Evaluating Exchange Rate Regimes: a Natural Experiment?

Amol Amol
Ashoka University
amol@ashoka.edu.in

Constantino Hevia
Universidad Torcuato Di Tella
chevia@utdt.edu

Juan Pablo Nicolini
Federal Reserve Bank of Minneapolis and Universidad Torcuato Di Tella
juanpa@minneapolisfed.org

September 8, 2026

Preliminary and Incomplete

Abstract

We develop a model for small open economies with primary commodity production and Calvo price and wage frictions, and ask how far these two standard, minimal ingredients can go in accounting for business cycle behavior in commodity-exporting economies. We discipline the quantitative exercise by comparing two countries in opposite corners of the spectrum of exchange rate policies. Chile opted for a fully floating exchange rate regime, and Ecuador opted to dollarize its economy, the strongest possible form of exchange rate control. Calibrated using only steady-state targets, with no reference to exchange rate data, the model reproduces several business cycle facts for Chile, and we show that in a model that exhibits both price and wage frictions, some form of dirty floating is optimal. The same ingredients, however, fail to account for the behavior of the real exchange rate in Ecuador, a puzzle that survives across a wide range of nominal rigidities and fiscal regimes, and that we also document for Saudi Arabia, another fixed exchange rate commodity exporter.

1 Introduction

A salient feature of emerging economies (EE) is the prominent role that exchange rate developments have in the monetary policy debate. A likely reason is the substantially larger pass-through of exchange rates to prices, relative to developed economies, as documented by Burstein and Gopinath [2014] among others. This observation rationalizes the reluctance of EE policy makers, relative to their peers in developed economies, to let their currencies fluctuate, as emphasized by Calvo and Reinhart [2002], who coined the by now famous "fear of floating" term.

The purpose of this paper is two-fold. First, we develop a standard quantitative model and study the extent to which it can reproduce salient business cycle facts in small open economies. Two main features are added to an otherwise totally standard neoclassical model. First, we explicitly model production and international trade of primary commodities. Second, we add standard Calvo frictions to the setting of both prices and wages. We pay special attention to the volatility of the real exchange rate and how it correlates with other main macro variables. Our main source of fluctuations is exogenous - to the small open economy - shocks to foreign prices of exportable commodities. To evaluate the model, we compare two countries that have many similarities but represent two extremes in terms of exchange rate policy: Ecuador, that adopted an extreme form of a peg since 2000, and Chile, that adopted an almost pure form of floating in 2000. Both countries produce and export sizable quantities of primary commodities.

Second, equipped with a model calibrated to reproduce the data, we evaluate the optimal degree of exchange rate flexibility. In doing so, we follow the literature and assume frictions in the setting of prices and wages. The existence of both frictions make the question of the optimal monetary policy regime a non-trivial one. Indeed, the literature has shown that in the presence of only one of the frictions, the optimal monetary policy is the one that stabilizes the prices that are subject to the friction.[1] In many circumstances, a way to decentralize the optimal allocation is to stabilize prices - through some form of inflation targeting - and let the nominal exchange rate to freely float. An interesting feature of this policy prescription is that at the optimum, policy is independent of the degree of the price frictions.[2] However, as shown in Hevia and Nicolini [2015], such a simple rule ceases to be optimal in the presence of both price and wage rigidities.[3]

[1] See Galí and Monacelli [2005], Hevia and Nicolini [2013], Schmitt-Grohe and Uribe [2016] and references therein.

[2] See Correia et al. [2008] for a formal proof in the context of a closed economy and Hevia and Nicolini [2013] for an extension to a small open economy.

[3] See also Galí and Monacelli [2016] for a related argument.

