{a_\sigma u^{NK}}{D_\sigma}, \qquad \chi_q^{NK} = -\frac{d_\sigma u^{NK}}{D_\sigma}, \tag{D.24}$$

$$\chi_\pi^{UIP} = \frac{\tilde{\kappa} u^{UIP}}{D_\sigma}, \qquad \chi_q^{UIP} = \frac{c_\sigma u^{UIP}}{D_\sigma}, \qquad D_\sigma = c_\sigma a_\sigma + \tilde{\kappa} d_\sigma, \tag{D.25}$$

where $\tilde{\kappa} = \kappa_{mc}(1 + \varphi\theta_q)$ and $d_\sigma = \phi_\pi - \rho_\sigma > 0$ under the corollary's hypothesis $\phi_\pi > \rho_\sigma$. With $K \equiv (\epsilon - 1)(1 + \varphi\theta_q)$, we have $\tilde{\kappa} = K/\kappa$.

With only the NKPC wedge,

$$\chi_q^{NK} = -\frac{u^{NK}}{\tilde{\kappa} + \frac{c_\sigma a_\sigma}{d_\sigma}} = -\frac{\kappa\, u^{NK}}{K + \kappa\frac{c_\sigma a_\sigma}{d_\sigma}}, \tag{D.26}$$

$$\chi_\pi^{NK} = \frac{a_\sigma}{d_\sigma}\frac{u^{NK}}{\tilde{\kappa} + \frac{c_\sigma a_\sigma}{d_\sigma}} = \frac{a_\sigma}{d_\sigma}\frac{\kappa\, u^{NK}}{K + \kappa\frac{c_\sigma a_\sigma}{d_\sigma}}. \tag{D.27}$$

With only the UIP wedge,

$$\chi_q^{UIP} = \frac{c_\sigma u^{UIP}}{\tilde{\kappa} d_\sigma + c_\sigma a_\sigma} = \frac{c_\sigma \kappa\, u^{UIP}}{K d_\sigma + \kappa c_\sigma a_\sigma}, \tag{D.28}$$

$$\chi_\pi^{UIP} = \frac{\tilde{\kappa} u^{UIP}}{\tilde{\kappa} d_\sigma + c_\sigma a_\sigma} = \frac{K\, u^{UIP}}{K d_\sigma + \kappa c_\sigma a_\sigma}. \tag{D.29}$$

Let $\kappa \to 0^+$. If $u^{NK}(\kappa)$ stays bounded, (D.26)–(D.27) imply $\chi_q^{NK} \to 0$ and $\chi_\pi^{NK} \to 0$. If $u^{UIP}(\kappa)$ stays bounded, (D.28) implies $\chi_q^{UIP} \to 0$; if in addition $u^{UIP}(\kappa) \to u^{UIP,0}$, (D.29) implies $\chi_\pi^{UIP} \to u^{UIP,0}/d_\sigma$.

The corollary assumes that the wedges stay bounded; we now give primitive sufficient conditions. In this paragraph only, $O(\cdot)$ and $\Theta(\cdot)$ refer to $\kappa \to 0^+$, not to the perturbation order. Recall $u^{NK} = \bar{\sigma}_\tau^2\beta(B_\pi + 3A_\pi^2 + 2A_\pi s_1)$ with $s_1 = (\theta_q - 1)A_q - \theta_\tau$, and $u^{UIP} = \bar{\sigma}_\tau^2((B_\pi + B_q) - (A_\pi + A_q)^2)$. The first-order coefficients solve $\mathbf{M}(\rho_\tau)(A_\pi, A_q)' = b_\tau$ with $b_\tau = (-\kappa_{mc}\Gamma_\tau, \phi_y\theta_\tau)'$ and $\kappa_{mc} = (\epsilon - 1)/\kappa$. Suppose $\phi_\pi \notin \{\rho_\tau, \rho_\tau^2\}$. The determinant formula in (D.11) then gives $\det \mathbf{M}(\rho_\tau) = \Theta(1/\kappa)$ and $\det \mathbf{M}(\rho_\tau^2) = \Theta(1/\kappa)$. The Cramer numerators in (D.12)–(D.13) are $O(1/\kappa)$, so $(A_\pi, A_q) = O(1)$. In (D.17), the first component of the right-hand side is $O(1/\kappa)$ and the second is $O(1)$, so the same Cramer formulas at ρ_τ^2 give $(B_\pi, B_q) = O(1)$. Hence u^{NK} and u^{UIP} remain bounded. If $\phi_\pi = \rho_\tau$ or $\phi_\pi = \rho_\tau^2$, the

corresponding determinant loses its $1/\kappa$ leading term, and boundedness must be checked directly.

End of proof: Corollary D.2.

D.6 Fully substituted primitive closed form for (χ_π, χ_q)

This subsection eliminates the matrix shorthand and the coefficients (A_π, A_q, B_π, B_q), expressing (χ_π, χ_q) directly in model primitives. Define

$$\Delta(\rho) \equiv (1-\beta\rho)\big(1-\rho+\phi_y\theta_q\big) + \kappa_{mc}\big(1+\varphi\theta_q\big)(\phi_\pi-\rho), \tag{D.30}$$

where $\kappa_{mc} \equiv (\epsilon-1)/\kappa$, 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. Define

$$N_\pi \equiv \theta_\tau\kappa_{mc}\left[-\varphi(1-\rho_\tau+\phi_y\theta_q)+(1+\varphi\theta_q)\phi_y\right], \tag{D.31}$$
$$N_q \equiv \theta_\tau\left[(1-\beta\rho_\tau)\phi_y+\kappa_{mc}\varphi(\phi_\pi-\rho_\tau)\right]. \tag{D.32}$$

Define

$$N_C \equiv y_{qq}N_q^2 + 2y_{q\tau}N_q\Delta_\tau + y_{\tau\tau}\Delta_\tau^2 + y_{\pi\pi}N_\pi^2, \tag{D.33}$$
$$N_m \equiv (1+\varphi\theta_q)N_q - \varphi\theta_\tau\Delta_\tau, \qquad N_s \equiv (\theta_q-1)N_q - \theta_\tau\Delta_\tau, \tag{D.34}$$

and

$$N_{D_1} \equiv \kappa_{mc}\big(\varphi N_C + N_m^2\big) + 2\beta\rho_\tau(\rho_\tau-1)N_\pi N_s - 3(1-\beta\rho_\tau^2)N_\pi^2, \tag{D.35}$$
$$N_{D_2} \equiv -\phi_y N_C. \tag{D.36}$$

Then

$$N_{B_\pi} \equiv N_{D_1}(1-\rho_\tau^2+\phi_y\theta_q) + \kappa_{mc}(1+\varphi\theta_q)N_{D_2}, \tag{D.37}$$
$$N_{B_q} \equiv (1-\beta\rho_\tau^2)N_{D_2} - (\phi_\pi-\rho_\tau^2)N_{D_1}. \tag{D.38}$$

Finally, define the composite blocks

$$N_1 \equiv N_{B_\pi} + \Delta_{\tau^2}\big(3N_\pi^2 + 2N_\pi N_s\big), \tag{D.39}$$
$$N_2 \equiv N_{B_\pi} + N_{B_q} - \Delta_{\tau^2}(N_\pi+N_q)^2. \tag{D.40}$$

Then the uncertainty coefficients from (D.22) are

$$\chi_\pi = \frac{\bar{\sigma}_\tau^2}{\Delta_\sigma \, \Delta_\tau^2 \, \Delta_{\tau^2}} \left[\beta(1 - \rho_\sigma + \phi_y \theta_q) N_1 + \kappa_{mc}(1 + \varphi \theta_q) N_2\right], \tag{D.41}$$

$$\chi_q = \frac{\bar{\sigma}_\tau^2}{\Delta_\sigma \, \Delta_\tau^2 \, \Delta_{\tau^2}} \left[(1 - \beta \rho_\sigma) N_2 - \beta(\phi_\pi - \rho_\sigma) N_1\right]. \tag{D.42}$$

Equations (D.2)–(D.42) therefore give (χ_π, χ_q) as an explicit, recursively defined function of the primitives

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

E A Two-Channel First-Order Representation

This section shows that, under complete markets, the Taylor rule and a common-persistence restriction, the third-order uncertainty responses of PPI inflation, the terms of trade and the variables recovered from them without a direct productivity term are reproduced, at every horizon and up to $O(4)$, by a constant-weight combination of two first-order disturbances. The representation holds equation by equation: a productivity disturbance reproduces the NKPC wedge and an additive UIP disturbance ψ_t reproduces the UIP wedge.

Assumption E.1 (Common persistence)**.** The tariff level, the volatility state and the two comparison disturbances are AR(1) with a common persistence $\rho \in (0, 1)$: $\rho_\tau = \rho_\sigma = \rho_A = \rho_\psi = \rho$. If persistence differs, a weighted combination of the two comparison disturbances can still match the impact uncertainty responses of (π_H, q), but constant weights generally do not reproduce every horizon.

Proposition E.1 (Two-channel representation)**.** *Consider the volatility experiment in Definition 1 under complete markets and the Taylor rule* (16)*. Impose Assumption E.1, suppose* $\det \mathbf{M}(\rho) \neq 0$*, and suppose the two first-order comparison impact vectors are linearly independent. Define the impact vectors of* (π_H, q) *by*

$$v^{\text{unc}} \equiv \begin{pmatrix} \Delta\pi_0 \\ \Delta q_0 \end{pmatrix}, \qquad v^{\text{Prod}} \equiv \begin{pmatrix} \text{IRF}_0^{\text{Prod}}(\pi) \\ \text{IRF}_0^{\text{Prod}}(q) \end{pmatrix}, \qquad v^{\psi} \equiv \begin{pmatrix} \text{IRF}_0^{\psi}(\pi) \\ \text{IRF}_0^{\psi}(q) \end{pmatrix}.$$

The unique weights $(\lambda_A, \lambda_\psi)$ *are defined by*

$$\begin{bmatrix} v^{\text{Prod}} & v^{\psi} \end{bmatrix} \begin{pmatrix} \lambda_A \\ \lambda_\psi \end{pmatrix} = v^{\text{unc}}.$$

Let x be PPI inflation, the terms of trade, or any variable whose first-order recovery equation has no direct productivity term and depends on the endogenous response only through (π_H, q). If $\mathrm{IRF}_t^{\mathrm{unc}}(x)$ denotes its third-order uncertainty response, then the same weights satisfy

$$\mathrm{IRF}_t^{\mathrm{unc}}(x) = \lambda_A \mathrm{IRF}_t^{\mathrm{Prod}}(x) + \lambda_\psi \mathrm{IRF}_t^{\psi}(x) + O(4) \quad \textit{for every } t \geq 0. \tag{E.1}$$

Roadmap. Steps 1–2 show that the volatility innovation shifts the linearized UIP condition and NKPC by additive wedges Δ_t^{UIP} and Δ_t^{NK}. Step 3 reproduces these shifts in the first-order model with a productivity disturbance and an additive UIP disturbance. Step 4 uses Assumption E.1, under which both comparison disturbances share the time profile of the volatility state, to show that one pair of weights applies at every horizon, and derives the weights in closed form.

Proof of Proposition E.1. Work under Assumption E.1 and consider the third-order perturbation around the deterministic steady state.

Step 1 (UIP wedge). Let

$$\mathcal{D}_{t+1} \equiv (q_{t+1} - q_t) + \pi_{t+1}$$

denote the change in the log terms of trade plus PPI inflation, which equals nominal depreciation when Foreign prices are constant. Taking logs of (24) yields

$$\phi_\pi \pi_t + \phi_y y_t + \psi_t = -\log \mathbb{E}_t\left[e^{-\mathcal{D}_{t+1}}\right]. \tag{E.2}$$

By a cumulant expansion around the deterministic steady state,

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

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, $\kappa_{3,t}(\mathcal{D}_{t+1}) = O(4)$, so the third-order corrections in (E.3) come from conditional second moments.

Under the volatility experiment (Definition 1), the time-t conditional dispersion of $\mathcal{D}_{t+1}$ comes from the next tariff innovation $\varepsilon_{\tau,t+1}$, scaled by $\bar{\sigma}_\tau e^{\sigma_{T,t}}$. Let $\xi_{\mathcal{D}}$ and $\zeta_{\mathcal{D}}$ denote the steady-state loadings of $\mathcal{D}_{t+1}$ on $\varepsilon_{\tau,t+1}$ and $\varepsilon^2_{\tau,t+1}$:

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

Under the innovation symmetry in Assumption 1, the order-three covariance between the

linear innovation term and the centered quadratic term is proportional to $\mathbb{E}[\varepsilon^3_{\tau,t+1}]$ and therefore vanishes, and the variance of the quadratic term is order four. Hence $\mathrm{Var}_t(\mathcal{D}_{t+1}) = \xi^2_{\mathcal{D}}\bar{\sigma}^2_\tau e^{2\sigma_{T,t}} + O(4)$, and the baseline-subtracted variance (Jensen) component is

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

The conditional mean $\mathbb{E}_t[\mathcal{D}_{t+1}]$ in (E.3) contains a curvature component from the quadratic term in this expansion. Since $\mathbb{E}[\varepsilon^2_{\tau,t+1}] = 1$, its baseline-subtracted value is

$$\Delta^{UIP,curv}_t \equiv \mathbb{E}_t[\mathcal{D}_{t+1}] - \mathbb{E}_t[\mathcal{D}_{t+1}]\Big|_{\sigma_{T,t}=0} = \frac{1}{2}\,\zeta_{\mathcal{D}}\,\bar{\sigma}^2_\tau\big(e^{2\sigma_{T,t}} - 1\big) + O(4). \tag{E.5}$$

Combining (E.4)–(E.5), the UIP wedge, which shifts the *linearized* log UIP restriction, is

$$\Delta^{UIP}_t \equiv \Delta^{UIP,curv}_t + \widetilde{\Delta}^{UIP}_t = \frac{1}{2}\,(\zeta_{\mathcal{D}} - \xi^2_{\mathcal{D}})\,\bar{\sigma}^2_\tau\big(e^{2\sigma_{T,t}} - 1\big) + O(4). \tag{E.6}$$

In the notation of (29), $\xi_{\mathcal{D}} = A_\pi + A_q$ and $\zeta_{\mathcal{D}} = B_\pi + B_q$. Since $\bar{\sigma}^2_\tau(e^{2\sigma_{T,t}} - 1) = 2\bar{\sigma}^2_\tau\sigma_{T,t} + O(4)$, (E.6) gives $\Delta^{UIP}_t = u^{UIP}\,\sigma_{T,t} + O(4)$, with u^{UIP} the UIP-wedge loading in (32). Because $\sigma_{T,t}$ follows (19), Δ^{UIP}_t is AR(1) with persistence ρ up to $O(4)$ under Assumption E.1.

Step 2 (NKPC wedge). Rewrite the nonlinear NKPC (23) in the form

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

where $\Omega_{t+1} \equiv \exp\big((y_{t+1} - y_t) + (q_t - q_{t+1})\big)$ coincides with the intertemporal term in (23). Expanding around the deterministic steady state and collecting third-order terms yields the linear NKPC augmented by an additive wedge:

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + \tilde{\kappa}q_t - c_\tau\tau_t - c_a a_t + \Delta^{NK}_t + O(4), \tag{E.7}$$

where $(\tilde{\kappa}, c_\tau, c_a)$ are defined in (D.9) and Δ^{NK}_t collects the baseline-subtracted conditional second-moment terms generated by the expectation in the Rotemberg NKPC. Under the volatility experiment, the only time-varying component of these moments again comes from $\varepsilon_{\tau,t+1}$, scaled by $\bar{\sigma}_\tau e^{\sigma_{T,t}}$, so Δ^{NK}_t is proportional to $\bar{\sigma}^2_\tau(e^{2\sigma_{T,t}} - 1)$ up to $O(4)$. Expanding the forward-looking Rotemberg term as in Section 3.2 gives $\Delta^{NK}_t = u^{NK}\,\sigma_{T,t} + O(4)$, with u^{NK} the NKPC-wedge loading in (32). Hence Δ^{NK}_t is also AR(1) with persistence ρ up to $O(4)$ under Assumption E.1.

Step 3 (Representation by two first-order disturbances). Consider the first-order linearized system for (π_t, q_t), given by (D.9)–(D.10). In that system, productivity a_t enters

the NKPC through $-c_a a_t$, while the UIP disturbance ψ_t enters the UIP condition additively. We can therefore reproduce the two additive terms $(\Delta_t^{NK}, \Delta_t^{UIP})$ constructed in Steps 1–2 with the comparison disturbances

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

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, up to $O(4)$ terms, for variables whose recovery equations satisfy the restriction in the proposition. It does not generally reproduce labor or marginal cost: in the comparison economy, those variables load directly on productivity even after (π_t, q_t) have been matched.

Step 4 (Time-invariant weights). Linearity of the first-order equilibrium mapping implies that the response to the combined disturbance (a_t^{eq}, ψ_t^{eq}) is the sum of the responses to its two components. Under Assumption E.1, the equivalent processes (a_t^{eq}, ψ_t^{eq}) and the comparison processes (a_t, ψ_t) share the same persistence ρ, so the responses of (π_t, q_t) and every recovery equation covered by the proposition have a common geometric time profile. It follows that there exist constants $(\lambda_A, \lambda_\psi)$ such that (E.1) holds for all horizons.

For closed-form weights, solve the linearized 2×2 system for one-standard-deviation comparison innovations:

$$v^{\text{Prod}} = \frac{\sigma_A c_a}{D_\rho} \begin{pmatrix} -(\phi_y \theta_q + 1 - \rho) \\ \phi_\pi - \rho \end{pmatrix}, \qquad v^{\psi} = -\frac{\sigma_\psi}{D_\rho} \begin{pmatrix} \tilde{\kappa} \\ 1 - \beta\rho \end{pmatrix},$$

where $D_\rho \equiv (1 - \beta\rho)(\phi_y \theta_q + 1 - \rho) + \tilde{\kappa}(\phi_\pi - \rho)$, $\tilde{\kappa} \equiv \kappa_{mc}(1 + \varphi\theta_q)$ 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 of the two comparison disturbances in (18)–(20). Matching impact responses delivers the linear system

$$\begin{pmatrix} v_\pi^{\text{Prod}} & v_\pi^{\psi} \\ v_q^{\text{Prod}} & v_q^{\psi} \end{pmatrix} \begin{pmatrix} \lambda_A \\ \lambda_\psi \end{pmatrix} = \begin{pmatrix} \eta_\sigma \chi_\pi \\ \eta_\sigma \chi_q \end{pmatrix}. \tag{E.8}$$

Since $\det[v^{\text{Prod}}, v^{\psi}] = \sigma_A c_a \sigma_\psi / D_\rho$, with $c_a > 0$ and $D_\rho = \det \mathbf{M}(\rho) \neq 0$ by hypothesis, the system has a unique solution, given by the closed-form weights

$$\begin{aligned} \lambda_A &= \eta_\sigma \, \frac{\tilde{\kappa}\chi_q - (1 - \beta\rho)\chi_\pi}{\sigma_A c_a}, \\ \lambda_\psi &= \eta_\sigma \, \frac{-\big(\phi_y \theta_q + 1 - \rho\big)\chi_q - (\phi_\pi - \rho)\chi_\pi}{\sigma_\psi}. \end{aligned} \tag{E.9}$$

In the quantitative illustrations we normalize $\sigma_A = \sigma_\psi = 0.01$ (1% innovations).

End of proof: Proposition E.1. □

Figure E.1 illustrates why both channels are needed. We choose the one-disturbance comparison by least squares across the displayed variables, after normalizing each response by the root-mean-square magnitude of its uncertainty response so that every variable receives equal weight. A productivity disturbance is the best one-disturbance comparison: it matches the real responses and the terms of trade almost exactly, but it essentially misses the PPI deflation and the policy-rate cut. Adding the UIP disturbance reproduces all displayed responses up to the $O(4)$ remainder established in Proposition E.1.

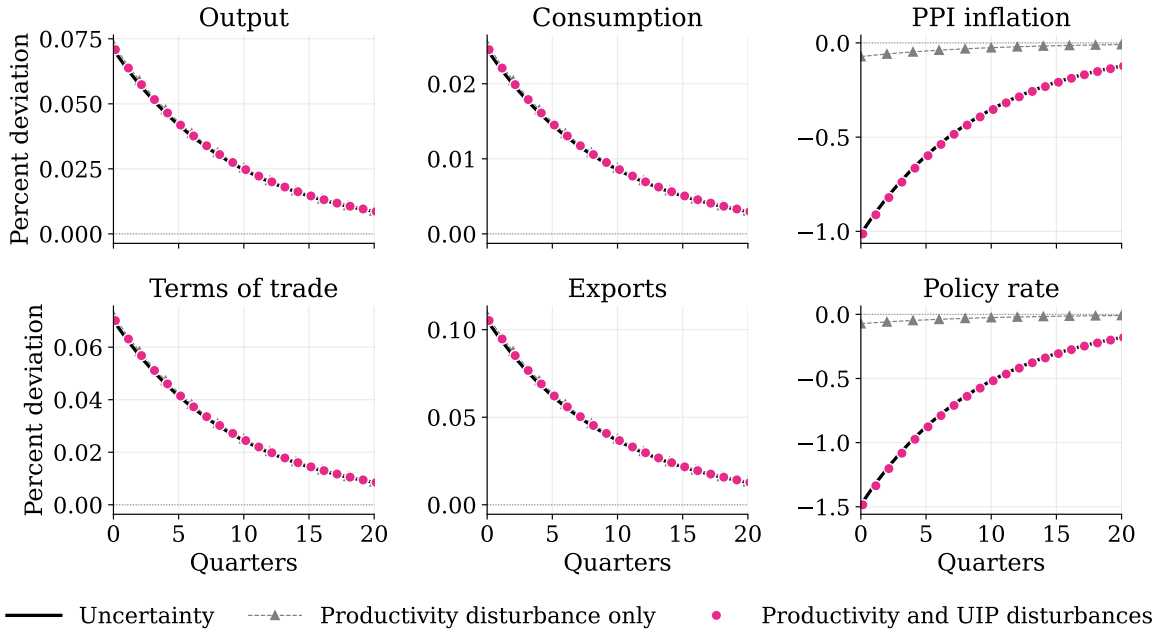


Figure E.1: A two-channel representation of the complete-markets uncertainty response.

Notes: The experiment is a one-standard-deviation volatility innovation at $t = 0$ under the baseline calibration in Table B.1. The solid line is the third-order uncertainty response. The gray dashed line with triangle markers is the best one-disturbance comparison when each displayed variable receives equal weight. The pink circles combine the first-order responses to productivity and UIP disturbances using the constant weights in Proposition E.1. Real variables and the terms of trade are percent deviations from the deterministic steady state; PPI inflation and the policy rate are percentage-point deviations.

The weights $(\lambda_A, \lambda_\psi)$ are not structural parameters. Because they solve the impact-matching system (E.8), they depend both on the uncertainty coefficients (χ_π, χ_q), which encode the wedges characterized in Section 3, and on the linearized responses to the two comparison disturbances. Figure E.2 shows that, in the plotted calibration, trade openness tends to scale up both weights, while nominal rigidity and policy feedback reshape their

relative importance. These patterns parallel the wedge comparative statics in main-text Figure 3, but they are not simple rescalings of them, because the weights also reflect the equilibrium transmission of each disturbance.

In the openness and trade-elasticity panels of the plotted calibration, higher α and higher γ amplify the uncertainty wedges (Section 3.2) and 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 disturbance carries most of the adjustment, while at higher openness the productivity disturbance 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 shifts the weighted combination toward the productivity disturbance: λ_A rises sharply before peaking at high rigidity. 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 panel, raising ϕ_π changes both the target vector v^{unc} and the two comparison vectors in (E.8). In the plotted calibration the target contracts relative to the comparison responses, so both weights fall. This is not a general monotonicity result: stronger feedback lowers the weights only when the target response contracts more than the unit-shock comparison responses.

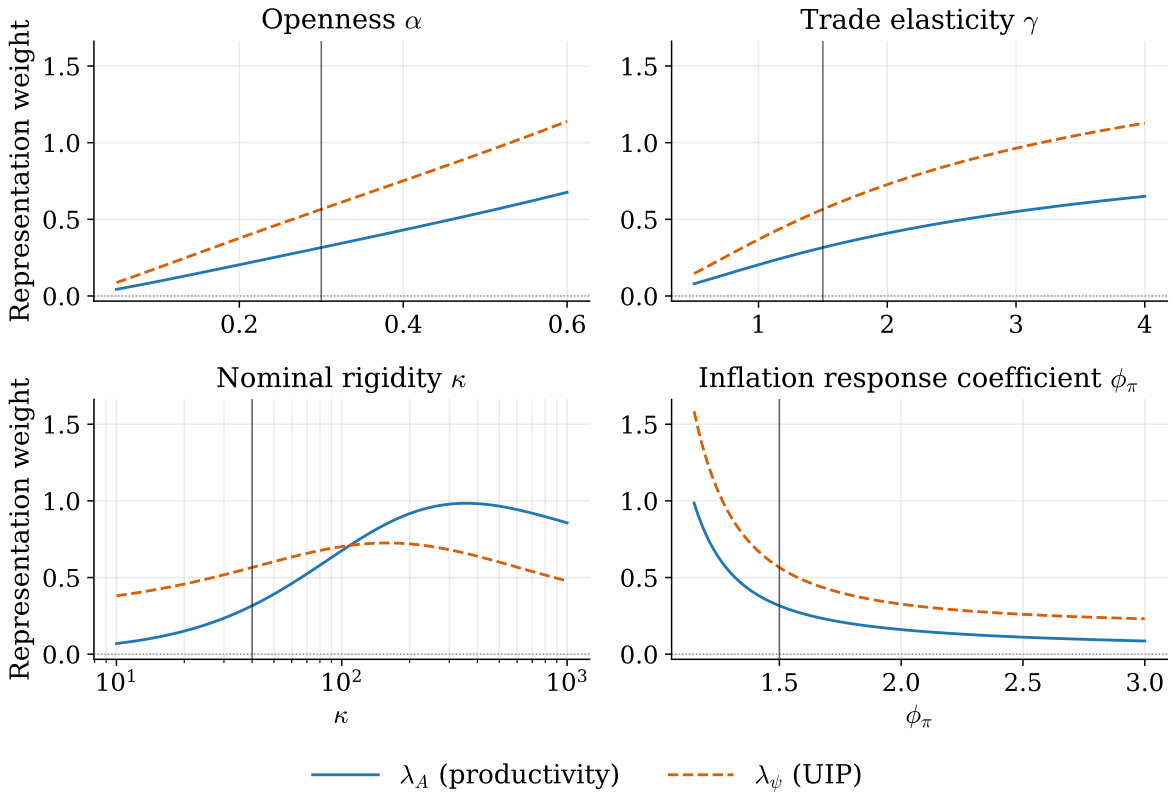


Figure E.2: Sensitivity of the two representation weights

Notes: The figure reports the constant weights $(\lambda_A, \lambda_\psi)$ in Proposition E.1 under the volatility experiment in Definition 1, varying one parameter at a time around the baseline calibration in Table B.1 (vertical lines). The productivity and UIP comparison shocks are normalized to 1% innovations, so the weights have the same percent scale (e.g. $\lambda_\psi = 0.30$ means a 0.30% UIP-disturbance innovation). Persistence is set to $\rho = \rho_\sigma$ to satisfy Assumption E.1.

F Complete-Markets Optimal Policy: Exact Conditions and Closed-Form Tariff-Uncertainty Responses

This appendix derives the optimal-policy results in Section 4 from preferences, market clearing, and firms' exact price-setting equation. It answers four questions. When does strict PPI targeting implement the efficient allocation after every tariff history (Appendices F.2–F.3)? What is the full Ramsey plan under the paper's fixed production correction (Appendices F.4–F.6)? How does that plan change if the central bank must follow the LQ Level-Ramsey targeting rule after date 0 (Appendix F.8)? And how costly is it to set to zero only the incremental date-0 PPI response to the uncertainty innovation while imposing the Ramsey first-order conditions thereafter (Appendix F.9)? Appendix F.10 reports welfare and approximation

orders, and Appendix F.11 the numerical verification. Appendix F.7 proves Proposition 4, and Appendix F.12 explains why the baseline Taylor and Ramsey real paths nearly overlap.

F.1 Policy experiment and contingent histories

Let h^t collect all innovations through date t. The export tariff and its log-volatility state s_t (denoted $\sigma_{T,t}$ in the main text) obey

$$\tau_t = \rho_\tau \tau_{t-1} + \bar{\sigma}_\tau e^{s_{t-1}} \varepsilon_{\tau,t}, \tag{F.1}$$

$$s_t = \rho_\sigma s_{t-1} + \eta_\sigma \varepsilon_{\sigma,t}. \tag{F.2}$$

The date-0 experiment raises tariff uncertainty without changing the tariff itself: $\varepsilon_{\sigma,0} = 1$ and $\varepsilon_{\tau,0} = 0$. The plotted impulse response sets all later innovations to zero, so the tariff stays at its baseline value along the displayed path even though the date-0 innovation widens the distribution of future tariffs:

$$\mathrm{Var}_t(\tau_{t+1}) = \bar{\sigma}_\tau^2 e^{2s_t}. \tag{F.3}$$

The displayed path is therefore one history among many. The Ramsey plan must also specify policy on histories in which the tariff does change, and expected welfare averages over all of them, so a zero tariff realization on the displayed path does not remove tariff risk from the policy problem. A response averaged over later realized innovations coincides with the displayed response at date 0 under this timing convention but not, in general, at later dates.

F.2 Primitive static maps and the constrained-efficient target

Set the Home import tariff to one and write $Q = e^q$ and $T^* = e^\tau$. The CPI–PPI ratio and complete-markets consumption allocation are

$$G(Q) = \left[(1-\alpha) + \alpha Q^{1-\gamma}\right]^{1/(1-\gamma)}, \tag{F.4}$$

$$C(Q) = \left(\frac{G(Q)}{Q}\right)^{-1/\sigma}. \tag{F.5}$$

Home absorption and exports are

$$C_H(Q) = (1-\alpha) G(Q)^\gamma C(Q), \tag{F.6}$$

$$X(Q,\tau) = \alpha \exp\{\gamma(q-\tau)\}. \tag{F.7}$$

Before price-adjustment resources, output is $Y^a(Q,\tau) = C_H(Q) + X(Q,\tau)$. Rotemberg adjustment costs use the fraction $\kappa(\Pi_H - 1)^2/2$ of output, so exact goods clearing gives

$$Y(q,\pi,\tau) = \frac{Y^a(e^q,\tau)}{1 - \frac{\kappa}{2}(e^\pi - 1)^2}. \tag{F.8}$$

To derive the real target, temporarily set $\pi = 0$. With productivity normalized to one, consider a static relaxation of the full Ramsey problem. It retains the complete-markets consumption map and goods clearing, $Y = Y^a$, but drops firms' Rotemberg price-setting equation; when PPI inflation is zero at every history, that equation reduces to the decentralized flexible-price condition introduced below. The nominal Euler equation only recovers the unconstrained rate that implements the chosen allocation, so it adds no static restriction. The relaxed planner can therefore choose Q directly to maximize

$$\mathcal{W}(Q;\tau) = \frac{C(Q)^{1-\sigma}}{1-\sigma} - \chi\frac{Y^a(Q,\tau)^{1+\varphi}}{1+\varphi}. \tag{F.9}$$

The exact first-order condition is

$$C^{-\sigma}C_Q = \chi(Y^a)^\varphi Y^a_Q. \tag{F.10}$$

At $\pi = 0$, household labor supply and $L = Y^a$ imply that raw marginal cost in Home-PPI units is

$$m^r(Q,\tau) \equiv \chi G(Q)C(Q)^\sigma \left[Y^a(Q,\tau)\right]^\varphi. \tag{F.11}$$

Multiplying both sides by GC^σ/Y^a_Q gives

$$\chi G C^\sigma (Y^a)^\varphi = G\frac{C_Q}{Y^a_Q}. \tag{F.12}$$

By (F.11), the left side is $m^r(Q,\tau)$. Hence the planner's exact marginal-cost target is

$$\mathcal{M}^e(Q;\tau) \equiv G(Q)\frac{C_Q(Q)}{Y^a_Q(Q,\tau)}. \tag{F.13}$$

This target follows from preferences, risk sharing, and resource feasibility. The planner chooses the relative price Q so that the consumption gain from changing Q equals the additional labor cost of producing the associated Home absorption and exports. The tariff enters the target through export demand, and hence through $Y^a_Q(Q,\tau)$; this is the dependence that a fixed production correction generally cannot offset after every tariff realization.

Let $s^p(\tau)$ be the production correction. The firm's symmetric flexible-price condition

$(1 - s^p)m^r = (\epsilon - 1)/\epsilon$ defines the raw marginal-cost target

$$\upsilon(\tau) \equiv \frac{\epsilon - 1}{\epsilon[1 - s^p(\tau)]}. \tag{F.14}$$

At flexible prices, firms require $m^r(Q, \tau) = \upsilon(\tau)$, whereas the relaxed planner's first-order condition requires $m^r(Q, \tau) = \mathcal{M}^e(Q; \tau)$. Suppose that, for every reachable τ, $\mathcal{W}(\cdot; \tau)$ has a unique global maximizer $Q^e(\tau)$ and the decentralized flexible-price equilibrium condition has a unique interior solution $Q^f(\tau)$. Then

$$Q^f(\tau) = Q^e(\tau) \quad \Longleftrightarrow \quad \upsilon(\tau) = \mathcal{M}^e(Q^e(\tau); \tau). \tag{F.15}$$

At $Q = 1$ and $\tau = 0$, the planner's target is

$$\mathcal{M}^e(1; 0) = \frac{C_Q(1)}{Y_Q^a(1, 0)} = \frac{(1 - \alpha)/\sigma}{\theta_q}. \tag{F.16}$$

The baseline calibration sets $\chi = \mathcal{M}^e(1; 0) = 35/101$. Together with $1 - s^p = (\epsilon - 1)/(\epsilon\bar{\upsilon})$ and $\bar{\upsilon} = m^r_{ss} = \chi$, this makes the private and planner conditions coincide at the deterministic steady state. The production correction alone does not deliver this equality, which also requires the stated calibration of χ, and it does not in general extend the equality to other tariff states.

F.2.1 When parameters make the target constant

The two tariff experiments give one case in which a constant production correction suffices and one in which it does not. For a Home import tariff T_F with the export tariff fixed at one, suppose $\gamma\sigma = 1$. Then $1/\sigma = \gamma$ and the static maps reduce to

$$C_H = (1 - \alpha)Q^\gamma, \qquad X = \alpha Q^\gamma, \qquad Y^a = Q^\gamma. \tag{F.17}$$

Moreover,

$$GC_Q = \gamma(1 - \alpha)Q^{\gamma-1}, \qquad Y_Q^a = \gamma Q^{\gamma-1}, \tag{F.18}$$

so

$$\mathcal{M}^e_F(Q, T_F) = 1 - \alpha. \tag{F.19}$$

A constant production correction can therefore align all import-tariff states.

For a Foreign tariff on Home exports, the same restriction gives instead

$$\mathcal{M}_X^e(T^*) = \frac{1-\alpha}{(1-\alpha)+\alpha/(T^*)^\gamma}. \tag{F.20}$$

It varies with T^* whenever $0 < \alpha < 1$ and $\gamma > 0$. Thus $\gamma\sigma = 1$ is an exact strict-PPI case for the import-tariff experiment, not for the paper's export-tariff experiment. For export tariffs, exact strict PPI optimality requires a state-contingent correction or another restriction that satisfies (F.15) on the entire support.

F.3 Strict PPI: implementation versus optimality

Proof of Proposition 2. Under strict PPI, $\Pi_{H,t} = 1$ at every history. Rotemberg costs vanish and the exact price-setting equation reduces to the decentralized flexible-price condition. Consequently strict PPI implements Q_t^f. It implements Q_t^e only if (F.15) holds.

When matching holds at every history, flexible prices deliver the best feasible static allocation in every state and strict PPI eliminates all price-adjustment costs. Because every history then attains its static optimum with zero adjustment cost, no feasible policy does better. The nominal rate that implements this allocation follows from the exact Euler equation:

$$1 = \beta R_t \mathbb{E}_t \left[\frac{Q_t}{Q_{t+1}\Pi_{H,t+1}} \right], \tag{F.21}$$

so

$$R_t^{PPI} = \left\{ \beta \mathbb{E}_t \left[\frac{Q_t^e}{Q_{t+1}^e} \right] \right\}^{-1}. \tag{F.22}$$

Thus strict PPI is globally optimal when flexible prices are efficient in every state. The argument does not depend on which tariff varies: if the Home import tariff T_F is the tariff state instead, the same steps go through with the static maps above evaluated at the current T_F.

End of proof: Proposition 2. □

This conclusion fails when the production correction is fixed at its no-tariff value. A log-volatility innovation does not change today's static target: current utility and resources in (F.9) depend on today's tariff level, which has not moved. But firms set prices today partly in anticipation of tomorrow. If the fixed production correction fails to make flexible prices efficient when a future tariff is realized, the planner can improve expected welfare by allowing a small price adjustment today and promising appropriate state-contingent policy later. That intertemporal consideration produces the nonzero third-order Ramsey response even on the plotted history where the tariff never changes.

F.4 Exact nonlinear commitment problem

This subsection states the exact Ramsey commitment problem under the production correction fixed at its no-tariff value, which, together with the calibration of χ, makes the no-tariff steady state constrained efficient, and the conditions under which its first-order system identifies a local optimum. Define normalized marginal cost and the current-date functions

$$mc(q,\pi,\tau) = e^{q}Y(q,\pi,\tau)^{\varphi}, \tag{F.23}$$

$$f(\pi) = (e^{\pi}-1)e^{\pi}, \qquad \mathcal{D}(q,\pi,\tau) = \frac{Y(q,\pi,\tau)}{e^{q}}, \tag{F.24}$$

$$\mathcal{A}(q,\pi,\tau) = \kappa\mathcal{D}(q,\pi,\tau)f(\pi), \tag{F.25}$$

$$\mathcal{B}(q,\pi,\tau) = (\epsilon-1)\mathcal{D}(q,\pi,\tau)\big(mc(q,\pi,\tau)-1\big). \tag{F.26}$$

Complete-market risk sharing implies $C_t^{-\sigma} = G_t/Q_t$. Multiplying the exact Rotemberg NKPC by $\mathcal{D}_t$ and using this identity removes the consumption-based factor from the expected future pricing term, leaving

$$\mathcal{F}_t \equiv \mathcal{A}_t - \mathcal{B}_t - \beta\mathbb{E}_t\mathcal{A}_{t+1} = 0. \tag{F.27}$$

Let

$$\mathcal{V}(q,\pi;\tau) = \frac{C(e^{q})^{1-\sigma}}{1-\sigma} - \chi\frac{Y(q,\pi,\tau)^{1+\varphi}}{1+\varphi}. \tag{F.28}$$

The central bank chooses the complete contingent plan $\mathcal{P}_0 = \{q_t(h^t), \pi_t(h^t)\}_{t,h^t}$:

$$\max_{\mathcal{P}_0} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \mathcal{V}_t \quad \text{subject to} \quad \mathcal{F}_t(h^t) = 0 \quad \forall (t,h^t). \tag{F.29}$$

Let ν_t multiply the date-t constraint (F.27). At date t the multiplier ν_{t-1} is predetermined: it was chosen one period earlier, enters the current history through the forward pricing term, and records the price-setting promise inherited at date t. The current multiplier ν_t becomes inherited only at date $t+1$. Commitment from scratch sets $\nu_{-1} = 0$. The Lagrangian is

$$\mathcal{J}_0(\nu_{-1}) = -\nu_{-1}\mathcal{A}_0 + \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\mathcal{V}_t + \nu_t(\mathcal{A}_t - \mathcal{B}_t - \beta\mathbb{E}_t\mathcal{A}_{t+1})\right]. \tag{F.30}$$

Reindexing the forward term requires keeping track of the upper endpoint. Truncating the

sum in (F.30) at T gives

$$\mathcal{J}_{0,T}(\nu_{-1}) = \mathbb{E}_0 \sum_{t=0}^{T} \beta^t \left[\mathcal{V}_t + (\nu_t - \nu_{t-1})\mathcal{A}_t - \nu_t \mathcal{B}_t\right] - \beta^{T+1}\mathbb{E}_0[\nu_T \mathcal{A}_{T+1}]. \tag{F.31}$$

We impose the discounted terminal-pairing condition

$$\lim_{T\to\infty} \beta^{T+1}\mathbb{E}_0[\nu_T \mathcal{A}_{T+1}] = 0. \tag{F.32}$$

For differentiation, its first variation must also vanish for every admissible local perturbation:

$$\lim_{T\to\infty} \beta^{T+1}\mathbb{E}_0[\delta\nu_T \mathcal{A}_{T+1} + \nu_T D\mathcal{A}_{T+1}\delta(q_{T+1}, \pi_{T+1})] = 0. \tag{F.33}$$

Equations (F.32)–(F.33) are the transversality restrictions associated with the forward pricing constraint; bounded controls alone do not imply them. Under (F.32), taking $T \to \infty$ gives

$$\mathcal{J}_0(\nu_{-1}) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\mathcal{V}_t + (\nu_t - \nu_{t-1})\mathcal{A}_t - \nu_t \mathcal{B}_t\right]. \tag{F.34}$$

The discount factor in the NKPC cancels the one-period discounting difference; there is no additional β multiplying ν_{t-1}. Differentiation gives the exact stationarity equations

$$\mathcal{S}_t^q \equiv \mathcal{V}_{q,t} + (\nu_t - \nu_{t-1})\mathcal{A}_{q,t} - \nu_t \mathcal{B}_{q,t} = 0, \tag{F.35}$$

$$\mathcal{S}_t^\pi \equiv \mathcal{V}_{\pi,t} + (\nu_t - \nu_{t-1})\mathcal{A}_{\pi,t} - \nu_t \mathcal{B}_{\pi,t} = 0. \tag{F.36}$$

Stack the two controls and the current NKPC multiplier as

$$\mathbf{k}_t = (q_t, \pi_t, \nu_t)', \qquad \mathcal{R}_t = (\mathcal{F}_t, \mathcal{S}_t^q, \mathcal{S}_t^\pi)'. \tag{F.37}$$

Define the selectors $e_q \equiv (1,0,0)'$, $e_\pi \equiv (0,1,0)'$, and $e_\nu \equiv (0,0,1)'$. They recover the terms-of-trade, inflation, and multiplier components of $\mathbf{k}_t$. Equations (F.27), (F.35), and (F.36), with $\nu_{-1} = 0$ and the no-explosive-mode condition, form the exact local Ramsey first-order, or Karush–Kuhn–Tucker (KKT), system. When unconstrained, the interest rate is recovered from (F.21) and adds no allocation trade-off.

These first-order conditions identify a candidate allocation, not by themselves an optimum. Let $\mathfrak{F}(\mathcal{P}_0) \equiv \{\mathcal{F}_t(h^t)\}_{t,h^t}$ collect the equilibrium constraints. The candidate is a locally unique Ramsey allocation under three maintained conditions, which are the local regularity

conditions invoked in Proposition 3. First, $D\mathfrak{F}$ is onto in the discounted sequence space, so the equality constraints satisfy a local constraint qualification. Second, the linearization of the KKT equations, the boundary $\nu_{-1} = 0$, and the no-explosive-mode condition is invertible; the implicit function theorem then selects a unique differentiable local branch. Third, the second variation of the Lagrangian is uniformly negative on every nonzero feasible tangent $\delta\mathcal{P}$:

$$D^2_{\mathcal{P}\mathcal{P}}\mathcal{J}_0[\delta\mathcal{P}, \delta\mathcal{P}] \leq -c\|\delta\mathcal{P}\|^2_\beta, \qquad c > 0, \qquad D\mathfrak{F}[\delta\mathcal{P}] = 0,$$

where $\|\delta\mathcal{P}\|^2_\beta = \mathbb{E}_0 \sum_{t\geq 0} \beta^t[(\delta q_t)^2 + (\delta\pi_t)^2]$, and the derivatives are evaluated at the candidate allocation and its multiplier sequence. The constrained second-order sufficient condition makes the selected branch a strict local maximum, and invertibility makes it locally unique. Without these conditions, the derivations below characterize only a non-explosive solution of the KKT system (a KKT branch); a root count and small KKT residuals do not establish optimality or global uniqueness.

F.5 Closed-form tariff-level continuation

Proof begins: Proposition 3.

The proof has two stages. Steps 1–4 (this subsection) derive the event-contingent KKT branch after a tariff innovation is realized; under the local regularity conditions of Appendix F.4, it is the local Ramsey allocation. Steps 5–6 (Appendix F.6) average these allocations over a change in tariff risk and impose the date-0 boundary. The order matters: even when the tariff never changes along the displayed path, the planner must first specify policy in the states in which it does change.

The differentiation below follows the standard perturbation approach of Schmitt-Grohé and Uribe (2004). We display the substitutions because the uncertainty response depends on both the slope and the curvature of the tariff-level continuation.

Step 1: write the residual as an explicit lead–lag map. On a deterministic tariff continuation, the conditional expectation in (F.27) equals its realized future value. Substituting the three exact conditions gives the smooth lead–lag map

$$\widetilde{\mathcal{R}}(\mathbf{k}_-, \mathbf{k}, \mathbf{k}_+; \tau, \tau_+) \equiv \begin{pmatrix} \mathcal{A} - \mathcal{B} - \beta\mathcal{A}_+ \\ \mathcal{V}_q + (\nu - \nu_-)\mathcal{A}_q - \nu\mathcal{B}_q \\ \mathcal{V}_\pi + (\nu - \nu_-)\mathcal{A}_\pi - \nu\mathcal{B}_\pi \end{pmatrix} = 0. \tag{F.38}$$

Here undated functions are evaluated at (q, π, τ), the plus subscript is evaluated at (q_+, π_+, τ_+), and $\mathbf{k}_- = (q_-, \pi_-, \nu_-)'$. Only its inherited multiplier ν_- enters the residual (in the last two rows); keeping the full vector makes the lead–lag notation uniform. All derivatives below are evaluated at $\bar{\mathbf{k}} = (0, 0, 0)'$ and $\bar{\tau} = 0$. Define

$$J_- \equiv D_{\mathbf{k}_{t-1}}\widetilde{\mathcal{R}}, \qquad J_0 \equiv D_{\mathbf{k}_t}\widetilde{\mathcal{R}}, \qquad J_+ \equiv D_{\mathbf{k}_{t+1}}\widetilde{\mathcal{R}}, \tag{F.39}$$

$$d_1(r) \equiv D_{\tau_t}\widetilde{\mathcal{R}} + rD_{\tau_{t+1}}\widetilde{\mathcal{R}}. \tag{F.40}$$

These objects follow directly from the primitives. Because $f(0) = 0$, $f'(0) = 1$, the production correction is fixed at its no-tariff value, which, together with the calibration of χ, makes the no-tariff steady state constrained efficient, and price-adjustment costs are locally quadratic,

$$\mathcal{A}_q = \mathcal{A}_\tau = \mathcal{B}_\pi = \mathcal{V}_{q\pi} = \mathcal{V}_{\pi\tau} = 0 \quad \text{at the steady state.} \tag{F.41}$$

Define $\theta_q \equiv \partial \log Y/\partial q|_{ss} > 0$. Because $mc = e^q Y^\varphi$ and steady-state marginal cost is one, $mc_q = 1 + \varphi\theta_q > 0$. The local pricing and curvature coefficients are

$$\begin{aligned} &a_0 \equiv \mathcal{A}_\pi = \kappa\bar{\mathcal{D}} > 0, \quad b_0 \equiv \mathcal{B}_q = (\epsilon - 1)\bar{\mathcal{D}}(1 + \varphi\theta_q) > 0, \\ &h_q \equiv -\mathcal{V}_{qq}, \qquad\qquad h_\pi \equiv -\mathcal{V}_{\pi\pi} > 0. \end{aligned} \tag{F.42}$$

The primitives make a_0, b_0, and h_π positive. $h_q > 0$ holds at the baseline ($h_q \approx 0.79$), and we maintain it. Direct differentiation of (F.38) then gives

$$\begin{aligned} J_- &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -a_0 \end{pmatrix}, \qquad J_0 = \begin{pmatrix} -b_0 & a_0 & 0 \\ -h_q & 0 & -b_0 \\ 0 & -h_\pi & a_0 \end{pmatrix}, \\ J_+ &= \begin{pmatrix} 0 & -\beta a_0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \qquad d_1(r) = \begin{pmatrix} -\mathcal{B}_\tau \\ \mathcal{V}_{q\tau} \\ 0 \end{pmatrix}. \end{aligned} \tag{F.43}$$

The generic definition (F.40) allows a direct future-tariff term. In this model, $\mathcal{A}_\tau = 0$ at the steady state, so $d_1(r)$ itself does not depend on r; tariff persistence still matters through the matrix $\mathcal{M}(r)$ below.