Our study is particularly motivated by recent concerns in Ecuador regarding the suitability of the choice to adopt the dollar as their currency made two decades ago. There is a recent literature that has criticized the adoption of hard fixed exchange rate regimes, a dollarization being an extreme special case. Another notable example in the region of a hard fixed exchange rate regime, is the currency board adopted by Argentina during the 1990s, that ended in the most dramatic economic crisis in that country since the Great Depression of the 1930s. Motivated by that example, a "fear of fixed exchange rate regimes" has developed in the literature, the most notable example being Schmitt-Grohe and Uribe [2016] (SGU). A first draft Schmitt-Grohe and Uribe [2011] of that paper's running title was 'Pegs and Pain", which quite clearly indicates the paper's conclusions. Their model only exhibits frictions in the setting of wages so a fixed exchange regime is clearly inefficient. Thus, the suboptimality of a peg is assumed, rather than established.

Our paper provides a complementary step towards the question of the optimal policy design in small open economies. By comparing the behavior of Chile and Ecuador during periods of large swings in their export prices, we let the differential behavior of the data in the two countries discipline the computation of the welfare costs (or benefits) of the two regimes.

It is worth emphasizing that the large swings in the prices of the exportable commodities are the key ingredients, since it is the interaction of these shocks, together with the frictions in the setting of wages, that generates the propagation effects that are behind the mechanism in SGU. And it is precisely that interaction that worries policy makers in the region. An important difference between our approach and that of SGU is that they assume shocks to the endowments as a reduced form of more structural shocks that generate movements in the exchange rate. In contrast, we provide an endogenous theory of exchange rate fluctuations, based on shocks to the primary commodity prices, and discipline our calibration using independent data on those prices, never using data on the comovement of the model's endogenous variables to do so. As we show, the resulting model reproduces several features of the data reasonably well, particularly in the case of Chile, in spite of embedding only two standard nominal frictions and nothing beyond them.

A main contribution of the paper is to compare the experiences of Ecuador and Chile to discipline the quantitative exercise. As mentioned above, Ecuador dollarized its economy over two decades ago. This is the most extreme form of pegging the exchange rate. Chile adopted a pure inflation targeting regime around the same time that Ecuador dollarized, and adopted the nominal interest rate as the policy instrument, letting the exchange rate to freely float explicitly. There were only two instances that lasted only months, in which the Central Bank decided to sell foreign reserves: once in April 2008 and again in January 2011.

On both occasions, the justification for the intervention was essentially the same: the terms of trade were too high and the nominal exchange rate too low. Thus, while the Chilean example is not exactly the one of a pure floater, it is, in our view, the cleanest among the "dirty floaters" in the region. Both countries also share other critical features: both produce and export primary commodities, with a high degree of concentration on a few of them (Copper in Chile and Oil in Ecuador), with exports that are a large share of total output, as we document below.

The attractiveness to evaluate and eventually estimate models using these two extreme cases is that both countries faced very large and persistent changes in the prices of their main exports during these two decades. There are no natural experiments in macro. Our subtitle reflects our belief that the comparison between Chile and Ecuador is the closest we can get to a natural experiment to evaluate and test our theories.

The model we build for this purpose is deliberately parsimonious. It is a purely neoclassical framework — standard preferences, competitive commodity and final-good production, monopolistic competition in intermediate goods — augmented with just two nominal frictions, Calvo price and wage rigidity. Every technology parameter is pinned down by a steady-state accounting identity rather than by matching a business-cycle moment, and every shock process is estimated independently on the corresponding exogenous series in the data; at no point does the calibration make use of data on the comovement of the endogenous variables. The question we ask is how far such a minimal set of ingredients can take us in accounting for the joint behavior of the real exchange rate, output, and commodity prices, before turning to a more comprehensive quantitative exercise. The paper should be read as a first step in that direction, and the exercise that follows is very much work in progress.

We first confront the model with the Chilean data. As emphasized above, none of the model's parameters are chosen to match any moment of GDP, the real exchange rate, or their comovement with commodity prices. Judged against that standard, the results are encouraging: the calibrated model reproduces the sign, and in several cases the approximate magnitude, of the volatility of the real exchange rate and of its correlation with commodity prices and GDP, without ever being asked to. We view this as evidence that the mechanism we emphasize — price stability translating commodity-price shocks into exchange-rate movements — is quantitatively relevant, not as a claim that the model offers a complete quantitative account of the Chilean business cycle. For the calibration, we assume a fully floating regime, so as to be consistent with the Chilean experience.