Introduce an innovation amplitude ζ. A tariff innovation at date 0 implies $\tau_t(\zeta) = \zeta\rho_\tau^t$ for $t \geq 0$. Write the associated Ramsey path as

$$\mathbf{k}_t(\zeta) = \bar{\mathbf{k}} + \zeta\mathbf{k}_{1,t} + \frac{\zeta^2}{2}\mathbf{k}_{2,t} + O(\zeta^3). \tag{F.44}$$

Differentiating (F.38) once with respect to ζ gives

$$J_{-}\mathbf{k}_{1,t-1} + J_0\mathbf{k}_{1,t} + J_{+}\mathbf{k}_{1,t+1} + d_1(\rho_\tau)\rho_\tau^t = 0. \tag{F.45}$$

Step 2: solve the part that follows the realized tariff. Try a particular solution $\mathbf{k}_{1,t}^p = p_\tau \rho_\tau^t$. Substitution into (F.45) yields

$$\left(\rho_\tau^{-1}J_{-} + J_0 + \rho_\tau J_{+}\right) p_\tau \rho_\tau^t + d_1(\rho_\tau)\rho_\tau^t = 0. \tag{F.46}$$

Define the matrix polynomial

$$\mathcal{M}(r) \equiv r^{-1}J_{-} + J_0 + rJ_{+}. \tag{F.47}$$

After dividing (F.46) by ρ_τ^t, the coefficient is

$$p_\tau = -\mathcal{M}(\rho_\tau)^{-1} d_1(\rho_\tau). \tag{F.48}$$

The coefficient p_τ is therefore not fitted: it is obtained by evaluating the derivative matrices of the exact Ramsey system at the tariff persistence ρ_τ.

Step 3: impose boundedness and commitment from scratch. A homogeneous geometric component wr^t satisfies

$$\mathcal{M}(r)w = 0 \quad \Longleftrightarrow \quad (J_{-} + rJ_0 + r^2 J_{+})w = 0. \tag{F.49}$$

The matrices in (F.43) make the endogenous commitment dynamics explicit. The three homogeneous rows are

$$-b_0 q_t + a_0\pi_t - \beta a_0 \pi_{t+1} = 0, \tag{F.50}$$

$$-h_q q_t - b_0 \nu_t = 0, \tag{F.51}$$

$$-h_\pi \pi_t + a_0 \nu_t - a_0 \nu_{t-1} = 0. \tag{F.52}$$

The second and third equations imply $q_t = -(b_0/h_q)\nu_t$ and $\pi_t = (a_0/h_\pi)(\nu_t - \nu_{t-1})$. Substituting both expressions into (F.50) leaves one scalar recurrence:

$$-\nu_{t-1} + (1 + \beta + \lambda_R)\nu_t - \beta\nu_{t+1} = 0, \qquad \lambda_R \equiv \frac{b_0^2 h_\pi}{a_0^2 h_q} > 0. \tag{F.53}$$

For a component $\nu_t = r^t$, the two nonzero roots solve

$$\beta r^2 - (1 + \beta + \lambda_R)r + 1 = 0. \tag{F.54}$$

Because $\lambda_R > 0$, the smaller root lies strictly between zero and one and the other root exceeds one. Boundedness eliminates the latter. The stable root and a corresponding eigenvector, normalized so that its multiplier component equals one, are therefore

$$\rho_R = \frac{1 + \beta + \lambda_R - \sqrt{(1 + \beta + \lambda_R)^2 - 4\beta}}{2\beta} \in (0, 1), \tag{F.55}$$

$$w_R = \begin{pmatrix} -b_0/h_q \\ a_0(\rho_R - 1)/(h_\pi \rho_R) \\ 1 \end{pmatrix}. \tag{F.56}$$

Thus the bounded first-order path is

$$\mathbf{k}_{1,t} = p_\tau \rho_\tau^t + c_1 w_R \rho_R^t. \tag{F.57}$$

The date-0 planner inherits no price-setting promise, so $\nu_{-1} = 0$. Taking the multiplier component and extrapolating each geometric component one period back gives

$$0 = e_\nu' \mathbf{k}_{1,-1} = \rho_\tau^{-1} e_\nu' p_\tau + c_1 \rho_R^{-1} e_\nu' w_R. \tag{F.58}$$

Solving this scalar equation gives

$$c_1 = -\frac{\rho_R}{\rho_\tau} \frac{e_\nu' p_\tau}{e_\nu' w_R}. \tag{F.59}$$

The backward extrapolation in (F.58) applies only to the inherited multiplier component; it does not assert that the tariff innovation occurred at date -1. The formula requires $\det \mathcal{M}(\rho_\tau) \neq 0$ and $e_\nu' w_R \neq 0$; both conditions hold at the baseline calibration.

Step 4: derive the curvature of the continuation. Let

$$\xi_{1,t} \equiv \left(\mathbf{k}_{1,t-1}', \mathbf{k}_{1,t}', \mathbf{k}_{1,t+1}', \rho_\tau^t, \rho_\tau^{t+1}\right)' \tag{F.60}$$

be the first derivative of every argument of (F.38). Let $\mathcal{H} \equiv D^2 \widetilde{\mathcal{R}}$ denote the vector-valued Hessian: the jth entry of $\mathcal{H}[a, b]$ is the Hessian of residual j applied to directions (a, b).

Differentiating the residual a second time gives

$$J_{-}\mathbf{k}_{2,t-1} + J_0\mathbf{k}_{2,t} + J_{+}\mathbf{k}_{2,t+1} + n_{2,t} = 0, \qquad n_{2,t} \equiv \mathcal{H}[\xi_{1,t}, \xi_{1,t}]. \tag{F.61}$$

Thus $n_{2,t}$ is the Hessian contribution to the second-order KKT residual. It collects every quadratic product of the first-order movements in the lagged, current, and future Ramsey variables and in the current and future tariff.

To make its time dependence explicit, define the two augmented modal directions

$$z_\tau \equiv (\rho_\tau^{-1}p_\tau', p_\tau', \rho_\tau p_\tau', 1, \rho_\tau)', \tag{F.62}$$

$$z_R \equiv (\rho_R^{-1}c_1 w_R', c_1 w_R', \rho_R c_1 w_R', 0, 0)'. \tag{F.63}$$

Equation (F.57) implies $\xi_{1,t} = z_\tau\rho_\tau^t + z_R\rho_R^t$. Bilinearity of the Hessian then gives

$$\begin{aligned} n_{2,t} = {} & \underbrace{\mathcal{H}[z_\tau, z_\tau]}_{n_2(\rho_\tau^2)}(\rho_\tau^2)^t + \underbrace{2\mathcal{H}[z_\tau, z_R]}_{n_2(\rho_\tau\rho_R)}(\rho_\tau\rho_R)^t \\ & + \underbrace{\mathcal{H}[z_R, z_R]}_{n_2(\rho_R^2)}(\rho_R^2)^t. \end{aligned} \tag{F.64}$$

Thus each $n_2(r)$ is the coefficient on one geometric component of the Hessian contribution $n_{2,t}$; it is generated by the curvature of the equilibrium and Ramsey conditions, not by a new exogenous disturbance. For each $r \in \mathcal{P}_2 \equiv \{\rho_\tau^2, \rho_\tau\rho_R, \rho_R^2\}$, substitution of $p_2(r)r^t$ into (F.61) gives

$$p_2(r) = -\mathcal{M}(r)^{-1}n_2(r). \tag{F.65}$$

The complete bounded second-order path is

$$\mathbf{k}_{2,t} = \sum_{r\in\mathcal{P}_2} p_2(r)r^t + c_2 w_R\rho_R^t, \qquad c_2 = -\frac{\rho_R}{e_\nu' w_R}\sum_{r\in\mathcal{P}_2} r^{-1}e_\nu' p_2(r). \tag{F.66}$$

The expression for c_2 follows by setting $e_\nu'\mathbf{k}_{2,-1} = 0$, exactly as in (F.58). The displayed geometric formulas require

$$\rho_\tau\rho_\sigma \neq 0, \qquad \det\mathcal{M}(r) \neq 0 \quad \text{for all } r \in \{\rho_\tau, \rho_\tau^2, \rho_\tau\rho_R, \rho_R^2, \rho_\sigma\}. \tag{F.67}$$

These conditions exclude resonance between an exogenous or quadratic mode and the endogenous commitment root. If resonance occurs, a solution can still be constructed, but the corresponding term generally takes the form $t\rho_R^t$ rather than a constant times ρ_R^t.

This completes the tariff-level calculation. Although the uncertainty response is third

order, it requires only the first two tariff-level derivatives. For a mean-zero future tariff innovation u, the first-order term $\mathbf{k}_{1,0}u$ has zero conditional mean, whereas the second-order term $\mathbf{k}_{2,0}u^2/2$ contributes $\mathbf{k}_{2,0}\mathbb{E}_t u^2/2$; the nonlinearity of the pricing residual adds a contribution quadratic in $\mathbf{k}_{1,0}$ (Step 5). Because the symmetric tariff innovation also satisfies $\mathbb{E}_t u^3 = 0$, no third tariff-level derivative enters the retained third-order response. Imposing the linear LQ targeting rule on $\mathbf{k}_{2,t}$ would omit the Hessian terms in (F.64) and so define a restricted continuation rather than the nonlinear Ramsey plan.

F.6 Third-order risk shift and closed-form response

This subsection converts the tariff-level continuation into the response to a volatility innovation, linking Appendix F.5 to the pure tariff-uncertainty experiment.

Step 5: integrate the continuation over the wider tariff distribution. Along the conditional history on which realized tariff innovations are zero, the innovation realized at date $t+1$ is

$$u_{t+1} = \bar{\sigma}_\tau e^{s_t}\varepsilon_{\tau,t+1}, \qquad \mathbb{E}_t\varepsilon_{\tau,t+1} = 0, \qquad \mathbb{E}_t\varepsilon^2_{\tau,t+1} = 1. \tag{F.68}$$

If u_{t+1} is realized, it becomes the date-0 tariff amplitude for the continuation derived above. Hence, at its impact date,

$$\begin{pmatrix} \mathbf{k}_{t+1}(u_{t+1}) \\ \tau_{t+1} \end{pmatrix} = \begin{pmatrix} \bar{\mathbf{k}} \\ 0 \end{pmatrix} + \underbrace{\begin{pmatrix} \mathbf{k}_{1,0} \\ 1 \end{pmatrix}}_{\ell_1} u_{t+1} + \frac{1}{2}\underbrace{\begin{pmatrix} \mathbf{k}_{2,0} \\ 0 \end{pmatrix}}_{\ell_2} u^2_{t+1} + O(u^3_{t+1}). \tag{F.69}$$

Only the pricing residual contains a conditional expectation after the Lagrangian is reindexed. Define its contribution from the future state by

$$\mathcal{C}(\mathbf{k}_+, \tau_+) \equiv (-\beta\mathcal{A}(q_+, \pi_+, \tau_+), 0, 0)'. \tag{F.70}$$

Taylor expanding this primitive continuation term before taking expectations gives

$$\begin{aligned} \mathcal{C}\{z(u)\} = \mathcal{C}(\bar{z}) &+ D\mathcal{C}\,\ell_1 u \\ &+ \frac{1}{2}\left[D\mathcal{C}\,\ell_2 + D^2\mathcal{C}[\ell_1, \ell_1]\right] u^2 + O(u^3). \end{aligned} \tag{F.71}$$

Every derivative in this equation is evaluated at the deterministic steady state. For order accounting, introduce a perturbation parameter δ that scales both the tariff-innovation standard deviation and the volatility-state innovation. Then $\bar{\sigma}_\tau = O(\delta)$ and $s_t = O(\delta)$, so

$\bar{\sigma}_\tau^2 s_t = O(\delta^3)$. The first-order term has zero conditional expectation. The change in the second moment is

$$\begin{aligned}\Delta_s \mathbb{E}_t u_{t+1}^2 &= \bar{\sigma}_\tau^2 \left(e^{2s_t} - 1\right) \\ &= 2\bar{\sigma}_\tau^2 s_t + O(\delta^4).\end{aligned} \tag{F.72}$$

Define the local risk state

$$v_t \equiv \bar{\sigma}_\tau^2 s_t = v_0 \rho_\sigma^t. \tag{F.73}$$

It equals the leading change in one-half of the conditional variance. Combining (F.71) and (F.72) shows that the variance change adds $g^v v_t$ to the KKT residual, where

$$g^v \equiv D\mathcal{C}\,\ell_2 + D^2\mathcal{C}[\ell_1, \ell_1]. \tag{F.74}$$

This vector can also be written entirely in the primitive derivatives defined above. Let $e_F = (1, 0, 0)'$ select the pricing residual, define $q_{1,0} = e_q' \mathbf{k}_{1,0}$ and $\pi_{j,0} = e_\pi' \mathbf{k}_{j,0}$ for $j \in \{1, 2\}$, and evaluate $\mathcal{D}$, $\mathcal{D}_q$, and $\mathcal{D}_\tau$ at the steady state. Since $\mathcal{C} = -\beta\kappa e_F \mathcal{D} f(\pi)$, $f'(0) = 1$, and $f''(0) = 3$, direct differentiation gives

$$g^v = -\beta\kappa e_F \left[\bar{\mathcal{D}}\pi_{2,0} + 3\bar{\mathcal{D}}\pi_{1,0}^2 + 2\mathcal{D}_q q_{1,0}\pi_{1,0} + 2\mathcal{D}_\tau \pi_{1,0}\right]. \tag{F.75}$$

The two terms in (F.74) have distinct roles. $D\mathcal{C}\,\ell_2$ maps the curvature of the tariff-contingent continuation policy into the conditional mean of the future pricing residual. $D^2\mathcal{C}[\ell_1, \ell_1]$ instead evaluates the curvature of that residual along the first-order tariff response, producing the Jensen and covariance effects of the tariff-level slope. In the primitive expression (F.75), $\bar{\mathcal{D}}\pi_{2,0}$ is the policy-curvature contribution, $3\bar{\mathcal{D}}\pi_{1,0}^2$ comes from the curvature of Rotemberg price adjustment, $2\mathcal{D}_q q_{1,0}\pi_{1,0}$ records the interaction between the terms-of-trade and inflation responses, and $2\mathcal{D}_\tau \pi_{1,0}$ the interaction between the tariff and inflation responses. Equation (F.74) therefore maps, at the retained order, the tariff-level continuation into the uncertainty response; the same higher-order logic underlies analyses of stochastic-volatility shocks in Fernández-Villaverde, Guerrón-Quintana, Rubio-Ramírez and Uribe (2011).

Equation (F.75) also proves the strict-PPI special case directly. If flexible prices implement the efficient allocation after every tariff realization, Ramsey sets $\pi_{1,0} = \pi_{2,0} = 0$. It follows that $g^v = 0$, and hence $p_v = c_v = 0$: a pure volatility innovation has no allocation effect under the state-by-state efficiency condition.

Step 6: solve the risk response and impose the date-0 boundary. For any contingent-plan variable x_t, define the baseline-subtracted local uncertainty response by

$$\Delta^U x_t \equiv x_t(h_t^U) - x_t(h_t^0). \tag{F.76}$$

Both histories start from $\tau_{-1} = s_{-1} = 0$ and the commitment-from-scratch boundary $\nu_{-1} = 0$. On h_t^U, $\varepsilon_{\sigma,0} = 1$; on the baseline history h_t^0, $\varepsilon_{\sigma,0} = 0$. Both histories set $\varepsilon_{\tau,j} = 0$ for every $j \geq 0$ and $\varepsilon_{\sigma,j} = 0$ for $j \geq 1$. The baseline still retains $\bar{\sigma}_\tau > 0$ in the tariff process, and conditional expectations continue to integrate over future off-path innovations. Thus Δ^U removes only the date-0 volatility innovation; it does not compare the stochastic economy with a zero-risk economy. Set $\mathbf{k}_t^v \equiv \Delta^U \mathbf{k}_t$. Because v_t is already third order, $\mathbf{k}_t^v$ solves the first-order system with $g^v v_t$ as its forcing term:

$$J_- \mathbf{k}_{t-1}^v + J_0 \mathbf{k}_t^v + J_+ \mathbf{k}_{t+1}^v + g^v v_0 \rho_\sigma^t = 0. \tag{F.77}$$

A particular solution $v_0 p_v \rho_\sigma^t$ gives

$$\mathcal{M}(\rho_\sigma) p_v + g^v = 0, \qquad p_v = -\mathcal{M}(\rho_\sigma)^{-1} g^v. \tag{F.78}$$

The only bounded homogeneous component is again $w_R \rho_R^t$, so $\mathbf{k}_t^v = v_0[p_v \rho_\sigma^t + c_v w_R \rho_R^t]$ for a scalar c_v. Commitment from scratch requires

$$0 = e_\nu' \mathbf{k}_{-1}^v = v_0 \left[\rho_\sigma^{-1} e_\nu' p_v + c_v \rho_R^{-1} e_\nu' w_R \right]. \tag{F.79}$$

Solving the scalar boundary equation produces the closed form

$$\mathbf{k}_t^v = v_0 \left[p_v \rho_\sigma^t + c_v w_R \rho_R^t \right], \qquad p_v = -\mathcal{M}(\rho_\sigma)^{-1} g^v, \qquad c_v = -\frac{\rho_R}{\rho_\sigma} \frac{e_\nu' p_v}{e_\nu' w_R}. \tag{F.80}$$

Existence requires $\det \mathcal{M}(\rho_\sigma) \neq 0$ and $e_\nu' w_R \neq 0$, the rank conditions of Step 3 with ρ_σ in place of ρ_τ.

If $\rho_\sigma = \rho_R$, the inverse in (F.78) is unavailable because the exogenous risk state and the stable Ramsey component have the same root. The solution remains closed form. Let ℓ_R be a left-null vector and retain w_R as the right-null vector, so $\ell_R' \mathcal{M}(\rho_R) = 0$ and $\mathcal{M}(\rho_R) w_R = 0$. For a simple root satisfying

$$\ell_R' \rho_R \mathcal{M}'(\rho_R) w_R \neq 0, \tag{F.81}$$

define

$$a_R = -\frac{\ell'_R g^v}{\ell'_R \rho_R \mathcal{M}'(\rho_R) w_R}, \qquad \mathcal{M}(\rho_R) p_R + a_R \rho_R \mathcal{M}'(\rho_R) w_R + g^v = 0. \tag{F.82}$$

The second equation is solvable because the definition of a_R makes its right-hand side orthogonal to the left null space. Any solution p_R can be used; adding a multiple of w_R only changes the homogeneous coefficient below. Direct substitution into (F.77) gives

$$\mathbf{k}^v_t = v_0 \rho_R^t \left[p_R + (a_R t + c_R) w_R \right], \qquad c_R = a_R - \frac{e'_\nu p_R}{e'_\nu w_R}. \tag{F.83}$$

The expression for c_R imposes $e'_\nu \mathbf{k}^v_{-1} = 0$. Thus resonance replaces the two distinct geometric terms by a polynomial–geometric term. The scalar reduction below yields the $(t+1)\rho_R^t$ limit reported in the main text.

To obtain the scalar representation in Proposition 3, let $g_v \equiv e'_F g^v$ and write the commitment polynomial (F.54) as

$$P_R(r) \equiv \beta r^2 - (1 + \beta + \lambda_R) r + 1. \tag{F.84}$$

Using the explicit matrices in (F.43), direct inversion gives the terms-of-trade loading of the particular solution,

$$e'_q p_v = -\frac{g_v b_0 h_\pi \rho_\sigma}{a_0^2 h_q P_R(\rho_\sigma)}. \tag{F.85}$$

The date-0 boundary (F.79) implies that the loading of the stable homogeneous mode is

$$e'_q (c_v w_R) = \frac{g_v b_0 h_\pi \rho_R}{a_0^2 h_q P_R(\rho_\sigma)}. \tag{F.86}$$

Define

$$\Gamma_v \equiv \frac{g_v b_0 h_\pi (\rho_R - \rho_\sigma)}{a_0^2 h_q P_R(\rho_\sigma)}. \tag{F.87}$$

Combining (F.85) and (F.86) yields

$$\Delta^U q_t = v_0 \Gamma_v \frac{\rho_R^{t+1} - \rho_\sigma^{t+1}}{\rho_R - \rho_\sigma} + O(\delta^4). \tag{F.88}$$

Tariff variance does not shift the second and third KKT rows directly, so they still imply $q^v_t = -(b_0/h_q)\nu^v_t$ and $\pi^v_t = (a_0/h_\pi)(\nu^v_t - \nu^v_{t-1})$. Eliminating the multiplier gives

$$\Delta^U \pi_t = -\frac{a_0 h_q}{b_0 h_\pi} \left(\Delta^U q_t - \Delta^U q_{t-1} \right) + O(\delta^4), \qquad \Delta^U q_{-1} = 0. \tag{F.89}$$

Equations (F.88) and (F.89) are (39)–(40), with $\omega_R = a_0 h_q/(b_0 h_\pi)$ and $\sum_{j=0}^{t} \rho_R^{t-j}\rho_\sigma^j = (\rho_R^{t+1} - \rho_\sigma^{t+1})/(\rho_R - \rho_\sigma)$. Consumption and output follow by applying the steady-state Jacobian of the exact static maps (F.5) and (F.8); because the response is already third order, higher powers of $\mathbf{k}_t^v$ are beyond the retained order.

End of proof: Proposition 3.

Implementing interest rate. Let $d_{t+1} = q_{t+1} - q_t + \pi_{t+1}$. Taking logs of (F.21) and applying a cumulant expansion gives

$$\widehat{r}_t = -\log \mathbb{E}_t e^{-d_{t+1}} = \mathbb{E}_t d_{t+1} - \frac{1}{2}\operatorname{Var}_t(d_{t+1}) + \frac{1}{6}\operatorname{Cum}_{3,t}(d_{t+1}) + O(\delta^4). \tag{F.90}$$

Here $\operatorname{Cum}_{3,t}(d_{t+1})$ is the conditional third cumulant of d_{t+1} given date-t information. The volatility-induced change in the third cumulant is beyond third order under the symmetric innovation used here. If a realized-tariff continuation has $d(u) = A_D u + B_D u^2/2 + O(u^3)$, the direct risk term in the rate is

$$(B_D - A_D^2)v_t, \qquad A_D = (e_q + e_\pi)'\mathbf{k}_{1,0}, \qquad B_D = (e_q + e_\pi)'\mathbf{k}_{2,0}. \tag{F.91}$$

The full rate response follows by substituting (F.80) into (F.90).

F.7 Proof of the Ramsey targeting rule with tariff cost-push

Proof begins: Proposition 4.

The proof rewrites the local KKT rows of Appendices F.5–F.6 in terms of the gap $\widetilde{q}_t \equiv q_t - q^e(\tau_t)$. No new approximation is introduced. Case (L) uses the first-order system (F.45), whose remainder is $O(\zeta^2)$ (or $O(\delta^2)$ under the joint perturbation scaling). Case (U) uses the third-order risk system (F.77), whose remainder is $O(\delta^4)$. All derivatives are evaluated at the deterministic steady state, as in (F.43).

Step 1: the two static slopes. At $\pi = 0$, $Y(q, 0, \tau) = Y^a(e^q, \tau)$, so $\mathcal{V}(q, 0; \tau) = \mathcal{W}(e^q; \tau)$ in (F.9). The constrained-efficient terms of trade therefore solve $\mathcal{V}_q(q^e(\tau), 0, \tau) = 0$, which is the static condition (F.10). The steady state satisfies it under the calibration of χ. Implicit differentiation at $\tau = 0$ gives

$$q_1^e = -\frac{\mathcal{V}_{q\tau}}{\mathcal{V}_{qq}} = \frac{\mathcal{V}_{q\tau}}{h_q}, \tag{F.92}$$

which is the slope q_1^e of (F.111). Under strict PPI, $\pi = 0$ after every history, so $\mathcal{A} \equiv 0$, and the scaled NKPC (F.27) reduces to $\mathcal{B}(q^f(\tau), 0, \tau) = 0$, that is, $mc(q^f(\tau), 0, \tau) = 1$. Differentiating,

and using $\mathcal{B}_q = b_0$, gives

$$q_1^f = -\frac{\mathcal{B}_\tau}{b_0} = -\frac{mc_\tau}{mc_q}, \qquad mc_q = 1 + \varphi\theta_q. \tag{F.93}$$

The second equality uses $\mathcal{B}_\tau = (\epsilon - 1)\bar{\mathcal{D}}\, mc_\tau$ and $b_0 = (\epsilon - 1)\bar{\mathcal{D}}\, mc_q$, both evaluated where $mc = 1$. The two slopes coincide exactly when the fixed correction keeps flexible prices efficient to first order in the tariff.

Step 2: case (L), the realized tariff. Write q_t, π_t, and ν_t for the first-order terms $\zeta e_q' \mathbf{k}_{1,t}$, $\zeta e_\pi' \mathbf{k}_{1,t}$, and $\zeta e_\nu' \mathbf{k}_{1,t}$, and let $\tau_t = \zeta \rho_\tau^t$. With the matrices in (F.43) and $d_1 = (-\mathcal{B}_\tau, \mathcal{V}_{q\tau}, 0)'$, the three rows of (F.45) are

$$a_0\pi_t - \beta a_0 \pi_{t+1} - b_0 q_t - \mathcal{B}_\tau \tau_t = 0, \tag{F.94}$$

$$-h_q q_t - b_0 \nu_t + \mathcal{V}_{q\tau}\tau_t = 0, \tag{F.95}$$

$$-h_\pi \pi_t + a_0(\nu_t - \nu_{t-1}) = 0, \qquad \nu_{-1} = 0. \tag{F.96}$$

These are (F.50)–(F.52) with the tariff forcing restored. To first order, $q_t = \widetilde{q}_t + q_1^e \tau_t$. By (F.92), the tariff terms in (F.95) cancel exactly, leaving

$$\widetilde{q}_t = -\frac{b_0}{h_q}\, \nu_t. \tag{F.97}$$

Combining (F.97) with (F.96) gives $\pi_t = (a_0/h_\pi)(\nu_t - \nu_{t-1}) = -\omega_R(\widetilde{q}_t - \widetilde{q}_{t-1})$, with $\omega_R = a_0 h_q/(b_0 h_\pi)$ and $\widetilde{q}_{-1} \equiv -(b_0/h_q)\nu_{-1} = 0$. As in Appendix F.8, $\widetilde{q}_{-1} = 0$ records the absence of an inherited promise. It does not impose a date -1 gap on the economy. Dividing (F.94) by a_0 and substituting $q_t = \widetilde{q}_t + q_1^e \tau_t$ and (F.93) gives

$$\pi_t = \beta\pi_{t+1} + \frac{b_0}{a_0}\widetilde{q}_t + \frac{b_0}{a_0}\Big(q_1^e + \frac{\mathcal{B}_\tau}{b_0}\Big)\tau_t = \beta\pi_{t+1} + \frac{b_0}{a_0}\widetilde{q}_t + \frac{b_0}{a_0}\big(q_1^e - q_1^f\big)\tau_t. \tag{F.98}$$

This is (48) with the cost-push term μ_t^L of (49). Its coefficient is exact at first order. Because $b_0/a_0 = \kappa_{mc}\, mc_q$, it equals $\kappa_{mc}(mc_q q_1^e + mc_\tau)\tau_t$, the first-order term of the LQ pricing discrepancy z_t in (F.102). The term μ_t^L is therefore the KKT counterpart of the LQ pricing discrepancy z_t.

The reduction can be reversed. Given $(\pi_t, \widetilde{q}_t)$ that satisfy the two equations, define $q_t = \widetilde{q}_t + q_1^e \tau_t$ and $\nu_t = -(h_q/b_0)\widetilde{q}_t$. Then (F.94)–(F.96) hold, and boundedness of $\widetilde{q}_t$ is equivalent to boundedness of ν_t. The two-equation system is thus equivalent to the first-order KKT branch with $\nu_{-1} = 0$ and the no-explosive-mode condition.

At first order, certainty equivalence implies the same two equations hold on every history once π_{t+1} is replaced by $\mathbb{E}_t\pi_{t+1}$. The response to an innovation at any date $s \geq 0$ is the date-0 response shifted by s periods, because ν_{s-1} is predetermined with respect to that innovation.

Step 3: case (U), tariff uncertainty. Let (q_t, π_t, ν_t) now denote the components of $\mathbf{k}_t^v = \Delta^U \mathbf{k}_t$. The forcing in (F.77) is $g^v v_t = g_v e_F v_t$, because only the pricing residual contains a conditional expectation (F.70). Hence

$$\begin{aligned} a_0\pi_t - \beta a_0\pi_{t+1} - b_0 q_t + g_v v_t &= 0, \\ -h_q q_t - b_0\nu_t &= 0, \\ -h_\pi\pi_t + a_0(\nu_t - \nu_{t-1}) = 0, \qquad \nu_{-1} &= 0, \end{aligned}$$

up to $O(\delta^4)$. The shocked and no-innovation histories both have $\tau_t = 0$, so $\Delta^U q^e(\tau_t) = 0$ and $\Delta^U \widetilde{q}_t = q_t$. The algebra of Step 2 then gives (48) with $\mu_t^U = -(g_v/a_0)v_t = -(g_v/a_0)v_0\rho_\sigma^t$.

The two cost-push terms differ in nature. The term μ^L depends only on the static slopes (F.92)–(F.93), and the first-order pricing row is firms' Phillips curve, which holds under any policy. The term μ^U is endogenous to Ramsey policy instead: by (F.75), g_v is built from the Ramsey continuation coefficients $q_{1,0}$, $\pi_{1,0}$, and $\pi_{2,0}$. Under a different continuation policy the pricing shift is different. Under strict PPI it is zero, because $\pi \equiv 0$ in every continuation.

Step 4: solve the scalar system. Substitute the targeting rule into the gap Phillips curve. Because $(b_0/a_0)/\omega_R = b_0^2 h_\pi/(a_0^2 h_q) = \lambda_R$, this gives

$$\beta\widetilde{q}_{t+1} - (1 + \beta + \lambda_R)\widetilde{q}_t + \widetilde{q}_{t-1} = \frac{\mu_t}{\omega_R}, \qquad t \geq 0, \quad \widetilde{q}_{-1} = 0, \tag{F.99}$$

whose characteristic polynomial is P_R in (F.84). Its roots are $\rho_R \in (0,1)$ and $(\beta\rho_R)^{-1} > 1$, so $\beta\rho_R + \rho_R^{-1} = 1 + \beta + \lambda_R$. With L and F the lag and lead operators,

$$(1-\rho_R L)(1-\beta\rho_R F)\widetilde{q}_t = (1+\beta\rho_R^2)\widetilde{q}_t - \beta\rho_R\widetilde{q}_{t+1} - \rho_R\widetilde{q}_{t-1} = -\rho_R\big[\beta\widetilde{q}_{t+1} - (1+\beta+\lambda_R)\widetilde{q}_t + \widetilde{q}_{t-1}\big],$$

so (F.99) is equivalent to $(1 - \rho_R L)(1 - \beta\rho_R F)\widetilde{q}_t = -(\rho_R/\omega_R)\mu_t$. Let $u_t \equiv \widetilde{q}_t - \rho_R\widetilde{q}_{t-1}$. A bounded u_t requires forward solution of $u_t - \beta\rho_R u_{t+1} = -(\rho_R/\omega_R)\mu_t$; any other solution contains the explosive mode $(\beta\rho_R)^{-t}$. For $\mu_t = \mu_0\rho^t$ with $0 \leq \rho < 1$,

$$u_t = -\frac{\rho_R}{\omega_R}\sum_{k\geq 0}(\beta\rho_R)^k\mu_{t+k} = -\frac{\rho_R\,\mu_0}{\omega_R(1-\beta\rho_R\rho)}\rho^t.$$

Iterating $\widetilde{q}_t = \rho_R \widetilde{q}_{t-1} + u_t$ from $\widetilde{q}_{-1} = 0$ gives the finite sum in (50), and $\pi_0 = -\omega_R \widetilde{q}_0 = \rho_R \mu_0/(1 - \beta\rho_R\rho)$. The solution is unique because the boundary fixes the coefficient on ρ_R^t and boundedness removes the other root. The finite sum remains valid when $\rho = \rho_R$, where it equals $(t+1)\rho_R^t$, so the scalar form needs no separate resonant case.

Step 5: cumulative inflation. Summing the targeting rule telescopes:

$$\sum_{s=0}^{t} \pi_s = -\omega_R\big(\widetilde{q}_t - \widetilde{q}_{-1}\big) = -\omega_R \widetilde{q}_t. \tag{F.100}$$

The cumulative change in the log PPI is therefore proportional to the current gap. Because $\rho_R < 1$ and $\rho < 1$, $\widetilde{q}_t \to 0$, so the PPI level returns to its initial value. Commitment thus works like a price-level target: any PPI inflation that Ramsey accepts on impact is later reversed by inflation of the opposite sign as the gap closes. The identity is exact for the first-order level coefficients and for the leading third-order uncertainty coefficients. It uses the date-0 boundary $\widetilde{q}_{-1} = 0$, and the return occurs at the slower of ρ_R and ρ. The proposition makes no such claim for the second-order tariff-level coefficients $\mathbf{k}_{2,t}$. Their Hessian term $n_{2,t}$ in (F.61) enters the inflation row, so $\pi_{2,t}$ is not proportional to $\nu_{2,t} - \nu_{2,t-1}$ in general.

Step 6: consequences. (i) If $q_1^e = q_1^f$, then $\mu^L \equiv 0$ and the unique bounded solution is $\widetilde{q}_t = \pi_t = 0$, so $q_t = q_1^f \tau_t$, the strict-PPI allocation. If $q_1^e \neq q_1^f$, then $\pi_0 \neq 0$ because $\rho_R > 0$, so strict PPI is not the first-order Ramsey response. Setting $\pi \equiv 0$ in the policy-invariant Phillips curve (F.98) gives the strict-PPI gap $\widetilde{q}_t^{PPI} = -(q_1^e - q_1^f)\tau_t$. Since $P_R(\rho_R) = 0$ implies $\rho_R \lambda_R = (1 - \rho_R)(1 - \beta\rho_R)$,

$$\frac{\widetilde{q}_0}{\widetilde{q}_0^{PPI}} = \frac{\lambda_R \rho_R}{1 - \beta\rho_R\rho_\tau} = \frac{(1-\rho_R)(1-\beta\rho_R)}{1 - \beta\rho_R\rho_\tau} \in (0, 1).$$

The ratio is below one because $(1-\rho_R)(1-\beta\rho_R) < 1 - \beta\rho_R \leq 1 - \beta\rho_R\rho_\tau$: on impact, Ramsey closes part of the terms-of-trade gap that strict PPI leaves open, about one fifth at the baseline (ratio 0.797; see the numerical check below). The comparison is limited to case (L), because μ^U is not policy invariant (Step 3).

(ii) Every term in (F.75) is proportional to $\pi_{1,0}$ or $\pi_{2,0}$, so $g_v = 0$ and $\mu^U \equiv 0$ when $\pi_{1,0} = \pi_{2,0} = 0$. First-order alignment $q_1^e = q_1^f$ gives $\pi_{1,0} = 0$ but not, in general, $\pi_{2,0} = 0$. It then leaves $g_v = -\beta\kappa\bar{\mathcal{D}}\pi_{2,0}$, which reflects misalignment at second order in the tariff. The state-by-state efficiency condition of Proposition 2 removes both terms.

(iii) Consistency with Proposition 3. With $\mu_0 = -(g_v/a_0)v_0$ and $\rho = \rho_\sigma$, the solution has impact loading $g_v\rho_R/[a_0\omega_R(1 - \beta\rho_R\rho_\sigma)]$. Because $P_R(r) = \beta(r - \rho_R)\big(r - (\beta\rho_R)^{-1}\big)$,

$P_R(\rho_\sigma) = (\rho_R - \rho_\sigma)(1 - \beta\rho_R\rho_\sigma)/\rho_R$, and the loading equals Γ_v in (F.87), including its limit at $\rho_\sigma = \rho_R$.

Scope. Because both systems are linear, simultaneous date-0 level and volatility innovations produce the sum of the two leading responses, with $\mu_t = \mu_t^L + \mu_t^U$. That sum is not the complete third-order response to the combined shock, which also contains tariff-level terms of order ζ^2, ζ^3, and $\zeta\bar{\sigma}_\tau^2$. Those terms involve the Hessian of the KKT system and do not satisfy the linear rule.

Numerical check. At the baseline calibration, the primitive coefficients give $b_0/a_0 = 0.45225$, $q_1^e = 0.49723$, $q_1^f = 0.22388$, and $\mu^L = 0.12362$ per unit log tariff. The primitive formula (F.75) gives $g_v = 1.84468$, matching Table F.2, and hence $\mu^U = -0.046117$ per unit v. The scalar solution reproduces $\pi_{1,0} = 0.027365$, which is 0.274 percentage points of impact PPI inflation after a 0.10 log-point tariff at first order. It also reproduces $\Gamma_v = 0.0812965$ and an impact gap equal to 0.797 of strict PPI's.

As an independent check, the three-row KKT system was solved directly in sequence space over 600 periods, with $\nu_{-1} = 0$ and a zero terminal condition. Its paths match (50) to within 4×10^{-16} per unit log tariff for the level cost-push term and to a relative error below 2×10^{-15} for the uncertainty cost-push term at the baseline v_0, and also in the resonant case $\rho = \rho_R$. The residuals of the gap Phillips curve and of the targeting rule are below 2×10^{-17}, and cumulative inflation summed over the first 300 periods is below 10^{-15} in absolute value.

End of proof: Proposition 4.

F.8 The LQ Level-Ramsey approximation and the optimal date-0 choice

This subsection derives the Level-Ramsey benchmark of main-text Subsection 4.3: the linear-quadratic (LQ) targeting rule that is optimal for small realized tariff changes, and the optimal date-0 response to a volatility innovation when that rule governs policy from date 1 onward. Let y_t^a denote output used for Home absorption and exports, let y_t^e denote its constrained-efficient target, and define $\widetilde{y}_t^a = y_t^a - y_t^e$. To first order, the Phillips curve can be written

$$\pi_t = \kappa_y \widetilde{y}_t^a + \beta \mathbb{E}_t \pi_{t+1} + z_t, \tag{F.101}$$

where, with $\kappa_{mc} = (\epsilon - 1)/\kappa$,

$$\omega_y = \frac{1 + \varphi\theta_q}{\theta_q}, \qquad \kappa_y = \kappa_{mc}\omega_y, \qquad z_t = \kappa_{mc}\zeta_t^e. \tag{F.102}$$

These objects follow from the first-order expansion of normalized marginal cost around the efficient target, $\widehat{mc}_t = \omega_y \widetilde{y}^a_t + \zeta^e_t + O(2)$. Here $\zeta^e_t \equiv \log mc(q^e_t, 0, \tau_t)$ is log normalized marginal cost evaluated at the constrained-efficient allocation. Thus z_t is the pricing discrepancy induced by a realized tariff when the fixed production correction fails to implement that allocation; it is zero under state-by-state efficiency and is not an independent policy shock.

Divide the second-order utility loss by χ and omit terms independent of monetary policy. The policy-relevant quadratic loss is

$$\frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\lambda_\pi \pi_t^2 + \lambda_y (\widetilde{y}^a_t)^2). \tag{F.103}$$

The corresponding unnormalized utility loss is χ times (F.103), up to third-order terms. The Rotemberg resource cost gives $\lambda_\pi = \kappa$. The real-gap weight is not generically $\sigma + \varphi$ because consumption and labor do not move one-for-one with the open-economy output gap. Let

$$\theta_q = \left.\frac{\partial y}{\partial q}\right|_{ss}, \qquad c_q = \frac{1-\alpha}{\sigma}, \qquad g_2 = (1-\gamma)\alpha(1-\alpha), \tag{F.104}$$

and let y^a_{qq} be the second derivative of the log static output map. Direct differentiation of utility along the feasible map gives

$$\lambda_y = \frac{(\sigma-1)c_q^2 + g_2/\sigma + \chi[y^a_{qq} + (1+\varphi)\theta_q^2]}{\chi \theta_q^2}. \tag{F.105}$$

The strict local second-order condition requires $\lambda_y > 0$, which holds at the baseline calibration.

Let ϕ_t multiply (F.101). The Lagrangian is

$$\mathcal{J} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{1}{2}\left(\lambda_\pi \pi_t^2 + \lambda_y (\widetilde{y}^a_t)^2\right) + \phi_t \left(\pi_t - \kappa_y \widetilde{y}^a_t - \beta \mathbb{E}_t \pi_{t+1} - z_t\right) \right\}, \tag{F.106}$$

where commitment begins at date 0. By the law of iterated expectations, the expected-inflation term can be reindexed as

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(-\beta \phi_t \mathbb{E}_t \pi_{t+1}\right) = \mathbb{E}_0 \sum_{t=1}^{\infty} \beta^t \left(-\phi_{t-1}\pi_t\right),$$

so date-t inflation enters the Lagrangian through $\phi_t - \phi_{t-1}$. Differentiating with respect to $\widetilde{y}^a_t$ and π_t gives the first-order conditions

$$\lambda_y \widetilde{y}^a_t - \kappa_y \phi_t = 0, \qquad \lambda_\pi \pi_t + \phi_t - \phi_{t-1} = 0, \qquad \phi_{-1} = 0. \tag{F.107}$$

For $t \geq 1$, eliminating ϕ_t and ϕ_{t-1} gives the continuation targeting criterion for this LQ problem:

$$\lambda_\pi \pi_t + \frac{\lambda_y}{\kappa_y}(\widetilde{y}^a_t - \widetilde{y}^a_{t-1}) = 0. \tag{F.108}$$

The criterion in (F.108) applies from date 1 onward. At date 0, commitment starts from scratch, so the economic boundary is instead $\phi_{-1} = 0$. The date-0 condition is therefore $\lambda_\pi \pi_0 + (\lambda_y/\kappa_y)\widetilde{y}^a_0 = 0$. Setting $\widetilde{y}^a_{-1} = 0$ merely writes this condition in the same difference form as (F.108); it does not impose a date--1 output gap on the economy. When $z_t = 0$ in every state, $\pi_t = \widetilde{y}^a_t = 0$ satisfies these conditions and minimizes the loss; this is the LQ form of the divine coincidence, whose exact counterpart is proved in Appendix F.3.

The rule is linear in the *gap*, but the economic object measured by that gap is not a straight line in the tariff. Let

$$y^a(q,\tau) \equiv \log\left(\frac{C_H(e^q) + X(e^q,\tau)}{Y_{ss}}\right), \qquad y^e(\tau) \equiv y^a(q^e(\tau),\tau), \tag{F.109}$$

where $q^e(\tau)$ is the solution to the static planner condition (F.10). The gap used by the targeting rule is

$$\widetilde{y}^a_t = y^a(q_t,\tau_t) - y^e(\tau_t). \tag{F.110}$$

To obtain its third-order representation from primitives, write the planner's static first-order condition as $H(q,\tau) = 0$ and define $q^e_j = d^j q^e/d\tau^j|_0$. Repeated implicit differentiation gives

$$\begin{aligned}
q^e_1 &= -\frac{H_\tau}{H_q},\\
q^e_2 &= -\frac{H_{\tau\tau} + 2H_{q\tau}q^e_1 + H_{qq}(q^e_1)^2}{H_q},\\
q^e_3 &= -\frac{H_{\tau\tau\tau} + 3H_{q\tau\tau}q^e_1 + 3H_{qq\tau}(q^e_1)^2 + H_{qqq}(q^e_1)^3 + 3(H_{q\tau} + H_{qq}q^e_1)q^e_2}{H_q}.
\end{aligned} \tag{F.111}$$

All derivatives are evaluated at the deterministic steady state. Applying the chain rule to $y^a(q^e(\tau),\tau)$ then yields

$$\begin{aligned}
A^e_Y &= y^a_q q^e_1 + y^a_\tau,\\
B^e_Y &= y^a_{qq}(q^e_1)^2 + 2y^a_{q\tau}q^e_1 + y^a_{\tau\tau} + y^a_q q^e_2,\\
C^e_Y &= y^a_{qqq}(q^e_1)^3 + 3y^a_{qq\tau}(q^e_1)^2 + 3y^a_{q\tau\tau}q^e_1 + y^a_{\tau\tau\tau} + 3(y^a_{qq}q^e_1 + y^a_{q\tau})q^e_2 + y^a_q q^e_3.
\end{aligned} \tag{F.112}$$

The targeting relation is exact only for the LQ problem (F.101)–(F.103). In the third-order model we impose that same linear relation on a nonlinear economic gap whose efficient target

is approximated through cubic order. Thus the Dynare implementation uses

$$\widetilde{y}_t^a = y^a(q_t, \tau_t) - A_Y^e \tau_t - \frac{1}{2} B_Y^e \tau_t^2 - \frac{1}{6} C_Y^e \tau_t^3 + O(4). \tag{F.113}$$

The tangent $\theta_q(q_t - q_1^e \tau_t)$ is sufficient for first-order tariff-level dynamics, but not for the variance-induced residual in a third-order tariff-uncertainty calculation. The numerical check below therefore compares both the tariff-level coefficients and the pure tariff-uncertainty coefficient with the independently solved model.

No quadratic-cubic targeting equation is used in the paper's Ramsey results. Such an equation would require the complete cubic expansion of utility and the complete second-order expansions of marginal cost, the intertemporal pricing kernel, and all state-dependent terms. The exact nonlinear KKT system in Appendix F.4 retains those objects before the local third-order coefficient is taken. The linear targeting relation (F.108) enters only as the imposed continuation of the Level-Ramsey benchmark.

Now impose (F.108) as the continuation policy from date 1 onward and ask for the optimal date-0 response to a pure tariff uncertainty shock. Let Δ^U denote the difference between a variable's path with and without the date-0 volatility innovation, and let $\Delta^U \mathbb{E}_0 \pi_1$ denote the resulting third-order contribution to expected date-1 PPI inflation under this continuation. Let $u_0^{NK,L}$ denote the remaining third-order conditional-mean and covariance contribution to the exact date-0 price-setting residual after that expected-inflation term is separated. The two objects are determined differently, and neither is a free policy input. The direct pricing shift $u_0^{NK,L}$ is computed by substituting the Level-Ramsey continuation's own first- and second-order tariff-level coefficients into these conditional-mean and covariance terms. Under the maintained AR(1) volatility process it obeys $u_t^{NK,L} = u_0^{NK,L} \rho_\sigma^t$; the superscript L distinguishes it from the Taylor-rule pricing wedge in main-text Section 3.2. Expected inflation $\Delta^U \mathbb{E}_0 \pi_1$ is instead endogenous: it is determined jointly with the date-0 gap by the continuation's history-dependent targeting recurrence, the date-0 boundary, and the nonexplosive root, as in (F.118)–(F.121) below. Define

$$\mathcal{Z}_0 \equiv \beta \Delta^U \mathbb{E}_0 \pi_1 + u_0^{NK,L} \tag{F.114}$$

as the complete date-0 pricing residual caused by the volatility innovation. It is an endogenous effect of tariff risk, not an assumed cost-push shock. Subtract the no-innovation path from (F.101), retain these third-order terms, and use the date-0 commitment boundary $\phi_{-1} = 0$.