Equipped with this model, we then turn to a welfare analysis, which we present as an illustration of a simple but important point: with two distinct nominal distortions — sticky prices and sticky wages — and a single policy instrument, the degree of exchange-

rate flexibility, a genuine trade-off emerges. We compare the two extreme cases for Chile and show that a pure floating regime dominates, in a welfare sense, a fixed exchange rate. More interestingly, the optimal policy is interior: some form of managed, or "dirty," floating dominates both extremes. This is, in our model, the formal sense in which a degree of "fear of floating" is optimal. The exact amount depends and is rather sensitive to the particular calibration, but the qualitative point — that a single instrument cannot simultaneously stabilize two distinct nominal distortions — does not.

We then subject the model to a considerably harder test: without changing a single parameter other than the policy rule, we ask whether switching to a peg — our model of Ecuador's dollarization — reproduces the Ecuadorian data with the same combination of ingredients that worked reasonably well for Chile. It does not. Most prominently, the correlation between the real exchange rate and the price of the exportable commodity and the impulse response of the real exchange rate to a commodity price shock have the opposite sign in the model and in the data. The exchange-rate regime and the price and wage frictions, on their own, cannot account for the differential behavior of Chile and Ecuador. We then add one further ingredient that is empirically relevant and common among oil exporters: a fiscal policy, in the form of an export levy, that stabilizes the domestic price of oil. Even this addition, while it moves some moments in the right direction, does not overturn the sign of the correlation between the real exchange rate and the price of the exportable that characterizes the Ecuadorian data. We view this as a genuine and informative failure: it tells us that whatever accounts for Ecuador's experience, it is not simply the interaction of the exchange-rate regime with nominal rigidities, nor the specific fiscal instrument oil producers commonly use. Identifying that additional mechanism is a natural next step, which we leave for future work.

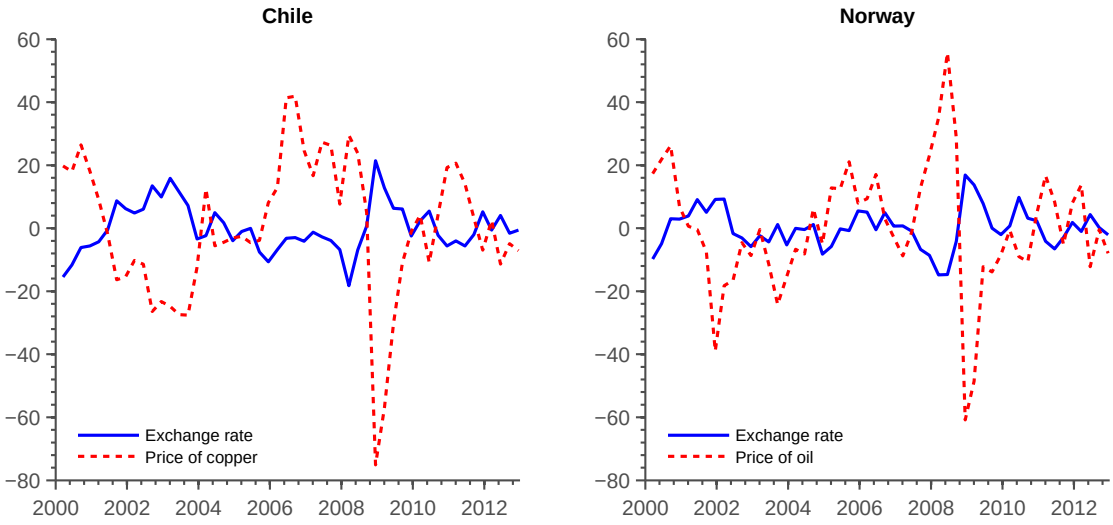
This paper is part of a research project that is motivated by the experiences of many small open economies that in the last two decades became—very effectively!—inflation targeters. Many of these economies are commodity producers, and the size of the commodity sector to GDP is very high.

For example, in the decade from 2000 to 2010, exports of copper and marine products for Chile were, on average, around 17 percent of GDP, while exports of oil and marine products in Norway accounted for 20 percent of GDP [Hevia and Nicolini, 2013]. During the same decade, the real price of copper and oil experienced changes of around 300%.