The date-0 Phillips curve and first-order condition are then

$$\Delta^U \pi_0 - \kappa_y \Delta^U \widetilde{y}_0^a = \mathcal{Z}_0, \tag{F.115}$$

$$\lambda_\pi \Delta^U \pi_0 + \frac{\lambda_y}{\kappa_y} \Delta^U \widetilde{y}_0^a = 0. \tag{F.116}$$

Therefore

$$\Delta^U \pi_0^{LQ} = \frac{\mathcal{Z}_0}{1+\chi_y}, \qquad \kappa_y \Delta^U \widetilde{y}_0^{a,LQ} = -\frac{\chi_y}{1+\chi_y}\mathcal{Z}_0, \qquad \chi_y \equiv \frac{\lambda_\pi \kappa_y^2}{\lambda_y}. \tag{F.117}$$

Because $\mathcal{Z}_0$ contains the endogenous term $\beta \Delta^U \mathbb{E}_0 \pi_1$, (F.117) is not yet a closed form in the direct pricing shift. To eliminate that term, impose the same pricing and targeting equations from date 1 onward. With χ_y as in (F.117), define

$$\mu_c \equiv \frac{\lambda_\pi \kappa_y}{\lambda_y}, \qquad \rho_L \equiv \frac{1+\beta+\chi_y - \sqrt{(1+\beta+\chi_y)^2 - 4\beta}}{2\beta}. \tag{F.118}$$

Combining the targeting rule with the pricing equation at every date yields

$$\beta \mathbb{E}_t \Delta^U \widetilde{y}_{t+1}^a - (1+\beta+\chi_y)\Delta^U \widetilde{y}_t^a + \Delta^U \widetilde{y}_{t-1}^a = \mu_c u_0^{NK,L} \rho_\sigma^t. \tag{F.119}$$

The root $\rho_L \in (0,1)$ selects the nonexplosive commitment path. With $\Delta^U \widetilde{y}_{-1}^a = 0$ and $d_L \equiv \rho_L^{-1} - \beta\rho_\sigma > 0$, that path is

$$\Delta^U \widetilde{y}_t^{a,LQ} = -\frac{\mu_c u_0^{NK,L}}{d_L} \sum_{j=0}^{t} \rho_L^{t-j} \rho_\sigma^j, \qquad \Delta^U \pi_t^{LQ} = -\frac{\lambda_y}{\lambda_\pi \kappa_y}\left(\Delta^U \widetilde{y}_t^{a,LQ} - \Delta^U \widetilde{y}_{t-1}^{a,LQ}\right). \tag{F.120}$$

Evaluating the path at date 0 gives the impact response in closed form:

$$\begin{aligned} \Delta^U \pi_0^{LQ} &= \frac{u_0^{NK,L}}{d_L} = \frac{\mathcal{Z}_0}{1+\chi_y}, \\ \kappa_y \Delta^U \widetilde{y}_0^{a,LQ} &= -\frac{\chi_y u_0^{NK,L}}{d_L} = -\frac{\chi_y}{1+\chi_y}\mathcal{Z}_0, \\ \mathcal{Z}_0 &= \frac{1+\chi_y}{d_L} u_0^{NK,L}. \end{aligned} \tag{F.121}$$

At the baseline calibration, $u_0^{NK,L} < 0$, so the optimum accepts a small PPI deflation and a positive gap that offsets part of the deflationary pricing shift. Because tariff uncertainty operates through curvature, the implementation evaluates $y_t^e(\tau_t)$ through third order, as in (F.113), before applying the linear targeting rule. This is a genuine date-0 optimum

conditional on the imposed LQ continuation. It is not the full nonlinear Ramsey optimum because the continuation's higher-order coefficients are restricted by (F.108).

F.9 Incremental date-0 PPI stabilization with Ramsey conditions from date 1 onward

The second restricted benchmark sets to zero only the incremental date-0 PPI response to the volatility innovation and imposes the Ramsey conditions from date 1 onward. The restriction is incremental because baseline tariff risk can generate lower-order risk corrections to inflation even without a date-0 volatility innovation. Write impact inflation as

$$\pi_0 = \pi_0^{\text{base}} + e_\pi' \mathbf{k}_0^v + O(4), \qquad v_0 = \bar{\sigma}_\tau^2 \eta_\sigma, \tag{F.122}$$

where π_0^{base} is impact inflation in the same tariff-risk environment without the date-0 volatility innovation. The restriction is

$$e_\pi' \mathbf{k}_0^v = 0. \tag{F.123}$$

It restricts a single third-order coefficient of the uncertainty response; it does not require total date-0 inflation to be zero. We call the resulting policy, the zero-impact-PPI benchmark of main-text Section 4.3, the *incremental date-0 PPI plan*: it sets this coefficient to zero and imposes all Ramsey first-order conditions at every history from date 1 onward.

Let ξ_0^v be the multiplier on (F.123). At third order, the date-0 boundary equations are

$$\mathcal{F}_0^v = 0, \qquad (\mathcal{S}_0^q)^v = 0, \qquad e_\pi' \mathbf{k}_0^v = 0, \qquad (\mathcal{S}_0^\pi)^v + \xi_0^v = 0, \qquad \nu_{-1}^v = 0. \tag{F.124}$$

The multiplier ξ_0^v enters only date-0 inflation stationarity; the new-commitment boundary $\nu_{-1}^v = 0$ is unchanged. At every history from date 1 onward, all three Ramsey conditions hold and the continuation inherits ν_0 from date 0, so commitment is not restarted.

The continuation is nevertheless not the one selected by unrestricted Ramsey: the date-0 restriction changes ν_0, and the inherited ν_0 changes the coefficient on the bounded history-dependent component of the solution.

Because the equations from date 1 onward and the risk-induced residual are unchanged, only that coefficient differs from full Ramsey:

$$\mathbf{k}_t^{H,v} = v_0 \left[p_v \rho_\sigma^t + c_H w_R \rho_R^t \right], \qquad c_H = -\frac{e_\pi' p_v}{e_\pi' w_R}, \qquad e_\pi' w_R \neq 0. \tag{F.125}$$

Writing $\mathbf{k}_t^{R,v}$ for the full Ramsey response in (F.80), the difference is

$$\mathbf{k}_t^{H,v} - \mathbf{k}_t^{R,v} = v_0(c_H - c_v)w_R\rho_R^t. \tag{F.126}$$

Thus the incremental date-0 PPI plan differs from full Ramsey only through the coefficient on the bounded history-dependent component $w_R\rho_R^t$.

To give the welfare statement an exact local meaning, suppose that for all sufficiently small δ there is a unique differentiable feasible policy family $H(\delta)$ that imposes the date-0 restriction corresponding to (F.123) while satisfying the exact nonlinear equilibrium conditions at every history. Let $H(0)$ coincide with the deterministic Ramsey allocation. Because the Ramsey feasible set contains both this family and the exact strict-PPI family, on these local branches

$$W^R \geq W^H, \qquad W^R \geq W^{PPI}. \tag{F.127}$$

These inequalities do not rank the incremental date-0 PPI plan against strict PPI, because the two restrictions are not nested.

Let δ scale the two shock loadings according to $\bar{\sigma}_\tau(\delta) = \delta\bar{\sigma}_\tau$ and $\eta_\sigma(\delta) = \delta\eta_\sigma$, while standardized innovations remain of order one. Let $z^j(\delta)$ denote the complete contingent allocation on the corresponding local policy branch. The incremental date-0 PPI plan changes only a third-order date-0 coefficient: its lower-order coefficients coincide with Ramsey at every history, and every history from date 1 onward satisfies the Ramsey first-order conditions. Under local differentiability and uniqueness, therefore,

$$\|z^H(\delta) - z^R(\delta)\|_\beta = O(\delta^3), \tag{F.128}$$

where $\|\cdot\|_\beta$ is a discounted sequence norm over the event tree. Because the first variation of the Ramsey Lagrangian is zero along feasible directions,

$$0 \leq W^R(\delta) - W^H(\delta) = O(\delta^6). \tag{F.129}$$

If the second variation is nondegenerate in this direction, the leading loss is of order δ^6; otherwise it is of still higher order. Thus full Ramsey and the incremental date-0 PPI plan have identical welfare through fourth order. This does not make the welfare effect of the uncertainty shock itself sixth order: a volatility innovation changes conditional second moments and can affect expected welfare at order δ^3. Strict PPI, by contrast, can generate a loss at order δ^2 (Appendix F.10).

F.10 Welfare accounting and approximation orders

This subsection derives the order of each policy's welfare loss relative to Ramsey, shows why third-order policy functions suffice for the fourth-order comparison, and describes the estimator behind Table F.1. The order statements follow from the Ramsey envelope condition. Let z^R denote the complete contingent Ramsey allocation and multipliers, and let Δz^j be the feasible displacement induced by policy j. A local expansion of the Ramsey Lagrangian gives

$$\mathcal{L}(z^R + \Delta z^j) - \mathcal{L}(z^R) = \underbrace{D\mathcal{L}(z^R)[\Delta z^j]}_{=0} + \frac{1}{2} D^2\mathcal{L}(z^R)[\Delta z^j, \Delta z^j] + O(\|\Delta z^j\|^3). \tag{F.130}$$

The first variation is zero because z^R satisfies the equilibrium constraints, the commitment boundary, and the planner stationarity conditions. Therefore a policy displacement of order δ^k can generate no welfare loss before order δ^{2k}; a nondegenerate second variation makes that bound exact. Strict PPI and the Taylor rule change allocations at first order on histories in which a tariff is realized, so their first possible welfare loss is second order. The Level-Ramsey experiment agrees with Ramsey at first order but can differ in second-order continuation curvature, so its first possible loss is fourth order. The incremental date-0 PPI plan changes only a third-order coefficient and retains the Ramsey first-order conditions from date 1 onward, so its first possible loss is sixth order. These orders refer to the full event tree, not only to the displayed history without realized tariff innovations.

We next show that third-order policy functions suffice for the fourth-order welfare comparison under the new date-0 commitment used here. Assume that the local allocation and multiplier branch and the utility, constraint, and shock-law integrands are five times continuously differentiable in the perturbation scale, with an integrable dominating function that permits differentiation under expectations. These conditions justify the $O(\delta^5)$ remainders below; under four continuous derivatives, the corresponding remainder is $o(\delta^4)$. Around the deterministic constrained steady state, write any feasible allocation and multiplier sequence as

$$z(\delta) = \bar{z} + \delta z_1 + \frac{\delta^2}{2} z_2 + \frac{\delta^3}{6} z_3 + \frac{\delta^4}{24} z_4 + O(\delta^5). \tag{F.131}$$

The Ramsey Lagrangian also depends directly on the perturbation scale through the shock laws. Define the extended argument

$$x(\delta) \equiv (z(\delta), \delta), \qquad x_1 = (z_1, 1), \qquad x_k = (z_k, 0), \quad k = 2, 3, 4. \tag{F.132}$$

Let all derivatives below be evaluated at $(\bar{z}, 0)$ and taken with respect to this extended

argument. Direct differentiation gives the fourth-order coefficient

$$[\delta^4]\mathcal{L}(x(\delta)) = \frac{1}{24}D\mathcal{L}[x_4] + \frac{1}{6}D^2\mathcal{L}[x_1, x_3] + \frac{1}{8}D^2\mathcal{L}[x_2, x_2] + \frac{1}{4}D^3\mathcal{L}[x_1, x_1, x_2] + \frac{1}{24}D^4\mathcal{L}[x_1, x_1, x_1, x_1]. \tag{F.133}$$

Under new commitment, the inherited multiplier is zero. The Lagrangian includes all date-0 and later equilibrium constraints, so deterministic Ramsey stationarity implies $D_z\mathcal{L}[\widetilde{z}] = 0$ for every admissible allocation direction $\widetilde{z}$. Since $x_4 = (z_4, 0)$, the first term in (F.133) vanishes. The unknown fourth-order policy coefficient therefore drops out, while x_1, x_2, x_3 and derivatives of the objective, constraints, and shock laws through order four identify the fourth-order welfare coefficient. Hence the calculation combines third-order policy functions with a fourth-order expansion of the Ramsey Lagrangian. With inherited multipliers, an additional lower-boundary term must be evaluated; Appendix H.6 explains why that case cannot use this shortcut without further work.

The same argument gives a sharper restriction for the incremental uncertainty effect under the Level-Ramsey experiment. Let $\Delta z_B^L = O(\delta^2)$ denote its baseline displacement from full Ramsey; the first-order tariff-level policies coincide, and the first difference is in second-order continuation curvature. Let $\Delta z_V^L = O(\delta^3)$ denote the difference between the two policies' incremental responses to the date-0 volatility innovation. Expanding the quadratic term in (F.130), the change in the policy loss caused by that innovation is

$$\begin{aligned}&\frac{1}{2}D^2\mathcal{L}[\Delta z_B^L + \Delta z_V^L, \Delta z_B^L + \Delta z_V^L] - \frac{1}{2}D^2\mathcal{L}[\Delta z_B^L, \Delta z_B^L] \\ &\quad = D^2\mathcal{L}[\Delta z_B^L, \Delta z_V^L] + \frac{1}{2}D^2\mathcal{L}[\Delta z_V^L, \Delta z_V^L] = O(\delta^5).\end{aligned} \tag{F.134}$$

Even without a date-0 volatility innovation, imposing the Level-Ramsey continuation can lower welfare relative to full Ramsey at fourth order because the two policies differ after future realized tariff-level innovations. The date-0 volatility innovation changes their welfare difference only at fifth order, so it contributes nothing to their fourth-order comparison. The numerical procedure imposes this analytical result rather than treating a near-zero estimate as evidence for it.

Equation (F.130) also explains why a common fourth-order calculation ranks strict PPI and the Taylor rule relative to Ramsey but cannot separate Ramsey from the incremental date-0 PPI plan. A nonzero fourth-order loss for that plan would indicate that total inflation, rather than its incremental uncertainty coefficient, had been constrained, or that approximation orders had been mixed.

To express losses in consumption units, write the lifetime utility loss of policy j relative to Ramsey, as in main-text equation (51), as

$$D_j(\delta) \equiv W^R(\delta) - W^j(\delta) = d_2^j\delta^2 + d_3^j\delta^3 + d_4^j\delta^4 + O(\delta^5). \tag{F.135}$$

Normalize deterministic steady-state consumption to one. The permanent consumption-equivalent loss λ_j is the proportional increase in that consumption stream whose utility value equals D_j:

$$\frac{1}{1-\beta}\frac{(1+\lambda_j)^{1-\sigma}-1}{1-\sigma} = D_j. \tag{F.136}$$

For $\sigma = 1$, the definition is

$$\frac{\log(1+\lambda_j)}{1-\beta} = D_j, \tag{F.137}$$

and the expansion below follows by continuity. Re-expanding this definition, rather than applying a nonlinear transformation to a truncated welfare number, gives

$$\begin{aligned}\lambda_j(\delta) = {} & (1-\beta)d_2^j\delta^2 + (1-\beta)d_3^j\delta^3 \\ & + \left[(1-\beta)d_4^j + \frac{\sigma}{2}(1-\beta)^2(d_2^j)^2\right]\delta^4 + O(\delta^5).\end{aligned} \tag{F.138}$$

For a *log*-consumption equivalent, the coefficient $\sigma/2$ would be replaced by $(\sigma-1)/2$; the tables report the level consumption equivalent in (F.138).

F.10.1 Numerical estimator for the fourth-order coefficients

We estimate the coefficients in (F.135) directly, rather than evaluating the unscaled third-order decision rule once and treating the resulting utility number as a fourth-order welfare approximation. Introduce a formal scale a and set

$$\bar{\sigma}_\tau(a) = a\bar{\sigma}_\tau, \qquad \eta_\sigma(a) = a\eta_\sigma. \tag{F.139}$$

Setting $a = 1$ recovers the calibrated shock loadings. The scale a implements the δ scaling above: standardized innovations stay fixed while both loadings change. Negative values of a are a formal device for separating odd from even coefficients, not negative standard deviations.

For each estimated policy and each scale, we solve the policy functions through third order in Dynare and reconstruct consumption, absorption, exports, Rotemberg resource use, output, and period utility from the exact nonlinear static maps. This direct-utility reconstruction is valid here for a reason specific to the complete-markets calibration. After substituting the

static allocation, the independent endogenous arguments of period utility are q and π_H. The steady-state matching condition in (F.15) implies $W_q(\bar{z}) = 0$, while the quadratic Rotemberg resource cost implies $W_{\pi_H}(\bar{z}) = 0$. Hence $DW(\bar{z})[z_4] = 0$ for every policy, and the unknown fourth-order allocation does not enter direct utility. Without these zero-gradient conditions, the Ramsey-Lagrangian envelope would justify the Ramsey coefficient but would not justify evaluating a fixed rule from only a third-order allocation.

All policies use the same covariance matrix for tariff and volatility innovations. Each of eight independent Matoušek–Affine–Owen scrambles supplies 128 Sobol normal vectors and their 128 exact negatives, for 256 antithetic paths per scramble. The same standardized innovations are used across policies and across the baseline and uncertainty states. The realized tariff paths can nevertheless differ across the two states because the date-0 volatility innovation changes the loading on later standardized tariff innovations. The welfare gap is formed path by path before averaging.

Each scramble is fitted separately. The reported coefficient is the mean of the eight replicate estimates, and its quasi-Monte Carlo standard error is their sample standard deviation divided by $\sqrt{8}$. The main horizon is $H = 160$ quarters. If $\bar{u}_{H,32}$ is mean flow utility over the final 32 quarters, the tail correction is $\beta^H \bar{u}_{H,32}/(1-\beta)$. The same paths are reevaluated with $H = 120$ and a 24-quarter tail. The parentheses in Table 2 and Table F.1 report these replicate standard errors.

Let $D_j^B(a)$ be the expected utility gap between Ramsey and policy j when there is no date-0 innovation but shocks can occur from date 1 onward. Let $D_j^V(a)$ be the additional gap caused by setting the date-0 volatility innovation to one while leaving the future draws unchanged. The signed-scale regressions impose the order restrictions

$$D_j^B(a) = b_{2j}a^2 + b_{4j}a^4 + b_{6j}a^6, \tag{F.140}$$

$$D_j^V(a) = w_{3j}a^3 + w_{4j}a^4 + w_{5j}a^5 + w_{6j}a^6. \tag{F.141}$$

Only orders two through four enter the reported welfare measure. Orders five and six are nuisance terms that prevent finite-scale curvature from loading onto the retained coefficients. For the Level-Ramsey experiment, (F.134) permits the sharper fits

$$D_L^B(a) = b_{4L}a^4 + b_{6L}a^6 + b_{8L}a^8, \qquad D_L^V(a) = w_{5L}a^5 + w_{6L}a^6 + w_{7L}a^7 + w_{8L}a^8. \tag{F.142}$$

We estimate b_{4L}, but the estimate fails the precision and horizon checks described below and is therefore not reported. The fourth-order coefficient of the loss added by the date-0 volatility innovation is set to its analytical value of zero, and the higher powers again serve only as finite-scale nuisance terms. The absolute scale grid is $\{0.025, 0.05, 0.075, 0.10, 0.125, 0.15\}$.

Each reported point estimate must pass: (i) a comparison of the six-point fit with the five smallest scales; (ii) a replicate-based simulation-error check for every retained coefficient; (iii) a relative fit-residual check; (iv) the even-symmetry restriction for baseline welfare; (v) the horizon-and-tail comparison; and (vi) a condition-number check for the scaled design matrix. The respective tolerances are 10 percent, 10 percent, 5 percent, 1 percent, 10 percent, and 10^6.

The subgrid statistic is the absolute difference between the six-scale and five-scale estimates divided by the larger absolute estimate. The fit statistic is the largest absolute polynomial residual divided by the largest absolute raw welfare gap. The symmetry statistic is the largest odd component of baseline welfare divided by its largest even component. The horizon statistic is the absolute 160-versus-120-quarter difference divided by the larger absolute estimate. These relative tests are meaningful only for a nonzero point estimate; they are never used to turn an unresolved near-zero coefficient into a reported loss.

Two further requirements concern policy definitions rather than statistical fit. First, the Level-Ramsey decision rule must use (F.113), face the same shock distribution as the other policies, and reproduce the analytical impact uncertainty response (Appendix F.11). Second, the incremental date-0 PPI plan is not estimated as a feedback rule: its fourth-order loss is zero by (F.129).

The retained coefficients deliver the main-text welfare ranking: strict PPI produces a small loss relative to Ramsey and the Taylor rule a much larger one. Both conclusions survive the scale-grid, simulation-error, symmetry, conditioning, and horizon checks. Because all policies face one common exogenous shock law, the comparison does not confound policy effects with different stochastic baselines.

Table F.1 separates each loss into its no-innovation baseline and the increment added by the date-0 volatility innovation. Strict PPI loses 3.954 basis points relative to Ramsey without the innovation and 4.072 basis points with it, so the innovation adds 0.118 basis points. The corresponding Taylor-rule losses are 322.759 and 331.168 basis points, an increment of 8.409 basis points. Most of each loss therefore reflects the policy's response to tariff risk that is present even without the date-0 innovation. The increment is the change in the Ramsey-relative gap, not the standalone welfare cost of tariff uncertainty under either policy.

For the Level-Ramsey experiment, the continuation restriction can lower welfare relative to Ramsey at fourth order on future histories with realized tariffs, but this gap is below the resolution of the simulation, and a point estimate would attach economic meaning to simulation error. The added uncertainty loss is zero through fourth order by (F.134).

Table F.1: Fourth-order complete-markets welfare losses relative to Ramsey: baseline and uncertainty increment

Policy design	Leading loss order	Baseline	Uncertainty state	Added uncertainty loss
Ramsey commitment	–	0.000	0.000	0.000
Zero-impact-PPI benchmark	Sixth	0.000†	0.000†	0.000†
Level-Ramsey benchmark	Fourth	Not resolved§	Not resolved§	0.000‡
Strict PPI	Second	3.954 (0.056)	4.072 (0.056)	0.118 (0.001)
Taylor rule	Second	322.759 (5.757)	331.168 (5.753)	8.409 (0.092)

Notes: Entries are fourth-order approximations to $10^4\lambda_j$, the permanent percentage increase in deterministic steady-state consumption, in basis points, that compensates for the Ramsey-relative utility gap; λ_j is defined in (F.136) and expanded in (F.138). "Baseline" sets the date-0 volatility innovation to zero while shocks occur from date 1 onward. "Uncertainty state" sets that innovation to one standard deviation and keeps the same future draws. "Added uncertainty loss" is the difference between the two. "Leading loss order" is the lowest power of the shock scale δ at which the loss can be nonzero. Parentheses are quasi-Monte Carlo standard errors across eight scrambled replicates. Ramsey is zero by construction. † The restricted plan differs from Ramsey at third order, so its loss is $O(\delta^6)$ by (F.129) and zero in the fourth-order approximation. ‡ The added loss is zero through fourth order by (F.134). § Not resolved: the estimate is below the resolution of the simulation and fails the precision and horizon checks, so no point estimate is reported.

F.11 Numerical verification and economic interpretation

We check the closed form in (F.80) against a separately coded third-order perturbation solution in Dynare, which is not used to construct the analytical coefficients. Over quarters 0–23, the largest absolute difference between the two paths for PPI inflation, output, the terms of trade, and the policy rate is 3.8×10^{-16}, and the largest residual in the finite-mode coefficient equations is 1.0×10^{-13}. Table F.2 reports the calibrated coefficients that enter Proposition 3. As a zero-variance check, setting the tariff-innovation variance to zero eliminates v_t, the variance-induced residual $g^v v_t$, and the uncertainty response. Together, these checks verify the analytical derivatives, the commitment boundary, and the propagation formula.

For the imposed Level-Ramsey continuation, the third-order numerical solution of (F.113) reproduces the analytically derived Level-Ramsey impact response at several shock scales after dividing by the cube of the scale. By contrast, replacing $y^e(\tau)$ with only its tangent gives the correct first-order tariff response but misses the third-order uncertainty response, so the agreement does not come merely from imposing the targeting rule's first-order coefficients.

In Figure F.1, the Ramsey formula and the numerical markers overlap. The Level-Ramsey path is close to Ramsey because it preserves the correct first-order response to realized tariffs, but it does not coincide with Ramsey because the uncertainty response also depends on the curvature of policy across future tariff states, which a targeting relation derived from the

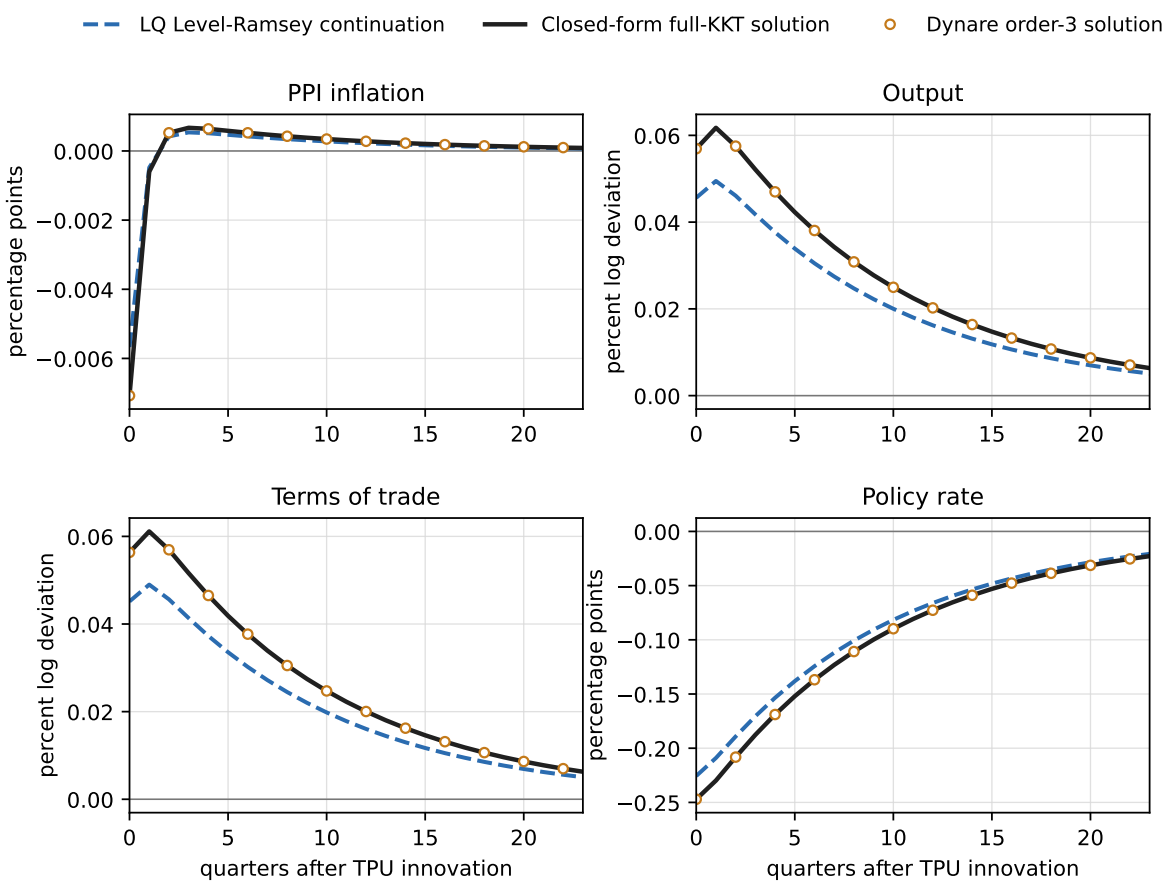


Figure F.1: The closed-form Ramsey response matches the numerical solution

Notes: All curves are local third-order responses to a one-standard-deviation date-0 log-volatility innovation, $\varepsilon_{\sigma,0} = 1$ with $\eta_\sigma = \log 2$, calibrated to double the exact conditional standard deviation of next quarter's tariff innovation from 0.10 to 0.20. Later realized tariff innovations are zero. The solid black line ("closed-form full-KKT solution") is the analytical Ramsey formula in (F.80); the hollow circles ("Dynare order-3 solution") are the separately computed third-order numerical solution described in the text. Their overlap verifies that both solve the same local system. The dashed blue line ("LQ Level-Ramsey continuation") shows the LQ-optimal date-0 response conditional on imposing the Level-Ramsey targeting relation (F.108) from date 1 onward. At date 0, commitment starts from the multiplier boundary $\phi_{-1} = 0$, which the difference form of the targeting criterion represents by $\widetilde{y}^a_{-1} = 0$. The continuation evaluates the gap around the cubic approximation to the efficient tariff allocation in (F.113); it is not reoptimized under the nonlinear Ramsey conditions. Real variables are percent log deviations; inflation and the quarterly policy rate are percentage-point deviations.

Table F.2: Calibrated coefficients of the closed-form Ramsey uncertainty response

Coefficient	Economic role	Calibrated value
ρ_R	Stable commitment root	0.184895201
ω_R	Inflation–terms-of-trade targeting coefficient	0.125572139
Γ_v	Impact loading from the variance shift to the terms of trade	0.081296531
$q_{1,0}$	Impact slope of the tariff-level Ramsey allocation	0.279303735
$\pi_{1,0}$	Impact inflation slope of the tariff-level allocation	0.027365043
$\pi_{2,0}$	Impact inflation curvature of the tariff-level allocation	−0.024353631
$g_{v,\text{mean}}$	Conditional-mean contribution to the risk loading	0.964403797
$g_{v,\text{cov}}$	Conditional-covariance contribution to the risk loading	0.880273908
g_v	Total risk loading	1.844677706

Notes: Coefficients are evaluated at the baseline calibration. ρ_R and ω_R are defined in main-text equation (42), and Γ_v in (F.87). The impact coefficients $q_{1,0}$, $\pi_{1,0}$, and $\pi_{2,0}$ of the tariff-level Ramsey allocation are defined before (F.75). The total risk loading g_v is the pricing-row entry of g^v in (F.74); its conditional-mean and conditional-covariance contributions are the entries of $D\mathcal{C}\,\ell_2$ and $D^2\mathcal{C}[\ell_1,\ell_1]$, respectively.

local linear-quadratic problem omits.

The experiment that sets only the incremental date-0 PPI response to zero is closer still: it changes one third-order boundary coefficient while leaving all continuation histories subject to the Ramsey conditions. Strict PPI at every history is a different policy. It also constrains the central bank after tariffs are realized and therefore changes allocations at lower orders. This distinction explains why the incremental date-0 PPI plan closely approximates the Ramsey uncertainty response even though strict PPI is not an exact solution to the full state-contingent Ramsey problem.

F.12 Why the baseline Taylor and Ramsey real paths nearly overlap

Main-text Figure 5 shows similar output and terms-of-trade responses under Ramsey policy and the Taylor rule. The closed forms show that this similarity is a feature of the baseline calibration, not a general implication of the model. Along the history with no realized tariff innovations, the two policy regimes imply

$$\Delta^T q_t = \eta_\sigma \chi_q \rho_\sigma^t + O(4), \tag{F.143}$$

$$\Delta^R q_t = v_0 \Gamma_v \frac{\rho_R^{t+1} - \rho_\sigma^{t+1}}{\rho_R - \rho_\sigma} + O(4), \tag{F.144}$$

$$\Delta^j y_t = \theta_q \Delta^j q_t + O(4), \qquad j \in \{T, R\}. \tag{F.145}$$

Equation (F.143) follows from Proposition 1; in particular,

$$\chi_q = \frac{(1-\beta\rho_\sigma)u^{UIP} - (\phi_\pi - \rho_\sigma)u^{NK}}{\det \mathbf{M}(\rho_\sigma)}. \tag{F.146}$$

Equation (F.144) restates (F.88). On impact the two responses differ ($v_0\Gamma_v = 0.056$ against $\eta_\sigma\chi_q = 0.070$ percent). Once the commitment term ρ_R^{t+1} has faded, (F.144) is approximately $v_0\Gamma_v\rho_\sigma(\rho_\sigma - \rho_R)^{-1}\rho_\sigma^t$, so similar real responses after the first quarter or two require the numerical relation $\eta_\sigma\chi_q \simeq v_0\Gamma_v\rho_\sigma/(\rho_\sigma - \rho_R)$. No equilibrium condition imposes this relation: χ_q is determined by the Taylor rule and the two uncertainty wedges, whereas Γ_v is determined by the planner's first-order conditions and the date-0 commitment boundary. At later dates, the Ramsey path also contains the stable commitment component ρ_R^t. This component fades quickly in the baseline ($\rho_R \approx 0.18$; Table F.2), leaving both responses to inherit much of the persistence of the volatility state ($\rho_\sigma = 0.9$). These two numerical features account for the apparent overlap.

Figure F.2 changes one parameter at a time and recomputes the third-order solution under both policies. The middle row sets $\rho_\sigma = 0.76$. Cho, Han, Oh and Rogantini Picco (2021) use this value for the persistence of productivity uncertainty; we use it as a literature-based alternative calibration for the persistence of the tariff-volatility state, not as an estimate of tariff volatility. This change leaves the slopes and curvatures of the allocation with respect to realized tariffs, and hence the NKPC and UIP wedges, unchanged; it changes their relative weights $1 - \beta\rho_\sigma$ and $\phi_\pi - \rho_\sigma$ in (F.146). The numerator of (F.146), which at the baseline is a partial cancellation between the two weighted wedges, then changes sign. The Taylor-rule terms-of-trade loading χ_q therefore changes sign, and Taylor and Ramsey predict real responses of opposite signs. The bottom row sets the Rotemberg adjustment-cost parameter to $\kappa = 80$. This value implies an NKPC slope of 0.1125 and, under the slope matching in Section J, a Calvo non-reset probability of about 0.72. Greater price rigidity flattens the NKPC and raises the welfare cost of inflation. It reduces the Taylor-rule terms-of-trade loading, whereas Ramsey, facing costlier inflation, shifts more of the adjustment to real variables, whose response becomes larger and more gradual. The peak Ramsey response is more than four times the Taylor-rule peak in this experiment.

Both policies are simulated in the same third-order Dynare model (with pruning) as the baseline comparison, and every comparison begins from the same deterministic steady-state values of all predetermined physical variables. Under every displayed calibration, (F.143)–(F.144) reproduce the numerical paths to machine precision. At the baseline, the uncertainty response under each policy is zero at first and second order, and removing tariff risk eliminates the third-order Ramsey response. The divergence in Figure F.2 therefore

reflects policy transmission, not a change in shock normalization, initial conditions, or approximation method: the baseline similarity of output and the terms of trade is not a structural equivalence between Ramsey policy and the Taylor rule.

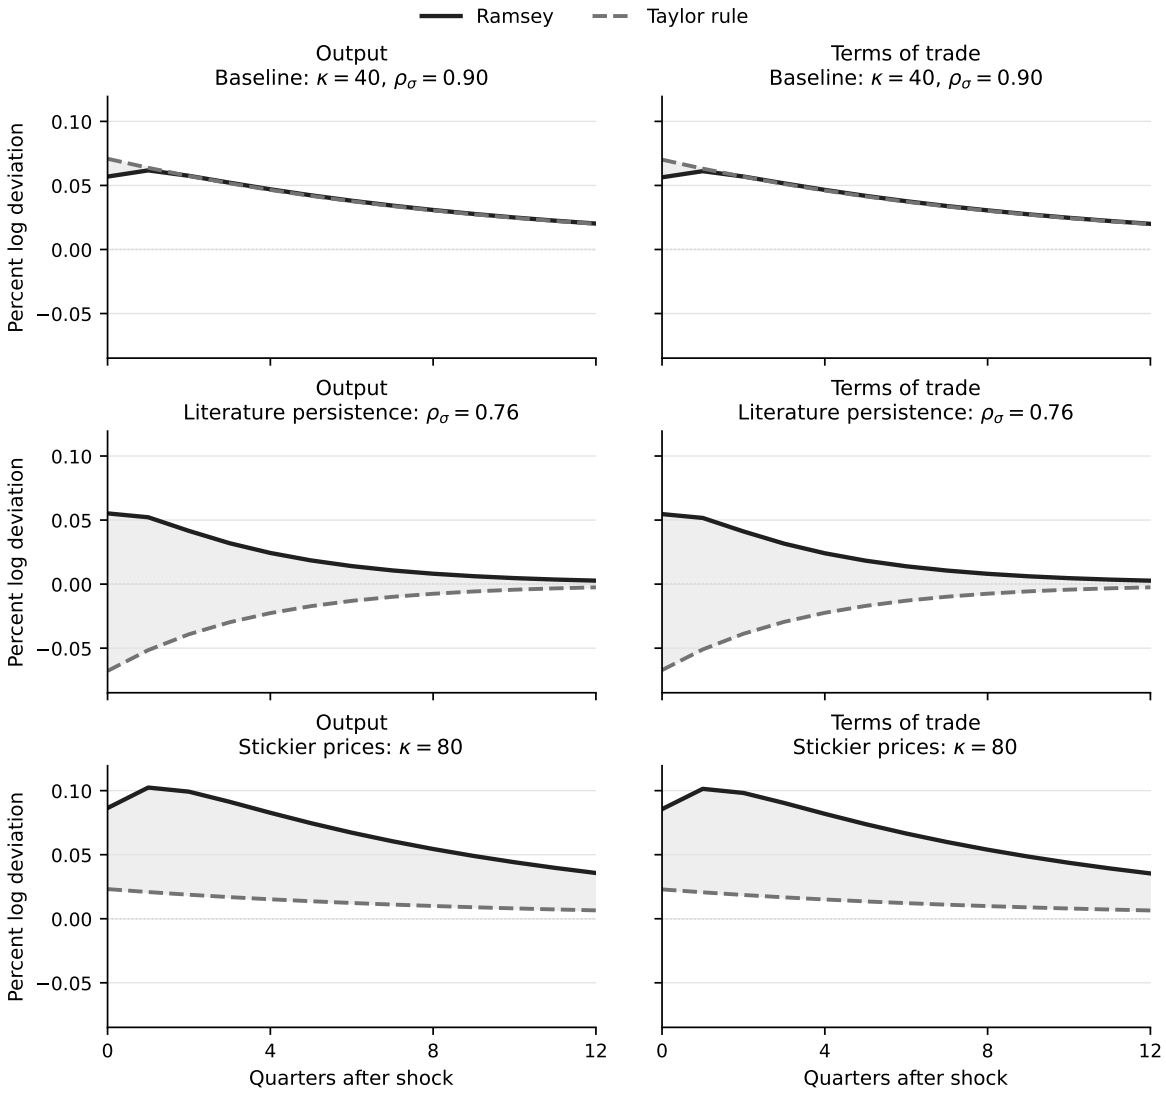


Figure F.2: Taylor and Ramsey real responses under alternative calibrations

Notes: The top row reproduces the calibration in main-text Figure 5. The middle row changes volatility persistence to $\rho_\sigma = 0.76$; the bottom row changes the Rotemberg adjustment-cost parameter to $\kappa = 80$. All other parameters, the date-0 volatility innovation, and the deterministic steady-state values of all predetermined physical variables are held fixed. Ramsey begins a new commitment at date 0; the Taylor-rule coefficients remain $(\phi_\pi, \phi_y) = (1.5, 0.5)$. Solid black lines denote Ramsey policy and dashed gray lines the Taylor rule; shading marks the distance between them. Output and the terms of trade are percent log deviations and share one vertical scale across all six panels. A rise in the terms of trade is a Home real depreciation. Each curve subtracts the no-innovation path under the same policy. Later tariff innovations are zero, while agents continue to account for future tariff risk.

G Wedge Decomposition, Special Cases, and Comparative Statics

This appendix supplements the closed-form analysis of Section 3. It shows why uncertainty responses are third order (Section G.1), decomposes the NKPC and UIP wedges (Section G.2), reports special cases of the closed-form solution (Section G.3) and one-dimensional comparative statics of impact responses (Section G.4), and checks the analytical recovery formulas for selected variables against the numerical solution (Section G.5).

G.1 Why uncertainty responses are third order

Under the volatility experiment in Definition 1, the tariff *level* stays at its steady state, so the shock acts only through time-varying conditional variances and the associated Jensen corrections. In a perturbation around the deterministic steady state, impulse responses to the volatility innovation are therefore zero at first *and* second order (second-order terms capture stochastic-mean shifts from baseline risk), and the first nonzero responses appear at third order through terms proportional to $\bar{\sigma}_\tau^2 \sigma_{T,t}$. Section D.4.1 gives the order-counting argument, and Figure G.1 illustrates it.

G.2 Decomposition of the uncertainty wedges

The NKPC (pricing) wedge $u^{NK} = \beta \bar{\sigma}_\tau^2 (3A_\pi^2 + 2A_\pi s_1 + B_\pi)$ in (32) combines three terms: a nonnegative Rotemberg adjustment-cost term; a covariance term, which arises because high-tariff states are low-sales states, so pricing mistakes there matter less for profits; and the curvature of the inflation response to the tariff. The UIP (asset-market) wedge $u^{UIP} = \bar{\sigma}_\tau^2[(B_\pi + B_q) - (A_\pi + A_q)^2]$ combines a nonpositive variance term, the Jensen correction to the Foreign-bond payoff measured in Home consumption units, with the curvature of the depreciation response to the tariff. Main-text Figure 3 plots these components as openness and the two Taylor-rule coefficients vary. Both wedges are negative at the baseline and over almost all of the displayed ranges; the NKPC wedge becomes slightly positive only when output feedback is close to zero.

G.3 Special cases of the closed-form perturbation solution

The baseline mechanism combines the NKPC and UIP wedges, a policy rule, and persistent dynamics. Figure G.2 separates these ingredients: the pricing-channel case sets $u^{UIP} = 0$, the

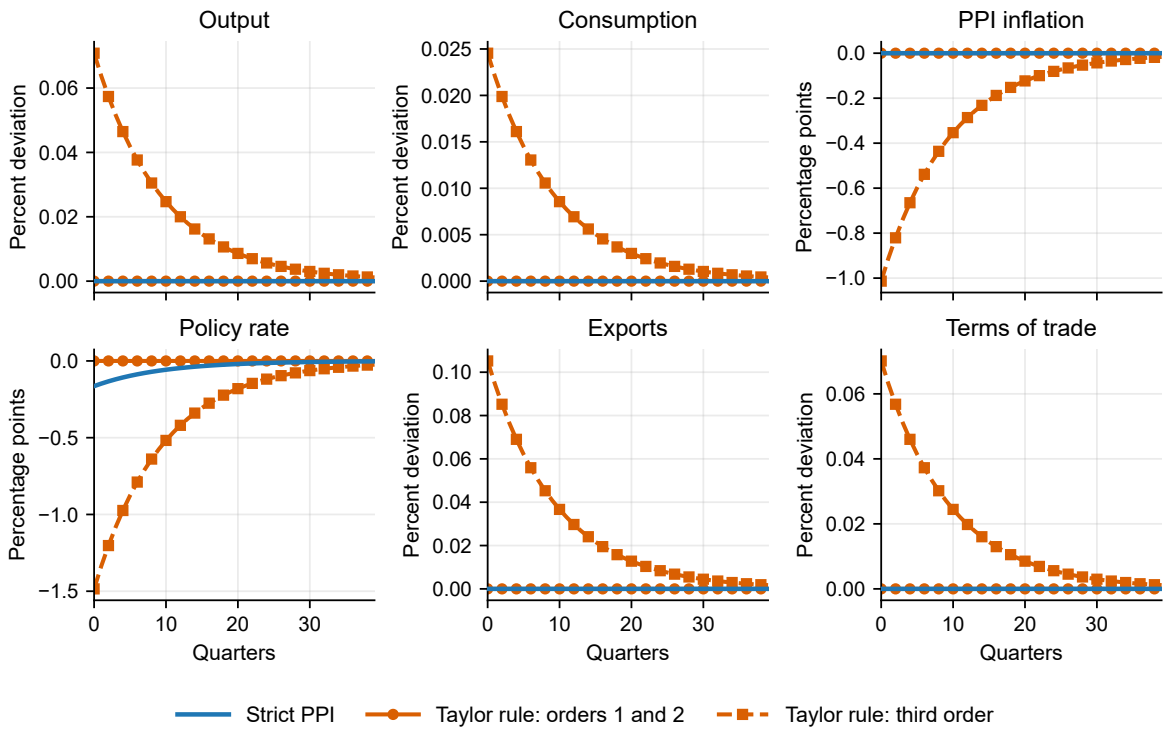


Figure G.1: Responses to a tariff uncertainty shock by approximation order

Notes: The experiment is the one-standard-deviation date-0 log-volatility innovation in Definition 1, with $\eta_\sigma = \log 2$, all realized tariff innovations set to zero, and all later realized volatility innovations set to zero. Future tariff and volatility risk remains in agents' conditional expectations. A volatility innovation affects equilibrium through conditional-variance (Jensen) terms and therefore has no first- or second-order response. The Taylor-rule line labelled "orders 1 and 2" is zero at both orders; the nonzero response first appears at third order. Under strict PPI in the complete-markets economy, output, consumption, exports, the terms of trade, and PPI inflation have zero responses on this no-realization history. The implementing policy rate need not be zero: the Euler equation can require it to move when conditional tariff risk changes.

asset-market-channel case sets $u^{NK} = 0$, the no-output-feedback case sets $\phi_y = 0$, and the remaining cases vary volatility persistence and the Rotemberg adjustment-cost parameter κ.

When the volatility state has no persistence ($\rho_\sigma = 0$), agents know that the tariff variance returns to its unconditional level after the impact period, so $\Delta\pi_t = \Delta q_t = 0$ for $t \geq 1$ (Corollary D.1) and every variable that is a contemporaneous function of (π_t, q_t) responds only on impact. Persistence is also quantitatively decisive on impact, because it changes how the two wedges offset: with $\rho_\sigma = 0$ their weights in the numerator of χ_q become 1 and ϕ_π, rather than $1 - \beta\rho_\sigma$ and $\phi_\pi - \rho_\sigma$ in (F.146).

Figure G.2 also reports lower-rigidity variants of the single-channel cases in Corollary D.2. More flexible prices strongly attenuate the pricing-channel real and inflation responses. They also move the asset-market-channel terms-of-trade and real responses toward zero while leaving a nominal response, as implied by the flexible-price limit in the same corollary. Changing κ moves both the wedges and, through $\tilde{\kappa}$, the propagation matrix.

The special cases yield four lessons. First, at the baseline wedge signs, the pricing channel alone produces a real depreciation and an output expansion, whereas the asset-market channel alone produces a real appreciation and a contraction. Second, more flexible prices attenuate the real effects of either channel rather than reversing their signs. Third, persistence turns an impact disturbance into a sustained path: when $\rho_\sigma = 0$, the responses end after the impact period. Fourth, under the baseline calibration, removing output feedback reverses the sign of the real response, producing a real appreciation and an output contraction. These implications follow directly from Corollary D.2, Corollary D.1, and the $\phi_y = 0$ case in Section D.5.

G.4 Comparative statics of the uncertainty responses

Figure G.3 uses the closed-form third-order solution in Proposition 1 to trace how impact responses vary with one primitive parameter at a time. A parameter moves each response through two objects: the NKPC and UIP wedges and the propagation matrix $\mathbf{M}(\rho_\sigma)^{-1}$. The figure complements the wedge decomposition in main-text Figure 3.