The size of these shocks is orders of magnitude above any business cycle shock we have ever seen. In addition, they have direct effects on the monetary side of these economies. For example, the correlation between the HP-filtered price of the exportable commodity and the HP-filtered nominal exchange rate—to concentrate our analysis at the business cycle

frequency—ranges from -50% to -70% for both Norway and Chile, depending on whether one includes the last three periods of the decade in the sample. That period of time witnessed very large changes in commodity prices, comoving with the peso in Chile and the krone in Norway.[4]

Figure 1: **Nominal exchange rate and commodity prices**



Note: Quarterly data. HP-filtered with a smoothing parameter of 1600.

In Figure 1, we plot the HP-filtered data for the nominal exchange rate and the relevant commodity price for both countries, where the correlation is strikingly clear. Clearly, this correlation obviously depends on the policy regime. In a dollarized economy like Ecuador, that correlation is naturally zero! But in a successful inflation-targeting regimes, like the one implemented in Chile, the nominal exchange rate freely floats, so the correlation is determined by market forces. According to the model that follows, the negative correlation is a direct implication of the inflation-targeting regime. Indeed, in the model, following a large increase in the price of the exportable, and given price stability, the nominal exchange rate—and the real, given price stability—suffers a strong appreciation, generating the traditional Dutch disease effect, which is the optimal response of prices and quantities to a relative price shock. But in a dollarized economy the adjustment ought to be different.

We present the model in Section 2. Section 3 presents the calibration for Chile and evaluates the model's positive implications, comparing simulated and empirical business-cycle moments and impulse responses. Section 4 turns to the policy analysis and evaluates the optimal degree of exchange rate flexibility. Section 5 confronts the model with the

[4]The series are first logged and then HP-filtered with a smoothing parameter of 1600.

Ecuadorian experience under a fixed exchange rate regime and takes stock of what the exercise has, and has not, been able to account for.

2 The model

We study a discrete-time small open economy inhabited by households, the government, competitive firms that produce a home primary commodity, competitive firms that produce a single non-traded final good, and a continuum of firms that produce differentiated non-traded intermediate goods. The final good is used for consumption, investment, and government purchases; it is not exported, and there is no competing imported final good. Openness operates entirely through the commodity market: production of intermediate goods requires the home commodity, which is exported, and an imported commodity, both of which are freely traded and priced in world markets. Capital is accumulated subject to adjustment costs, and the economy borrows from and lends to the rest of the world through a single non-contingent bond, whose interest rate carries a debt-elastic premium. Monetary policy follows a general rule, indexed by a single parameter ν, that nests a fully flexible exchange rate and a fixed exchange rate (a peg, or dollarization in the limit) as the two extremes of a continuum.

Throughout, we work directly with the detrended, stationary equilibrium system. Real variables are denoted in lower case: $w_t = W_t/P_t$ is the real wage, $s_t = S_t P_t^*/P_t$ is the real exchange rate (the nominal rate S_t deflated by the ratio of a foreign to the domestic price level), and P_t^{x*}, P_t^{m*} are the foreign-currency prices of the home and imported commodities. D_t denotes external debt ($D_t > 0$: the country is a net debtor), $\pi_t = P_t/P_{t-1}$ is domestic CPI inflation, and γ_g is the government spending share (the subscript distinguishes it from the coefficient of relative risk aversion, γ). Derivations that are routine but tedious are collected in Appendix A; here we lay out the primitives and the resulting first-order conditions, and comment on their economic content.