Three patterns stand out. First, openness and the trade elasticity (α, γ) amplify both the output and the inflation responses because they strengthen exposure to tariff risk. Second, the policy coefficients (ϕ_π, ϕ_y) can amplify the response or reverse the sign of output by changing both wedges and the propagation matrix; responses are especially large when $\mathbf{M}(\rho_\sigma)$ is close to singular. Third, nominal rigidity κ and labor-supply curvature φ alter the relative strength of the pricing and asset-market channels, producing hump-shaped or sign-changing output responses even when inflation remains deflationary. The baseline combination of higher output and lower inflation is therefore a property of the baseline calibration, not a

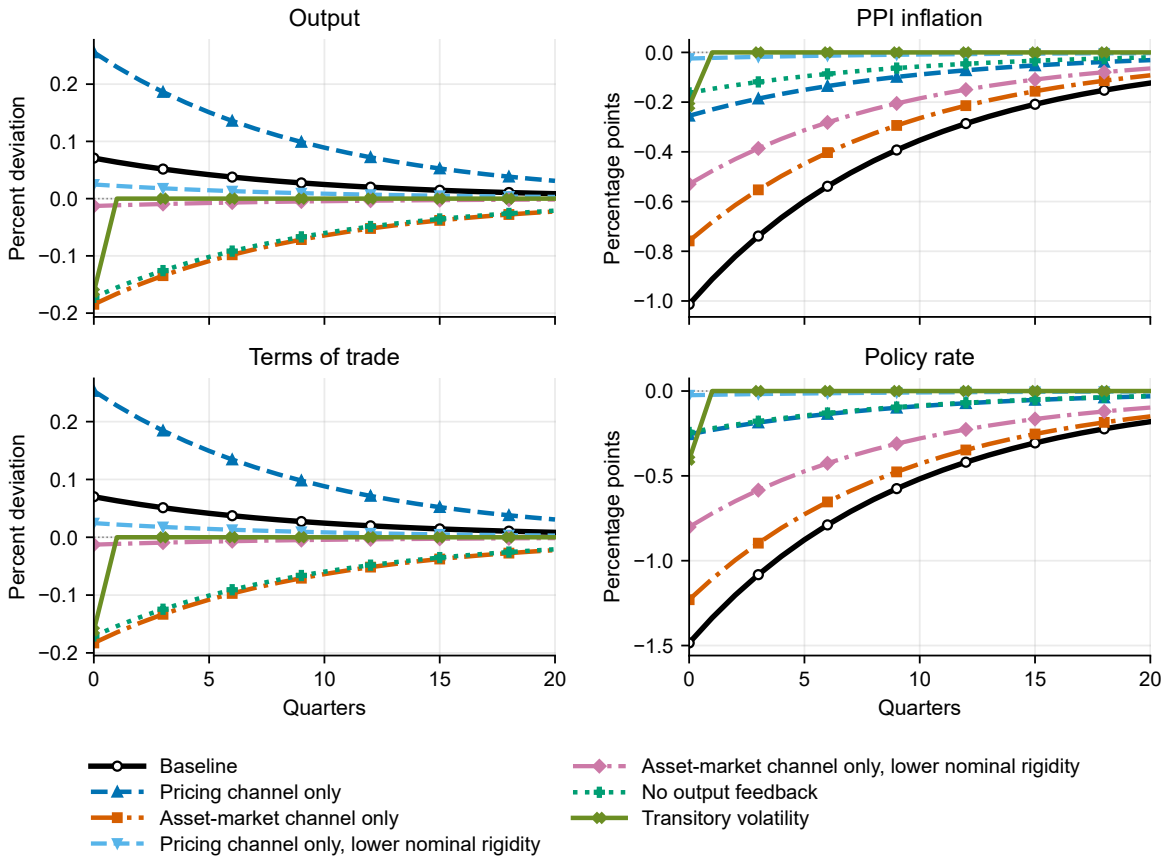


Figure G.2: Impulse responses under special cases

Notes: Closed-form third-order responses to the tariff-volatility innovation in Definition 1. The legend identifies the baseline and each restriction defined in the text. The lower-rigidity cases reduce the Rotemberg adjustment-cost parameter from its baseline value. Real variables and the terms of trade are percent deviations from steady state; inflation and the policy rate are percentage-point deviations.

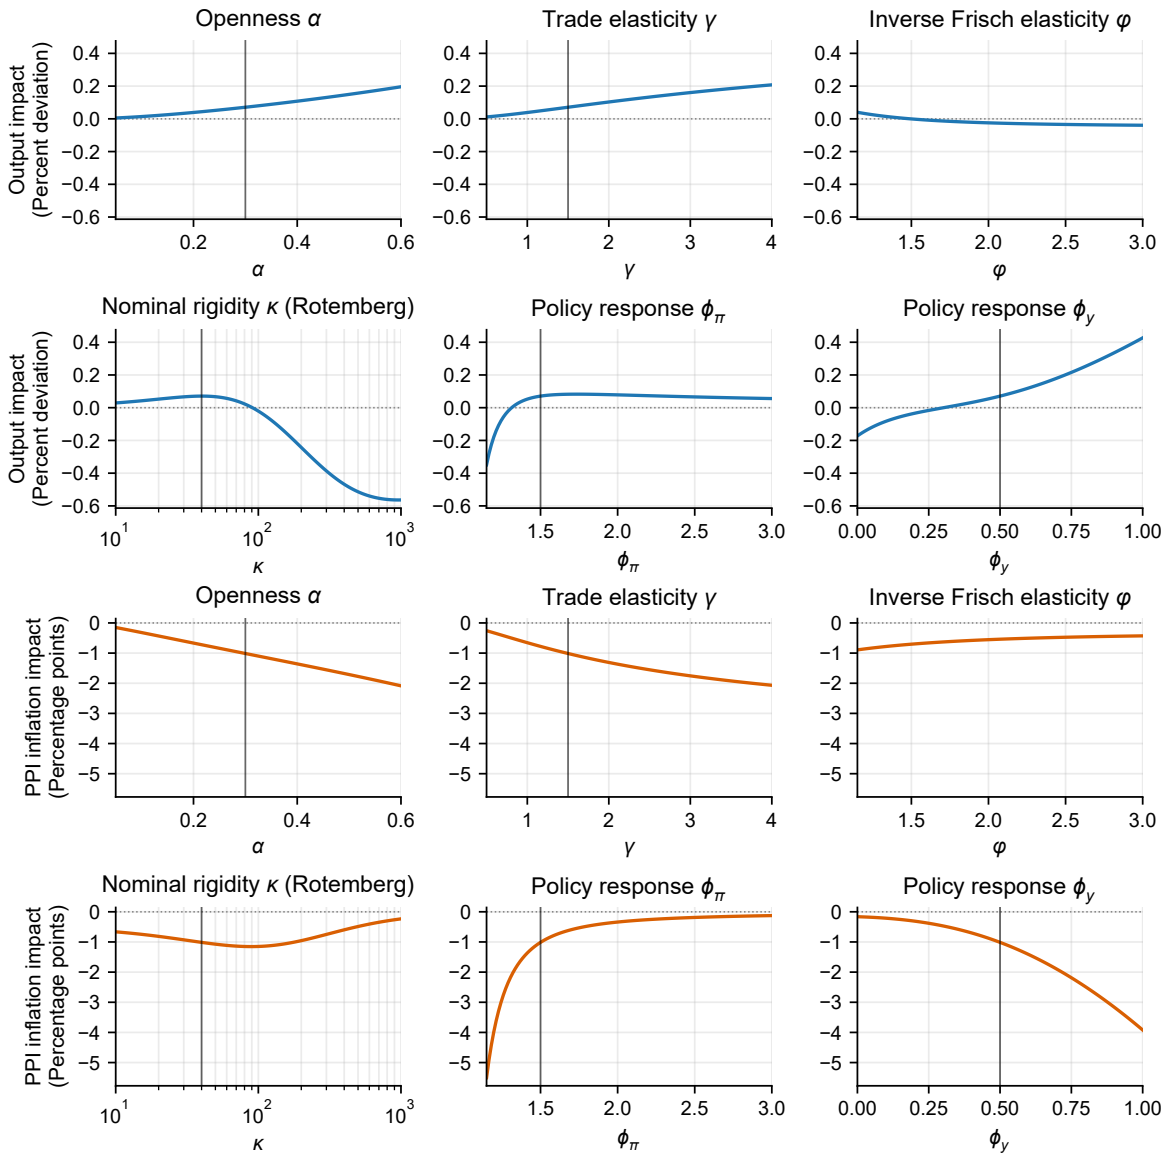


Figure G.3: Comparative statics of impact responses to a tariff uncertainty shock

Notes: Each panel reports the impact response to the one-standard-deviation date-0 log-volatility innovation in Definition 1, with $\eta_\sigma = \log 2$, while varying one parameter around the baseline in Table B.1, marked by the vertical line. The top two rows show output in percent; the bottom two show PPI inflation in percentage points. Responses use the closed-form third-order solution in Proposition 1; the κ axis is logarithmic.

sign restriction imposed by the model.

G.5 Selected equilibrium responses and numerical check

Figure G.4 reports six responses that summarize the real, nominal, and exchange-rate adjustment under the PPI-based Taylor rule and complete markets. Their close agreement with the separately computed third-order numerical solution checks both the closed-form (π_t, q_t) response and the static recovery formulas used for the remaining variables.

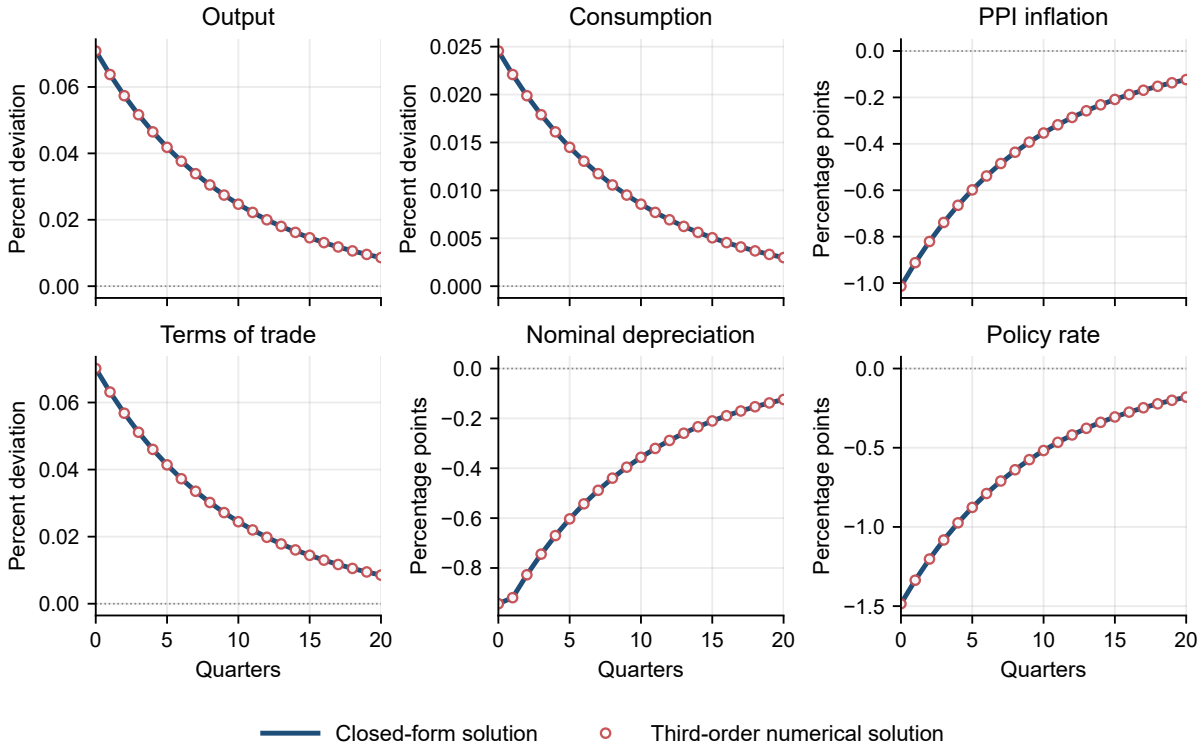


Figure G.4: Selected complete-markets responses under the Taylor rule

Notes: Third-order responses to the one-standard-deviation date-0 log-volatility innovation in Definition 1, with $\eta_\sigma = \log 2$, under the PPI Taylor rule. Later realized tariff innovations are zero. Solid lines are closed-form responses; circles are the separately computed third-order numerical solution. Output, consumption, and the terms of trade are percent deviations; PPI inflation, nominal depreciation, and the policy rate are percentage-point deviations. The nominal depreciation rate is $\Delta e_t = (q_t - q_{t-1}) + \pi_{H,t}$.

Output and consumption rise as the NKPC wedge induces monetary easing and a real depreciation, while PPI inflation and the policy rate fall. The terms of trade rise (a Home real depreciation), yet nominal depreciation is negative on impact. The two signs are consistent: with Foreign inflation fixed, $\Delta e_t = \Delta q_t + \pi_{H,t}$, and Home-price deflation is initially strong enough to produce a nominal appreciation despite the real depreciation. The nominal exchange rate therefore cannot be inferred from the terms of trade alone.

Part III

Extensions and Robustness

The following appendices ask how the policy conclusions change when the baseline complete-markets, producer-currency-pricing environment is relaxed. We first add external saving through a Foreign bond and then reverse the tariff incidence from exports to imports. Two further extensions replace Rotemberg pricing with Calvo adjustment and producer-currency pricing with local-currency pricing. In each case, the central question is whether strict PPI remains a useful approximation to Ramsey and why any departure arises. Appendix L then collects the answers across all computed environments.

H Incomplete Markets: Exact Closure and Optimal Policy

This appendix derives the exact bond-economy closure and the third-order Ramsey response to an export-tariff uncertainty shock in the small-open-economy (SOE) model of Section 5. Three blocks play distinct roles: the domestic Euler equation only recovers the implementing policy rate; the reduced Foreign Euler equation and the net-foreign-asset (NFA) law restrict the real allocation; and the planner-stationarity conditions select the Ramsey allocation among implementable paths. Section H.5 derives the response in closed form and verifies it numerically, and Section H.6 states what a welfare ranking would additionally require.

H.1 Experiment, approximation order, and normalization

Let $\tau_t \equiv \log(T_t^*/\bar{T}^*)$ be the export-tariff state and let s_t be its log volatility (denoted $\sigma_{T,t}$ in the main text). The timing is

$$\tau_{t+1} = \rho_\tau \tau_t + \bar{\sigma}_\tau e^{s_t} \varepsilon_{\tau,t+1}, \qquad \mathbb{E}_t \varepsilon_{\tau,t+1} = 0, \qquad \mathbb{E}_t \varepsilon^2_{\tau,t+1} = 1, \tag{H.1}$$

$$s_t = \rho_\sigma s_{t-1} + \eta_\sigma \varepsilon_{\sigma,t}. \tag{H.2}$$

As in Assumption 1, tariff innovations are independent of volatility innovations and conditionally symmetric; symmetry is used below when tariff histories are integrated out. The pure-uncertainty experiment sets $\varepsilon_{\sigma,0} = 1$, $\eta_\sigma = \log 2$, and the date-zero tariff innovation to zero, so the conditional standard deviation of next quarter's tariff innovation rises from 0.10 to 0.20. Along the displayed path all later realized tariff innovations are zero; we call this the

no-realization path. Responses start from the common deterministic physical state and are measured relative to the path from that state without the volatility innovation; future tariff risk remains in agents' expectations along both paths.

For order accounting, multiply both the tariff innovation scale and the volatility-state deviation by a common perturbation parameter ϑ (the main-text δ; the $\delta \equiv 1-\alpha$ defined below and the risk-sharing constant ϑ in (B.23) are unrelated). The baseline-subtracted half-variance state is

$$v_t(\vartheta) \equiv \frac{1}{2}\vartheta^2\bar{\sigma}_\tau^2\left(e^{2\vartheta s_t}-1\right) = \vartheta^3\bar{\sigma}_\tau^2 s_t + \mathcal{O}(\vartheta^4), \qquad \mathbb{E}_t v_{t+1} = \rho_\sigma v_t + \mathcal{O}(\vartheta^4). \tag{H.3}$$

This order statement concerns conditional moments and the no-realization path; it does not imply that the two full adapted tariff processes coincide through second order. To see the distinction, let u_{t+1}^ι denote the tariff innovation in economy $\iota \in \{0,1\}$. Coupling the two economies with the same standardized shock gives

$$\begin{aligned} u_{t+1}^1 - u_{t+1}^0 &= \vartheta\bar{\sigma}_\tau\left(e^{\vartheta s_t}-1\right)\varepsilon_{\tau,t+1} \\ &= \vartheta^2\bar{\sigma}_\tau s_t\varepsilon_{\tau,t+1} + \frac{1}{2}\vartheta^3\bar{\sigma}_\tau s_t^2\varepsilon_{\tau,t+1} + \mathcal{O}(\vartheta^4), \\ \tau_{t+1}^1 - \tau_{t+1}^0 &= \rho_\tau(\tau_t^1-\tau_t^0) + (u_{t+1}^1 - u_{t+1}^0). \end{aligned} \tag{H.4}$$

Let $\Delta u_{t+1}^{[2]}$ and $\Delta\tau_t^{[2]}$ denote the coefficients on ϑ^2 in the two-economy differences. Then

$$\mathbb{E}_t\Delta u_{t+1}^{[2]} = 0, \qquad \mathbb{E}_t\Delta\tau_{t+1}^{[2]} = \rho_\tau\Delta\tau_t^{[2]}. \tag{H.5}$$

After an off-path tariff realization the inherited difference $\Delta\tau_t^{[2]}$ is known at date t, so the tariff-state difference need not have zero conditional mean. It has zero mean at the common initial node and after averaging tariff histories conditional only on the volatility history, and $\Delta\tau_t^{[2]} = 0$ at every date on the no-realization path. The order-two innovation difference nevertheless changes the conditional half variance at order three, as in (H.3). Hence the baseline-subtracted response shown in the figures is third order, even though off-path realized allocations can differ at second order. We set $\vartheta = 1$ after collecting coefficients. Because $\eta_\sigma = \log 2$ is large for a local expansion, the paths are third-order approximations: the finite variance doubling fixes the shock normalization but does not enlarge the approximation's local domain.

The policy object is the continuation of a Ramsey plan chosen before date 0. At date 0 the planner inherits both the deterministic physical state and the lagged multipliers of that earlier plan, so promises made before the shock are kept. A new date-0 commitment

would instead reset those multipliers to zero and would generally select a different impact response. We call the solution with inherited multipliers the *inherited Ramsey continuation*, or simply the Ramsey continuation. It should not be confused with the complete-markets Level-Ramsey continuation of Appendix F.8. Strict PPI means $\Pi_{H,t} = 1$ after every history under the same physical initial state and the same fixed production correction.

H.2 Exact private equilibrium and bond closure

Let $G_t \equiv P_t/P_{H,t}$ and let a rise in Q_t denote a Home real depreciation. With the import tariff fixed at its normalized value, the static block is

$$G_t^{1-\gamma} = (1-\alpha) + \alpha Q_t^{1-\gamma}, \tag{H.6}$$

$$C_{H,t} = (1-\alpha) G_t^{\gamma} C_t, \qquad X_t = \alpha \left(\frac{Q_t}{T_t^*}\right)^{\gamma}, \tag{H.7}$$

$$Y_t = \frac{C_{H,t} + X_t}{1 - \frac{\kappa}{2}(\Pi_{H,t} - 1)^2}, \qquad tb_t = \frac{C_{H,t} + X_t}{G_t} - C_t, \tag{H.8}$$

$$\Pi_{C,t} = \Pi_{H,t} \frac{G_t}{G_{t-1}}. \tag{H.9}$$

Raw and normalized real marginal cost are, respectively,

$$m_t^r \equiv \chi G_t C_t^{\sigma} Y_t^{\varphi} A_t^{-(1+\varphi)}, \qquad \bar{\upsilon} \equiv \chi, \qquad mc_t \equiv \frac{m_t^r}{\bar{\upsilon}}. \tag{H.10}$$

The fixed production correction satisfies

$$1 - s^p = \frac{\epsilon - 1}{\epsilon \bar{\upsilon}}. \tag{H.11}$$

Fixed at its no-tariff value, the correction normalizes deterministic steady-state marginal cost to one in the NKPC,

$$\begin{aligned} \kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) \\ + \beta\kappa \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \left(\frac{C_t}{C_{t+1}}\right)^{\sigma} \frac{G_t}{G_{t+1}} \right]. \end{aligned} \tag{H.12}$$

Because the correction is fixed across states, the flexible-price allocation need not coincide with the constrained-efficient allocation in every tariff state.

Households in the SOE trade a domestic nominal bond B_t in zero net supply and a Foreign-currency nominal bond B_t^*. Before market clearing, the household's nominal budget

constraint is

$$\begin{aligned} P_t C_t + B_t + \mathcal{E}_t B_t^* = W_t L_t + R_{t-1} B_{t-1} + \mathcal{E}_t R_{t-1}^* \Theta_{t-1} B_{t-1}^* \\ + \mathcal{P}_t + T_t^{gov}, \end{aligned} \tag{H.13}$$

where $\mathcal{P}_t$ denotes profits, T_t^{gov} collects net lump-sum transfers, and

$$\Theta_t \equiv e^{-\psi_b b_t}. \tag{H.14}$$

The debt-elastic premium makes aggregate NFA mean revert. Each atomistic household takes aggregate NFA b_t, and therefore Θ_t, as given when it chooses its own bond holding. The planner later internalizes how aggregate NFA enters the equilibrium restrictions. If Λ_t^P is the multiplier on (H.13), differentiating with respect to consumption, the Home bond, and the Foreign bond gives

$$\begin{aligned} 0 &= C_t^{-\sigma} - \Lambda_t^P P_t, & \Lambda_t^P &= \frac{C_t^{-\sigma}}{P_t}, \\ 0 &= -\Lambda_t^P + \beta R_t \mathbb{E}_t \Lambda_{t+1}^P, & \Lambda_t^P &= \beta R_t \mathbb{E}_t \Lambda_{t+1}^P, \\ 0 &= -\Lambda_t^P \mathcal{E}_t + \beta R_t^* \Theta_t \mathbb{E}_t [\Lambda_{t+1}^P \mathcal{E}_{t+1}], & & \\ & & \Lambda_t^P \mathcal{E}_t &= \beta R_t^* \Theta_t \mathbb{E}_t [\Lambda_{t+1}^P \mathcal{E}_{t+1}]. \end{aligned} \tag{H.15}$$

Because the household takes Θ_t as given, differentiating its own Foreign-bond position introduces no factor $1 - \psi_b b_t$. Normalizing real quantities by steady-state output, define aggregate NFA as

$$b_t \equiv \frac{\mathcal{E}_t B_t^*}{P_t \bar{Y}}, \qquad \bar{Y} = 1, \tag{H.16}$$

so $b_t > 0$ denotes net assets; external debt is $D_t \equiv -b_t$. The effective Foreign gross return is $R_t^* e^{-\psi_b b_t}$, where $\psi_b > 0$ raises the borrowing cost when NFA falls.

Proposition H.1 (Exact bond-economy closure)**.** *Under incomplete international asset markets, the exact asset-market block is*

$$1 = \beta R_t \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \right], \tag{H.17}$$

$$1 = \beta R_t^* e^{-\psi_b b_t} \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{C,t+1}^{-1} \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t} \right], \tag{H.18}$$

$$b_t = R_{t-1}^* e^{-\psi_b b_{t-1}} \frac{\mathcal{E}_t}{\mathcal{E}_{t-1}} \frac{b_{t-1}}{\Pi_{C,t}} + tb_t. \tag{H.19}$$

Using $\mathcal{E}_{t+1}/\mathcal{E}_t = (Q_{t+1}/Q_t)\Pi_{H,t+1}$ *and* (H.9)*, the two allocation restrictions are equivalently*

$$1 = \beta R_t^* e^{-\psi_b b_t} \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{Q_{t+1}}{Q_t} \frac{G_t}{G_{t+1}} \right], \tag{H.20}$$

$$b_t = R_{t-1}^* e^{-\psi_b b_{t-1}} \frac{Q_t}{Q_{t-1}} \frac{G_{t-1}}{G_t} b_{t-1} + tb_t. \tag{H.21}$$

Conditional on an admissible real allocation, the domestic Euler equation recovers the implementing nominal rate from

$$R_t^{-1} = \beta \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \Pi_{H,t+1}^{-1} \frac{G_t}{G_{t+1}} \right]. \tag{H.22}$$

There is no separate exact nonlinear UIP equation.

Proof. Substituting $\Lambda_t^P = C_t^{-\sigma}/P_t$ into the two bond conditions in (H.15) gives (H.17)–(H.18). In particular, $P_t/P_{t+1} = \Pi_{C,t+1}^{-1}$ converts each nominal payoff into Home consumption units. Aggregate the household budget, impose zero net supply of the Home bond, and use goods-market clearing. Dividing the inherited Foreign payoff by $P_t \bar{Y}$ gives

$$R_{t-1}^* e^{-\psi_b b_{t-1}} \frac{\mathcal{E}_t}{\mathcal{E}_{t-1}} \frac{P_{t-1}}{P_t} b_{t-1},$$

while the remaining net export value is tb_t. This proves (H.19). The two price identities then give (H.20)–(H.21). Finally, substituting CPI inflation into the domestic Euler equation gives (H.22). A fall in b_t raises the effective Foreign return in the household Euler equation; the resulting movements in q_t and c_t change the trade balance, and the Foreign Euler equation and NFA law jointly generate mean reversion. Around zero NFA, the premium does not enter the first-order inherited-payoff term directly because its derivative is multiplied by steady-state NFA, which is zero. □

H.3 Third-order risk terms in private and planner conditions

This subsection derives the third-order risk terms in the private equilibrium rows and in the planner-stationarity rows. The domestic Euler equation isolates the conditional-variance component of the implementing rate. Let lower-case letters denote log deviations and define

$$X_{t+1}^H \equiv -\sigma(c_{t+1} - c_t) + r_t - \pi_{C,t+1}. \tag{H.23}$$

Equation (H.17) implies $0 = \log \mathbb{E}_t[e^{X^H_{t+1}}]$. A third-order cumulant expansion and subtraction of the same risky baseline give

$$c^v_t = \mathbb{E}_t c^v_{t+1} - \frac{1}{\sigma}\left(r^v_t - \mathbb{E}_t \pi^v_{C,t+1}\right) - \frac{1}{\sigma}\Delta^E_t + \mathcal{O}(4), \tag{H.24}$$

$$\Delta^E_t = \frac{1}{2}(\sigma\xi_c + \xi_{\pi_C})^2 \bar{\sigma}^2_\tau (e^{2s_t} - 1) + \mathcal{O}(4) = (\sigma\xi_c + \xi_{\pi_C})^2 v_t + \mathcal{O}(4), \tag{H.25}$$

where ξ_c and ξ_{π_C} are the first-order loadings on next period's tariff innovation. Equation (H.25) isolates only the Jensen (conditional-variance) contribution; the complete order-three rate recovery, (H.38) below, also includes the change in the conditional mean. After a positive volatility innovation the Jensen term is locally nonnegative and, holding the expected future allocation and real rate fixed, lowers current consumption. With an unconstrained nominal rate, however, policy can offset this term, so its presence alone does not imply that tariff uncertainty moves the real allocation.

Let (ℓ_t, μ_t, ν_t) denote the Ramsey multipliers on the normalized NKPC, reduced Foreign Euler equation, and NFA law, respectively, under the residual sign convention in (H.44)–(H.46). Write $\lambda_t \equiv (\ell_t, \mu_t, \nu_t)'$ and use bars for deterministic stationary values. The NKPC multiplier, denoted ν_t in Appendix F.4, is ℓ_t here because ν_t multiplies the NFA law. Likewise, the multiplier μ_t here is unrelated to the tariff cost-push term μ_t of main-text Proposition 4. The additive risk terms in the real-allocation equations use the first- and second-order tariff-level response of the same Ramsey problem, written as deviations from the deterministic steady state:

$$\widetilde{x}_{t+1}(u_{t+1}) \equiv x_{t+1}(u_{t+1}) - \bar{x} = \xi_x u_{t+1} + \frac{1}{2} B_x u^2_{t+1} + \mathcal{O}(3), \qquad x \in \{\pi, q, c, \nu\}, \tag{H.26}$$

where u_{t+1} is the scaled tariff innovation, $\xi_x = \partial x_{t+1}/\partial u_{t+1}|_{u=0}$, and $B_x = \partial^2 x_{t+1}/\partial u^2_{t+1}|_{u=0}$. In particular, the multiplier expansion is for $\nu_{t+1} - \bar{\nu}$, not for the level of ν_{t+1}. Define

$$\delta \equiv 1 - \alpha, \quad y_q \equiv \alpha\gamma(2 - \alpha), \quad t_q \equiv y_q - \alpha, \quad g_{qq} \equiv \alpha\delta(1 - \gamma), \tag{H.27}$$

and

$$m_q \equiv \alpha + \varphi y_q, \quad m_c \equiv \sigma + \varphi\delta, \quad k_q \equiv \frac{\epsilon - 1}{\kappa} m_q, \quad k_c \equiv \frac{\epsilon - 1}{\kappa} m_c. \tag{H.28}$$

At the steady state, reduced utility has derivatives

$$U_q = -\chi y_q, \qquad U_c = 1 - \chi\delta. \tag{H.29}$$

The NFA stationarity row gives $\bar{\mu} = 0$. The remaining q and c stationarity rows are

$$-\chi y_q - k_q \bar{\ell} - t_q \bar{\nu} = 0, \tag{H.30}$$

$$1 - \chi\delta - k_c \bar{\ell} + \alpha \bar{\nu} = 0. \tag{H.31}$$

In the q condition, $-\chi y_q$ is the marginal welfare effect through labor needed to produce output, $-k_q \bar{\ell}$ is the contribution of the pricing constraint, and $-t_q \bar{\nu}$ is the contribution of the trade balance to NFA accumulation. In the c condition, $1 - \chi\delta$ is the direct consumption benefit net of the labor required to meet Home consumption demand, $-k_c \bar{\ell}$ again comes from the pricing constraint, and $\alpha\bar{\nu}$ is the NFA-law contribution. Defining $\Delta_\lambda \equiv k_q \alpha + k_c t_q$, Cramer's rule gives

$$\bar{\ell} = \frac{-\chi y_q \alpha - t_q(\chi\delta - 1)}{\Delta_\lambda}, \qquad \bar{\nu} = \frac{k_q(\chi\delta - 1) - k_c \chi y_q}{\Delta_\lambda}. \tag{H.32}$$

At the baseline calibration $\Delta_\lambda \neq 0$, so both multipliers are uniquely determined. Both depend on the calibrated utility weight χ, the same weight that enters the planner objective.

The pricing and reduced Foreign-Euler risk coefficients are

$$\xi_\Gamma \equiv t_q \xi_q + (\delta - \sigma)\xi_c - \alpha\gamma, \tag{H.33}$$

$$d_{NK} \equiv \beta \left(B_\pi + 3\xi_\pi^2 + 2\xi_\pi \xi_\Gamma \right), \tag{H.34}$$

$$d_B \equiv \underbrace{-\sigma B_c + \delta B_q - g_{qq}\xi_q^2}_{\text{conditional-mean curvature}} + \underbrace{(\delta\xi_q - \sigma\xi_c)^2}_{\text{conditional variance}}. \tag{H.35}$$

Keeping only the squared term in (H.35) would omit the second-order conditional-mean response.

Here $\xi_{\pi_C} = \xi_\pi + \alpha\xi_q$, and the Home-Euler conditional-mean and conditional-variance coefficients are

$$d_M^H \equiv \sigma B_c + B_\pi + \alpha B_q + g_{qq}\xi_q^2, \tag{H.36}$$

$$d_V^H \equiv (\sigma\xi_c + \xi_\pi + \alpha\xi_q)^2. \tag{H.37}$$

Combining both terms gives the complete third-order implementing rate:

$$\begin{aligned} r_t^v = {} & \sigma(\mathbb{E}_t c_{t+1}^v - c_t^v) + \mathbb{E}_t \pi_{t+1}^v + \alpha(\mathbb{E}_t q_{t+1}^v - q_t^v) \\ & + (d_M^H - d_V^H) v_t. \end{aligned} \tag{H.38}$$

With this multiplier convention, the four planner-stationarity risk entries are

$$f \equiv \begin{pmatrix} 0 \\ \bar{\ell} t_q d_{NK} \\ \bar{\ell}(\delta - \sigma) d_{NK} \\ -\left\{ B_\nu + \bar{\nu}[\delta B_q + (\delta^2 - g_{qq})\xi_q^2] + 2\delta \xi_\nu \xi_q \right\} \end{pmatrix}. \tag{H.39}$$

The last entry is the risk contribution to stationarity with respect to NFA. It combines the conditional-mean change in the NFA shadow value, the curvature of the effective Foreign-return term, and the covariance between the two.

Proposition H.2 (Third-order tariff-variance terms)**.** *Let $\mathcal{R}$ collect the conditional residuals of the exact private and planner Karush–Kuhn–Tucker (KKT) system. Let k and $h^{(2)}$ collect the first and second derivatives in* (H.26)*, including multiplier derivatives. The coefficient on the half-variance state is*

$$g^v = D\mathcal{R}\, h^{(2)} + D^2\mathcal{R}[k, k]. \tag{H.40}$$

Its private allocation rows reduce to

$$\pi_t^v - \beta \mathbb{E}_t \pi_{t+1}^v - k_q q_t^v - k_c c_t^v = d_{NK} v_t, \tag{H.41}$$

$$\delta(q_t^v - \mathbb{E}_t q_{t+1}^v) - \sigma(c_t^v - \mathbb{E}_t c_{t+1}^v) + \psi_b b_t^v = d_B v_t, \tag{H.42}$$

$$b_t^v - \beta^{-1} b_{t-1}^v - t_q q_t^v + \alpha c_t^v = 0. \tag{H.43}$$

The domestic Euler risk term is absent from these allocation rows because (H.22) *absorbs it in the implementing rate.*

Proof. Let $\mathcal{S}_{t+1}(u)$ collect the private variables and multipliers that enter the continuation residuals, and write its lower-order expansion as $\mathcal{S}_{t+1}(u) = \bar{\mathcal{S}} + ku + \frac{1}{2}h^{(2)}u^2 + \mathcal{O}(u^3)$. The volatility innovation changes its conditional mean by $h^{(2)}v_t$ and its conditional covariance by $2kk'v_t$. The factor two cancels the one-half in the second-order Taylor term of $\mathcal{R}$, giving (H.40). Direct differentiation of (H.12), (H.20), and (H.21) gives (H.34), (H.35), and (H.41)–(H.43); differentiating their multiplier-weighted terms gives (H.39).

The relevant pure-volatility interaction is $u^2 s$: two powers of the tariff innovation and one of its volatility state. Although this interaction is third degree in the exogenous perturbations, it enters the planner's first-order condition by differentiating a fourth-degree term $xu^2 s$ in the Ramsey Lagrangian with respect to the allocation or policy variation x. Differentiation removes x and leaves $u^2 s$. Therefore, deriving the third-order policy condition requires retaining the relevant fourth-degree terms in utility and the equilibrium constraints. Only

terms containing this tariff-variance interaction contribute here because volatility changes the innovation scale but has no direct first-order effect on the allocation. □

H.4 Ramsey conditions and the strict-PPI test

Let $z_t \equiv (\pi_t, q_t, c_t, b_t)'$ and let $F_t \equiv (F_t^N, F_t^E, F_t^B)'$ collect the normalized NKPC, Foreign-Euler, and NFA residuals. Their signs are fixed below.

$$F_t^N \equiv (\Pi_{H,t} - 1)\Pi_{H,t} - \frac{\epsilon - 1}{\kappa}(mc_t - 1) - \beta \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \left(\frac{C_t}{C_{t+1}} \right)^{\sigma} \frac{G_t}{G_{t+1}} \right], \tag{H.44}$$

$$F_t^E \equiv 1 - \beta R_t^* e^{-\psi_b b_t} \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{Q_{t+1}}{Q_t} \frac{G_t}{G_{t+1}} \right], \tag{H.45}$$

$$F_t^B \equiv b_t - R_{t-1}^* e^{-\psi_b b_{t-1}} \frac{Q_t}{Q_{t-1}} \frac{G_{t-1}}{G_t} b_{t-1} - t b_t. \tag{H.46}$$

Each residual equals zero in equilibrium. After the nominal rate is recovered from the domestic Euler equation, a planner that can recommit at date 0, holding the predetermined physical state fixed but carrying no promises from an earlier plan, solves

$$\max_{\{z_t\}_{t \geq 0}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t, L_t) \quad \text{subject to} \quad F_t = 0, \tag{H.47}$$

where $U(C, L) = C^{1-\sigma}/(1-\sigma) - \chi L^{1+\varphi}/(1+\varphi)$. Equation (H.47) is the new-commitment problem: its date-zero first-order conditions have no lower-endpoint multiplier, or equivalently set $\lambda_{-1} = 0$. This appendix instead studies the inherited Ramsey continuation (Section H.1). Fix a realized date-zero history h^0. Conditional on the inherited multiplier λ_{-1} fixed by the earlier plan, the nodewise boundary-correct Lagrangian is

$$\mathcal{J}_0^{R,\text{inh}}(h^0) = \beta^{-1} \lambda_{-1}' \mathfrak{f}_{-1}^{\text{lead}}(z_0(h^0); h^0) + \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \{ U_t + \lambda_t' F_t \}. \tag{H.48}$$

Here $\mathbb{E}_0$ is already conditional on h^0. The lower-endpoint vector contains the date--1 lead integrands evaluated at that node:

$$\mathfrak{f}_{-1}^{\text{lead}}(z_0(h^0);h^0) \equiv \begin{pmatrix} -\beta(\Pi_{H,0}-1)\Pi_{H,0}\dfrac{Y_0}{Y_{-1}}\left(\dfrac{C_{-1}}{C_0}\right)^{\sigma}\dfrac{G_{-1}}{G_0} \\ -\beta R^*_{-1}e^{-\psi_b b_{-1}}\left(\dfrac{C_0}{C_{-1}}\right)^{-\sigma}\dfrac{Q_0}{Q_{-1}}\dfrac{G_{-1}}{G_0} \\ 0 \end{pmatrix}. \tag{H.49}$$

Every date-zero object in this vector is evaluated at h^0. To see why the date--1 expectation disappears, consider a discrete event tree. The derivative of the ex-ante Lagrangian at node h^0 contains its conditional probability under both the date--1 restriction and the date-zero continuation objective. Conditioning on h^0 divides the first-order condition by that common probability, leaving the integrand in (H.49). For a continuous shock, the same result follows from disintegration and localized variations under a regular conditional density; it does not rely on dividing by a zero-probability event. No additional likelihood ratio appears here because both lead restrictions are linear conditional expectations. The first two entries are, respectively, the lead terms in F^N_{-1} and F^E_{-1} under the residual signs in (H.44) and (H.45); F^B_{-1} contains no date-zero lead and therefore contributes the zero in the third entry. All date--1 objects are fixed, while $(\Pi_{H,0}, Y_0, C_0, Q_0, G_0)$ are the date-zero objects implied by $z_0(h^0)$ and the static block. Terms in F_{-1} that do not depend on z_0 are omitted because they do not enter the date-zero first-order conditions. The boundary term is not part of household welfare: it records how the date--1 equilibrium restriction containing z_0 enters the first-order conditions of the earlier commitment. Setting $\lambda_{-1} = 0$ recovers the new-commitment problem (H.47).

The Hessian blocks used below follow from direct differentiation of the static block:

$$y_{cc} = \alpha\delta, \qquad y_{qc} = -\alpha\delta^2\gamma, \qquad y_{qq} = \alpha\delta^2\gamma(1-\alpha\gamma), \tag{H.50}$$

$$m_{cc} = \varphi\alpha\delta, \qquad m_{qc} = \varphi y_{qc}, \qquad m_{qq} = g_{qq} + \varphi y_{qq}. \tag{H.51}$$

Because the NKPC contains $mc - 1 = e^m - 1$, define $M_{ij} \equiv m_{ij} + m_i m_j$ for $i, j \in \{q, c\}$. The trade-balance curvatures are

$$t^{(2)}_{qq} = (y_{qq} - g_{qq}) + t_q^2, \qquad t^{(2)}_{qc} = y_{qc} + \delta t_q, \qquad t^{(2)}_{cc} = -\alpha. \tag{H.52}$$

The direct-utility Hessian is

$$U_{\pi\pi} = -\chi\kappa, \qquad U_{qq} = -\chi\{y_{qq} + (1+\varphi)y_q^2\}, \tag{H.53}$$

$$U_{qc} = -\chi\{y_{qc} + (1+\varphi)y_q\delta\}, \qquad U_{cc} = 1 - \sigma - \chi\{\alpha\delta + (1+\varphi)\delta^2\}. \tag{H.54}$$

Combining these terms with the steady multipliers in (H.32) yields

$$H_{\pi\pi} = -\chi\kappa - \bar{\ell}(\epsilon-1)\varphi, \qquad H_{\pi q} = -\bar{\ell}t_q, \qquad H_{\pi c} = -\bar{\ell}(\delta - \sigma), \tag{H.55}$$

$$H_{qq} = U_{qq} - \bar{\ell}\frac{\epsilon-1}{\kappa}M_{qq} - \bar{\nu}t^{(2)}_{qq}, \quad H_{qc} = U_{qc} - \bar{\ell}\frac{\epsilon-1}{\kappa}M_{qc} - \bar{\nu}t^{(2)}_{qc}, \tag{H.56}$$

$$H_{cc} = U_{cc} - \bar{\ell}\frac{\epsilon-1}{\kappa}M_{cc} - \bar{\nu}t^{(2)}_{cc}, \quad H_{qb} = \bar{\nu}\delta, \qquad H_{bb} = 2\psi_b\bar{\nu}. \tag{H.57}$$

Therefore the current and lead blocks of the transformed Lagrangian are

$$Q_0 = \begin{pmatrix} H_{\pi\pi} & H_{\pi q} & H_{\pi c} & 0 \\ H_{\pi q} & H_{qq} & H_{qc} & H_{qb} \\ H_{\pi c} & H_{qc} & H_{cc} & 0 \\ 0 & H_{qb} & 0 & H_{bb} \end{pmatrix}, \tag{H.58}$$

$$Q_+ = \begin{pmatrix} 0 & 0 & 0 & 0 \\ \beta\bar{\ell}t_q & 0 & 0 & 0 \\ \beta\bar{\ell}(\delta-\sigma) & 0 & 0 & 0 \\ 0 & -\bar{\nu}\delta & 0 & 0 \end{pmatrix}, \qquad Q_- = \beta^{-1}Q_+', \tag{H.59}$$

where the last identity is discounted adjointness. Evaluated at the baseline calibration, these blocks enter the root and boundary calculations below. The four linearized stationarity rows are

$$\begin{aligned} 0 = {} & Q_- z^v_{t-1} + Q_0 z^v_t + Q_+ \mathbb{E}_t z^v_{t+1} \\ & + \begin{pmatrix} \ell^v_t - \ell^v_{t-1} \\ -k_q\ell^v_t + \delta\mu^v_t - t_q\nu^v_t - \delta\beta^{-1}\mu^v_{t-1} \\ -k_c\ell^v_t - \sigma\mu^v_t + \alpha\nu^v_t + \sigma\beta^{-1}\mu^v_{t-1} \\ \psi_b\mu^v_t + \nu^v_t - \mathbb{E}_t\nu^v_{t+1} \end{pmatrix} + fv_t. \end{aligned} \tag{H.60}$$

Equations (H.41)–(H.43) and (H.60) form seven equations for four allocation paths and three multiplier paths. The private rows characterize implementability; the four stationarity rows select the planner's preferred implementable path.

H.4.1 Conditions for local Ramsey optimality

The seven equations above are first-order conditions. Interpreting their bounded solution as a locally optimal Ramsey plan also requires a valid multiplier representation, transversality, local second-order sufficiency, and local uniqueness. This subsection states these conditions as maintained assumptions. Lemma H.1 collects them, Proposition H.3 uses them to characterize when strict PPI coincides with Ramsey, and a finite-support curvature check closes the subsection. None of this material is used to construct the response in Section H.5; readers interested only in the response can skip to that section.

Fix a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$ generated by the shock history. Condition on the fixed date-zero history and denote the resulting regular conditional probability measure by $\mathbb{P}_0$; equivalently, take $\mathcal{F}_0$ to be trivial. For any finite dimension m, define

$$\mathbb{H}^2_\beta(\mathbb{R}^m) \equiv \left\{ z : z_t \in L^2(\Omega, \mathcal{F}_t; \mathbb{R}^m), \quad \mathbb{E}_{\mathbb{P}_0} \sum_{t=0}^{\infty} \beta^t \|z_t\|^2 < \infty \right\}, \qquad \|z\|^2_{\mathbb{H}^2_\beta} \equiv \mathbb{E}_{\mathbb{P}_0} \sum_{t=0}^{\infty} \beta^t \|z_t\|^2. \tag{H.61}$$

Here n_z and n_g denote the finite numbers of allocation coordinates and equilibrium residuals per history. Square integrability alone does not guarantee that the model's exponential maps are well defined or Fréchet differentiable. We therefore fix a Banach space $\mathbb{X}_\beta \subset \mathbb{H}^2_\beta(\mathbb{R}^{n_z})$ of adapted processes on which the exact objective and equilibrium maps are well defined. For bounded histories, one concrete choice is

$$\mathbb{X}_\beta = \mathbb{H}^2_\beta \cap L^\infty_{\text{ad}}, \qquad \|z\|_{\mathbb{X}_\beta} = \|z\|_{\mathbb{H}^2_\beta} + \sup_{t\geq 0} \|z_t\|_{L^\infty}. \tag{H.62}$$

For an unbounded shock law, the fixed domain must instead impose the exponential integrability required by the nonlinear maps; finite moments alone are not sufficient. Let $\mathbb{Y}_\beta = \mathbb{H}^2_\beta(\mathbb{R}^{n_g})$ be the residual space. The weighted inner product identifies its continuous dual with the same Hilbert space. We therefore take $\mathbb{M}_\beta = \mathbb{H}^2_\beta(\mathbb{R}^{n_g})$, with its $\mathbb{H}^2_\beta$ norm, and let an adapted multiplier process act through

$$\langle \Lambda, g \rangle_\beta \equiv \mathbb{E}_{\mathbb{P}_0} \sum_{t=0}^{\infty} \beta^t \Lambda_t' g_t. \tag{H.63}$$

Cauchy–Schwarz makes this pairing continuous.

Let $\mathcal{W}(z; \vartheta, \iota)$ be exact conditional welfare and let $\mathcal{G}(z; \vartheta, \iota) \in \mathbb{Y}_\beta$ collect the exact private-equilibrium restrictions, where $\iota \in \{0, 1\}$ denotes absence or presence of the normalized volatility innovation. Lead, lag, and conditional-expectation operators are assumed bounded between the stated spaces. Fix one closed neighborhood $\mathcal{N}$ of the deterministic steady state

in $\mathbb{X}_\beta \times \Theta$, where Θ is an open interval around zero. The process maps $\mathcal{W}$ and $\mathcal{G}$ are assumed four times continuously Fréchet differentiable on $\mathcal{N}$ for both values of ι, with a common integrable envelope that permits four differentiations under conditional expectations. These process-level conditions, including the required exponential integrability under unbounded shocks, are maintained assumptions.

The exact restrictions have at most one lead. If $D_+\mathcal{G}_T$ denotes the derivative of the date-T residual with respect to its date-$T+1$ argument, the upper endpoint left by the truncated first variation is

$$\mathcal{T}_T^\iota(z, \Lambda, \vartheta; \Delta z) \equiv \mathbb{E}_{\mathbb{P}_0}\left[\beta^T \Lambda_T' D_+ \mathcal{G}_T(z; \vartheta, \iota) \Delta z_{T+1}\right]. \tag{H.64}$$

For $\Lambda \in \mathbb{M}_\beta$, define the fixed endpoint seminorm

$$e_T(\Lambda) \equiv \sup_{\substack{(z,\vartheta)\in\mathcal{N},\ \iota\in\{0,1\} \\ \|\Delta z\|_{\mathbb{X}_\beta}\le 1}} \left|\mathcal{T}_T^\iota(z, \Lambda, \vartheta; \Delta z)\right|. \tag{H.65}$$

Let $\mathfrak{I}$ denote the index set in the supremum in (H.65), and let $\mathbb{B} \equiv \ell^\infty(\mathfrak{I})$ with its supremum norm. For each T, define the bounded-function-valued linear map

$$(\mathcal{E}_T\Lambda)(z, \vartheta, \iota, \Delta z) \equiv \mathcal{T}_T^\iota(z, \Lambda, \vartheta; \Delta z), \qquad \|\mathcal{E}_T\Lambda\|_{\mathbb{B}} = e_T(\Lambda).$$

The endpoint operator is $\mathcal{E}\Lambda \equiv (\mathcal{E}_T\Lambda)_{T\ge 0}$, with values in $\ell^\infty(\mathbb{N}_0; \mathbb{B})$ on its graph domain. Define $c_0(\mathbb{N}_0; \mathbb{B})$ as the closed subspace of sequences whose $\mathbb{B}$-norm converges to zero. We maintain that $\mathcal{E}$ is closed and define its graph domain and the transversality subspace by

$$\begin{aligned} D(\mathcal{E}) &\equiv \{\Lambda \in \mathbb{M}_\beta : \mathcal{E}\Lambda \in \ell^\infty(\mathbb{N}_0; \mathbb{B})\}, \\ \mathbb{M}_\beta^{\mathrm{tvc}} &\equiv \{\Lambda \in D(\mathcal{E}) : \mathcal{E}\Lambda \in c_0(\mathbb{N}_0; \mathbb{B})\}, \\ \|\Lambda\|_{\mathbb{M}_\beta^{\mathrm{tvc}}} &\equiv \|\Lambda\|_{\mathbb{H}_\beta^2} + \sup_{T\ge 0} e_T(\Lambda). \end{aligned} \tag{H.66}$$

The closed-graph assumption makes $D(\mathcal{E})$ complete in this graph norm; because c_0 is closed, $\mathbb{M}_\beta^{\mathrm{tvc}}$ is a fixed closed Banach subspace independent of the branch. Transversality is therefore an admissibility restriction on multipliers, not an extra finite-date Euler equation. For the necessity statement below, we separately maintain that the dual multiplier delivered by the equality-constraint theorem has an adapted Riesz representative in $\mathbb{M}_\beta^{\mathrm{tvc}}$; constraint qualification alone does not prove that representation. The lower endpoint is the inherited-commitment term, with the date--1 multiplier fixed as data.

Let $\bar{d}_{-1}$ collect the predetermined physical coordinates entering the date-zero private restrictions and the inherited multipliers entering the lower endpoint of (H.48). These objects

are fixed parameters, not date-zero choice coordinates or additional residual rows.

For all sufficiently small $|\vartheta|$, assume that $D_z\mathcal{G}(z^\iota(\vartheta);\vartheta,\iota;\bar{d}_{-1})$ is surjective for both ι, with right inverses bounded uniformly in ϑ and ι. Also assume that one $\eta>0$ satisfies

$$\begin{aligned} D^2_{zz}\mathcal{J}_0^{R,\mathrm{inh}}(z^\iota(\vartheta),\Lambda^\iota(\vartheta);\vartheta,\iota;\bar{d}_{-1})[\Delta z,\Delta z] \\ \leq -\eta\|\Delta z\|^2_{\mathbb{H}^2_\beta} \quad \text{for every } \Delta z\in\ker D_z\mathcal{G}(z^\iota(\vartheta);\vartheta,\iota;\bar{d}_{-1}), \end{aligned} \tag{H.67}$$

with objective and constraint Taylor remainders uniformly $o(\|\Delta z\|^2_{\mathbb{H}^2_\beta})$. Here $\mathcal{J}_0^{R,\mathrm{inh}}$ is the exact boundary-correct Lagrangian for the inherited Ramsey continuation, formed from household welfare, the lower-endpoint functional in (H.49), and the multiplier pairing in (H.63). For exact feasible perturbations in the maintained $\mathbb{X}_\beta$ neighborhood, these conditions give strict local maximality and quadratic separation measured by $\|\Delta z\|^2_{\mathbb{H}^2_\beta}$. They do not establish coercivity in the stronger $\mathbb{X}_\beta$ norm. The finite-dimensional calculations below do not establish either claim.

With this dependence displayed, define the exact operator

$$\mathfrak{K}^\iota((z,\Lambda),\vartheta;\bar{d}_{-1})\equiv\begin{pmatrix} D_z\mathcal{J}_0^{R,\mathrm{inh}}(z,\Lambda;\vartheta,\iota;\bar{d}_{-1}) \\ \mathcal{G}(z;\vartheta,\iota;\bar{d}_{-1})\end{pmatrix}:\mathbb{X}_\beta\times\mathbb{M}^{\mathrm{tvc}}_\beta\times\Theta\longrightarrow\mathbb{X}^*_\beta\times\mathbb{Y}_\beta. \tag{H.68}$$

Assume each $\mathfrak{K}^\iota$ is three times continuously Fréchet differentiable and its derivative with respect to (z,Λ) at the common deterministic steady state is a bounded isomorphism. The Banach-space implicit-function theorem gives unique local branches $Z^\iota(\vartheta)=(z^\iota(\vartheta),\Lambda^\iota(\vartheta))$.

The lower-order cancellation applies to the response reported in the paper, not to every off-path node of the adapted branch. We maintain that the local solution branch admits continuous Markov policy-function representatives on a fixed neighborhood of the deterministic state. On those representatives, let $\mathfrak{P}_{\mathrm{nr}}$ evaluate the process along the history on which all realized tariff innovations from date zero onward are zero, and define $Z^\iota_{\mathrm{nr}}\equiv\mathfrak{P}_{\mathrm{nr}}Z^\iota$. Equation (H.4) shows why this projection is needed: the full adapted branches can differ at order two at histories with a nonzero tariff realization. The sup-norm component of (H.62) makes evaluation of the continuous representative bounded; point evaluation would not be well defined on unqualified L^2 equivalence classes.

At second order, the difference between the full branch equations has the form $A\Delta Z^{[2]}=-r_2$, where A is the common linearized KKT operator and r_2 is proportional to the order-two innovation in (H.4). Thus r_2 is odd in tariff innovations and is generally nonzero. Symmetry, uniqueness, and preservation of parity by A^{-1} imply that $\Delta Z^{[2]}$ is also odd under sign reversal of the tariff-innovation history. Let $\mathcal{F}^\sigma_t$ be the sigma-field generated by the volatility history.

Conditional sign symmetry implies

$$\mathbb{E}[\Delta Z_t^{[2]} \mid \mathcal{F}_t^{\sigma}] = 0, \qquad \mathfrak{P}_{\mathrm{nr}} \Delta Z^{[2]} = 0. \tag{H.69}$$

This does not imply a zero one-step conditional mean at an arbitrary off-path node, because the realized odd tariff state is then part of the full information set. At a no-realization node, however, the next-period odd continuation again has zero conditional mean. The first difference that survives the reported projection is the order-three half-variance term in (H.3). Hence

$$Z_{\mathrm{nr}}^{1}(0) = Z_{\mathrm{nr}}^{0}(0), \qquad Z_{\mathrm{nr}}^{1\prime}(0) = Z_{\mathrm{nr}}^{0\prime}(0), \qquad Z_{\mathrm{nr}}^{1\prime\prime}(0) = Z_{\mathrm{nr}}^{0\prime\prime}(0). \tag{H.70}$$

Date-zero means and means conditional only on the volatility history obey the same lower-order cancellation under the stated symmetry condition; the figures report the no-realization projection. Applying Taylor's theorem to the projected branches gives the normalization

$$\begin{aligned} Z_{\mathrm{nr}}^{1}(\vartheta) - Z_{\mathrm{nr}}^{0}(\vartheta) &= \vartheta^3 \bar{\sigma}_{\tau}^2 s_0 Z^v + o(\vartheta^3), \\ Z^v &= \frac{1}{6\bar{\sigma}_{\tau}^2 s_0} \frac{d^3}{d\vartheta^3} \big[Z_{\mathrm{nr}}^{1}(\vartheta) - Z_{\mathrm{nr}}^{0}(\vartheta) \big] \Big|_{\vartheta=0}, \end{aligned} \tag{H.71}$$

where $s_0 \neq 0$ is the date-zero volatility-state coefficient. Thus the order-three object below is the baseline-subtracted no-realization-path coefficient, not the full adapted-process difference or the unscaled third derivative. Strong regularity and the uniform sufficiency conditions are assumed; finite-horizon curvature or root checks do not prove them.

Lemma H.1 (Existence of a local Ramsey branch)**.** *Under the function-space and differentiability conditions above, suppose the constraint derivative is surjective, the multiplier has the maintained adapted transversality representation, the second-order condition* (H.67) *and its remainder condition hold, and the exact KKT map* (H.68) *is strongly regular. Then the private restrictions, planner stationarity equations, inherited-multiplier boundary, and transversality condition characterize a unique strict local Ramsey maximum.*

Proof. Surjectivity and the multiplier representation make the first-order conditions necessary. The coercive second variation and uniform remainder give a strict local constrained maximum with quadratic separation in the $\mathbb{H}_{\beta}^2$ norm. Bounded invertibility of the derivative of (H.68) then gives a unique local branch by the implicit-function theorem. □

Proposition H.3 (When strict PPI coincides with Ramsey)**.** *Maintain the conditions of Lemma H.1. Let z^{PPI} be the strict-PPI allocation from the same physical initial state and shock process. Strict PPI coincides locally with Ramsey if and only if z^{PPI} can be paired with multiplier paths satisfying the Ramsey stationarity equations, inherited boundary, and*

transversality condition. A nonzero PPI-inflation coefficient in the Ramsey response therefore rejects strict PPI through the retained order. A zero coefficient alone is inconclusive because the remaining allocation and multiplier conditions must also hold.

Proof. Strict PPI imposes $\pi_t^v = 0$ after every retained history. If the Ramsey coefficient is nonzero at any date or history, the allocations differ. If it is zero, only the inflation restriction has been checked. If strict PPI coincides locally with Ramsey, then z^{PPI}, paired with the Ramsey multiplier paths, satisfies every stationarity, boundary, and transversality condition. Conversely, z^{PPI} satisfies the private restrictions by construction; if it can be paired with multiplier paths satisfying these remaining conditions, the local uniqueness in Lemma H.1 implies that it is the Ramsey branch. □

Finite-support curvature check. A simple numerical check restricts the infinite-horizon second variation to feasible perturbations with finite support. For terminal date T, stack

$$Z_T = (z_0', \ldots, z_T')'$$

and extend the perturbation by zero for every $t > T$. Substitute the fixed inherited deviation $b_{-1} = 0$ into the date-zero rows; it is data and does not add a row to the finite matrix. Let $A_T Z_T = 0$ collect the resulting homogeneous linearized private restrictions (H.41)–(H.43) for dates 0 through T and the NFA restriction at date $T + 1$. Because the perturbation is zero from $T + 1$ onward, the last restriction reduces to $-\beta^{-1} b_T = 0$ and therefore imposes $b_T = 0$. The additive risk term is irrelevant because feasible directions satisfy the homogeneous restrictions.

Let H_T denote the discounted block-tridiagonal Hessian with blocks

$$(H_T)_{tt} = \beta^t Q_0, \qquad (H_T)_{t,t+1} = \beta^t Q_+, \qquad (H_T)_{t+1,t} = \beta^t Q_+'. \tag{H.72}$$

Thus H_T is the principal block, for dates 0 through T, of the infinite-horizon Lagrangian Hessian evaluated at the deterministic steady state. If the columns of N_T form an orthonormal basis for $\ker(A_T)$, every deterministic coefficient-path direction represented by A_T can be written $Z_T = N_T d$. Its second variation is $d'(N_T' H_T N_T) d$.

At every finite horizon we compute, the constraint Jacobian has full row rank and the largest eigenvalue of the reduced Hessian is negative, so the second variation is negative for every deterministic finite-support feasible direction at those horizons. As the horizon grows, however, that eigenvalue moves toward zero. The check does not establish the uniform infinite-horizon bound in (H.67), transversality, or local isolation, and it does not span every history-dependent perturbation in $\mathbb{H}_\beta^2$. We therefore call the computed object the solution to

the Ramsey first-order conditions; interpreting it as the local Ramsey optimum requires the conditions in Lemma H.1.

Strict PPI sets normalized $mc_t = 1$ in (H.12) and removes Rotemberg adjustment costs, but it does not make the Foreign Euler equation, the NFA shadow value, or the remaining stationarity rows redundant. As noted after (H.12), the fixed production correction does not make the flexible-price allocation efficient in every tariff state.

H.5 Closed-form response and numerical verification

This subsection derives the third-order response in closed form and verifies it against the numerical solution. The analytical solution takes as given only the order-one and order-two tariff-level derivatives (ξ, B) in (H.26); no order-three response is an input. The response is the sum of a component that decays at the exogenous persistence of tariff risk and three endogenous decaying components. At the baseline calibration, a fast-decaying real root adjusts the inherited price-setting commitment, while a slowly decaying complex-conjugate pair carries NFA and its shadow value. The derivation obtains these components and their weights from the lower-order derivatives and the inherited Ramsey boundary.

To define the root count before any normalization, let $\mathcal{S}_t = (z_t', \ell_t, \mu_t, \nu_t)'$ and let $\mathcal{R}_t^h = 0$ stack the seven homogeneous private and stationarity rows. Define

$$\mathcal{A}_j \equiv D_{\mathcal{S}_{t+j}} \mathcal{R}_t^h \big|_{\bar{\mathcal{S}}}, \qquad \widetilde{\mathcal{P}}(r) \equiv r^{-1}\mathcal{A}_{-1} + \mathcal{A}_0 + r\mathcal{A}_{+1}, \qquad \mathcal{P}(r) \equiv r\widetilde{\mathcal{P}}(r). \tag{H.73}$$

Thus $\mathcal{P}$ is the seven-equation quadratic matrix polynomial associated with a solution component proportional to r^t. At the baseline, direct determinant evaluation gives

$$\det \mathcal{P}(r) = Cr^4 P_R(r), \qquad C \neq 0, \tag{H.74}$$

where P_R is the degree-six polynomial displayed below and $P_R(0) = \beta^{-3} \neq 0$. Direct evaluation of the coefficient matrices also gives

$$\operatorname{rank} \mathcal{A}_{-1} = \operatorname{rank} \mathcal{A}_{+1} = 3. \tag{H.75}$$

Thus zero has algebraic and geometric multiplicity four, so its elementary divisors are four copies of r. The grade-two reversal

$$\operatorname{rev}_2 \mathcal{P}(s) \equiv \mathcal{A}_{+1} + s\mathcal{A}_0 + s^2\mathcal{A}_{-1} \tag{H.76}$$

also has four simple zero elementary divisors. Equivalently, $\mathcal{P}$ has four simple infinite

elementary divisors. These structural divisors are not finite nonzero roots of the original Laurent system and therefore do not enter the displayed r^t modal expansion. This classification does not by itself establish endpoint uniqueness. To connect the divisor count to admissible paths, invoke the maintained bounded-isomorphism assumption for the derivative of the full KKT map in (H.68), on the same bounded path space and with the same inherited boundary and transversality conditions. Injectivity then rules out a nonzero homogeneous component supported only at the lower endpoint. In a regular strong linearization, the zero and infinite blocks are therefore algebraic blocks rather than additional free propagating modes.

For the six finite nonzero roots, assume that none satisfies $|r| = 1$, all are simple, the volatility persistence is distinct from them, and the three-by-three matrix formed by (H.106)–(H.108) is nonsingular when evaluated on the three roots with $|r| < 1$. Those three roots generate bounded paths that satisfy transversality; the other three are their discounted reciprocals and are excluded because their paths grow without bound. At the baseline, the computed finite roots solve the characteristic equation, pair as discounted reciprocals, and give a nonsingular boundary matrix. The ratio normalizations used below also require the denominator conditions in (H.102). Under the full-operator injectivity, finite-root, endpoint-rank, and denominator conditions, the scalar equation derived below contains every admissible finite nonzero endogenous component and no extraneous one.

A repeated root with a complete eigenbasis requires additional independent eigenvectors but no polynomial factor. A defective root requires terms of the form $t^k r^t$. If the exogenous risk persistence equals an endogenous root, the particular solution also requires a term proportional to $t^k r^t$ unless the risk vector satisfies the relevant range condition. Taking the limit of the complete sum as roots approach one another can recover these terms, but direct substitution into the simple-root Cramer formulas cannot. These cases lie outside the displayed closed form.

Consider a component proportional to r^t and let h_x denote its coefficient for $x \in \{\pi, q, c, b\}$. Define

$$\zeta(r) \equiv 1 - (\beta r)^{-1}, \tag{H.77}$$

$$E_q(r) \equiv \delta(1-r) + \frac{\psi_b t_q}{\zeta(r)}, \qquad E_c(r) \equiv -\sigma(1-r) - \frac{\psi_b \alpha}{\zeta(r)}, \tag{H.78}$$

$$\Delta(r) \equiv k_q E_c(r) - k_c E_q(r), \qquad S(r) \equiv (1-\beta r)h_\pi - d_{NK}. \tag{H.79}$$

Substitution into the three private rows gives

$$k_q h_q + k_c h_c = S(r), \qquad E_q(r)h_q + E_c(r)h_c = d_B, \tag{H.80}$$

and the NFA law then recovers h_b. These formulas require $r \neq 0$, $\zeta(r) \neq 0$, and $\Delta(r) \neq 0$. Under those conditions, Cramer's rule yields

$$h_q = \frac{E_c(r)S(r) - k_c d_B}{\Delta(r)}, \tag{H.81}$$

$$h_c = \frac{k_q d_B - E_q(r)S(r)}{\Delta(r)}, \tag{H.82}$$

$$h_b = \frac{t_q h_q - \alpha h_c}{\zeta(r)}. \tag{H.83}$$

Thus the allocation coefficients for every privately feasible component can be written

$$h_z(r) = z^0(r) + n(r)h_\pi, \tag{H.84}$$

where z^0 sets $h_\pi = 0$ in (H.81)–(H.83), and

$$n(r) = \begin{pmatrix} 1 \\ (1-\beta r)E_c(r)/\Delta(r) \\ -(1-\beta r)E_q(r)/\Delta(r) \\ [t_q n_q(r) - \alpha n_c(r)]/\zeta(r) \end{pmatrix}. \tag{H.85}$$

Write $z^0 = d_{NK} z_N + d_B z_B$, where z_N and z_B are obtained from the same Cramer formulas after setting $(d_{NK}, d_B) = (1, 0)$ and $(0, 1)$, respectively.

For a component proportional to r^t, the matrix multiplying the three constraint multipliers in (H.60) is

$$A_\lambda(r) = \begin{pmatrix} 1 - r^{-1} & 0 & 0 \\ -k_q & \delta\zeta(r) & -t_q \\ -k_c & -\sigma\zeta(r) & \alpha \\ 0 & \psi_b & 1 - r \end{pmatrix}. \tag{H.86}$$

Direct multiplication using (H.85) verifies that its normalized left-null vector is

$$w(r) = n\big((\beta r)^{-1}\big), \qquad w(r)' A_\lambda(r) = 0. \tag{H.87}$$

Define the allocation Hessian kernel

$$\mathcal{Q}(r) \equiv r^{-1} Q_- + Q_0 + r Q_+. \tag{H.88}$$

Left-multiplying the four stationarity rows by $w(r)'$ removes all three multiplier coefficients.

The remaining scalar projection condition is

$$w'\mathcal{Q}z^0 + w'\mathcal{Q}n\,h_\pi + w'f = 0. \tag{H.89}$$

Consequently,

$$h_\pi(r) = \frac{\omega_{NK}(r)d_{NK} + \omega_B(r)d_B + w(r)'f}{\Lambda_R(r)}, \tag{H.90}$$

where

$$\Lambda_R \equiv -w'\mathcal{Q}n, \qquad \omega_{NK} \equiv w'\mathcal{Q}z_N, \qquad \omega_B \equiv w'\mathcal{Q}z_B. \tag{H.91}$$

The projection is well defined whenever $\Lambda_R(r) \neq 0$. Its sign is not implied by the finite-support curvature check above. At the baseline, $\Lambda_R(\rho_\sigma)$ is positive, so (H.90) determines a finite inflation response. Equation (H.90) shows why the Foreign-Euler risk term alone does not determine Ramsey policy: the planner-stationarity term $w'f$ enters at the same order.

For a homogeneous component the risk vector is zero and $z_t = n(r)r^t$. Applying the same projection gives the scalar characteristic equation

$$\mathcal{K}(r) \equiv n\big((\beta r)^{-1}\big)'\,\mathcal{Q}(r)n(r) = 0. \tag{H.92}$$

Let $P_R(r)$ denote the degree-six numerator obtained after clearing the denominators in (H.92). Discounted adjointness implies the reciprocal factorization

$$P_R(r) = \prod_{j=0}^{2}\left(r^2 - x_j r + \beta^{-1}\right), \tag{H.93}$$

where x_j solve

$$\begin{aligned} &x^3 - S_1x^2 + S_2x - S_3 = 0, \\ &S_1 = \sum_{j=0}^{2} x_j, \qquad S_2 = \sum_{0\le j<k\le 2} x_jx_k, \qquad S_3 = \prod_{j=0}^{2} x_j. \end{aligned} \tag{H.94}$$

The values of (S_1, S_2, S_3) are read directly from the coefficients of P_R after matching (H.93); no numerical root is needed for this reduction. For an explicit radical representation, define

$$p_C = S_2 - \frac{S_1^2}{3}, \qquad q_C = -\frac{2S_1^3}{27} + \frac{S_1S_2}{3} - S_3, \qquad \Delta_C = \left(\frac{q_C}{2}\right)^2 + \left(\frac{p_C}{3}\right)^3. \tag{H.95}$$

Let $\omega_C = e^{2\pi i/3}$ and choose the cube-root branches to satisfy

$$A_C^3 = -\frac{q_C}{2} + \sqrt{\Delta_C}, \qquad B_C^3 = -\frac{q_C}{2} - \sqrt{\Delta_C}, \qquad A_C B_C = -\frac{p_C}{3}. \tag{H.96}$$

The product condition is the branch restriction required by Cardano's formula; degenerate cases are obtained by continuity. These branches give

$$x_j = \frac{S_1}{3} + \omega_C^j A_C + \omega_C^{-j} B_C, \qquad j = 0, 1, 2. \tag{H.97}$$

Each pair of Ramsey roots is then

$$r_{j,\pm} = \frac{x_j \pm \sqrt{x_j^2 - 4/\beta}}{2}. \tag{H.98}$$

At the baseline, define the three selected roots by

$$\{r_1, r_2, r_3\} \equiv \{r_{j,\pm} : j = 0, 1, 2,\ |r_{j,\pm}| < 1\}, \qquad r_1 \in \mathbb{R}, \quad r_3 = \overline{r_2}. \tag{H.99}$$

Thus r_1 is the real root and (r_2, r_3) is the slowly decaying complex-conjugate pair. For each selected root r_j, the other root in its quadratic pair is $1/(\beta r_j)$ and has modulus greater than one, so its associated path is excluded by boundedness. The numerical calculation evaluates (H.92) at full precision.

The multiplier coefficients are also explicit. For a homogeneous component, let $s(r) \equiv -\mathcal{Q}(r)n(r)$. The remaining coefficients are

$$\ell(r) = \frac{s_\pi(r)}{1 - r^{-1}}, \qquad A_\mu(r) = \delta\zeta(r), \tag{H.100}$$

$$H_\mu(r) = A_\mu(r)(1 - r) + \psi_b t_q, \qquad R_q(r) = s_q(r) + k_q \ell(r). \tag{H.101}$$

The modal formulas require the following regularity conditions. Let $\mathcal{D}_R \equiv \{\rho_\sigma, r_1, r_2, r_3\}$. With $0 < \beta < 1$ and $0 < \rho_\sigma < 1$, require, for every $r \in \mathcal{D}_R$,

$$r\,\zeta(r)\,\Delta(r)\,\zeta\big((\beta r)^{-1}\big)\,\Delta\big((\beta r)^{-1}\big)\,(1 - r^{-1})H_\mu(r) \neq 0, \qquad \Lambda_R(\rho_\sigma) \neq 0. \tag{H.102}$$

The factors involving $(\beta r)^{-1}$ are needed because the left-null vector $w(r)$ evaluates n at the discounted reciprocal. The last condition selects a finite particular response to the risk state. All conditions hold at the baseline. If any factor in (H.102) vanishes, the corresponding ratio normalization is invalid: the component must instead be obtained from the nullspace or forced linear system of the full pencil $\widetilde{\mathcal{P}}(r)$. Repeated, defective, and resonant cases then

use the generalized modes described above rather than the simple-root formulas. The q- and b-stationarity rows give

$$\begin{aligned}\mu(r) &= \frac{R_q(r)(1-r) + t_q s_b(r)}{H_\mu(r)},\\ \nu(r) &= \frac{A_\mu(r) s_b(r) - \psi_b R_q(r)}{H_\mu(r)}.\end{aligned} \tag{H.103}$$

Let $e(r) = (n(r)', \ell(r), \mu(r), \nu(r))'$. For the component driven directly by the risk state, replace n by h_z, replace s by $s^h = -\mathcal{Q}h_z - f$, and apply (H.100)–(H.103); denote the resulting vector by $h = (h_z', h_\ell, h_\mu, h_\nu)'$.

It remains to impose the inherited Ramsey boundary. The deviations of the inherited histories satisfy

$$q^v_{-1} = c^v_{-1} = b^v_{-1} = 0, \qquad \ell^v_{-1} = \mu^v_{-1} = 0. \tag{H.104}$$

Because q_0 and c_0 can adjust at date 0, the three independent conditions that determine the component weights are inherited b_{-1}, inherited μ_{-1}, and the actual date-zero inflation stationarity row

$$e_\pi'\{Q_0 z_0^v + Q_+ \mathbb{E}_0 z_1^v\} + \ell_0^v + f_\pi v_0 = 0, \qquad e_\pi = (1,0,0,0)'. \tag{H.105}$$

Let $e_j = e(r_j)$, with the roots ordered as in (H.99). Define the three-by-three matrix M and vector d by

$$M_{1j} = e_{j,b}/r_j, \qquad d_1 = -h_b/\rho_\sigma, \tag{H.106}$$

$$M_{2j} = e_{j,\mu}/r_j, \qquad d_2 = -h_\mu/\rho_\sigma, \tag{H.107}$$

$$M_{3j} = e_\pi'(Q_0 + r_j Q_+)e_{j,z} + e_{j,\ell}, \qquad d_3 = -\{e_\pi'(Q_0 + \rho_\sigma Q_+)h_z + h_\ell + f_\pi\}. \tag{H.108}$$

Cramer's rule requires $\det(M) \neq 0$ and gives $a_j = \det(M_j)/\det(M)$, with M_j formed by replacing column j of M by d. The baseline boundary matrix is nonsingular, so these three weights are uniquely determined. The order-three path that excludes the three explosive roots and satisfies the Ramsey first-order conditions and inherited boundary is therefore

$$\mathcal{S}_t^v = v_0 \left[h\rho_\sigma^t + \sum_{j=1}^{3} a_j e_j r_j^t \right]. \tag{H.109}$$

Replacing (H.105) by a backward extrapolation of an impact multiplier would preserve the roots but solve a different date-zero problem.

We verify the closed form in four steps. First, a Dynare implementation of the Ramsey problem solves the exact nonlinear private-equilibrium and planner-stationarity equations

with R_t as the instrument. Second, its order-one and order-two tariff-level solution supplies (ξ, B) for (H.34)–(H.39). Third, the analytical target, characteristic roots, and boundary equations generate (H.109). Fourth, a separately computed third-order solution of the same first-order-condition system produces the pure-volatility response. The two calculations share the model and its first- and second-order derivatives, but the numerical third-order response is not used to choose the analytical target, roots, or weights.

Proposition H.4 (Analytical formula matches the numerical response)**.** *For the baseline calibration and the experiment in Section H.1, the analytical path matches, at numerical precision over the reported horizon, the separately computed third-order path that satisfies the same Ramsey first-order conditions and inherited boundary. The same agreement holds for exports recovered from the static block. This comparison verifies the computed response, not the infinite-horizon curvature, transversality, and local-isolation conditions required for a local Ramsey optimum; that interpretation additionally requires the conditions in Lemma H.1.*

Proof. We recompute d_{NK}, d_B, and f from the lower-order solution, solve the characteristic and boundary equations, and evaluate (H.109). The characteristic equation and inherited boundary hold at numerical precision. Direct comparison with the numerical solution of the same first-order-condition system then gives machine-precision agreement for every displayed series and for exports recovered from the static equations. □

Figure H.1 also illustrates the economics of the modal decomposition. The terms of trade adjust quickly to the risk innovation, while NFA accumulates and then unwinds slowly. Consumption inherits both movements. The policy rate implements this real allocation, and PPI inflation remains small but nonzero. Under complete markets, the fixed production correction aligns the flexible-price allocation with the constrained-efficient one at the deterministic steady state but not, in general, in other tariff states. Because the correction does not adjust with the tariff, the Ramsey allocation can already depart from strict PPI away from that state. Incomplete markets add two further margins: households adjust intertemporal saving, and the resulting NFA position changes external wealth and future trade balances. These additional margins help explain why the Ramsey continuation need not satisfy strict PPI, even though its PPI-inflation response remains small.

H.6 Requirements for a fourth-order welfare comparison

The agreement in Figure H.1 concerns the local conditional response; it does not rank policies by expected lifetime utility. A welfare comparison must integrate household utility over the complete future tariff event tree and retain every term of total order four. This subsection

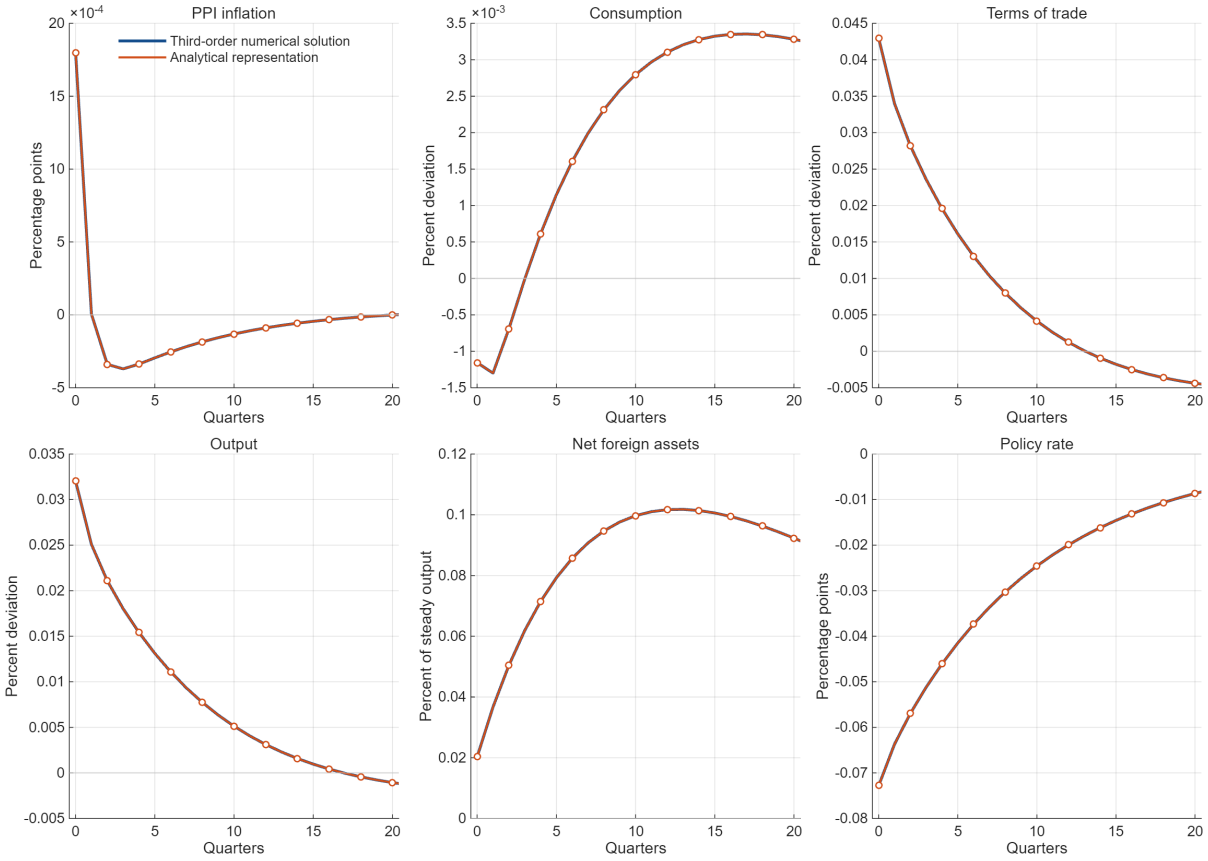


Figure H.1: Analytical and numerical incomplete-markets Ramsey responses

Notes: The experiment is a one-standard-deviation date-0 log-volatility innovation, $\varepsilon_{\sigma,0} = 1$ with $\eta_\sigma = \log 2$; later realized tariff innovations are zero while future risk remains in expectations. Each curve subtracts the Ramsey continuation's no-volatility path and inherits its pre-date-0 planner multipliers. Solid blue lines are the separately computed third-order numerical responses; orange lines and markers are the analytical formula in (H.109). The formula uses only first- and second-order tariff-level derivatives and the inherited Ramsey boundary. Inflation and the policy rate are percentage-point deviations; real variables are percent deviations; net foreign assets are a percent of steady-state output, with a positive response denoting a larger Home creditor position.

states what that calculation requires and why the third-order response is not, by itself, sufficient.

Let δ_s be a common perturbation scale for the tariff-level and volatility shocks, while the standardized innovations remain fixed. For a Ramsey commitment boundary $\mathfrak{c} \in \{\text{new}, \text{inh}\}$, let $W^R_{\mathfrak{c}}$ denote Ramsey welfare and let W^j denote welfare under policy j from the same physical state and shock process. For the pure-risk experiment, the calculation below assumes that the relevant allocation branches and utility and constraint integrands are five times continuously differentiable in a neighborhood of $\delta_s = 0$, with an integrable dominating function that permits differentiation under the expectation. With only four continuous derivatives, each $O(\delta_s^5)$ remainder below must instead be written $o(\delta_s^4)$. The common deterministic allocation and shock symmetry remove the constant and first-order terms, so the signed lifetime-utility gap can be written

$$\Delta W_{j,\mathfrak{c}}(\delta_s) \equiv W^R_{\mathfrak{c}}(\delta_s) - W^j(\delta_s) = d^j_{2,\mathfrak{c}}\delta_s^2 + d^j_{3,\mathfrak{c}}\delta_s^3 + d^j_{4,\mathfrak{c}}\delta_s^4 + O(\delta_s^5). \tag{H.110}$$

The gap need not be positive when $\mathfrak{c} = \text{inh}$: continuing a past plan is not the same problem as maximizing welfare afresh at date 0. The coefficient $d^j_{4,\mathfrak{c}}$ depends on fourth-order conditional moments even when policy functions are solved only through third order. For example, if $\varepsilon_{\tau,i}$ and $\varepsilon_{\tau,k}$ are independent standard-normal tariff innovations at two future dates, as in the numerical welfare calculation, then

$$\mathbb{E}\varepsilon^4_{\tau,i} = 3, \qquad \mathbb{E}(\varepsilon^2_{\tau,i}\varepsilon^2_{\tau,k}) = 1 \quad (i \neq k). \tag{H.111}$$

These moments capture both the fourth power of a tariff innovation and mixed products of squared tariff innovations at different dates. Under the weaker symmetry assumption used for the analytical results, the first value is instead the unspecified fourth moment $\mu_4 \equiv \mathbb{E}\varepsilon^4_{\tau,i}$.

A complete calculation must also integrate future volatility innovations. To see why they enter at fourth order, suppress time subscripts and expand the scaled conditional variance:

$$\delta_s^2\bar{\sigma}^2_\tau e^{2\delta_s s} = \bar{\sigma}^2_\tau\left(\delta_s^2 + 2\delta_s^3 s + 2\delta_s^4 s^2\right) + O(\delta_s^5). \tag{H.112}$$

The s^2 term depends on the variance of future volatility innovations. Conditioning those innovations to zero would define a different welfare object from integrating over the joint tariff–volatility process.

A one-shock-at-a-time rule reproduces first and second moments but not the fourth

moments required here. To see why, define

$$\mathcal{I}_2[h] \equiv h(0) + \frac{1}{2}\sum_i \{h(e_i) + h(-e_i) - 2h(0)\}, \tag{H.113}$$

where e_i places a unit innovation at date i and zero innovations at every other date. This rule reproduces expected values for polynomials of total degree at most three under independent symmetric shocks. It is not fourth-order accurate: $\mathcal{I}_2[\varepsilon_{\tau,i}^4] = 1$ rather than 3, and $\mathcal{I}_2[\varepsilon_{\tau,i}^2\varepsilon_{\tau,k}^2] = 0$ rather than 1. Hence it omits pure fourth moments and cross-date fourth-order terms. Increasing the horizon cannot correct this failure because the rule never evaluates joint shocks at two dates and assigns the wrong fourth moment to a single shock.

There is a second issue. Let z_j denote the complete state-contingent allocation under policy j. If it is expanded as

$$z_j(\delta_s) = \bar{z} + \delta_s z_{1,j} + \frac{\delta_s^2}{2} z_{2,j} + \frac{\delta_s^3}{6} z_{3,j} + \frac{\delta_s^4}{24} z_{4,j} + O(\delta_s^5), \tag{H.114}$$

then a direct fourth-order expansion of household utility contains the term $DW(\bar{z})[z_{4,j}]/24$. Because household utility has a nonzero gradient at the steady state (see (H.29)), a direct utility expansion generally depends on the policy's unknown fourth-order allocation.

Two procedures can deliver a valid $d_{4,\mathfrak{c}}^j$. The first is to solve every policy's allocation through fourth order and integrate the resulting utility expansion over the joint shock process with a rule that reproduces the required fourth moments. For Ramsey, this route requires a five-times differentiable KKT branch and therefore stronger smoothness than the third-order response proof above.

The second procedure applies when the Ramsey benchmark is a new date-0 commitment. Add every equilibrium constraint with its Ramsey multiplier and expand that Lagrangian through fourth order, treating the perturbation scale as an explicit argument. Ramsey stationarity then eliminates the unknown fourth-order allocation, leaving policy responses through third order and fourth derivatives of utility and the equilibrium conditions. Section F.10 derives this new-commitment envelope calculation for complete markets. Its incomplete-markets counterpart must include the two Euler equations, the debt-elastic premium, and the net-foreign-asset law.

The inherited continuation requires an additional boundary calculation. Its Lagrangian in (H.48) contains

$$B(z_0(h^0); h^0) \equiv \beta^{-1}\lambda_{-1}' \mathbf{f}_{-1}^{\text{lead}}(z_0(h^0); h^0), \tag{H.115}$$

where λ_{-1} is fixed by the pre-date-0 plan. For a feasible allocation, $\mathcal{J}_0^{R,\text{inh}} = W + B$, so stationarity removes the unknown fourth-order allocation from $\mathcal{J}_0^{R,\text{inh}}$ but not, in general,

from B. A third-order policy solution is sufficient only if one also computes the fourth-order boundary contribution or formulates an explicit common promise-keeping constraint that makes it cancel across policies. We take neither step for this economy.

The consumption-equivalent conversion must be expanded to the same order. With deterministic steady-state consumption normalized to one, let $\ell_{j,\mathfrak{c}}$ be the permanent proportional increase in consumption whose utility value equals $\Delta W_{j,\mathfrak{c}}$:

$$\frac{1}{1-\beta}\frac{(1+\ell_{j,\mathfrak{c}})^{1-\sigma}-1}{1-\sigma} = \Delta W_{j,\mathfrak{c}}. \tag{H.116}$$

For $\sigma = 1$, the left side is $\log(1+\ell_{j,\mathfrak{c}})/(1-\beta)$. The expansion below is valid in that limit as well. Expanding and inverting this expression gives

$$\begin{aligned}\ell_{j,\mathfrak{c}}(\delta_s) = {} & (1-\beta)d^j_{2,\mathfrak{c}}\delta_s^2 + (1-\beta)d^j_{3,\mathfrak{c}}\delta_s^3 \\ & + \left[(1-\beta)d^j_{4,\mathfrak{c}} + \frac{\sigma}{2}(1-\beta)^2(d^j_{2,\mathfrak{c}})^2\right]\delta_s^4 + O(\delta_s^5).\end{aligned} \tag{H.117}$$

The final term is essential: multiplying a welfare difference by $1-\beta$ does not produce a fourth-order consumption equivalent when $d^j_{2,\mathfrak{c}}$ is nonzero. Equation (H.117) is the incomplete-markets counterpart of (F.138).

Finally, all policies must use the same predetermined physical state and joint shock process, and the Ramsey benchmark must state how promises made before date 0 are treated. The response in this appendix retains the planner multipliers inherited from the pre-date-0 plan. A new date-0 commitment sets the inherited multiplier term to zero and requires the planner problem to be solved again. Fixed rules do not carry planner multipliers, so an inherited-Ramsey comparison must define the promise-keeping restriction or evaluate the boundary term in (H.115); merely starting from the same physical state is not enough.

Because the available one-shock-at-a-time integration and consumption-equivalent conversion omit fourth-order terms, and the fourth-order inherited-boundary contribution is not computed, we report no export-tariff incomplete-markets welfare ranking among Ramsey, strict PPI targeting, and the Taylor rule. Nor do we compute a new-commitment Ramsey solution with welfare integrated over the joint tariff and volatility shock process for this economy. The import-tariff appendix instead uses a new-commitment Ramsey reference and integrates future Gaussian innovations jointly (Appendix I.5), which is why a welfare ranking is reported there.

I Import Tariffs: Derivations and Policy Comparisons

This appendix derives the import-tariff results from primitives. It keeps four distinctions explicit: tariff levels versus tariff uncertainty, PPI versus CPI policy rules, flexible prices versus constrained efficiency, and complete versus incomplete international asset markets. Sections I.1 and I.2 define the experiments and the model. Section I.3 derives the complete-markets responses under Taylor rules, the flexible-price and constrained-efficient allocations, and Ramsey policy; Section I.4 repeats the analysis in a bond economy; and Section I.5 reports a fourth-order welfare comparison. The quantitative comparisons use the baseline calibration in Table B.1.

I.1 Experiments, units, and commitment boundary

Let $\tau_t^F = \log T_{F,t}$ and let s_t be its log volatility (the import-tariff counterpart of $\sigma_{T,t}$ in the main text). The shock laws are

$$\tau_t^F = \rho_\tau \tau_{t-1}^F + \bar{\sigma}_\tau e^{s_{t-1}} \varepsilon_{T,t}, \tag{I.1}$$

$$s_t = \rho_\sigma s_{t-1} + \eta_\sigma \varepsilon_{\sigma,t}. \tag{I.2}$$

The innovations $\varepsilon_{T,t}$ and $\varepsilon_{\sigma,t}$ are centered, independent, unit-variance Gaussian variables, independent across dates. In particular, $\mathbb{E}\varepsilon_{T,t}^4 = 3$ and $\mathbb{E}(\varepsilon_{T,t}^2 \varepsilon_{\sigma,t}^2) = 1$; these joint moments matter for a fourth-order welfare calculation. The level experiment sets $\varepsilon_{T,0} = 1$ and $\varepsilon_{\sigma,0} = 0$. Since $\bar{\sigma}_\tau = 0.10$, it gives $\tau_0^F = 0.10$ and $e^{0.10} - 1 = 10.52\%$. The uncertainty experiment instead sets $\varepsilon_{T,0} = 0$ and $\varepsilon_{\sigma,0} = 1$. With $\eta_\sigma = \log 2$, next quarter's conditional standard deviation doubles from 0.10 to 0.20, while $\tau_0^F = 0$. In both experiments, the realized continuation sets $\varepsilon_{T,t} = \varepsilon_{\sigma,t} = 0$ for every $t \geq 1$.

The conditional uncertainty response is the path after the volatility innovation minus the path without that innovation under the same policy, approximation order, physical initial state, and commitment boundary. Realized tariff innovations are zero at every date on both paths. This does not truncate the shock laws: agents continue to integrate over the conditional Gaussian distributions of future tariff and volatility innovations in (I.1)–(I.2). The response is therefore neither a no-risk counterfactual nor an average over future shock histories. Quantities and the pre-tariff terms of trade $q_t = \log(P_{F,t}/P_{H,t})$ are reported in percent; inflation and rates are quarterly percentage-point deviations; net foreign assets (NFA) and the trade balance are percentages of steady output. A negative q_t is a Home real appreciation.

Every Ramsey comparison below states its commitment boundary. The complete-markets

common-state and four-policy comparisons (Figures I.4 and I.5) and the welfare exercise (Section I.5) begin new date-0 commitments and set inherited planner multipliers to zero. The incomplete-markets response comparison (Figure I.7) continues a pre-date-0 plan and retains inherited multipliers. Within each figure, all policies share the same predetermined physical state; fixed rules carry no planner multipliers.

I.2 Primitives, fiscal ledger, and price setting

The representative household has utility and CES consumption

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{L_t^{1+\varphi}}{1+\varphi} \right], \tag{I.3}$$

$$C_t = \left[(1-\alpha)^{1/\gamma} C_{H,t}^{(\gamma-1)/\gamma} + \alpha^{1/\gamma} C_{F,t}^{(\gamma-1)/\gamma} \right]^{\gamma/(\gamma-1)}. \tag{I.4}$$

For $\sigma = 1$, the consumption term in (I.3) is understood as $\log C_t$; for $\gamma = 1$, (I.4) is understood as its Cobb–Douglas limit. Home production is $Y_t = A_t L_t$. Let $P_{F,t}$ be the Home-currency foreign producer price, $Q_t = P_{F,t}/P_{H,t}$ the pre-tariff terms of trade, and $G_t = P_t/P_{H,t}$ the CPI–PPI ratio. CES demand gives

$$G_t^{1-\gamma} = (1-\alpha) + \alpha (T_{F,t} Q_t)^{1-\gamma}, \tag{I.5}$$

$$C_{H,t} = (1-\alpha) G_t^{\gamma} C_t, \qquad C_{F,t} = \alpha \left(\frac{T_{F,t} Q_t}{G_t} \right)^{-\gamma} C_t, \tag{I.6}$$

$$X_t = \alpha Q_t^{\gamma} C_t^*, \qquad Y_t = \frac{C_{H,t} + X_t}{1 - \frac{\kappa}{2} (\Pi_{H,t} - 1)^2}. \tag{I.7}$$

The exact adding-up identity is $G_t C_t = C_{H,t} + T_{F,t} Q_t C_{F,t}$, and

$$\Pi_{C,t} = \Pi_{H,t} \frac{G_t}{G_{t-1}}. \tag{I.8}$$

The labor first-order condition implies raw PPI marginal cost

$$m_t^r \equiv \frac{MC_t}{P_{H,t}} = \chi G_t C_t^{\sigma} Y_t^{\varphi} A_t^{-(1+\varphi)}. \tag{I.9}$$

Define $\bar{\upsilon} = m_{ss}^r = \chi$ and normalized marginal cost

$$mc_t \equiv \frac{m_t^r}{\bar{\upsilon}}. \tag{I.10}$$

The fixed production correction s^p satisfies

$$1 - s^p = \frac{\epsilon - 1}{\epsilon \bar{v}}. \tag{I.11}$$

Its sign is not imposed. The government rebates tariff revenue, finances the production correction, and returns the residual lump sum:

$$T_t^{gov} = (T_{F,t} - 1) P_{F,t} C_{F,t} - s^p M C_t Y_t. \tag{I.12}$$

Neither transfer uses real resources. Substitution of (I.11) into the firm's first-order condition gives the normalized nonlinear NKPC

$$\begin{aligned} \kappa(\Pi_{H,t} - 1)\Pi_{H,t} = (\epsilon - 1)(mc_t - 1) \\ + \beta\kappa \mathbb{E}_t \left[(\Pi_{H,t+1} - 1)\Pi_{H,t+1} \frac{Y_{t+1}}{Y_t} \left(\frac{C_t}{C_{t+1}} \right)^{\sigma} \frac{G_t}{G_{t+1}} \right]. \end{aligned} \tag{I.13}$$

The distinction between m_t^r and mc_t is substantive: the former enters the fiscal correction and the latter enters the NKPC with steady state one.

I.3 Complete markets

Complete state-contingent claims imply

$$C_t = C_t^* \left(\frac{Q_t}{G_t} \right)^{1/\sigma}. \tag{I.14}$$

With $C_t^* = 1$, normalized marginal cost becomes $mc_t = Q_t Y_t^{\varphi} A_t^{-(1+\varphi)}$. The domestic Euler equation and risk sharing imply the exact asset-market restriction

$$1 = \beta R_t e^{\psi_t} \mathbb{E}_t \left[\frac{Q_t}{Q_{t+1} \Pi_{H,t+1}} \right]. \tag{I.15}$$

with $\psi_t = 0$. The auxiliary multiplicative disturbance ψ_t is a reduced-form shift in this uncovered-interest-parity relation; it serves only to isolate the asset-market channel in the two-channel comparison below and is zero in all baseline tariff experiments. Equations (I.13) and (I.15), together with the static block, are an exact two-equation system in $(\Pi_{H,t}, Q_t)$.

The PPI and CPI Taylor rules are, respectively,

$$r_t = \phi_\pi \pi_{H,t} + \phi_y y_t, \tag{I.16}$$

$$r_t = \phi_\pi \pi_{C,t} + \phi_y y_t, \qquad \pi_{C,t} = \pi_{H,t} + \Delta g_t. \tag{I.17}$$

The additional Δg_t term is why the two rules are not interchangeable for an import tariff.

I.3.1 Level and uncertainty coefficients under the PPI rule

For lower-case log deviations,

$$g_q = g_{\tau^F} = \alpha,$$

$$g_{qq} = g_{q\tau^F} = g_{\tau^F\tau^F} = \alpha(1-\alpha)(1-\gamma), \tag{I.18}$$

$$y_t = \theta_q q_t + \theta_\tau^F \tau_t^F + \frac{1}{2}\left(y_{qq} q_t^2 + 2 y_{q\tau} q_t \tau_t^F + y_{\tau\tau}(\tau_t^F)^2 + \kappa \pi_{H,t}^2\right) + O(3), \tag{I.19}$$

$$\theta_q = \alpha\gamma(2-\alpha) + \frac{(1-\alpha)^2}{\sigma},$$

$$\theta_\tau^F = \alpha(1-\alpha)\left(\gamma - \frac{1}{\sigma}\right). \tag{I.20}$$

Unlike θ_τ in the main text, which enters output with a minus sign, θ_τ^F enters with a plus sign, so its sign is that of $\gamma - 1/\sigma$. The condition $\gamma > 1/\sigma$ governs the sign of the decentralized flexible-price response to a tariff-level shock; it is not a constrained-efficiency condition and does not govern a pure uncertainty shock.

Let $\kappa_{mc} = (\epsilon - 1)/\kappa$, $\Gamma_q = 1 + \varphi\theta_q$, and $\widetilde{\kappa} = \kappa_{mc}\Gamma_q$. For any persistence ρ, define the following objects (the letters a, b, d are local to this subsection and differ from a_σ, d_σ in Appendix D.4.2):

$$a(\rho) = 1 - \beta\rho, \qquad b(\rho) = \phi_\pi - \rho, \qquad d(\rho) = 1 - \rho + \phi_y\theta_q,$$

$$M(\rho) = \begin{pmatrix} a(\rho) & -\widetilde{\kappa} \\ b(\rho) & d(\rho) \end{pmatrix}, \qquad D(\rho) = a(\rho)d(\rho) + \widetilde{\kappa} b(\rho). \tag{I.21}$$

Assume the fixed-rule equilibrium is locally determinate and $D(\rho_\tau) \neq 0$. The first-order level coefficients satisfy

$$M(\rho_\tau)\begin{pmatrix} A_\pi \\ A_q \end{pmatrix} = \begin{pmatrix} \kappa_{mc}\varphi\theta_\tau^F \\ -\phi_y\theta_\tau^F \end{pmatrix}, \tag{I.22}$$

and hence

$$A_\pi = \frac{\theta_\tau^F[\kappa_{mc}\varphi d(\rho_\tau) - \widetilde{\kappa}\phi_y]}{D(\rho_\tau)},$$
$$A_q = -\frac{\theta_\tau^F[a(\rho_\tau)\phi_y + \kappa_{mc}\varphi b(\rho_\tau)]}{D(\rho_\tau)}. \tag{I.23}$$

If, in addition, $D(\rho_\tau^2) \neq 0$, the second derivatives (B_π, B_q) solve

$$M(\rho_\tau^2)\begin{pmatrix} B_\pi \\ B_q \end{pmatrix} = \begin{pmatrix} H_\pi \\ H_q \end{pmatrix}, \tag{I.24}$$

where H_π, H_q are the exact Hessian contractions of the NKPC, asset equation, static allocation, and Taylor rule along the first-order tariff direction. Explicitly, with $\mathcal{C}_y = y_{qq}A_q^2 + 2y_{q\tau}A_q + y_{\tau\tau} + \kappa A_\pi^2$, $\mathfrak{m}_1 = \Gamma_q A_q + \varphi\theta_\tau^F$, and $s_1 = (\theta_q - 1)A_q + \theta_\tau^F$,

$$H_\pi = \kappa_{mc}(\varphi\mathcal{C}_y + \mathfrak{m}_1^2) + 2\beta\rho_\tau(\rho_\tau - 1)A_\pi s_1 - 3(1 - \beta\rho_\tau^2)A_\pi^2,$$
$$H_q = -\phi_y\mathcal{C}_y. \tag{I.25}$$

A log-volatility innovation uses these lower-order tariff derivatives. The two analytically derived coefficients in the third-order approximation are

$$u_F^{NK} = \beta\bar{\sigma}_\tau^2\left[B_\pi + 3A_\pi^2 + 2A_\pi s_1\right], \tag{I.26}$$
$$u_F^{AM} = \bar{\sigma}_\tau^2\left[(B_\pi + B_q) - (A_\pi + A_q)^2\right]. \tag{I.27}$$

We call u_F^{NK} the pricing wedge and u_F^{AM} the asset-market wedge; the superscript AM refers to the complete-markets asset equation (I.15). They are the import-tariff counterparts of the main-text wedges u^{NK} and u^{UIP}. For $D(\rho_\sigma) \neq 0$,

$$M(\rho_\sigma)\begin{pmatrix} \chi_\pi \\ \chi_q \end{pmatrix} = \begin{pmatrix} u_F^{NK} \\ u_F^{AM} \end{pmatrix}, \tag{I.28}$$
$$\chi_\pi = \frac{d(\rho_\sigma)u_F^{NK} + \widetilde{\kappa}u_F^{AM}}{D(\rho_\sigma)}, \qquad \chi_q = \frac{a(\rho_\sigma)u_F^{AM} - b(\rho_\sigma)u_F^{NK}}{D(\rho_\sigma)}. \tag{I.29}$$

Thus $\Delta\pi_{H,t} = \eta_\sigma\chi_\pi\rho_\sigma^t$, $\Delta q_t = \eta_\sigma\chi_q\rho_\sigma^t$, and $\Delta y_t = \eta_\sigma\theta_q\chi_q\rho_\sigma^t$ under the PPI Taylor rule. The CPI rule adds the predetermined g_{t-1} state through (I.8), so its path need not be a single geometric decay.

Figure I.1 applies (I.29) to separate the two sources of the import-tariff uncertainty response. At the baseline calibration, the pricing and asset-market wedges both lower PPI

inflation, so they reinforce one another on the nominal side, but their real effects oppose each other: the pricing wedge produces a real depreciation and raises output, whereas the asset-market wedge produces an appreciation and lowers output. The pricing wedge dominates the real response. This sign pattern is the import-tariff counterpart of the channel decomposition in the main text.

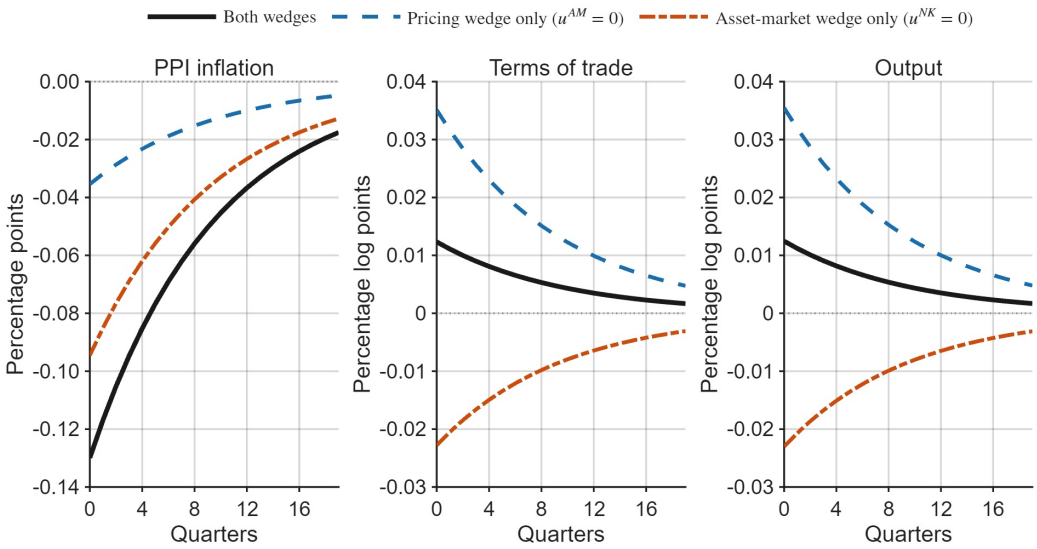


Figure I.1: The two channels of import-tariff uncertainty

Notes: Complete-markets third-order responses to a one-standard-deviation date-0 log-volatility innovation, which doubles the conditional standard deviation of next quarter's tariff innovation, under the PPI Taylor rule. The solid line includes both wedges in (I.28); the dashed lines set one wedge at a time to zero. Realized tariff innovations are zero along the displayed path. Inflation is in quarterly percentage points; the terms of trade and output are percentage log deviations.

Each wedge is itself a sum of economically distinct components. Figure I.2 varies openness and the Taylor-rule feedback coefficients around their baseline values. Price-adjustment convexity tends to offset the other components of the pricing wedge, while mean curvature and inverse-return variance jointly determine the asset-market wedge. Greater inflation feedback attenuates both wedges; greater output feedback generally strengthens them. Openness changes not only the exposure to tariff risk but also the covariance terms, so its effect need not be monotone. The figure therefore cautions against attributing the uncertainty response to a single Jensen term.

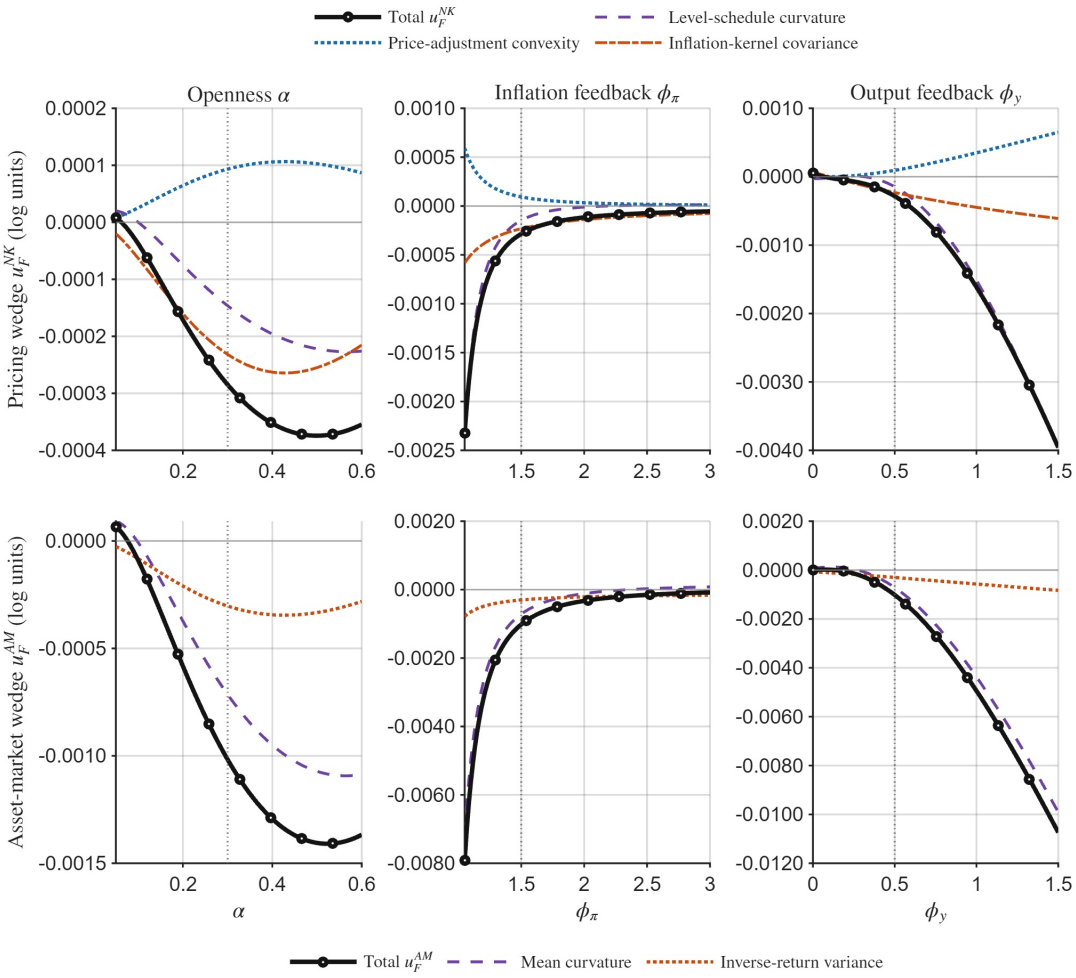


Figure I.2: What determines the import-tariff uncertainty wedges

Notes: The top row decomposes the pricing wedge u_F^{NK}; the bottom row decomposes the asset-market wedge u_F^{AM}. Each column varies one parameter while holding the others at their baseline values, marked by the vertical dotted line. All components are computed from the exact coefficients in (I.26)–(I.27).

Figure I.3 compares the first-order response to a realized import-tariff increase with the third-order response to greater tariff uncertainty. Both shocks lower PPI inflation under the PPI Taylor rule, but the output effects differ sharply in scale. A realized tariff immediately changes relative demand and the terms of trade. A volatility innovation leaves the current tariff unchanged and affects current decisions only through the distribution of future tariffs. Its output response is therefore much smaller and first appears at third order.

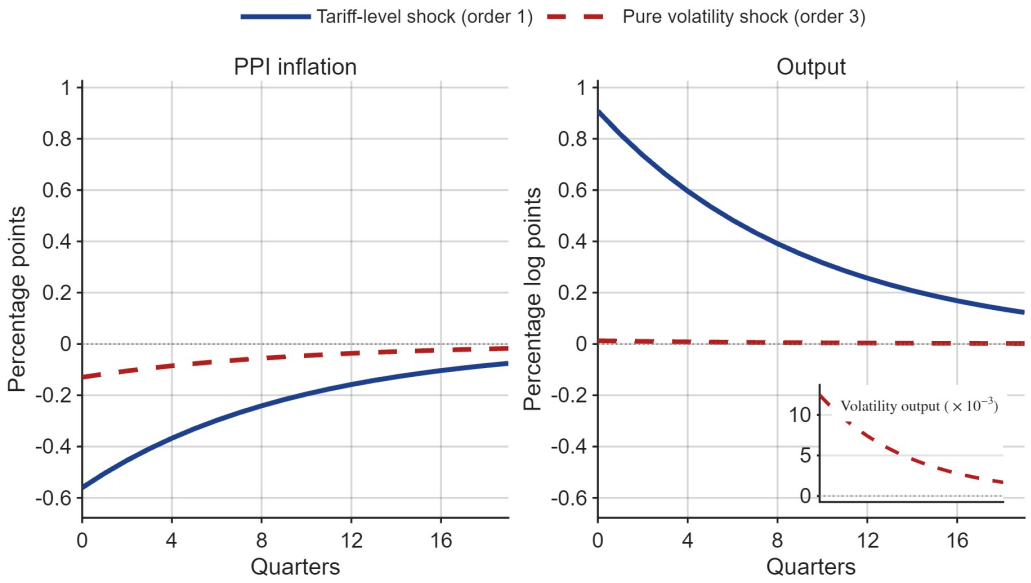


Figure I.3: Realized import tariffs and import-tariff uncertainty

Notes: Complete-markets responses under the PPI Taylor rule. The solid line is the first-order response to a one-standard-deviation tariff-level innovation. The dashed line is the local third-order approximation to the response after a log-volatility innovation that doubles the conditional standard deviation of next quarter's tariff innovation, with subsequent realized tariff innovations set to zero. It is not an exact nonlinear response to a finite variance doubling. Inflation is in quarterly percentage points and output is in percentage log deviations; the inset magnifies the uncertainty response of output.

I.3.2 Flexible prices and constrained efficiency

With adjustment costs removed, the Home planner chooses Q_t subject to the static allocation and risk sharing. Here C_Q and Y_Q denote derivatives of the substituted static consumption and output allocations with respect to Q, holding the tariff state fixed. The planner's exact interior condition is

$$C_t^{-\sigma} C_Q = \chi L_t^{\varphi} \frac{Y_Q}{A_t}. \tag{I.30}$$

The steady-normalized planner wedge is

$$\zeta^{e,F}(q, \tau^F) = \log \left[\frac{G(q, \tau^F) C_Q(q, \tau^F)}{Y_Q(q, \tau^F)} \right] - \log \left[\frac{\bar{G} \bar{C}_Q}{\bar{Y}_Q} \right]. \tag{I.31}$$

Write $\zeta_t^{e,F} = d_q^F q_t + d_\tau^F \tau_t^F + O(2)$ and $\Gamma_F = \varphi\theta_\tau^F$. The decentralized flexible-price allocation and the constrained-efficient allocation are, respectively,

$$q_t^f = -\frac{\Gamma_F}{\Gamma_q}\tau_t^F, \qquad y_t^f = \frac{\theta_\tau^F}{\Gamma_q}\tau_t^F, \tag{I.32}$$

$$q_t^e = \frac{d_\tau^F - \Gamma_F}{\Gamma_q - d_q^F}\tau_t^F, \qquad y_t^e = \theta_q q_t^e + \theta_\tau^F \tau_t^F. \tag{I.33}$$

Define the planner wedge evaluated along the constrained-efficient response by

$$\iota_\zeta^F \equiv d_q^F q_\tau^e + d_\tau^F, \qquad q_\tau^e \equiv \frac{d_\tau^F - \Gamma_F}{\Gamma_q - d_q^F}. \tag{I.34}$$

Direct subtraction of (I.32) and (I.33) gives

$$q_\tau^e - q_\tau^f = \frac{\iota_\zeta^F}{\Gamma_q}, \qquad y_\tau^e - y_\tau^f = \frac{\theta_q}{\Gamma_q}\iota_\zeta^F. \tag{I.35}$$

Thus ι_ζ^F measures the first-order conflict between the allocation implemented by flexible prices and the allocation preferred by the planner. Equation (I.35) also gives the sign interpretation: when $\iota_\zeta^F > 0$, constrained-efficient output responds more strongly to an import-tariff increase; when it is negative, the ranking is reversed. The two allocations coincide when the production correction neutralizes the planner wedge in every reachable tariff state. One such case is $\gamma = 1/\sigma$, for which $\theta_\tau^F = d_q^F = d_\tau^F = \iota_\zeta^F = 0$.

I.3.3 Complete-markets Ramsey policy

After substituting the static block, let $N_t = 0$ denote the exact NKPC and $A_t^R = 0$ the exact asset equation. Ramsey maximizes substituted utility $\widetilde{U}_t$ with Lagrangian

$$\mathcal{L} = \mathbb{E}_0 \sum_{t\geq 0} \beta^t [\widetilde{U}_t + \lambda_t N_t + \mu_t A_t^R]. \tag{I.36}$$

We maintain that the exact objective and constraints are four times continuously differentiable, the constraint derivative has full row rank, and the stable-root equations together with the stated commitment boundary have a nonsingular coefficient matrix. These conditions select a locally unique branch of the Ramsey KKT conditions. Calling that branch a local Ramsey optimum additionally requires the uniform second-order condition of Appendix F.4, applied to the present constraints, with the Taylor remainder uniformly $o(\|\Delta z\|_\beta^2)$ in the discounted sequence norm. The numerical calculations select the KKT branch and do not establish

a global optimum. For brevity, the text and figure legends call this branch "Ramsey." Stationarity with respect to the nominal rate gives $\mu_t(A_t^R)_R = 0$; since $(A_t^R)_R \neq 0$, $\mu_t = 0$. The rate therefore implements the chosen real allocation and is recovered from (I.15).

Let $Y_t^a \equiv C_{H,t} + X_t$ denote output available for private absorption before Rotemberg resource costs, and let $Y_t^{a,e}$ denote the same object at the constrained-efficient allocation. The welfare-relevant gap is $\widetilde{y}_t = \log Y_t^a - \log Y_t^{a,e}$. The second-order loss and linearized NKPC are

$$\frac{1}{2}\left(\lambda_\pi \pi_{H,t}^2 + \lambda_y \widetilde{y}_t^2\right), \qquad \lambda_\pi = \kappa, \quad \lambda_y = \frac{\Gamma_q - d_q^F}{\theta_q}, \tag{I.37}$$

$$\pi_{H,t} = \beta \mathbb{E}_t \pi_{H,t+1} + \kappa_y \widetilde{y}_t + \kappa_{mc} \iota_\zeta^F \tau_t^F, \quad \kappa_y = \kappa_{mc} \frac{\Gamma_q}{\theta_q}. \tag{I.38}$$

Under commitment, eliminating the NKPC multiplier gives the first-order target criterion

$$\lambda_\pi \pi_{H,t} + \frac{\lambda_y}{\kappa_y}(\widetilde{y}_t - \widetilde{y}_{t-1}) = 0. \tag{I.39}$$

This target is history dependent. It is exact within the linear-quadratic approximation, but it does not imply that every higher-order term in the full nonlinear Ramsey conditions vanishes.

Define

$$\vartheta = \frac{\lambda_y}{\lambda_\pi \kappa_y}, \qquad \chi_y = \frac{\lambda_\pi \kappa_y^2}{\lambda_y},$$
$$\rho_y = \frac{1 + \beta + \chi_y - \sqrt{(1+\beta+\chi_y)^2 - 4\beta}}{2\beta}. \tag{I.40}$$

The stable level solution is the convolution

$$\widetilde{y}_t^R = b_\zeta \iota_\zeta^F \tau_0^F \frac{\rho_\tau^{t+1} - \rho_y^{t+1}}{\rho_\tau - \rho_y}, \qquad b_\zeta = -\frac{(\lambda_\pi \kappa_y / \lambda_y)\kappa_{mc}}{\rho_y^{-1} - \beta \rho_\tau}, \tag{I.41}$$

$$\pi_{H,t}^R = -\vartheta(\widetilde{y}_t^R - \widetilde{y}_{t-1}^R), \qquad q_t^R = q_t^e + \widetilde{y}_t^R / \theta_q. \tag{I.42}$$

The denominators require $\rho_\tau \neq \rho_y$ and $\rho_y^{-1} \neq \beta \rho_\tau$; coincident-root cases use the continuous limits.

For curvature and pure risk, define the exact reduced Karush–Kuhn–Tucker (KKT) residual in $w_t = (\pi_{H,t}, q_t, \lambda_t)'$ by

$$\mathcal{R}_t = \begin{pmatrix} N_t \\ U_{z,t} + \lambda_t N_{z,t} + \beta^{-1} \lambda_{t-1} N_{z^+,t-1} \end{pmatrix} = 0, \qquad z_t = (\pi_{H,t}, q_t)'. \tag{I.43}$$

Let J_-, J_0, J_+ be its lag, current, and lead Jacobians and $K(r) = r^{-1} J_- + J_0 + r J_+$. For

an exogenous component proportional to r^t, let d_r be the coefficient on r^t in the first-order residual after the terms involving w_t have been moved to the left-hand side. Let $g_r^{[2]}$ be the corresponding second-order residual coefficient generated by the Hessian terms evaluated on the first-order solution. When r does not coincide with an endogenous KKT root, a nonsingular forced root satisfies

$$\begin{aligned} K(r)h_r &= -d_r, & h_r &= -K(r)^{-1}d_r, & &\det K(r) \neq 0, \\ K(r)h_r^{[2]} &= -g_r^{[2]}, & h_r^{[2]} &= -K(r)^{-1}g_r^{[2]}, & &\det K(r) \neq 0. \end{aligned} \tag{I.44}$$

The endogenous KKT roots are homogeneous components, not forced roots: they satisfy $\det K(r) = 0$, and if v_r is a right-null vector, they contribute $a_r v_r r^t$ to w_t. The applicable multiplier boundary determines a_r jointly with the other homogeneous coefficients: a new date-0 commitment sets the lagged multiplier to zero, whereas an inherited continuation carries its predetermined value into date 0. Denote the resulting first- and second-order tariff derivatives by (A_π^R, A_q^R) and (B_π^R, B_q^R). Define $s_1^R \equiv (\theta_q - 1)A_q^R + \theta_\tau^F$ as the first-order tariff loading of the exact NKPC continuation kernel along this Ramsey path. The direct term is $+\theta_\tau^F$ because output in (I.19) loads on $+\theta_\tau^F \tau_t^F$, matching the earlier fixed-rule definition of s_1. These objects imply

$$u_{NK}^R = \beta\bar{\sigma}_\tau^2[B_\pi^R + 3(A_\pi^R)^2 + 2A_\pi^R s_1^R], \tag{I.45}$$

$$u_{AM}^R = \bar{\sigma}_\tau^2[(B_\pi^R + B_q^R) - (A_\pi^R + A_q^R)^2]. \tag{I.46}$$

As in the fixed-rule case, u_{NK}^R and u_{AM}^R are the pricing and asset-market wedges, now evaluated along the Ramsey path. The first-order target criterion in (I.39) need not hold automatically for the third-order response. To measure the additional nonlinear terms, define the target residual on the uncertainty path by

$$\mathcal{E}_t^V \equiv \lambda_\pi \pi_{H,t}^{R,V} + \frac{\lambda_y}{\kappa_y}(\widetilde{y}_t^{R,V} - \widetilde{y}_{t-1}^{R,V}). \tag{I.47}$$

The numerical solution of the full third-order Ramsey conditions does not impose $\mathcal{E}_t^V = 0$.

Setting this residual to zero instead yields a transparent analytical benchmark:

$$\widetilde{y}_t^{R,V} = b_v u_{NK}^R \eta_\sigma \frac{\rho_\sigma^{t+1} - \rho_y^{t+1}}{\rho_\sigma - \rho_y}, \qquad b_v = -\frac{\lambda_\pi \kappa_y / \lambda_y}{\rho_y^{-1} - \beta \rho_\sigma}, \tag{I.48}$$

$$\pi_{H,t}^{R,V} = -\vartheta(\widetilde{y}_t^{R,V} - \widetilde{y}_{t-1}^{R,V}), \qquad q_t^{R,V} = \widetilde{y}_t^{R,V}/\theta_q,$$

$$r_t^{R,V} = \frac{\widetilde{y}_{t+1}^{R,V} - \widetilde{y}_t^{R,V}}{\theta_q} + \pi_{H,t+1}^{R,V} + u_{AM}^R \eta_\sigma \rho_\sigma^t. \tag{I.49}$$

The formula is therefore a benchmark, not an exact general closed form for nonlinear Ramsey policy. Its value is transparency: the same stable policy root that governs the tariff-level problem propagates the third-order pricing wedge, while the asset-market wedge determines part of the implementing interest-rate response. The numerical Ramsey branch used in the figures retains $\mathcal{E}_t^V$ and all other third-order terms; the benchmark illustrates the main propagation mechanism and does not supply an error bound.

Common-state policy comparison. Figure I.4 compares a new Ramsey commitment, strict PPI stabilization, and the PPI-based Taylor rule from the same predetermined physical state at date -1. For Ramsey, all lagged planner multipliers are set to zero; the other policies are fixed rules without multiplier states. Each curve subtracts that policy's no-volatility path from the same initial physical state.

Because the tariff itself remains at its steady-state value, the figure isolates adjustment to expected tariff risk. New-commitment Ramsey permits only small movements in PPI inflation, output, and the terms of trade. Strict PPI sets all three responses exactly to zero but implements that allocation with a temporary increase in the policy rate, opposite to Ramsey's small rate cut. The Taylor rule produces a much larger initial deflation and interest-rate cut. Strict PPI is therefore close to Ramsey in inflation and the plotted real variables, but not in the implementing interest-rate response.

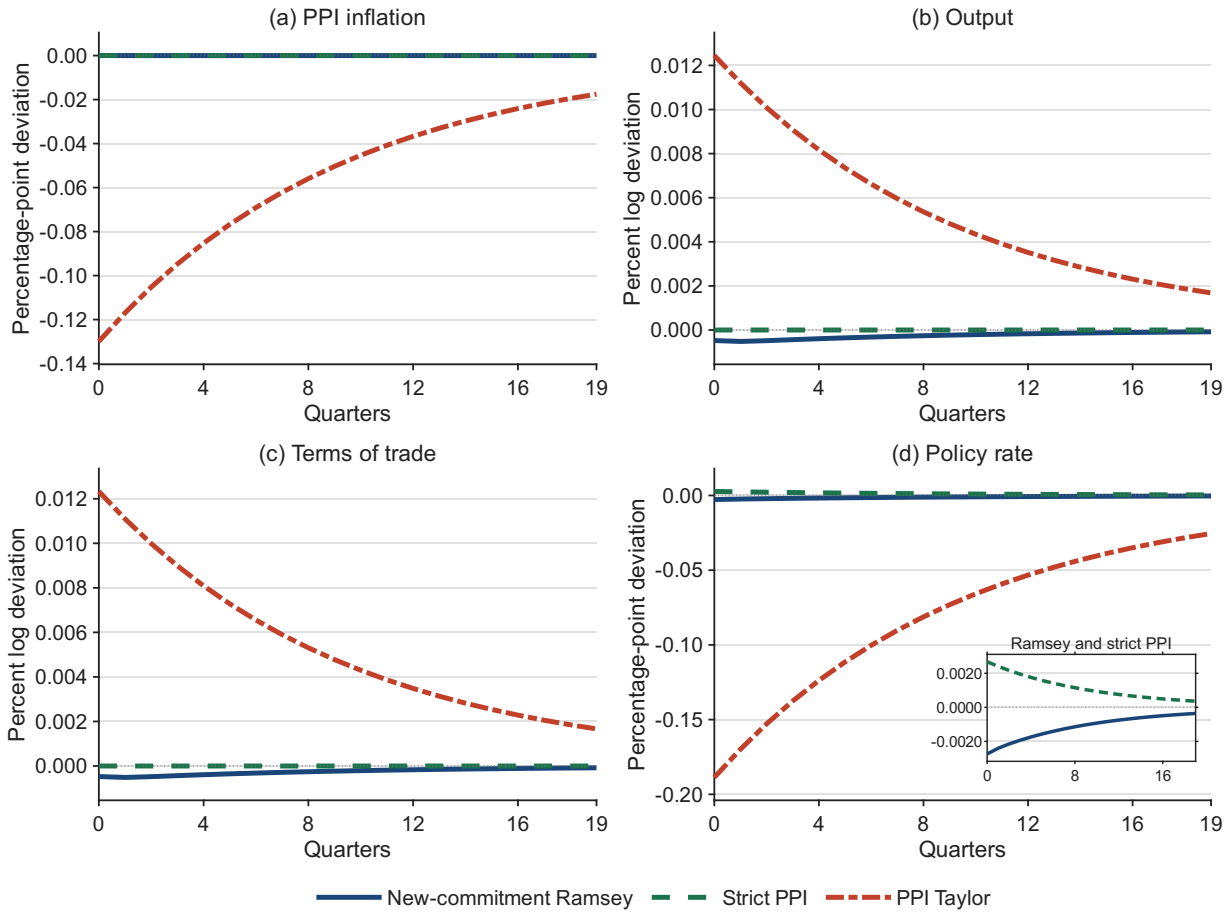


Figure I.4: Monetary policy under complete markets with import tariffs

Notes: Each panel reports the local third-order response to a one-standard-deviation date-0 innovation in log import-tariff volatility, which doubles the conditional standard deviation of next quarter's tariff innovation. The tariff remains at its deterministic steady-state level and realized tariff innovations are zero at all displayed dates; future tariff risk remains in expectations. All policies start from the same predetermined physical state at date -1, and each curve subtracts that policy's no-volatility path. New-commitment Ramsey sets all lagged planner multipliers to zero. The inset in panel (d) magnifies the Ramsey and strict-PPI policy-rate paths. Output and the terms of trade are percent log deviations; PPI inflation and the quarterly policy rate are percentage-point deviations.

Two restricted date-0 plans. The first restricted plan chooses the optimal date-0 response when policy must obey the linear-quadratic target criterion (I.39) from date 1 onward; it is the import-tariff counterpart of the main text's Level-Ramsey continuation. Write

$$v_t \equiv \frac{1}{2}\left[\mathrm{Var}_t(\tau^F_{t+1} \mid s_t) - \mathrm{Var}_t(\tau^F_{t+1} \mid s_t = 0)\right] = \frac{\bar{\sigma}^2_\tau}{2}(e^{2s_t} - 1). \tag{I.50}$$

Thus $v_t = \bar{\sigma}^2_\tau s_t + O(4)$ under the common perturbation scaling, and a unit date-0 volatility innovation gives $v_0 = \bar{\sigma}^2_\tau \eta_\sigma + O(4)$. This one-half-variance normalization makes a conditional-variance correction equal to v_t times the relevant second directional derivative.

To calculate that derivative, let δ be a date-1 log-tariff realization about the deterministic state and write the Level-Ramsey continuation as

$$\pi_{H,1}^{LQ}(\delta) = A_\pi^{LQ}\delta + \frac{1}{2}B_\pi^{LQ}\delta^2 + O(\delta^3),$$
$$q_1^{LQ}(\delta) = A_q^{LQ}\delta + O(\delta^2), \qquad s_1^{LQ} \equiv (\theta_q - 1)A_q^{LQ} + \theta_\tau^F. \tag{I.51}$$

The form of s_1^{LQ} follows from complete-market risk sharing: the log of the continuation kernel in the NKPC is $y_1 - q_1$ when the date-0 physical state is held at its deterministic value. Define

$$\mathcal{P}_1^{LQ}(\delta) \equiv [e^{\pi_{H,1}^{LQ}(\delta)} - 1]e^{\pi_{H,1}^{LQ}(\delta)}\frac{Y_1^{LQ}(\delta)}{Y_0}\left(\frac{C_0}{C_1^{LQ}(\delta)}\right)^\sigma \frac{G_0}{G_1^{LQ}(\delta)}. \tag{I.52}$$

After expected log inflation is separated from the exact pricing kernel, the remaining coefficient is the explicit derivative contraction

$$\begin{aligned}\bar{z}_{NK}^{LQ} &\equiv \beta\,\frac{d^2}{d\delta^2}[\mathcal{P}_1^{LQ}(\delta) - \pi_{H,1}^{LQ}(\delta)]\Big|_{\delta=0} \\ &= \beta\left[3(A_\pi^{LQ})^2 + 2A_\pi^{LQ}s_1^{LQ}\right].\end{aligned} \tag{I.53}$$

Indeed, the second derivative of $\mathcal{P}_1^{LQ}$ is $B_\pi^{LQ} + 3(A_\pi^{LQ})^2 + 2A_\pi^{LQ}s_1^{LQ}$, while B_π^{LQ} is already contained in the curvature contribution to $\Delta^V\mathbb{E}_0\pi_{H,1}^{LQ}$. Their sum therefore reproduces the bracket in (I.45), with the Level-Ramsey rather than the unrestricted-KKT level coefficients.

Under the common scaling parameter ε (the main-text δ; unrelated to the demand elasticity ϵ), define the coefficient on ε^3 by

$$\check{v}_0 \equiv [\varepsilon^3]v_0 = \bar{\sigma}_\tau^2\eta_\sigma, \qquad v_0(\varepsilon) = \varepsilon^3\check{v}_0 + O(\varepsilon^4).$$

Thus $\check{v}_0$ is the third-order coefficient, not the third derivative. Using this coefficient, write

$$\mathcal{Z}_0^F \equiv \beta\Delta^V\mathbb{E}_0\pi_{H,1}^{LQ} + z_{NK,0}^{LQ,V}, \qquad z_{NK,0}^{LQ,V} = \check{v}_0\bar{z}_{NK}^{LQ}, \tag{I.54}$$

where Δ^V denotes the response to the date-0 volatility innovation and $\bar{z}_{NK}^{LQ}$ is given exactly by (I.53). Thus $\mathcal{Z}_0^F$ combines the change in expected date-1 inflation with the remaining covariance terms in the nonlinear pricing kernel. It is not a free policy disturbance. Let ϕ_t denote the multiplier on the linearized Phillips curve. The LQ output-gap first-order condition gives $\phi_t = (\lambda_y/\kappa_y)\widetilde{y}_t$. Because commitment begins at date 0, its lower-endpoint condition is $\phi_{-1} = 0$, which we write as $\widetilde{y}_{-1} \equiv (\kappa_y/\lambda_y)\phi_{-1} = 0$ in the difference-form target criterion; as in Appendix F.8, this is a boundary convention, not a realized output gap at

date -1. The Phillips curve and the planner's date-0 condition are therefore

$$\pi_{H,0} - \kappa_y \widetilde{y}_0 = \mathcal{Z}_0^F, \qquad \lambda_\pi \pi_{H,0} + \frac{\lambda_y}{\kappa_y} \widetilde{y}_0 = 0. \tag{I.55}$$

Solving these two equations gives

$$\pi_{H,0}^{LQ} = \frac{\mathcal{Z}_0^F}{1+\chi_y}, \qquad \kappa_y \widetilde{y}_0^{LQ} = -\frac{\chi_y}{1+\chi_y} \mathcal{Z}_0^F, \qquad \chi_y = \frac{\lambda_\pi \kappa_y^2}{\lambda_y}. \tag{I.56}$$

Within the quadratic welfare approximation, this is the optimal date-0 response conditional on imposing the Level-Ramsey targeting criterion from date 1 onward. The criterion is fixed, not the continuation path: the allocation is re-solved after each future tariff history. Full Ramsey instead solves the complete third-order KKT system without imposing that LQ continuation criterion, so its higher-order continuation coefficients need not coincide with those implied by the LQ rule.

The second restricted plan sets only the incremental date-0 PPI-inflation response to zero and then honors the commitment selected under that restriction. To see why the continuation must change, let p_v be the particular loading at the volatility root ρ_σ, let w_R be the stable homogeneous Ramsey component associated with root ρ_R, and let e_π select PPI inflation from the KKT vector in (I.43). Any local third-order response has the form

$$h_t^V = \check{v}_0 [p_v \rho_\sigma^t + a_R \, w_R \rho_R^t]. \tag{I.57}$$

Both full Ramsey and the restricted plan begin a new commitment and impose the same lower-endpoint condition, $\lambda_{-1} = 0$, in the exact KKT notation of (I.43). Let $\mathcal{S}_0^\pi$ and $\mathcal{S}_0^q$ denote its date-0 stationarity rows for PPI inflation and the terms of trade. With inflation free, full Ramsey uses $\mathcal{S}_0^\pi = 0$ to select a_R. The restricted plan adds a multiplier ξ_0 on the incremental-inflation constraint

$$e_\pi' h_0^V = \Delta^V \pi_{H,0} = 0, \qquad a_R = -\frac{e_\pi' p_v}{e_\pi' w_R}, \qquad e_\pi' w_R \neq 0. \tag{I.58}$$

At date 0, it retains the Phillips curve $N_0 = 0$, terms-of-trade stationarity $\mathcal{S}_0^q = 0$, and $\lambda_{-1} = 0$. In the reduced allocation system, the incremental-PPI restriction replaces $\mathcal{S}_0^\pi = 0$ as the equation that selects a_R; in the full constrained KKT system, inflation stationarity is $\mathcal{S}_0^\pi + \xi_0 = 0$. From date 1 onward, the plan satisfies all unrestricted Ramsey first-order conditions. The new loading determines λ_0, which is carried into that continuation. Appending zero impact inflation to the original Ramsey continuation would therefore violate the date-0 Phillips curve; the restricted plan must be solved as a whole.

Figure I.5 compares the two restricted plans with full Ramsey and strict PPI. Imposing zero incremental PPI inflation produces a slightly larger date-0 decline in output and the terms of trade than full Ramsey, but its real and interest-rate paths converge rapidly thereafter. Thus Ramsey's small real response does not depend on its positive, but quantitatively tiny, impact inflation response. The LQ experiment generates a larger contraction because the date-1-onward response is governed by the Level-Ramsey criterion rather than reoptimized at third order. Strict PPI sets inflation to zero at every date and, under complete markets, eliminates the output and terms-of-trade responses. Its implementing interest-rate response has the opposite sign from Ramsey's. These are comparisons of conditional responses, not welfare rankings.

I.4 Incomplete markets: rebated bond economy

The household now trades domestic and foreign one-period bonds. Let b_t be end-of-period real net foreign assets, with $\bar{b} = 0$. The exact domestic and foreign Euler equations are

$$C_t^{-\sigma} = \beta R_t \mathbb{E}_t \left[\frac{C_{t+1}^{-\sigma}}{\Pi_{C,t+1}} \right], \tag{I.59}$$

$$1 = \beta R^* e^{-\psi_b (b_t - \bar{b}) + \psi_t} \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{Q_{t+1}}{Q_t} \frac{G_t}{G_{t+1}} \right]. \tag{I.60}$$

The return Home households face on foreign bonds includes a borrowing premium. Because $\psi_b > 0$, greater foreign borrowing (a lower b_t) raises this premium and prevents net foreign assets from drifting without bound. The government rebate implies the exact trade balance

$$\begin{aligned} TB_t &= \frac{X_t - Q_t C_{F,t}}{G_t} \\ &= \frac{C_{H,t} + X_t}{G_t} - C_t + (T_{F,t} - 1) \frac{Q_t C_{F,t}}{G_t}. \end{aligned} \tag{I.61}$$

The last term records tariff revenue returned to the household. The rebate is part of household income and therefore of external accounting. A model that omits it instead treats tariff revenue as a resource loss and lowers the direct first-order tariff coefficient by α. NFA evolves as

$$b_t = R^* e^{-\psi_b (b_{t-1} - \bar{b}) + \psi_{t-1}} \frac{Q_t}{Q_{t-1}} \frac{G_{t-1}}{G_t} b_{t-1} + TB_t. \tag{I.62}$$

After the real allocation is chosen, the domestic Euler equation recovers the implementing nominal rate. Real non-neutrality is instead governed by (I.60) and (I.62).

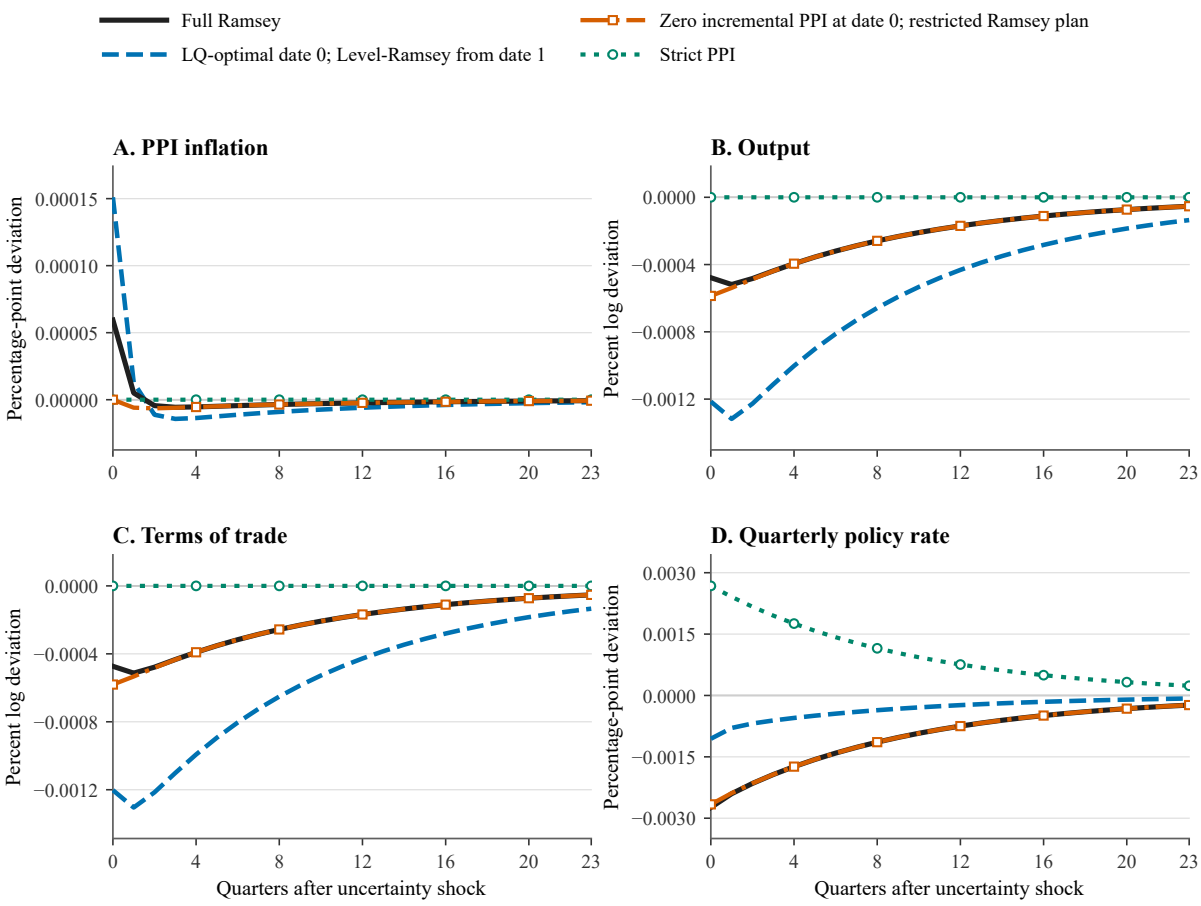


Figure I.5: Ramsey, strict PPI, and two restricted date-0 plans under import tariffs

Notes: Each panel reports the local third-order response to a one-standard-deviation date-0 innovation in log import-tariff volatility, which doubles the conditional standard deviation of next quarter's tariff innovation. The tariff remains at its deterministic steady-state level and realized tariff innovations are zero at all displayed dates; future risk remains in expectations. All policies start from the same predetermined physical state at date -1, and each curve subtracts its own no-volatility path. Full Ramsey and the date-0-PPI restricted plan impose the same new-commitment boundary $\lambda_{-1} = 0$; the LQ experiment uses the corresponding boundary $\phi_{-1} = 0$. The date-0-PPI plan re-solves the entire continuation under its impact restriction. Strict PPI is a fixed rule with no planner multiplier. Output and the terms of trade are percent log deviations; PPI inflation and the quarterly policy rate are percentage-point deviations.

At first order,

$$
\begin{aligned}
g_t &= \alpha(q_t + \tau_t^F), \\
y_t &= (1-\alpha)c_t + \alpha\gamma(2-\alpha)q_t + \alpha(1-\alpha)\gamma\tau_t^F, && \text{(I.63)} \\
tb_t &= -\alpha c_t + \alpha[\gamma(2-\alpha)-1]q_t + \alpha(1-\alpha)\gamma\tau_t^F. && \text{(I.64)}
\end{aligned}
$$

The real dynamic block is

$$
\begin{aligned}
\pi_{H,t} &= \kappa_{mc}[m_c c_t + m_q q_t + m_\tau \tau_t^F] + \beta\mathbb{E}_t\pi_{H,t+1}, && \text{(I.65)} \\
0 &= -\psi_b b_t + \sigma c_t - (1-\alpha)q_t + \alpha\tau_t^F + \mathbb{E}_t[-\sigma c_{t+1} + (1-\alpha)q_{t+1} - \alpha\tau_{t+1}^F], && \text{(I.66)} \\
b_t &= \beta^{-1}b_{t-1} + tb_t, && \text{(I.67)}
\end{aligned}
$$

where $m_c = \sigma + \varphi(1-\alpha)$, $m_q = \alpha + \varphi\alpha\gamma(2-\alpha)$, and $m_\tau = \alpha + \varphi\alpha(1-\alpha)\gamma$.

Table I.1 compares impact responses to the realized level experiment, $\tau_0^F = 0.10$, across the two financial structures. Under the PPI-based Taylor rule, incomplete markets reverse the sign of the PPI-inflation and policy-rate responses and essentially eliminate the output expansion. With complete markets, output rises by 0.94 percent and PPI inflation falls by 0.58 percentage point; with incomplete markets, output falls slightly, by 0.04 percent, and PPI inflation rises by 0.40 percentage point. The two economies differ because the Foreign Euler equation and NFA accumulation in (I.66)–(I.67) replace the risk-sharing condition (I.14) as the restrictions that determine consumption. This changes how the tariff enters marginal cost (I.9). Because labor supply depends on the consumption wage, marginal cost in PPI units rises with the CPI–PPI ratio G_t, which the tariff raises directly. Under complete markets, risk sharing with $C_t^* = 1$ gives $G_t C_t^\sigma = Q_t$, so consumption offsets G_t exactly and the ratio drops out of marginal cost. Under incomplete markets, no condition ties consumption to Q_t/G_t, so the tariff's direct effect on G_t passes into marginal cost: it is the first term, α, of m_τ in (I.65). This cost push, absent under complete markets, helps explain why, under the PPI-based Taylor rule, PPI inflation rises with incomplete markets but falls with complete markets.

The Ramsey rows show that the optimal response to a realized import tariff also departs from strict PPI targeting, in a direction that depends on the financial structure. Under complete markets, Ramsey policy is more expansionary than strict PPI: it accepts 0.10 percentage point of PPI inflation, lowers the policy rate by 0.06 percentage point where strict PPI raises it by 0.10, and raises output by 1.24 rather than 1.04 percent. Strict PPI targeting essentially reproduces the flexible-price allocation, but the constrained-efficient allocation calls for a larger output expansion: at the baseline calibration, the tariff increase in the

table would raise efficient output by about 2.1 percent, twice the 1.04 percent flexible-price response (the gap is given by (I.35)). Ramsey policy therefore trades a small amount of PPI inflation for moving output part of the way toward the constrained-efficient response. Under incomplete markets, Ramsey policy is more contractionary: PPI inflation falls by 0.06 percentage point, the policy rate rises by 0.22 rather than 0.10 percentage point, and output falls by 0.04 percent where it rises by 0.09 percent under strict PPI. When the level innovation coincides with an uncertainty innovation, the joint rows show that Ramsey's response is close to its level response under complete markets. Under incomplete markets, the PPI-based Taylor rule's level and uncertainty components offset, so its joint PPI response (-0.08 percentage point) lies close to Ramsey's (-0.10). Its policy-rate response (0.07 against 0.27 percentage point) and its welfare gap (Figure I.8) remain far from Ramsey's. A similar net inflation response therefore does not indicate that the rule is close to optimal.

Table I.1: Impact responses to an import-tariff level shock, alone and with an uncertainty shock

	Complete markets				Incomplete markets			
Policy	y	π_H	q	r	y	π_H	q	r
Level innovation								
Ramsey	1.24	0.10	−0.83	−0.06	−0.04	−0.06	−3.18	0.22
Strict PPI	1.04	0.00	−1.04	0.10	0.09	0.00	−3.05	0.10
PPI Taylor rule	0.94	−0.58	−1.22	−0.40	−0.04	0.40	−3.20	0.59
CPI Taylor rule	−0.57	−0.93	−2.89	1.48	−1.35	0.30	−4.52	2.21
Level and uncertainty innovations together								
Ramsey	1.24	0.10	−0.83	−0.06	0.24	−0.10	−2.52	0.27
Strict PPI	1.04	0.00	−1.04	0.11	0.45	0.00	−2.31	0.10
PPI Taylor rule	0.96	−0.71	−1.21	−0.59	0.38	−0.08	−2.36	0.07
CPI Taylor rule	−0.63	−1.09	−2.94	1.19	−1.03	−0.41	−3.75	1.66

Notes: Entries are order-three impact responses, measured relative to the path without the innovation under the same policy. The level experiment sets $\varepsilon_{T,0} = 1$ and $\varepsilon_{\sigma,0} = 0$, so that $\tau_0^F = 0.10$; the joint experiment sets $(\varepsilon_{T,0}, \varepsilon_{\sigma,0}) = (1, 1)$, where the volatility innovation doubles the conditional standard deviation of next quarter's tariff innovation. In the retained order-three solution the joint response equals the sum of the level and uncertainty responses. Output y and the terms of trade q are percent log deviations; PPI inflation π_H and the policy rate r are quarterly percentage-point deviations. Under complete markets, the new-commitment and inherited Ramsey boundaries give the same responses. Under incomplete markets, the Ramsey responses continue a plan chosen before date 0 and retain inherited planner multipliers, as in Figure I.7; the welfare comparison in Section I.5 instead begins a new date-0 commitment.

Figure I.6 holds the PPI-based Taylor rule coefficients, tariff process, common structural parameters, and predetermined physical state fixed while replacing complete international risk sharing with a noncontingent-bond economy. The bond economy also includes the

foreign-bond premium that keeps NFA from drifting. Without state-contingent insurance, greater tariff risk induces precautionary saving: consumption and import demand initially fall, demand shifts toward Home goods, and NFA rises. The larger movements in consumption and relative prices enter marginal cost through (I.65), producing a larger initial deflation. As the saving response unwinds, consumption, output, inflation, and the terms of trade reverse direction. Market incompleteness therefore changes both the size and the persistence of import-tariff uncertainty transmission even though the monetary rule is held fixed.

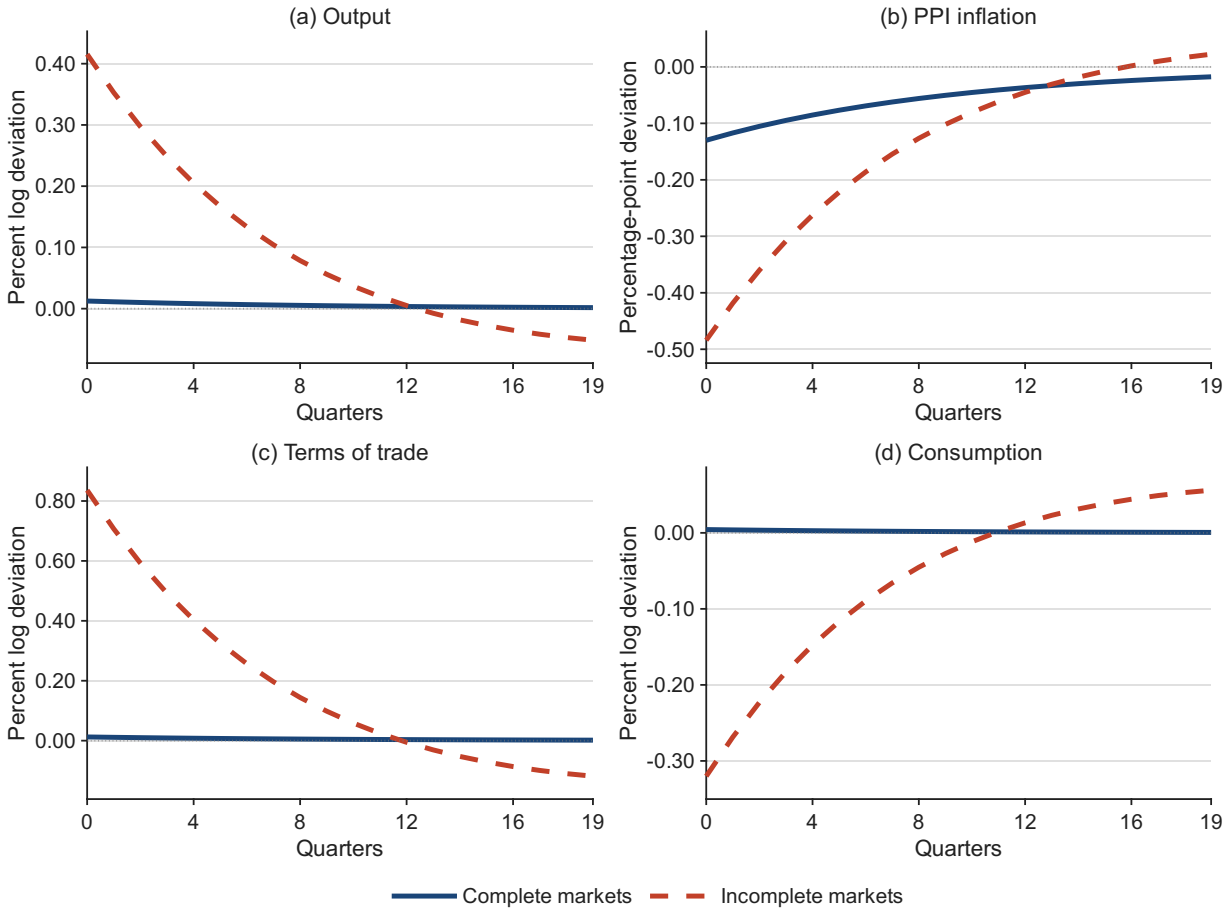


Figure I.6: Import-tariff uncertainty and market completeness under the Taylor rule

Notes: Each panel reports the local third-order response to a one-standard-deviation date-0 innovation in log import-tariff volatility, which doubles the conditional standard deviation of next quarter's tariff innovation. The comparison uses the same Taylor-rule coefficients, tariff process, and values for all parameters shared by the two economies. Both start from the same predetermined physical state at date -1; the bond economy starts with zero NFA and includes the foreign-bond premium. The tariff remains at its deterministic steady-state level and realized tariff innovations are zero at all displayed dates; future risk remains in expectations. Each curve subtracts the corresponding no-volatility path. Inflation is in quarterly percentage points; the other variables are percent log deviations.

I.4.1 Strict PPI and the real external block

Strict PPI sets $\pi_{H,t} = 0$, so the NKPC implies $c_t = -a_q q_t - a_\tau \tau_t^F$, with $a_q = m_q/m_c$ and $a_\tau = m_\tau/m_c$. Define

$$\begin{aligned} d &= \sigma a_q + 1 - \alpha, & e &= \sigma a_\tau - \alpha, \\ \ell_q &= \alpha[\gamma(2-\alpha) - 1] + \alpha a_q, & \ell_\tau &= \alpha(1-\alpha)\gamma + \alpha a_\tau. \end{aligned}$$

Eliminating NFA gives

$$q_{t+1} - \left(1 + R^* + \frac{\psi_b \ell_q}{d}\right) q_t + R^* q_{t-1} = K_\tau \tau_t^F, \tag{I.68}$$

where

$$K_\tau = \frac{\psi_b \ell_\tau}{d} + \frac{e}{d}(1-\rho_\tau)\left(1 - \frac{R^*}{\rho_\tau}\right). \tag{I.69}$$

The first term transmits the tariff through the trade balance and NFA; the second transmits the expected change in the tariff through the Foreign Euler equation. Using the steady-state relation $\beta R^* = 1$, the homogeneous roots are

$$\lambda_{s,u} = \frac{1 + R^* + \psi_b \ell_q/d \mp \sqrt{(1 + R^* + \psi_b \ell_q/d)^2 - 4R^*}}{2}. \tag{I.70}$$

For the stable solution, assume these roots are real and satisfy $|\lambda_s| < 1 < |\lambda_u|$. Excluding explosive paths removes the unstable component. Also assume $\rho_\tau \neq 0$ and that ρ_τ does not coincide with either root, so the denominator below is nonzero. A particular solution proportional to the tariff state and the remaining stable component have coefficients

$$\begin{aligned} A_\tau &= \frac{K_\tau}{\rho_\tau - (1 + R^* + \psi_b \ell_q/d) + R^*/\rho_\tau}, \\ C_s &= \frac{\ell_q A_\tau + \ell_\tau - (d/\psi_b) A_\tau(\rho_\tau - 1) - (e/\psi_b)(\rho_\tau - 1)}{(d/\psi_b)(\lambda_s - 1) - \ell_q}. \end{aligned} \tag{I.71}$$

The coefficient A_τ follows the persistent tariff state, while C_s enforces the initial asset position $b_{-1} = 0$. Hence

$$q_t^{PPI} = \tau_0^F [A_\tau \rho_\tau^t + C_s \lambda_s^t], \qquad c_t^{PPI} = -a_q q_t^{PPI} - a_\tau \tau_t^F. \tag{I.72}$$

The formula reproduces the first-order strict-PPI paths implied by the rebated trade-balance identity. It also shows why PPI stabilization need not leave the real allocation unchanged: the Foreign Euler equation and NFA accumulation still determine consumption and the terms

of trade.

For pure risk, use the half-variance state in (I.50) and let δ denote the date-$t+1$ log-tariff innovation. Around the deterministic date-t physical state, write $j(\delta) = j(0) + \xi_j \delta + B_j \delta^2/2 + O(\delta^3)$ for $j \in \{c, q\}$. Thus $\xi_z = (\xi_c, \xi_q)'$ and $B_z = (B_c, B_q)'$ are derivatives with respect to the realized log tariff, not its standardized innovation.

To fix signs, we take logs of the right-hand side of (I.60) and define

$$\begin{aligned} F_t &\equiv \log(\beta R^*) - \psi_b(b_t - \bar{b}) + \psi_t + \log \mathbb{E}_t e^{X_{t+1}}, \qquad && F_t = 0, \\ X_{t+1} &\equiv -\sigma(c_{t+1} - c_t) + (q_{t+1} - q_t) + (g_t - g_{t+1}). \end{aligned} \tag{I.73}$$

This convention gives a positive sign to a conditional-risk contribution in F_t; moving that contribution to the right-hand side of a linearized equation would reverse its sign. Along the level direction, $D_\delta X_{t+1}|_0 = (1-\alpha)\xi_q - \sigma\xi_c - \alpha$. With $g_{xx} = \alpha(1-\alpha)(1-\gamma)$, the second derivative of X_{t+1} and the square of its first derivative are

$$\begin{aligned} d_M^{F,imp} &= -\sigma B_c + (1-\alpha) B_q - g_{xx}(1+\xi_q)^2, \\ d_V^{F,imp} &= [(1-\alpha)\xi_q - \sigma\xi_c - \alpha]^2, \qquad && d_B^{F,imp} = d_M^{F,imp} + d_V^{F,imp}. \end{aligned} \tag{I.74}$$

A second-order cumulant expansion then gives

$$\Delta^V F_t = v_t d_B^{F,imp} + O(4), \tag{I.75}$$

where $d_M^{F,imp}$ is the conditional-mean curvature of X_{t+1} and $d_V^{F,imp}$ is its conditional-variance loading. In particular,

$$\Delta^V \operatorname{Var}_t(X_{t+1}) = 2 v_t d_V^{F,imp} + O(4). \tag{I.76}$$

Equation (I.75) is the positive residual coefficient used below; in the resolvent convention $\mathcal{K}(r) h_r = -d_r$, it is part of d_r and therefore appears with a minus sign in the solution $h_r = -\mathcal{K}(r)^{-1} d_r$. This is the real risk term that survives strict PPI; the domestic-Euler Jensen term affects the implementing rate but does not, by itself, imply a real response. To derive the trade-balance curvature, first define the second derivative of the log CPI–PPI ratio,

$$B_g = \alpha B_q + g_{xx}(1+\xi_q)^2. \tag{I.77}$$

Write the first and second derivatives of the export-revenue exponent as a_1, a_2 and those of

the import-expenditure exponent as m_1, m_2:

$$\begin{aligned} a_1 &= \gamma\xi_q - \alpha(1+\xi_q), & a_2 &= \gamma B_q - B_g, \\ m_1 &= \xi_c + (1-\gamma)\xi_q - \gamma + (\gamma-1)\alpha(1+\xi_q), & m_2 &= B_c + (1-\gamma)B_q + (\gamma-1)B_g. \end{aligned} \tag{I.78}$$

Since the rebated trade balance is exactly $TB = \alpha(e^a - e^m)$ and both exponents equal zero at the steady state, differentiating twice gives

$$B_{TB} = \alpha[(a_2 + a_1^2) - (m_2 + m_1^2)]. \tag{I.79}$$

The squared first derivatives are the curvature generated by exponentiation; the a_2 and m_2 terms capture curvature of the underlying allocations. Together with inherited-asset valuation, the Foreign-Euler risk term enters the second directional NFA law.

I.4.2 Incomplete-markets Ramsey system

After eliminating static reporting identities and the implementing rate, set

$$z_t = (\pi_{H,t}, q_t, c_t, b_t)'.$$

To fix the signs of the lower-endpoint terms, we orient the three exact residuals as

$$\begin{aligned} N_t &\equiv \kappa(\Pi_{H,t} - 1)\Pi_{H,t} - (\epsilon - 1)(mc_t - 1) - \beta\kappa\mathbb{E}_t\mathcal{P}_{t+1}, \\ \mathcal{P}_{t+1} &\equiv (\Pi_{H,t+1} - 1)\Pi_{H,t+1}\frac{Y_{t+1}}{Y_t}\left(\frac{C_t}{C_{t+1}}\right)^{\sigma}\frac{G_t}{G_{t+1}}, \end{aligned} \tag{I.80}$$

$$\mathcal{B}_t \equiv b_t - R^* e^{-\psi_b(b_{t-1}-\bar{b})+\psi_{t-1}}\frac{Q_t}{Q_{t-1}}\frac{G_{t-1}}{G_t}b_{t-1} - TB_t. \tag{I.81}$$

Thus $N_t = 0$ is (I.13), $F_t = 0$ is the log Foreign-Euler residual in (I.73), and $\mathcal{B}_t = 0$ is (I.62).

The response figure below continues a commitment chosen before date 0. Its date-0 stationarity conditions must therefore retain the marginal effects of the date--1 restrictions that contain z_0. Fix a realized date-zero history h^0 and define the inherited likelihood weight

$$\omega_0(h^0) \equiv \frac{e^{X_0(h^0)}}{\mathbb{E}_{-1}e^{X_0}}. \tag{I.82}$$

For any nodewise variation $\Delta z_0(h^0)$, the exact lower-endpoint contribution to the conditional

first-order condition is

$$D\mathfrak{B}_{-1}^{h^0}[\Delta z_0] \equiv \beta^{-1}\big\{ -\beta\kappa\lambda_{-1}^N D\mathcal{P}_0(h^0)[\Delta z_0] + \lambda_{-1}^F \omega_0(h^0) DX_0(h^0)[\Delta z_0]\big\}. \tag{I.83}$$

The date--1 NFA residual depends on z_{-1} and z_{-2} but not on z_0, so it contributes no lower-endpoint term. The date-zero node probability appears in both the ex-ante restriction and the continuation objective and cancels when the first-order condition is conditioned on h^0. The weight $\omega_0(h^0)$ remains because differentiating $\log \mathbb{E}_{-1} e^{X_0}$ changes measure toward histories with larger e^{X_0}. It satisfies $\mathbb{E}_{-1}\omega_0 = 1$ and equals one at the deterministic steady state. The exact nonlinear endpoint therefore requires the history-dependent inherited costate

$$\widehat{\lambda}_{-1}^F(h^0) \equiv \lambda_{-1}^F \omega_0(h^0). \tag{I.84}$$

Holding this inherited costate fixed in the conditional continuation, a scalar local representative of the endpoint is

$$\widehat{\mathfrak{B}}_{-1}^{h^0}(z_0) \equiv \beta^{-1}\left[-\beta\kappa\lambda_{-1}^N \mathcal{P}_0(h^0) + \widehat{\lambda}_{-1}^F(h^0) X_0(h^0)\right], \tag{I.85}$$

whose first variation is (I.83). This scalar expression is a first-variation device, not a second-order replacement for the inherited ex-ante restriction. For joint variations across date-zero sibling histories, define $\mathbb{E}_{-1}^{\omega_0} Y \equiv \mathbb{E}_{-1}[\omega_0 Y]$. The exact endpoint curvature is

$$D^2 \log \mathbb{E}_{-1} e^{X_0}[\Delta z_0]^2 = \mathbb{E}_{-1}^{\omega_0}\Big[D^2 X_0[\Delta z_0]^2 + \big(DX_0[\Delta z_0]\big)^2\Big] - \big(\mathbb{E}_{-1}^{\omega_0} DX_0[\Delta z_0]\big)^2. \tag{I.86}$$

On the feasible tangent to the date--1 Foreign-Euler restriction, the last term is zero, but the squared first derivative inside the expectation generally remains. It is omitted if the inherited likelihood weight is frozen node by node.

The boundary-correct conditional stationarity system is defined by

$$D\mathcal{L}_0^{\mathrm{inh}}(h^0)[\Delta z] \equiv D\mathfrak{B}_{-1}^{h^0}[\Delta z_0] + D\left\{\mathbb{E}_0 \sum_{t\geq 0} \beta^t [U_t + \lambda_t^N N_t + \lambda_t^F F_t + \lambda_t^B \mathcal{B}_t]\right\}[\Delta z] = 0. \tag{I.87}$$

Setting $\lambda_{-1}^N = \lambda_{-1}^F = 0$ instead defines a new date-0 commitment. For $t \geq 0$, the lagged derivatives below denote the analogous nodewise derivatives after conditioning on the date-t

history. In particular, $F_{t-1,j_t} = \omega_t D_{j_t} X_t$, where $\omega_t = e^{X_t}/\mathbb{E}_{t-1} e^{X_t}$. The four stationarity rows are

$$0 = U_{\pi,t} + \lambda_t^N N_{\pi,t} + \beta^{-1} \lambda_{t-1}^N N_{t-1,\pi_t}, \tag{I.88}$$

$$\begin{aligned} 0 = {} & U_{q,t} + \lambda_t^N N_{q,t} + \lambda_t^F F_{q,t} + \lambda_t^B \mathcal{B}_{q,t} \\ & + \beta^{-1}(\lambda_{t-1}^N N_{t-1,q_t} + \lambda_{t-1}^F F_{t-1,q_t}) + \beta \mathbb{E}_t[\lambda_{t+1}^B \mathcal{B}_{t+1,q_t}], \end{aligned} \tag{I.89}$$

$$\begin{aligned} 0 = {} & U_{c,t} + \lambda_t^N N_{c,t} + \lambda_t^F F_{c,t} + \lambda_t^B \mathcal{B}_{c,t} \\ & + \beta^{-1}(\lambda_{t-1}^N N_{t-1,c_t} + \lambda_{t-1}^F F_{t-1,c_t}), \end{aligned} \tag{I.90}$$

$$0 = \lambda_t^F F_{b,t} + \lambda_t^B \mathcal{B}_{b,t} + \beta \mathbb{E}_t[\lambda_{t+1}^B \mathcal{B}_{t+1,b_t}]. \tag{I.91}$$

Linearized around the deterministic steady state, the last row implies

$$\lambda_t^B = \mathbb{E}_t \lambda_{t+1}^B + \psi_b \lambda_t^F. \tag{I.92}$$

For the fixed history h^0, use the discounted norm

$$\|\Delta z\|_{\beta,h^0}^2 \equiv \sum_{t=0}^{\infty} \beta^t \mathbb{E}_0 \left[\|\Delta z_t\|^2 \mid h^0 \right]. \tag{I.93}$$

For variations across date-zero sibling histories, the corresponding ex-ante norm is

$$\|\Delta z\|_{\beta,-1}^2 \equiv \mathbb{E}_{-1} \|\Delta z\|_{\beta,h^0}^2. \tag{I.94}$$

Equations (I.88)–(I.91) characterize an inherited Ramsey continuation. As in Appendix H.4.1, we maintain four-times continuous differentiability of the exact objective and constraints, full row rank of the constraint derivative, and nonsingularity of the stable-root and endpoint system, which make this continuation locally unique. It is a strict local Ramsey optimum if, in addition, there is a $c > 0$ such that the second variation of the exact ex-ante Lagrangian, retaining $\lambda_{-1}^F \log \mathbb{E}_{-1} e^{X_0}$ jointly across date-zero sibling histories, is no greater than $-c\|\Delta z\|_{\beta,-1}^2$ on the ex-ante feasible tangent space and the Taylor remainder is uniformly $o(\|\Delta z\|_{\beta,-1}^2)$ in that discounted norm. Relative to the fixed-weight scalar representation, the exact endpoint adds

$$\beta^{-1} \lambda_{-1}^F \operatorname{Var}_{-1}^{\omega_0}(DX_0[\Delta z_0]) \tag{I.95}$$

to the quadratic form, where the variance is taken under the likelihood-weighted date-zero distribution. The modal calculation below verifies the KKT equations and selects their locally determinate branch; it does not establish global optimality.

Stack the four allocation variables and three multipliers in w_t. If J_-, J_0, J_+ are the exact lag, current, and lead Jacobians of the seven reduced rows, then an exogenous component proportional to r^t satisfies

$$\mathcal{K}(r)h_r = -d_r, \qquad \mathcal{K}(r) = r^{-1}J_- + J_0 + rJ_+. \tag{I.96}$$

Here d_r is the coefficient on r^t in the residual generated by the exogenous state after the endogenous linear terms have been moved to the left-hand side. If $\det \mathcal{K}(r) \neq 0$, then $h_r = -\mathcal{K}(r)^{-1}d_r$. At second order, the analogous coefficient $g_r^{[2]}$ collects the terms generated by the Hessians and the first-order solution, and $h_r^{[2]} = -\mathcal{K}(r)^{-1}g_r^{[2]}$ whenever $\mathcal{K}(r)$ is nonsingular. Endogenous KKT roots instead satisfy $\det \mathcal{K}(r) = 0$ and contribute right-null vectors of $\mathcal{K}(r)$, with their coefficients fixed by stability and the inherited multiplier conditions. Third-order residuals are constructed in the same way from Hessian and third-derivative terms generated by lower-order components. Because planner stationarity differentiates the full Lagrangian, a third-order Ramsey policy requires fourth derivatives of utility and every nonlinear implementability restriction.

For the modal representation, assume that the three endogenous stable roots are simple and nondefective, are distinct from the exogenous persistence ρ, and satisfy $\det \mathcal{K}(\rho) \neq 0$. Let V_s collect their stable eigenvectors and let B_{end} impose the three inherited-state and multiplier conditions. We also require $\operatorname{rank}(B_{\text{end}}V_s) = 3$. Under these conditions, the calibrated inherited Ramsey continuation contains the tariff or volatility root, one faster stable root, and a complex-conjugate pair with modulus $R_s < 1$ and angle ω. Each element of z_t can therefore be written as

$$z_t = A_\rho \rho^t + A_f \lambda_f^t + R_s^t[A_c \cos(\omega t) + A_s \sin(\omega t)]. \tag{I.97}$$

The roots are the eigenvalues of the state-transition matrix of the first-order Ramsey decision rule. For each variable, the four coefficients are then identified from four quarters of the computed path. Equation (I.97) is therefore closed form in time but not in primitive parameters: the roots and coefficients are obtained numerically from the seven-equation KKT system. At the baseline, the three computed endogenous roots are $\lambda_f = 0.2204685424$ and $\lambda_\pm = 0.9536589973 \pm i\,0.0003278969$. Because four quarters identify the coefficients, the remaining quarters of the computed 50-quarter level and uncertainty paths provide a consistency check on the four-mode representation: the maximum absolute discrepancy over these paths and over the strict-PPI formula (I.72) is 9.062×10^{-15}. Separately, the first-order tariff-level path has a projected KKT residual of 7.105×10^{-15}; this check does not test the third-order uncertainty residual. Under incomplete markets, the shadow values of foreign

assets rule out the one-equation reduction available under complete markets; the modal representation (I.97) takes its place.

Strict PPI removes domestic inflation, but it leaves the Foreign-Euler and NFA conditions in (I.66)– (I.67) as restrictions on the allocation. The Ramsey planner also assigns shadow values to those conditions, as (I.92) shows. Strict PPI therefore need not satisfy the full Ramsey stationarity system under incomplete markets. Whether it is quantitatively close to Ramsey is assessed below from a common deterministic physical state.

Figure I.7 holds the bond economy fixed and compares the inherited Ramsey continuation with strict PPI and the PPI-based Taylor rule. All three policies start from the same predetermined physical state at date -1 and zero NFA. The Ramsey curve carries forward from date -1 the multipliers that enter the date-0 conditions under the plan chosen before the volatility innovation; the central bank neither resets them nor chooses a new plan at date 0. Strict PPI and the Taylor rule are fixed rules without planner multipliers. Each curve is measured relative to that policy's no-volatility path.

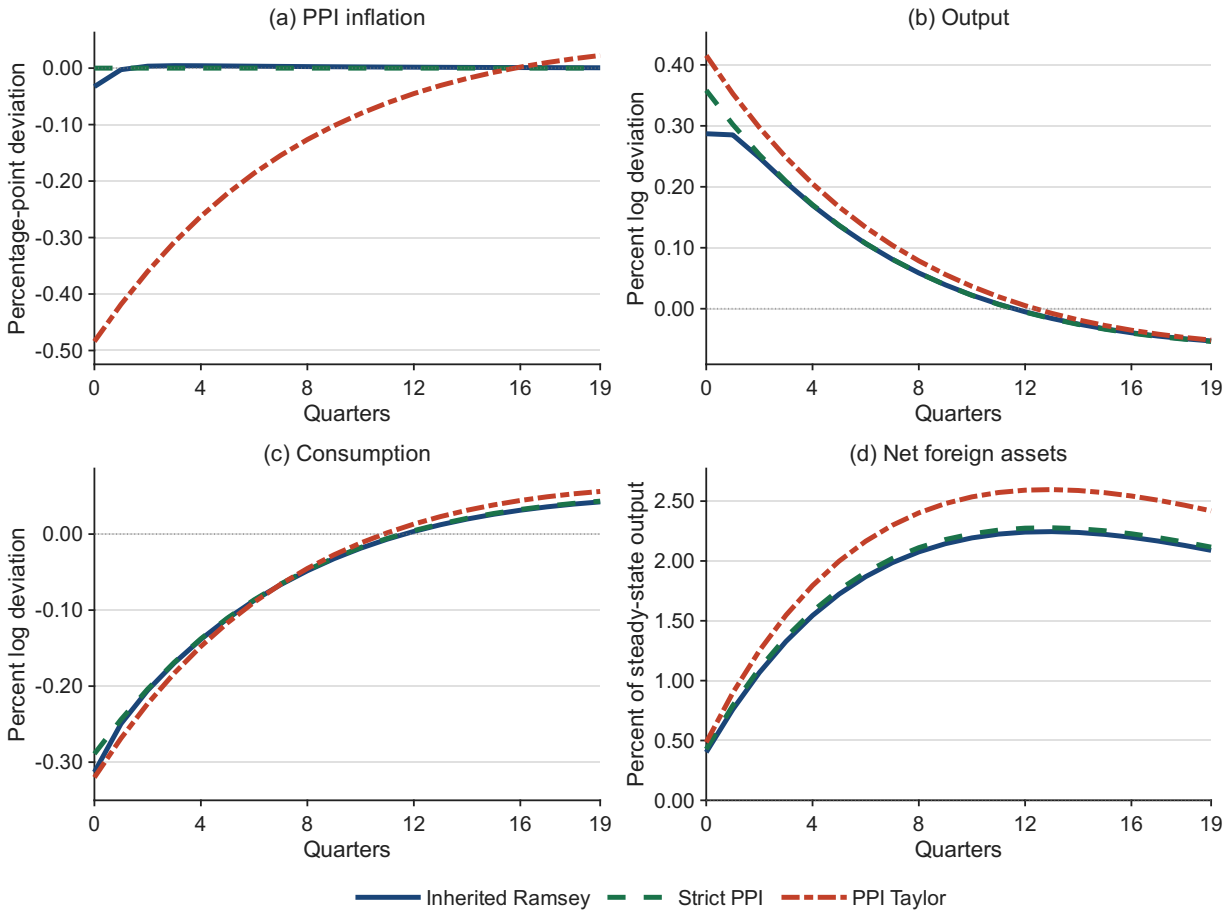


Figure I.7: Monetary policy under incomplete markets with import tariffs

Notes: Each panel reports the local third-order response to a one-standard-deviation date-0 innovation in log import-tariff volatility, which doubles the conditional standard deviation of next quarter's tariff innovation. The tariff remains at its deterministic steady-state level and realized tariff innovations are zero at all displayed dates; future risk remains in expectations. Each policy starts from the same predetermined physical state at date -1 and zero NFA, and each curve subtracts its own no-volatility path. The Ramsey curve continues the pre-date-0 plan, carrying forward from date -1 the multipliers that enter the date-0 conditions rather than resetting them; the other two curves are fixed-rule responses. Inflation is in quarterly percentage points; output and consumption are percent log deviations; NFA is a percentage of steady-state output.

The figure separates PPI-inflation stabilization from households' accumulation of foreign assets to smooth consumption across possible tariff states. Ramsey keeps PPI inflation close to zero, and strict PPI sets it exactly to zero. The two policies generate close consumption and NFA paths because both remain subject to the Foreign Euler equation and the NFA accumulation law. Ramsey chooses inflation jointly with consumption, the terms of trade, and NFA; the shadow values of those equilibrium restrictions enter the stationarity conditions (I.89)–(I.91). Strict PPI instead fixes inflation at zero and lets the same equilibrium restrictions determine the real adjustment. It produces a larger impact increase in output but converges quickly toward Ramsey's real paths. The Taylor rule produces a larger initial deflation

and larger adjustments in output and NFA. Strict PPI is therefore a close proxy for the displayed real responses, especially after impact, but the figure establishes neither exact Ramsey optimality nor welfare equivalence.

I.5 Approximation order and welfare

Scale $\bar{\sigma}_\tau$ and η_σ by a common perturbation parameter ε. A realized volatility innovation changes conditional tariff variance by

$$\bar{\sigma}_\tau^2(e^{2s_t} - 1) = 2\bar{\sigma}_\tau^2 s_t + O(\varepsilon^4) = O(\varepsilon^3). \tag{I.98}$$

Hence the conditional uncertainty response with zero realized tariff innovations is zero through order two and first appears at order three. Changing the ex-ante variance of the volatility innovation changes expected tariff variance only at order four. A Ramsey continuation computed with tariff-level innovations alone therefore coincides, through order three, with one that also allows future volatility innovations. This is not an exact nonlinear equivalence.

A third-order Ramsey policy must be evaluated with welfare through fourth order. Let $x = (h, \bar{\sigma}_\tau, \eta_\sigma)$ collect the policy-specific allocations and state variables h together with every shock-law argument that is scaled directly by ε. Write the scalar objective as $\mathcal{L}(x)$ and define $x^{[k]} = d^k x(\varepsilon)/d\varepsilon^k|_{\varepsilon=0}$. For $x(\varepsilon) = \bar{x} + \varepsilon x^{[1]} + \varepsilon^2 x^{[2]}/2 + \varepsilon^3 x^{[3]}/6 + \varepsilon^4 x^{[4]}/24 + O(\varepsilon^5)$,

$$\begin{aligned}\mathcal{L}^{[4]} = {} & \mathcal{L}_a x_a^{[4]} + 4\mathcal{L}_{ab} x_a^{[1]} x_b^{[3]} + 3\mathcal{L}_{ab} x_a^{[2]} x_b^{[2]} \\ & + 6\mathcal{L}_{abc} x_a^{[1]} x_b^{[1]} x_c^{[2]} + \mathcal{L}_{abcd} x_a^{[1]} x_b^{[1]} x_c^{[1]} x_d^{[1]}.\end{aligned} \tag{I.99}$$

Repeated component indices a, b, c, d are summed. This identity shows why third-order responses alone do not generally determine a fourth-order welfare ranking. Because planner stationarity differentiates the full Lagrangian, the fourth-order Ramsey solution used below requires the exact objective and constraints to be five times continuously differentiable, one more derivative than is maintained above for the third-order responses; with only four, the $O(\varepsilon^5)$ remainder becomes $o(\varepsilon^4)$. At a correctly initialized local Ramsey optimum, the envelope condition removes the unknown allocation component of $x^{[4]}$; it does not do so for a fixed rule. A valid comparison with strict PPI or a Taylor rule must therefore evaluate primitive household utility using a genuine fourth-order allocation for each rule. All policies face the same joint distribution of future tariff and volatility innovations and start from the same physical state. The Ramsey reference is a new date-0 commitment with inherited planner multipliers set to zero, distinct from the inherited Ramsey continuation shown in Figure I.7. The envelope argument is conditional on the local sufficiency conditions stated above; the

numerical calculation verifies the selected KKT branch, not global concavity.

Let $W^{(4)}_{R,\mathrm{KKT}}$ denote fourth-order lifetime utility on the selected new-commitment Ramsey-KKT branch and let $W^{(4)}_j$ denote the corresponding value under policy j. The signed utility difference is

$$\Delta W^{(4)}_j \equiv W^{(4)}_{R,\mathrm{KKT}} - W^{(4)}_j. \tag{I.100}$$

We solve a fourth-order decision rule for every policy without discarding higher-order state interactions. Reference utility is evaluated from the planner's fourth-order KKT solution; for each fixed rule, a recursive household-welfare equation is solved jointly with the allocation. Future Gaussian innovations are integrated through the fourth-order decision polynomial under a common joint shock distribution. This retains the mixed moment $\mathbb{E}(\varepsilon_T^2\varepsilon_\sigma^2) = 1$ rather than adding two separate one-shock calculations.

Figure I.8 reports (I.100) for strict PPI and the two Taylor rules. Strict PPI is the closest fixed rule to the new-commitment Ramsey reference under complete and incomplete markets and under each date-0 conditioning event. The Taylor-rule gaps are substantially larger, especially when households cannot trade state-contingent claims, and their ordering differs across financial structures: the CPI rule has the smaller gap under complete markets and the larger gap under incomplete markets. Strict PPI's gap is about 1.01 percent of the smaller Taylor-rule gap under complete markets and 0.28 percent under incomplete markets (Table L.1). Thus the import-tariff extension reinforces the paper's quantitative conclusion: strict PPI is not the exact Ramsey policy, but in this calibration its utility gap to Ramsey is much smaller than that of either Taylor rule. This ranking is not a proof of optimality over all feasible nonlinear policies.

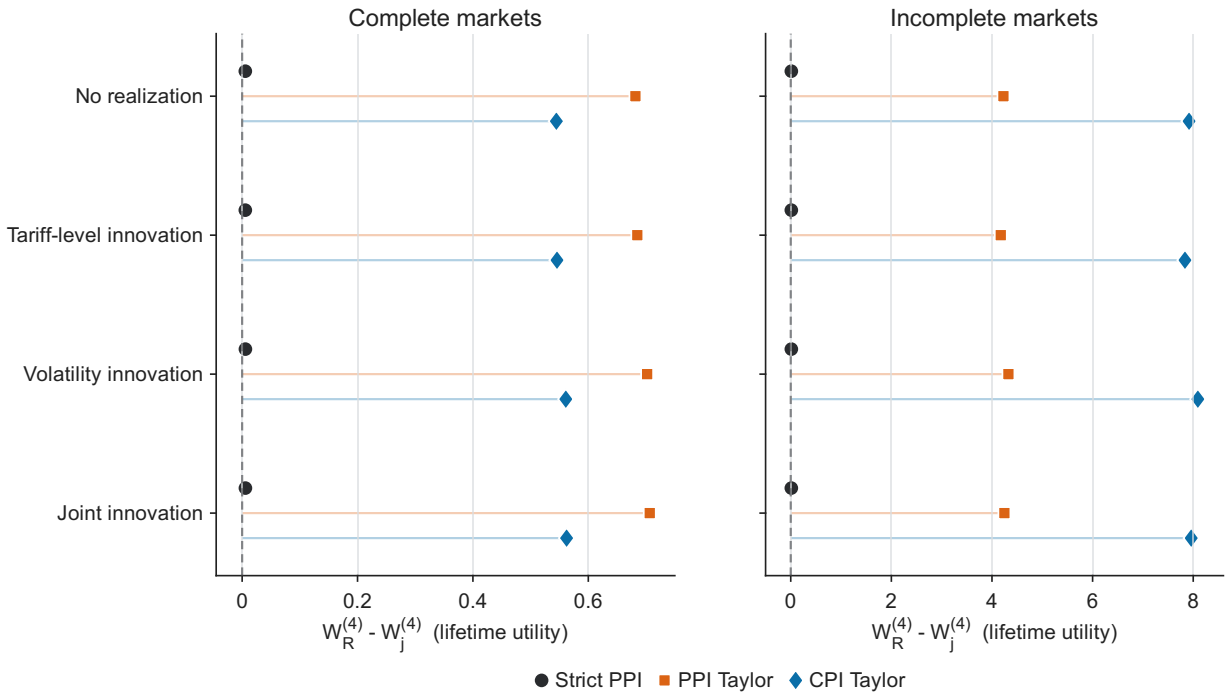


Figure I.8: Fourth-order utility differences relative to new-commitment Ramsey

Notes: Each marker reports $W^{(4)}_{R,\mathrm{KKT}} - W^{(4)}_{j}$ in lifetime household-utility units, where the reference is the selected new-commitment Ramsey KKT branch, which sets lagged multipliers to zero at date 0. The panels use separate horizontal scales because incomplete markets amplify the differences across policies. The four markers for each rule condition on, respectively, no date-0 realization, a one-standard-deviation tariff-level innovation, a one-standard-deviation log-volatility innovation that doubles the conditional standard deviation of next quarter's tariff innovation, and both innovations together. Future tariff and volatility innovations remain stochastic and are integrated jointly. All policies start from the same predetermined physical state. The comparison ranks the computed branches and does not establish global optimality.

The welfare gaps are measured in lifetime household-utility units. A valid fourth-order consumption equivalent would require an order-by-order re-expansion of the consumption scaling; applying an exact nonlinear power transform to truncated welfare would reintroduce uncomputed higher-order terms. We therefore report no consumption-equivalent number.

J Calvo Pricing

The main-text pricing block uses Rotemberg quadratic adjustment costs (Rotemberg, 1982). This appendix replaces it with Calvo time-dependent pricing (Calvo, 1983), choosing θ_C so that the linearized Calvo NKPC has the Rotemberg slope. Under a fixed policy rule, the two models then have identical linearized responses to a tariff-level shock. Their responses to tariff uncertainty can differ substantially, because Calvo reset prices and inherited price dispersion enter at higher orders. Under complete markets and the PPI Taylor rule, slope-matched

Calvo pricing reverses the baseline response to export-tariff uncertainty: on impact, PPI inflation rises by about 0.63 percentage points instead of falling by 1.01, and output falls by about 1.12 percent instead of rising by 0.07 percent. This reversal is consistent with the precautionary reset-price channel that makes uncertainty inflationary under Calvo but not under Rotemberg pricing (Oh, 2020). A Calvo firm that resets its price may have to keep it for several quarters, and because its profits fall faster when its price is too low than when it is too high, greater tariff risk pushes it toward a higher reset price. A Rotemberg firm can adjust its price every period, so this precautionary motive does not arise. Under incomplete markets, the Taylor-rule deflation survives: PPI inflation falls by about 0.09 percentage points under Calvo pricing and by about 0.10 under Rotemberg pricing.

Under a fixed monetary rule, the two technologies share the wedge logic of Section 3.2: a volatility innovation enters through conditional-moment corrections to pricing and asset valuation. Under Calvo pricing, however, these higher-order corrections must be solved in an augmented system that includes the reset values and inherited dispersion; under Ramsey policy, the system also includes the planner's stationarity conditions and inherited multipliers.

Roadmap. Section J.1 derives the slope-matched Calvo block and its first-order equivalence with Rotemberg pricing under a fixed rule. Section J.2 gives the augmented higher-order coefficient system for fixed rules. Section J.3 defines the inherited Ramsey continuation used in the numerical exercise, compares it with strict PPI and the PPI Taylor rule, and reports accuracy checks.

J.1 Setup

Reset-price problem. Each Home producer faces CES demand with elasticity $\epsilon > 1$, and a fraction $1 - \theta_C$ of firms can reset their nominal price each period. The household labor condition and technology imply raw marginal cost in Home-PPI units

$$m_t^r = \chi G_t L_t^{\varphi} C_t^{\sigma} / A_t, \qquad mc_t \equiv m_t^r / \bar{\upsilon}, \qquad \bar{\upsilon} = m_{ss}^r, \tag{J.1}$$

where mc_t is normalized to one at the deterministic steady state. The fixed production correction satisfies $1 - s^p = (\epsilon - 1)/(\epsilon \bar{\upsilon})$, so the reset-price first-order condition contains normalized mc_t rather than raw m_t^r. The correction is a fiscal transfer; Calvo price dispersion is the real resource distortion.

Define the relative reset price $p_t^{\#} \equiv P_t^{\#} / P_{H,t}$. Two auxiliary objects characterize the

reset-price first-order condition:

$$\mathcal{N}_t = mc_t Y_t + \theta_C \, \mathbb{E}_t\big[\mathcal{M}^H_{t,t+1} \, \Pi^{\epsilon}_{H,t+1} \, \mathcal{N}_{t+1}\big] \,, \tag{J.2}$$

$$\mathcal{D}_t = Y_t + \theta_C \, \mathbb{E}_t\big[\mathcal{M}^H_{t,t+1} \, \Pi^{\epsilon-1}_{H,t+1} \, \mathcal{D}_{t+1}\big] \,, \tag{J.3}$$

where $\mathcal{M}^H_{t,t+1} = \beta(C_t/C_{t+1})^{\sigma} G_t/G_{t+1}$ is the stochastic discount factor in Home-PPI units. With the baseline production correction above, the reset price satisfies

$$p^{\#}_t = \mathcal{N}_t/\mathcal{D}_t, \tag{J.4}$$

and the aggregate price index satisfies the exact Calvo restriction

$$1 = (1-\theta_C)\,(p^{\#}_t)^{1-\epsilon} + \theta_C \, \Pi^{\epsilon-1}_{H,t}. \tag{J.5}$$

A non-resetting firm keeps its previous nominal price, so the cross-section of firm-level relative prices is dispersed whenever inflation is nonzero. The demand-weighted dispersion state is $v^C_t \equiv \int_0^1 (P_{H,t}(j)/P_{H,t})^{-\epsilon} \mathrm{d}j$ and evolves as $v^C_t = (1-\theta_C)(p^{\#}_t)^{-\epsilon} + \theta_C \Pi^{\epsilon}_{H,t} \, v^C_{t-1}$. Aggregate labor satisfies $L_t = v^C_t Y_t / A_t$. Calvo pricing therefore replaces the Rotemberg adjustment-cost equation with reset-price, price-index, and dispersion restrictions. Price dispersion has no linear effect around the zero-inflation steady state, but it enters the higher-order reset-price equations.

Slope matching and calibration. Linearizing (J.2)–(J.5) around the zero-inflation steady state gives the standard Calvo NKPC

$$\pi_t = \lambda_C \, \widehat{mc}_t + \beta \, \mathbb{E}_t \pi_{t+1}, \qquad \lambda_C \equiv \frac{(1-\theta_C)(1-\beta\theta_C)}{\theta_C}.$$

Slope matching pins θ_C so that $\lambda_C = \kappa_{mc} = (\epsilon - 1)/\kappa$ from (22). At the baseline calibration of Table B.1, this equation uniquely determines the Calvo non-reset probability. With $\epsilon = 10$ and $\kappa = 40$, $\kappa_{mc} = 0.225$, so $\theta_C \approx 0.63$, an average price duration of about 2.7 quarters. All other model components are held at the main-text baseline: preferences, the static block, the tariff and volatility processes (18)–(19) and Definition 1, the asset-market closure (13), and the policy rules.

Because the linearized NKPC slope is common, under a fixed policy rule the first-order tariff-exposure coefficients (A_π, A_q) from (D.11)–(D.13), and hence the first-order responses to a tariff-*level* shock, are the same under Rotemberg and slope-matched Calvo pricing (Proposition J.1). The two technologies differ only in the higher-order Jensen terms that

drive a tariff-*uncertainty* shock.

Rotemberg pricing averages a risky path of *aggregate inflation*: firms can reset every period but pay a convex resource cost. Calvo averages a risky *reset-price problem*: the resetter chooses one nominal price $p_t^{\#}$ knowing it may still be in place when future tariff states arrive. The nonlinear price-index restriction and continuation-value recursions therefore enter the Calvo certainty equivalent directly. Both technologies generate an NKPC wedge, but its components have different economic interpretations; Section J.2 makes the mapping explicit.

Proposition J.1 (Slope-matched Calvo reduction)**.** *Consider the complete-markets PCP economy with the same static block, asset-market closure, exogenous processes, and policy rule as in the main text, but replace Rotemberg pricing with the Calvo block* (J.2)–(J.5) *and impose the slope match* $\lambda_C = \kappa_{mc}$*. Then the first-order equilibrium in* (π_H, q) *is identical to the Rotemberg equilibrium in Lemma C.1. Consequently, the first-order tariff-level responses and the first-order responses to the productivity shock and to the auxiliary UIP disturbance* ψ_t *are common across the two pricing technologies. The reset-price, price-index, and dispersion states do not alter the first-order* (π_H, q) *equilibrium; their distinct role in the volatility experiment appears through higher-order Jensen terms.*

Proof begins: Proposition J.1.

In the first-order pricing block, the only object from the sticky-price technology that enters the (π_H, q) equilibrium is the NKPC slope. The slope match imposes $\lambda_C = \kappa_{mc}$, while the static block, asset-market closure, exogenous laws, and policy rule are held fixed. Therefore the linearized Calvo and Rotemberg systems have the same coefficient matrix and the same disturbance vector for any first-order tariff or productivity shock or auxiliary UIP disturbance ψ_t. The reset-price, price-index, and dispersion recursions are normalized at the symmetric steady state and do not add an independent first-order restriction in (π_H, q); they enter the volatility experiment only through the higher-order conditional-moment terms described below.

End of proof: Proposition J.1.

J.2 Higher-order coefficient system under fixed rules

This subsection derives the higher-order coefficient system under a fixed rule. Calvo pricing adds reset-price continuation values and inherited dispersion, so we collect the local variables in the vector

$$d_t^C = \left(\pi_{H,t}, q_t, \widehat{\mathcal{N}}_t, \widehat{\mathcal{D}}_t, \widehat{p}_t^{\#}, \widehat{v}_t^C\right)'. \tag{J.6}$$

After substituting the exact static demand, risk-sharing, marginal-cost, and policy maps, write the reset recursions, price-index identity, dispersion law, and asset equation as the primitive residual system

$$\mathcal{F}^C(d^C_{t-1}, d^C_t, d^C_{t+1}; \tau_t, \sigma_{T,t}, \varepsilon_{\tau,t+1}) = 0, \tag{J.7}$$

where the equations containing $t+1$ variables are inside the time-t conditional expectation. All derivatives below are evaluated at the zero-inflation deterministic steady state.

Let F^C_-, F^C_0, and F^C_+ denote the Jacobians of (J.7) with respect to dated variables. The linearized residual has the form

$$F^C_- d^C_{t-1} + F^C_0 d^C_t + F^C_+ \mathbb{E}_t d^C_{t+1} + b^C_t = 0. \tag{J.8}$$

For a geometric component with root r, define the matrix polynomial

$$\mathcal{K}_C(r) = r^{-1} F^C_- + F^C_0 + r F^C_+. \tag{J.9}$$

A geometric inverse alone is not the date-0 solution because price dispersion is predetermined. Let (r^C_j, h^C_j) denote the stable generalized eigenpairs satisfying

$$\mathcal{K}_C(r^C_j) h^C_j = 0, \qquad |r^C_j| < 1,$$

and collect them in $H_C = (h^C_1, \ldots, h^C_{m_C})$ and $R_C = \mathrm{diag}(r^C_1, \ldots, r^C_{m_C})$. The selector

$$S_C = (0, 0, 0, 0, 0, 1)$$

extracts the inherited dispersion state from (J.6). Local determinacy requires one stable mode for this inherited state, so $m_C = 1$; we keep the matrix notation because it makes the boundary condition transparent.

For a residual term br^t that begins at date 0, suppose $\mathcal{K}_C(r)$ and $S_C H_C R_C^{-1}$ are nonsingular. The associated particular solution, boundary coefficient, and path are

$$p_C(b,r) \equiv -\mathcal{K}_C(r)^{-1} b, \tag{J.10}$$

$$c_C(b,r) \equiv -(S_C H_C R_C^{-1})^{-1} S_C p_C(b,r) r^{-1}, \tag{J.11}$$

$$\mathcal{P}_C(b,r;t) \equiv p_C(b,r) r^t + H_C R_C^t c_C(b,r). \tag{J.12}$$

The first term is the geometric particular solution. The second is the stable homogeneous component required by the inherited-state condition $S_C d_{-1}^C = 0$: substituting $t = -1$ into (J.12) gives zero exactly. Thus (J.12), rather than the particular inverse alone, is the response to a date-0 disturbance. These formulas require a regular matrix pencil, the locally determinate stable-root count, and no coincidence between the disturbance persistence and an endogenous root. Repeated roots require the corresponding generalized-eigenvector polynomial terms.

Let f_τ^C be the primitive first derivative of the residual with respect to the current tariff. The first-order response to an initial tariff level τ_0 is

$$d_t^{C,[1]} = \tau_0 \mathcal{P}_C(f_\tau^C, \rho_\tau; t). \tag{J.13}$$

At the zero-inflation steady state, the linearized price-index and dispersion equations imply $\widehat{v}_t^C = \theta_C \widehat{v}_{t-1}^C$. Starting from $\widehat{v}_{-1}^C = 0$, tariff innovations therefore have no first-order dispersion effect, so the boundary coefficient in (J.13) is zero. Slope matching then makes its (π_H, q) entries equal to their Rotemberg counterparts.

At second order, deterministic tariff curvature and conditional innovation variance are distinct objects and have different time profiles. Let $\mathfrak{g}_\tau^C$ denote the total first-order direction of every argument of $\mathcal{F}^C$ induced by one unit of the current tariff state, and let $\mathfrak{g}_\varepsilon^C$ denote the direction induced by one unit of the next tariff innovation, including its effect on next period's decision rule. Using the usual one-half Taylor convention, define the two primitive Hessian contractions

$$b_{\tau\tau}^C \equiv \frac{1}{2} D^2 \mathcal{F}^C[\mathfrak{g}_\tau^C, \mathfrak{g}_\tau^C], \qquad b_{\varepsilon\varepsilon}^C \equiv \frac{1}{2} \mathbb{E}_t D^2 \mathcal{F}^C[\mathfrak{g}_\varepsilon^C, \mathfrak{g}_\varepsilon^C]. \tag{J.14}$$

The deterministic term follows $\tau_t^2 = \rho_\tau^{2t} \tau_0^2$, whereas at $\sigma_{T,t} = 0$ the conditional innovation variance is the constant $\bar{\sigma}_\tau^2$. The second-order path is therefore

$$d_t^{C,[2]} = \tau_0^2 \mathcal{P}_C(b_{\tau\tau}^C, \rho_\tau^2; t) + \bar{\sigma}_\tau^2 \mathcal{P}_C(b_{\varepsilon\varepsilon}^C, 1; t). \tag{J.15}$$

The first term contains deterministic reset-price, price-index, and dispersion curvature.

The second is the stochastic-mean correction generated by baseline tariff risk. Propagating both terms at ρ_τ^2 would incorrectly assign the constant conditional-variance correction the persistence of tariff curvature.

Finally, let b_σ^C denote the full third-order residual coefficient on $\bar{\sigma}_\tau^2 \sigma_{T,t}$ obtained by differentiating the conditional innovation-variance contraction and substituting the lower-order rules in (J.13)–(J.15). Along the pure volatility experiment, the constant second-order stochastic-mean term is common to the shocked and baseline paths and cancels. The baseline-subtracted response is

$$\Delta^U d_t^{C,[3]} = \bar{\sigma}_\tau^2 \eta_\sigma \mathcal{P}_C(b_\sigma^C, \rho_\sigma; t) + O(4). \tag{J.16}$$

This expression includes the homogeneous stable component needed to keep the inherited dispersion state fixed at date -1. Numerically, the Jacobians, Hessian contractions, stable root, and boundary coefficient are evaluated from the primitive nonlinear equations.

Proposition J.2 (Slope-matched Calvo uncertainty system)**.** *Under complete markets and a fixed policy rule, suppose the Calvo matrix pencil is regular, the fixed-rule equilibrium is locally determinate with one stable mode corresponding to inherited dispersion, and the nonsingularity conditions stated above hold at $r \in \{\rho_\tau, \rho_\tau^2, 1, \rho_\sigma\}$. Then* (J.13) *gives the first-order tariff response,* (J.15) *separates deterministic tariff curvature from the stochastic-mean correction, and* (J.16) *gives the leading response to a pure log-volatility innovation on the displayed history with no realized tariff innovations. That conditional-path response is zero through order two and its order-three term contains reset-price, price-index, dispersion, continuation-value, and asset-pricing components.*

Proof begins: Proposition J.2.

Differentiate (J.7) along the tariff state. Substituting a geometric term r^t gives (J.9); adding the stable null-space mode and imposing $S_C d_{-1}^C = 0$ gives (J.12)–(J.13). At second order, deterministic tariff curvature has root ρ_τ^2, whereas the innovation-variance correction at the zero-volatility state is constant. Applying the same boundary-adjusted propagator gives (J.15). A pure volatility innovation leaves the realized tariff unchanged, so its first two coefficients vanish under the common shock scaling. Its variance change has root ρ_σ; applying (J.12) to the third-order contraction gives (J.16).

End of proof: Proposition J.2.

In this representation, the Calvo-specific risk channels of Section J.1—a reset price that may survive into future tariff states, the nonlinear price-index identity, and inherited price

dispersion—enter b_σ^C separately, even though the augmented matrix polynomial propagates them jointly.

J.3 Ramsey continuation and quantitative comparison

This subsection defines the Ramsey continuation under Calvo pricing and compares it with strict PPI and the PPI Taylor rule.

The Calvo policy line is not obtained by inserting an "optimal wedge" into the fixed-rule formula. We differentiate the primitive household objective and every nonlinear Calvo constraint and compute a third-order branch of the resulting Karush–Kuhn–Tucker (KKT) conditions. We call this branch the inherited Ramsey continuation, or simply the Ramsey continuation. Let (H_C^s, R_C^s) denote the stable deflating subspace of a regular strong linearization of the augmented KKT system after purely algebraic zero and infinite blocks are removed. Let B_C impose the n_C^b inherited-state and multiplier boundary conditions. We maintain hyperbolicity of the finite nonzero part and

$$\dim \operatorname{im} H_C^s = n_C^b, \qquad \operatorname{rank}\big[B_C H_C^s (R_C^s)^{-1}\big] = n_C^b. \tag{J.17}$$

Thus the endpoint conditions determine the stable-mode coefficients uniquely. A matching root count without the rank condition in (J.17) would make the system square but would not establish uniqueness. Under the stated conditions, the linearized boundary-value operator is invertible and the implicit-function theorem selects the reported local KKT branch. Its KKT residuals have the order required by a third-order solution. A strict local optimum additionally requires a constraint qualification and a negative-definite Hessian of the Lagrangian on the tangent space of the constraints. Those second-order sufficiency conditions, and global concavity, are not established here, so the figures report the locally determinate KKT branch and do not claim a global optimum. The path in Figure J.1 continues a pre-date-0 plan and inherits its steady-state planner multipliers. The main complete-markets exercise begins a new date-0 commitment and sets the inherited multiplier to zero.

Under strict PPI, $\Pi_{H,t} = 1$. The price-index identity then implies $p_t^{\#} = 1$; starting from unit dispersion, its law gives $v_t^C = 1$. The reset-price condition requires $mc_t = 1$, so strict PPI removes the domestic price-setting distortion and implements the decentralized flexible-price allocation. Exact Ramsey optimality requires one additional condition: that allocation must also be constrained efficient in every tariff state the economy may reach. The production correction is fixed at its no-tariff value, which, together with the calibration of χ, makes the no-tariff steady state constrained efficient, but it does not vary with the export tariff. Once a tariff is realized, the flexible-price allocation therefore need not solve the planner's allocation

problem. This is why strict PPI and the Ramsey continuation can differ even though their responses to tariff risk are quantitatively close.

Figures J.1 and J.2 compare the policies after a log-volatility innovation $\eta_\sigma = \log 2$. The innovation doubles the conditional standard deviation of next period's export-tariff innovation from 0.10 to 0.20. The displayed path conditions on zero realized tariff innovations, while future tariff risk remains in agents' expectations; these are conditional responses, not averages over future tariff histories.

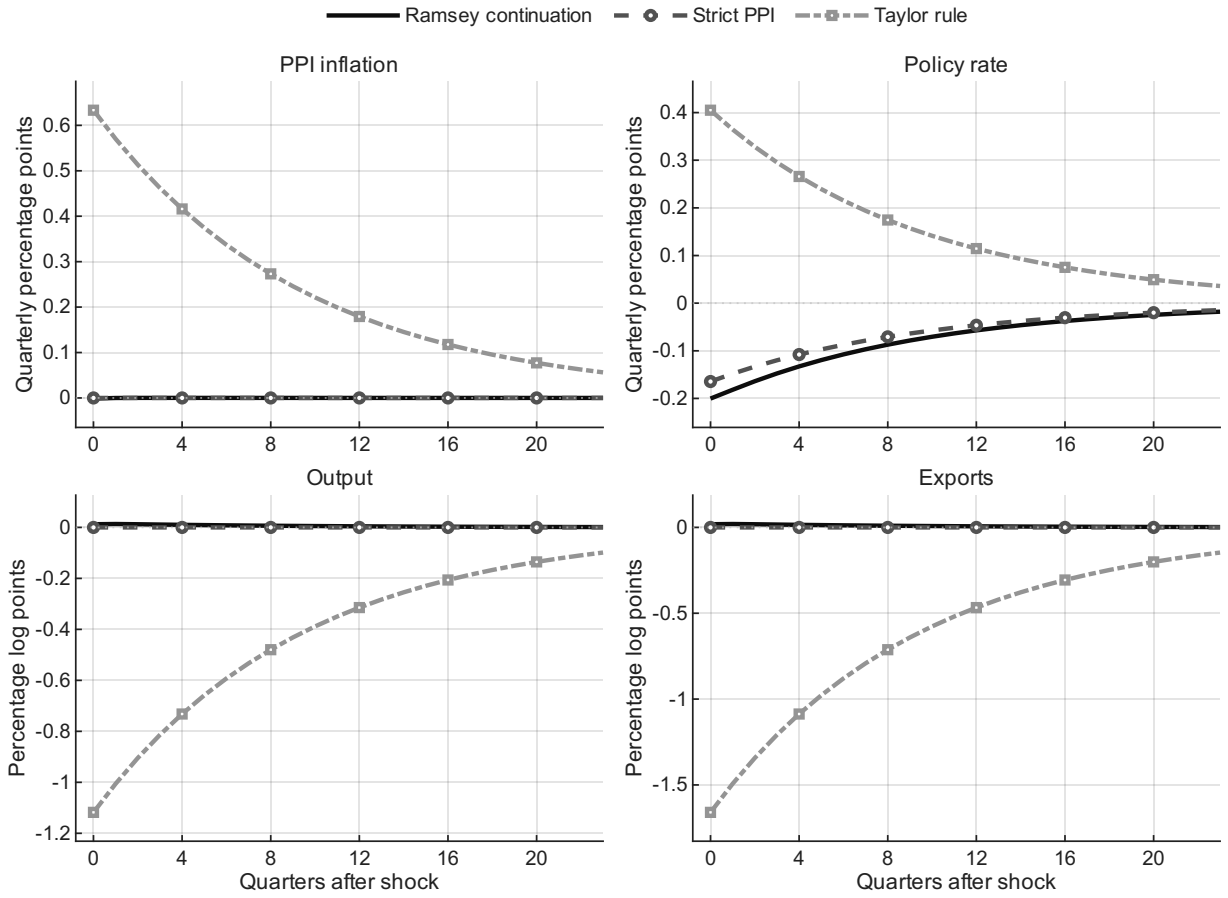


Figure J.1: Ramsey continuation and simple rules under Calvo pricing

Notes: Complete-markets third-order responses to the export-tariff log-volatility innovation $\eta_\sigma = \log 2$. Realized tariff innovations are set to zero, while future tariff risk remains in expectations. The Ramsey continuation inherits the promises made when the policy plan was chosen before date 0; strict PPI and the PPI Taylor rule are solved from the same physical state. Inflation and the policy rate are in quarterly percentage points (not annualized); output and exports are percent log deviations. Responses are local third-order approximations, not exact effects of a finite variance doubling.

Figure J.1 shows that the Taylor rule produces substantially larger nominal and real responses than either the Ramsey continuation or strict PPI: on impact, PPI inflation moves by 0.63 percentage point under the Taylor rule against -0.0014 under the Ramsey continuation

(Table L.1). The Ramsey and strict-PPI real allocations are almost indistinguishable at that scale, even though their interest-rate paths need not coincide: the interest rate is an implementation instrument and can respond to changes in expected future marginal utility and inflation while current output and trade move very little. Figure J.2 removes the Taylor line and magnifies output and exports; the remaining gap is visible but small. For nominal stabilization, the Calvo extension therefore preserves the main-text conclusion: strict PPI is a close approximation to the Ramsey continuation for this uncertainty experiment, whereas the Taylor rule is not. The real responses are not directly comparable with the main text, for two reasons: the Calvo Ramsey path is an inherited continuation rather than a new date-0 commitment, and under complete markets and Calvo pricing the Taylor rule reverses the sign of the baseline inflation and output responses.

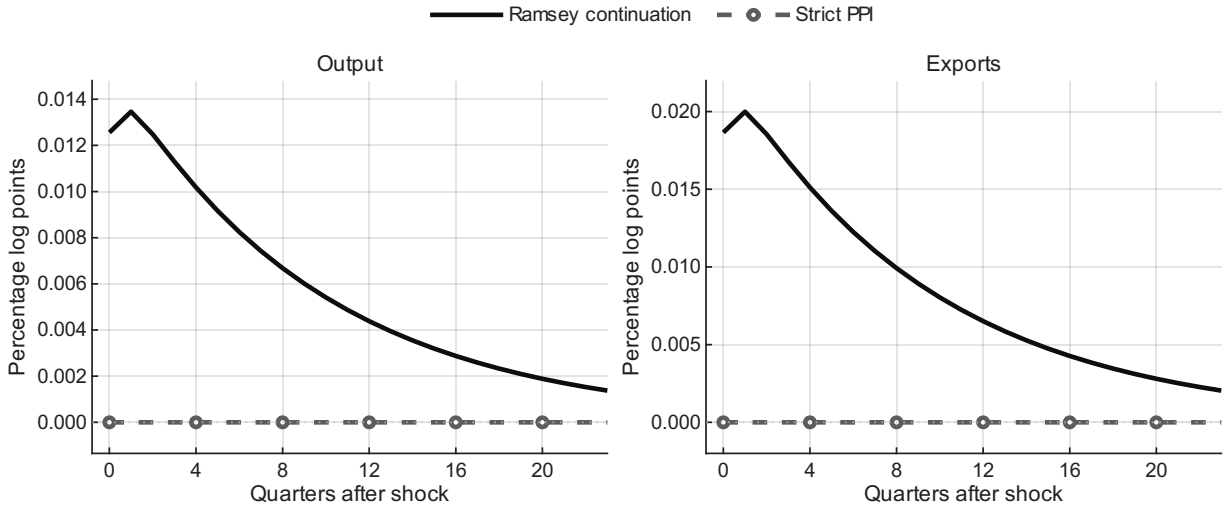


Figure J.2: The small Calvo Ramsey-continuation departure from strict PPI

Notes: The shock, conditioning event, units, and local third-order interpretation are the same as in Figure J.1. The panel-specific scales make the small differences in output and exports visible. Both policies are state contingent. Strict PPI imposes zero domestic inflation after every history, whereas the Ramsey continuation jointly solves for inflation and the real allocation subject to the full Calvo constraints.

The displayed no-realization response is zero through second order and first appears at third order, as Proposition J.2 predicts. When all shocks are rescaled toward the steady state, the conditional pricing and asset-equation residuals decline at approximately fourth order, as a third-order perturbation requires. As in the main text (see the discussion after Definition 1), the figures are local third-order approximations, not exact responses to a finite variance doubling. We have not computed a fourth-order Calvo welfare ranking, so the comparison concerns conditional responses, not welfare equivalence.

K Local-Currency Export Pricing

The main model assumes producer-currency pricing (PCP): Home exporters post their sticky price in Home currency, so the terms of trade governs both risk sharing and the Foreign-customer price. Under local-currency pricing (LCP), exporters instead post a separate sticky price in Foreign currency. This separates the customer price from the risk-sharing relative price and adds a predetermined export-price state to the policy problem. Appendix K.1 derives the export-pricing block from the exporter's problem, Appendix K.2 reduces the fixed-rule equilibrium to three equations with one predetermined state, and Appendix K.3 compares the inherited Ramsey continuation with strict Home-PPI targeting and the Home-PPI Taylor rule. In this calibration, strict PPI remains close to the Ramsey continuation for nominal stabilization and exports but is not distinctly closer than the Taylor rule for output on impact, and the sign of the Taylor-rule output response is reversed relative to producer-currency pricing.

K.1 The export-pricing block from primitives

Define the Foreign-currency export price and its Home-currency revenue value by

$$V_t \equiv \frac{P^*_{X,t}}{P^*_t}, \qquad S_{X,t} \equiv Q_t V_t, \qquad \Pi_{X,t} \equiv \frac{V_t}{V_{t-1}}. \tag{K.1}$$

The lag V_{t-1} is predetermined. Foreign customers pay the tariff-inclusive price $T^*_t V_t$, whereas Home exporters receive $S_{X,t}$ in Home-PPI units. CES demand therefore gives

$$X_t = \alpha (T^*_t V_t)^{-\gamma} C^*_t, \qquad x_t = -\gamma(\tau_t + v_t), \tag{K.2}$$

where $v_t \equiv \log V_t$. The elasticity γ governs substitution between the aggregate Home export good and Foreign goods. Within the Home export sector, varieties have the same elasticity ϵ used in domestic price setting:

$$X_t(i) = \left(\frac{P^*_{X,t}(i)}{P^*_{X,t}} \right)^{-\epsilon} X_t. \tag{K.3}$$

Hence aggregate export demand involves γ, whereas the individual exporter's markup condition involves $\epsilon - 1$.

For an individual exporter i, the export-market part of nominal profits is

$$\left[\mathcal{E}_t P^*_{X,t}(i) - (1 - s^p) MC_t \right] X_t(i) - \mathcal{E}_t P^*_{X,t} \frac{\kappa_X}{2} \left(\Pi_{X,t}(i) - 1 \right)^2 X_t. \tag{K.4}$$

The aggregate export price in the adjustment-cost term is taken as given by an atomistic firm. Differentiating (K.4), imposing symmetry, and using the household discount factor yields

$$\begin{aligned}\kappa_X(\Pi_{X,t}-1)\Pi_{X,t} = (\epsilon-1)\left(\frac{mc_t}{S_{X,t}}-1\right) \\ + \beta\kappa_X\mathbb{E}_t\left[(\Pi_{X,t+1}-1)\Pi_{X,t+1}\frac{X_{t+1}}{X_t}\left(\frac{C_t}{C_{t+1}}\right)^{\sigma}\frac{G_t}{G_{t+1}}\frac{S_{X,t+1}}{S_{X,t}}\right].\end{aligned} \tag{K.5}$$

Here $mc_t = m_t^r/\bar{v}$ is normalized marginal cost, so the markup term uses the same fixed production correction as the main model.

Because the adjustment expenditure in (K.4) is valued at the export-revenue price, dividing by $P_{H,t}$ gives the resource constraint

$$Y_t = C_{H,t} + X_t + \frac{\kappa_H}{2}(\Pi_{H,t}-1)^2 C_{H,t} + \frac{\kappa_X}{2}(\Pi_{X,t}-1)^2 S_{X,t}X_t. \tag{K.6}$$

The factor $S_{X,t}$ converts export-price adjustment expenditure into Home-PPI resource units. The same conversion produces the ratio $S_{X,t+1}/S_{X,t}$ in (K.5), so it enters both the pricing condition and aggregate feasibility.

The Ramsey planner maximizes the primitive household objective

$$\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\left[\frac{C_t^{1-\sigma}}{1-\sigma}-\chi\frac{L_t^{1+\varphi}}{1+\varphi}\right] \tag{K.7}$$

subject to Home and export pricing, (K.6), demand, the price index, complete-markets risk sharing, the tariff-level law, the log-volatility-state law, and the Euler-consistent implementation equation. The nominal rate is the instrument, and no Taylor rule or strict-PPI restriction is imposed.

K.2 A three-equation local representation

The first-order static block makes the extra state transparent. Define the appendix-specific static coefficients

$$\begin{aligned} h_q^L \equiv \frac{1-\alpha}{\sigma}+\gamma\alpha, \quad y_q \equiv \omega_C h_q^L, \quad y_v = y_\tau \equiv -\gamma\omega_X, \\ m_q \equiv 1+\varphi y_q, \qquad m_v \equiv \varphi y_v, \qquad m_\tau \equiv \varphi y_\tau, \end{aligned} \tag{K.8}$$

where $\omega_C = \bar{C}_H/\bar{Y}$ and $\omega_X = \bar{X}/\bar{Y}$. Then $y_t = y_q q_t + y_v v_t + y_\tau \tau_t$ and $\widehat{mc}_t = m_q q_t + m_v v_t + m_\tau \tau_t$ when productivity is fixed.

Let κ_H and κ_X denote the domestic and export-price adjustment cost coefficients, and

define the corresponding linearized Phillips-curve slopes

$$\lambda_H \equiv \frac{\epsilon - 1}{\kappa_H}, \qquad \lambda_X \equiv \frac{\epsilon - 1}{\kappa_X}. \tag{K.9}$$

The baseline sets $\kappa_H = \kappa_X = 40$, the main-text value of κ. For a Home-PPI Taylor rule, stack $z_t = (\pi_{H,t}, q_t, v_t)'$ and define $e_v = (0, 0, 1)'$. The entries $\pi_{H,t}$ and q_t are jump variables; v_{t-1} is the predetermined customer-price state carried into date t. The domestic NKPC, export NKPC, asset equation, and identity $\pi_{X,t} = v_t - v_{t-1}$ reduce to

$$A_0^L z_t = A_+^L \mathbb{E}_t z_{t+1} + a_-^L v_{t-1} + b_t^L. \tag{K.10}$$

The matrices in (K.10) are

$$A_0^L = \begin{pmatrix} 1 & -\lambda_H m_q & -\lambda_H m_v \\ 0 & -\lambda_X(m_q - 1) & 1 + \beta - \lambda_X(m_v - 1) \\ \phi_\pi & 1 + \phi_y y_q & \phi_y y_v \end{pmatrix}, \quad A_+^L = \begin{pmatrix} \beta & 0 & 0 \\ 0 & 0 & \beta \\ 1 & 1 & 0 \end{pmatrix}, \tag{K.11}$$

and $a_-^L = (0, 1, 0)'$. The vector of tariff and risk terms is

$$b_t^L = \begin{pmatrix} \mathcal{W}_{H,t} + \lambda_H m_\tau \tau_t \\ \mathcal{W}_{X,t} + \lambda_X m_\tau \tau_t \\ \mathcal{W}_{A,t} - \psi_t - \phi_y y_\tau \tau_t \end{pmatrix}. \tag{K.12}$$

Here $\mathcal{W}_{H,t}$, $\mathcal{W}_{X,t}$, and $\mathcal{W}_{A,t}$ collect the higher-order conditional-moment terms generated by the domestic pricing equation, the export pricing equation, and the asset equation, respectively. Each is signed so that a positive value raises the right-hand side of its equation in (K.10); the auxiliary UIP disturbance ψ_t enters the asset equation with the opposite sign. In the reported tariff-volatility experiment, $\psi_t = 0$ and the realized tariff path is fixed, so only the third-order conditional-moment terms remain. For a realized tariff-level shock, the τ_t terms enter and the three $\mathcal{W}$ terms are zero at first order. Thus $\mathcal{W}_{X,t}$ is the additional third-order term generated by export price setting.

For any disturbance $b_t^L = U_s s_t$ with $s_{t+1} = \rho_s s_t$, conjecture $z_t = K_v v_{t-1} + K_s s_t$ and define the endogenous state root $r_v \equiv e_v' K_v$. Substitution into (K.10) gives the coefficient equations

$$A_0^L K_v = A_+^L K_v e_v' K_v + a_-^L, \tag{K.13}$$

$$G_s^L \equiv A_0^L - A_+^L (K_v e_v' + \rho_s I), \tag{K.14}$$

$$K_s = (G_s^L)^{-1} U_s. \tag{K.15}$$

Proposition K.1 states the conditions under which these equations determine a unique local solution; nonsingularity of G_s^L rules out coincidence of the disturbance persistence ρ_s with an endogenous propagation root. Equations (K.13)–(K.15) show why a single geometric pair such as (χ_π, χ_q) in Proposition 1 is insufficient: the disturbance has its own persistence, while the customer price also inherits last period's export-price state. The tariff-level and uncertainty responses use the same propagation matrix; the uncertainty response loads it with risk coefficients derived from the second-order tariff rules.

Proposition K.1 (LCP local reduction)**.** *Under complete markets and LCP, the private equilibrium with a fixed policy rule reduces locally to* (K.10)*. Suppose the associated matrix pencil is regular, it has exactly one stable generalized root for the predetermined customer-price state,* (K.13) *selects the corresponding unique K_v with $|e_v' K_v| < 1$, and G_s^L is nonsingular. Then the locally determinate response to a geometrically persistent disturbance is characterized by* (K.13)–(K.15)*. Consequently, the tariff-volatility response has three conditional-moment loadings—on the domestic pricing, export pricing, and asset equations—and one inherited customer-price state.*

Proof begins: Proposition K.1.

Substituting the complete-markets static allocation, with coefficients (K.8), into the two linearized pricing equations and the asset equation, and replacing export-price inflation by $v_t - v_{t-1}$, yields (K.10)–(K.12). Substituting $z_t = K_v v_{t-1} + K_s s_t$ and matching the coefficients on v_{t-1} and s_t gives (K.13) and (K.15). Selecting the determinate stable root and using the assumed nonsingularity of G_s^L proves the result.

End of proof: Proposition K.1.

For the Ramsey continuation, the Taylor row is removed and the planner stationarity conditions for every primitive constraint are added. The augmented system keeps the three-variable structure above but also contains inherited commitment multipliers, and we do not derive a compact scalar third-order Ramsey formula for it. Instead, we differentiate the primitive Lagrangian numerically and solve the inherited Ramsey continuation to third order, which lets the two pricing blocks and their costates be chosen jointly.

K.3 Ramsey continuation and quantitative comparison

The Ramsey path is the KKT branch of (K.7) under the full nonlinear constraints. The policy plan was chosen before the volatility innovation, so date 0 inherits the lagged planner multipliers; a new date-0 commitment would instead set them to zero. We maintain the local-determinacy conditions used for the Calvo extension in (J.17), applied to the LCP

system: with (H_L^s, R_L^s) the stable deflating subspace of a regular strong linearization after purely algebraic zero and infinite blocks are removed, and B_L imposing the n_L^b inherited physical-state and multiplier restrictions, we require a hyperbolic finite nonzero part and

$$\dim \operatorname{im} H_L^s = n_L^b, \qquad \operatorname{rank}\left[B_L H_L^s (R_L^s)^{-1}\right] = n_L^b. \tag{K.16}$$

The root count makes the endpoint system square, and the rank condition makes it invertible, so the continuation is locally unique. As in the Calvo extension, a strict local optimum would additionally require a constraint qualification and a negative-definite reduced Hessian of the Lagrangian; we have not established these second-order conditions or global concavity, and we make no global optimality claim.

Figure K.1 shows that the no-realization uncertainty response is zero through second order and appears at third order. The Ramsey continuation permits a small movement in Home PPI inflation because stabilizing domestic prices alone does not eliminate the export-price distortion. Strict Home-PPI targeting therefore does not reproduce that continuation exactly. For exports, Home PPI inflation, and the policy rate, strict Home-PPI targeting nevertheless stays close to the Ramsey continuation: on impact, exports fall by 0.42 percent under strict PPI and by 0.43 percent under the Ramsey continuation, Home PPI inflation is zero against -0.01 percentage points, and the policy rate falls by 0.26 against 0.20 percentage points. The Taylor rule instead lowers Home PPI inflation by 0.98 percentage points and the policy rate by 1.55 percentage points on impact. Its output response, however, reverses sign relative to producer-currency pricing: output rises by about 0.071 percent on impact under producer-currency pricing but falls under local-currency pricing. Under producer-currency pricing, export demand depends on the terms of trade Q_t, so Taylor-rule easing brings a real depreciation that raises exports by about 0.11 percent. Under local-currency pricing, export demand (K.2) depends instead on the sticky Foreign-currency export price V_t, and exports fall by about 0.46 percent on impact under the Taylor rule, close to the declines under strict PPI and the Ramsey continuation. On impact, strict PPI and the Taylor rule bracket Ramsey output at similar distances: output falls by 0.080 percent under strict PPI, 0.110 percent under the Ramsey continuation, and 0.145 percent under the Taylor rule. In this calibration, strict Home-PPI targeting remains close to the Ramsey continuation for nominal stabilization and exports, but not distinctly closer than the Taylor rule for output on impact.

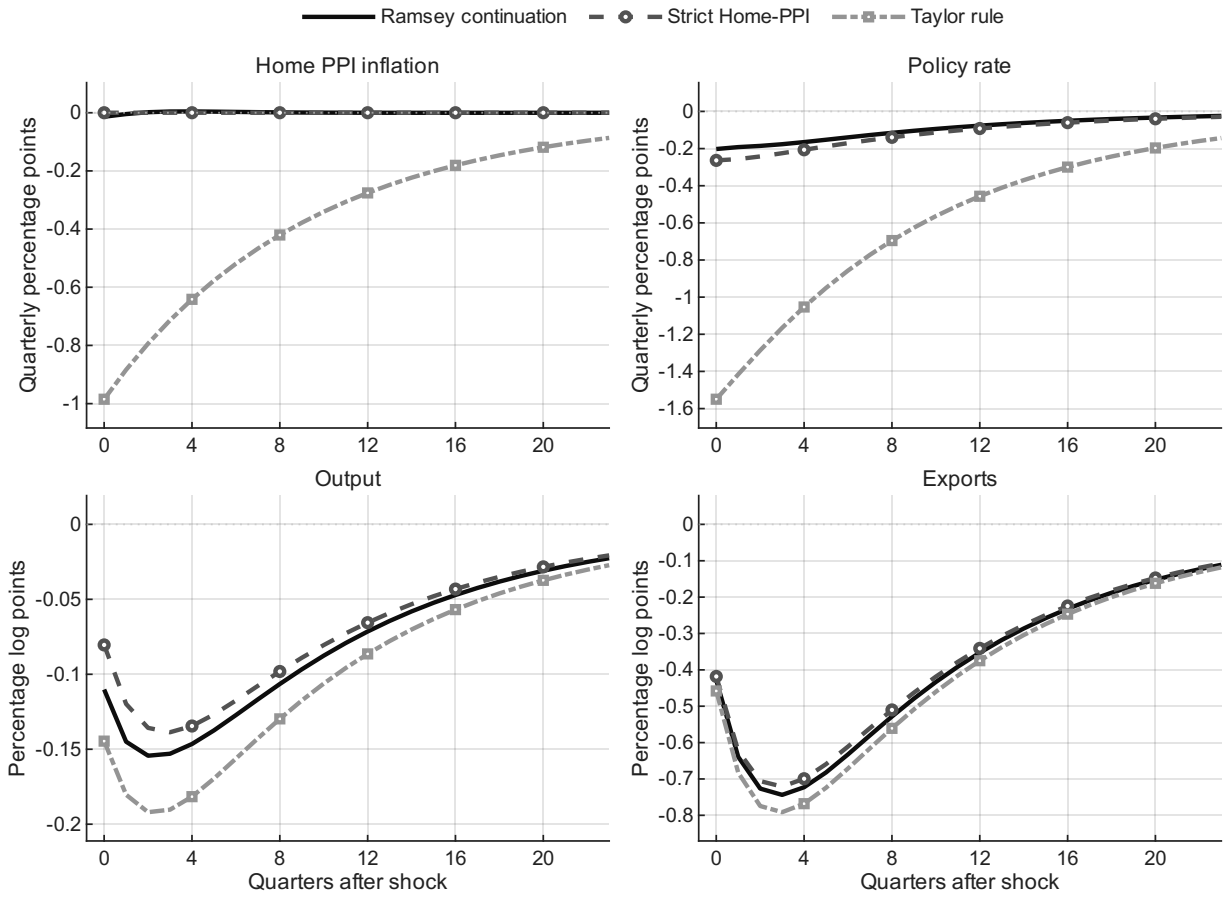


Figure K.1: Ramsey continuation and simple rules under local-currency export pricing

Notes: Complete-markets third-order responses to the log-volatility innovation $\eta_\sigma = \log 2$ in Definition 1. This innovation doubles the conditional standard deviation of the next export-tariff innovation from 0.10 to 0.20, while realized export-tariff innovations are set to zero. Future tariff risk remains in expectations, so the curves are conditional responses rather than averages over future shock histories. The Ramsey continuation inherits the promises made before date 0; strict Home-PPI targeting and the Home-PPI Taylor rule are solved from the same physical state. Inflation and the quarterly policy rate are in percentage points; output and exports are percentage log deviations. The third-order responses are local perturbation objects, not exact nonlinear responses to a finite variance doubling.

Expectations in the forward equations and the augmented KKT system are computed by Gauss–Hermite quadrature over next period's tariff and volatility innovations. When all shocks are rescaled toward the steady state, the conditional KKT residuals decline at approximately fourth order, as a third-order perturbation requires.

Strict Home-PPI stabilization closes only the domestic pricing margin. The export NKPC and the inherited V_t state remain active, which is why the gap between the Ramsey continuation and strict PPI is nonzero and why (K.12) contains the additional export-pricing term $\mathcal{W}_{X,t}$. Its small magnitude in the figure is a quantitative feature of the calibration. When $\mathcal{W}_{X,t} \neq 0$, the productivity and UIP disturbances that reproduce the uncertainty responses

under producer-currency pricing (Proposition E.1) cannot by themselves reproduce the LCP responses; a third, export-pricing disturbance is also required. We have not computed a fourth-order LCP welfare ranking, so Figure K.1 establishes an impulse-response comparison rather than a welfare-equivalence claim.

L Strict PPI Targeting across Environments

Table L.1 collects every environment in which we compute Ramsey policy and compares it with strict PPI targeting and the Taylor rule. The comparison concerns nominal stabilization and welfare. Strict PPI targeting holds PPI inflation at zero by construction, so in PPI inflation its distance from Ramsey policy equals the size of Ramsey's own response.

Two cases are exact, and the table omits them. The first restates Proposition 2: when the production correction aligns the flexible-price and constrained-efficient allocations in every tariff state, strict PPI targeting is the Ramsey policy for every allocation, nonlinearly and under any commitment boundary. The second considers a permanent export tariff of 5, 10, or 20 percent, with the correction fixed at its no-tariff value. In this deterministic economy the stationary Ramsey plan coincides with strict PPI targeting: the two terms-of-trade levels agree to machine precision, and Ramsey PPI inflation is zero. A new date-0 commitment can improve on this plan, but by at most 0.019 basis points of permanent consumption, at the 20 percent tariff. The exactness does not extend to tariff innovations. Around the permanent tariff, Ramsey responds to a realized tariff with about as much PPI inflation as it does at a zero tariff.

The table reports local third-order impact responses along the path with no tariff realizations after date 0; agents continue to account for future tariff risk. In the exact shock process, the uncertainty shock doubles the conditional standard deviation of next quarter's tariff innovation, from 0.10 to 0.20, and leaves the current tariff at zero. The level shock is a 0.10 log-point innovation in the tariff. A new commitment sets inherited planner multipliers to zero at date 0; an inherited commitment continues a plan chosen before date 0. Each row uses the commitment boundary of the exercise in which it is reported, and the table states that boundary.

Ramsey policy essentially looks through tariff uncertainty in producer prices. On impact, its PPI response to the uncertainty shock is at most 0.014 percentage point in every export-tariff environment and 0.033 percentage point under import tariffs with incomplete markets. The export-tariff rows include 27 complete-markets calibrations that combine $\kappa \in \{20, 40, 80\}$, $\rho_\tau \in \{0.5, 0.9, 0.95\}$, and $\rho_\sigma \in \{0.5, 0.9, 0.95\}$. In the same experiments the Taylor rule moves PPI inflation by 0.035 to 3.13 percentage points. Strict PPI targeting is therefore closer to

Ramsey than the Taylor rule in the PPI response to uncertainty in every row. Across the 27 calibrations, this ranking also holds for the discounted distance between the 20-quarter PPI paths, $\left(\sum_{t=0}^{19}\beta^t(\pi_t^j-\pi_t^R)^2\right)^{1/2}$ for policy j. Measured as minus the ratio of the uncertainty to the level PPI response, as in Table 1, uncertainty offsets 1.4 to 5.0 percent of Ramsey's level-shock PPI response across the 27 calibrations. Under the Taylor rule it offsets 24 to 106 percent in the 18 calibrations with $\rho_\tau \in \{0.9, 0.95\}$. With $\rho_\tau = 0.5$, the Taylor rule's level response is below 0.01 percentage point, so the ratio is uninformative.

The level-shock column shows that, outside the exact cases, Ramsey departs from strict PPI targeting in response to realized tariffs, and that the direction of the departure depends on the environment. Ramsey is more expansionary than strict PPI after a realized export tariff in all 27 calibrations: output is 0.30 to 1.05 percent higher and the policy rate 0.25 to 1.01 percentage points lower. The same holds under incomplete markets, where Ramsey cuts the rate by 0.40 percentage point against 0.22 under strict PPI, and after an import tariff under complete markets. After an import tariff under incomplete markets, Ramsey is more contractionary: PPI inflation falls by 0.06 percentage point, the rate rises by 0.22 percentage point against 0.10, and output falls slightly while it rises under strict PPI. Section 4.4 interprets this level-driven departure.

The welfare column is available where we compute a fourth-order welfare ranking against a new-commitment Ramsey plan. In each such case, strict PPI's welfare gap to Ramsey is a small fraction of the Taylor rule's gap: 1.23 percent for export tariffs under complete markets (4.07 against 331.17 basis points of permanent consumption; Table 2), 1.01 percent for import tariffs under complete markets, and 0.28 percent for import tariffs under incomplete markets. The import gaps are lifetime-utility differences without a valid consumption-equivalent conversion (Appendix I.5), so we report only their ratios. These ratios are stable across the four date-0 conditioning events in Figure I.8, and strict PPI is the closest fixed rule in each event. We compute no welfare ranking for the 27 calibrations, Calvo pricing, or local-currency pricing. For export tariffs under incomplete markets, the Ramsey path is an inherited continuation and we report no ranking (Appendix H.6).

Strict PPI's closeness to Ramsey concerns nominal stabilization and welfare, not real allocations. Strict PPI targeting does not reproduce Ramsey's small real adjustment. After the uncertainty shock, the Taylor rule's output path is closer to Ramsey's than strict PPI's in 4 of the 27 calibrations, including the baseline in Figure 5. In 15 of the 27, the Taylor rule's impact output response has the opposite sign from Ramsey's. As Appendix F.12 shows, the baseline overlap of the Ramsey and Taylor real paths reflects the calibration, not a structural equivalence between the two policies.

Table L.1: Ramsey policy, strict PPI targeting, and the Taylor rule across environments

		Impact PPI inflation (percentage points)			
		Uncertainty shock		Level shock	Welfare
Environment	Commitment	Ramsey	Taylor	Ramsey	ratio (%)
Export tariff, complete markets					
Baseline	New	−0.0071	−1.013	0.262	1.23
27 calibrations, range[a]	New	−0.0112 to −0.0039	−3.130 to −0.035	0.196 to 0.291	–
Calvo pricing	Inherited	−0.0014	0.633	–	–
Local-currency pricing	Inherited	−0.0137	−0.985	–	–
Export tariff, incomplete markets					
Baseline	Inherited	0.0018	−0.105	0.126[b]	–
Import tariff					
Complete markets	New[c]	0.0001	−0.130 / −0.161[d]	0.101	1.01
Incomplete markets	Inherited[e]	−0.0327	−0.484 / −0.706[d]	−0.062	0.28

Notes: The uncertainty columns report signed impact PPI inflation after a volatility innovation $\eta_\sigma = \log 2$, which in the exact shock process doubles the conditional standard deviation of next quarter's tariff innovation from 0.10 to 0.20 and leaves the current tariff at zero. The level column reports signed Ramsey impact PPI inflation after a 0.10 log-point innovation in the export or import tariff. Entries are local third-order impact responses, with pruning, along the path with no tariff realizations after date 0, in quarterly percentage points. Strict PPI targeting sets PPI inflation to zero in every row. "New" sets inherited planner multipliers to zero at date 0; "Inherited" continues a plan chosen before date 0. The Taylor rule is (16) with $(\phi_\pi, \phi_y) = (1.5, 0.5)$. The welfare ratio is strict PPI's fourth-order welfare gap to new-commitment Ramsey divided by the Taylor rule's gap (for import tariffs, the smaller of the PPI- and CPI-based rules' gaps). For export tariffs it uses consumption equivalents with the date-0 uncertainty innovation (Table 2); without the innovation it is also 1.23 percent. For import tariffs it uses lifetime-utility gaps (Figure I.8) and conditions on no date-0 realization; across the four conditioning events it lies between 1.008 and 1.009 percent under complete markets and between 0.276 and 0.283 percent under incomplete markets. A dash means that the object is not computed or not defined. [a] Range over $\kappa \in \{20, 40, 80\}$, $\rho_\tau \in \{0.5, 0.9, 0.95\}$, and $\rho_\sigma \in \{0.5, 0.9, 0.95\}$; the baseline is one of the 27 cells. [b] First-order response. [c] Under complete markets the new and inherited boundaries give the same import-tariff responses. [d] PPI-based / CPI-based Taylor rule; the CPI-based rule replaces PPI inflation with CPI inflation. [e] Responses continue a plan chosen before date 0; the welfare ratio uses a new date-0 commitment (Appendix I.1).

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.

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

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

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.

Kitamura, Kyla H., and Danielle M. Trachtenberg. 2026. "USMCA Joint Review: Process and Role of Congress." Congressional Research Service Report R48787.

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

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

Schmitt-Grohé, Stephanie, and Martín Uribe. 2004. "Solving dynamic general equilibrium models using a second-order approximation to the policy function." *Journal of Economic Dynamics and Control*, 28(4): 755–775.

USTR. 2026. "Ambassador Greer Issues Statement on the USMCA Joint Review." Office of the United States Trade Representative Press release July 1.