2.1 Households

A representative household is composed of a continuum of members indexed by $h \in [0, 1]$, each supplying a differentiated labor variety n_{ht} under monopolistic competition, exactly as in Erceg, Henderson and Levin (2000). We assume complete consumption insurance across household members, so that the household behaves as a single decision-maker with preferences

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\gamma}}{1-\gamma} - \varsigma \frac{(\bar{L} - L_t)^{1+\psi}}{1+\psi} \right], \tag{1}$$

where C_t is consumption of the (single, non-traded) final good, $L_t = \bar{L} - \int_0^1 n_{ht} dh$ is leisure, $N_t \equiv \int_0^1 n_{ht} dh$ is aggregate labor, γ is the coefficient of relative risk aversion, and ψ is the inverse Frisch elasticity of labor supply. Marginal utilities are

$$U_{1t} = C_t^{-\gamma}, \qquad U_{2t} = \varsigma(\bar{L} - L_t)^{\psi}. \tag{2}$$

Internationally, the household saves and borrows through a single, non-contingent bond denominated in tradable-goods units. Let r_t be its interest rate, augmented by a debt-elastic premium as in Schmitt-Grohé and Uribe (2003) over a stochastic world interest rate r_t^*,

$$r_t = r_t^* + \chi\left(e^{D_t - \bar{D}} - 1\right), \tag{3}$$

where r_t^* follows an AR(1) around its steady-state value $\bar{r}^*$, $\log\left(r_t^*/\bar{r}^*\right) = \rho_{r^*} \log\left(r_{t-1}^*/\bar{r}^*\right) + \varepsilon_{r^*t}$. The debt-elastic term $\chi\left(e^{D_t - \bar{D}} - 1\right)$ pins down a unique, stationary steady-state debt level $\bar{D}$ and closes the small-open-economy model in the standard way. Combining the household's Euler equation for this bond with the definition of the real exchange rate s_t (Appendix A, Section A.1) gives, in terms of the co-state $\lambda_t \equiv U_{1t} s_t$,

$$\lambda_t = U_{1t} s_t, \qquad \lambda_t = (1 + r_t)\, \beta E_t \lambda_{t+1}. \tag{4}$$

The role of s_t here is purely one of units: since the bond is denominated in tradable-goods units while utility is a function of the domestic consumption bundle, s_t converts the marginal utility of domestic consumption into the corresponding value in tradable-goods units, so that (4) is a standard consumption Euler equation expressed in a common numeraire.

Labor supply by household member h is set under Calvo (1983) wage rigidity: with probability $1-\alpha^w$, member h is free to reoptimize her wage in a given period; with probability α^w, her previous wage carries over. Members supply labor inelastically at whatever wage they are quoted, and the household pools their income. Because average labor per member exceeds the amount that would be supplied absent wage dispersion, we define wage dispersion $\Delta_t^w \geq 1$ by

$$\bar{L} - L_t = \Delta_t^w N_t. \tag{5}$$

Aggregating the reset wages of workers who get to reoptimize (a standard Calvo recursion;

full derivation in Appendix A.2) delivers the wage-setting block

$$\Delta_t^w = (1-\alpha^w)\left[\frac{1-\alpha^w\left(\pi_t^w\right)^{\theta^w-1}}{1-\alpha^w}\right]^{\frac{\theta^w}{\theta^w-1}} + \alpha^w\left(\pi_t^w\right)^{\theta^w}\Delta_{t-1}^w, \tag{6}$$

$$\frac{\Psi_t^w}{\Phi_t^w} = (1-\tau^n)\frac{\theta^w-1}{\theta^w}\left[\frac{1-\alpha^w\left(\pi_t^w\right)^{\theta^w-1}}{1-\alpha^w}\right]^{\frac{1}{1-\theta^w}}, \tag{7}$$

$$\Psi_t^w = U_{2t}N_t + \beta\alpha^w E_t\left[\left(\pi_{t+1}^w\right)^{\theta^w}\Psi_{t+1}^w\right], \tag{8}$$

$$\Phi_t^w = U_{1t}w_tN_t + \beta\alpha^w E_t\left[\left(\pi_{t+1}^w\right)^{\theta^w-1}\Phi_{t+1}^w\right], \tag{9}$$

where $\pi_t^w = (w_t/w_{t-1})\,\pi_t$ is wage inflation, θ^w is the elasticity of substitution across labor varieties, and the labor subsidy τ^n eliminates the steady-state wage markup (below). Equations (6)–(9) are the direct analogue of the price-setting block below, with labor playing the role of the differentiated good.

2.2 Investment and capital accumulation

The household also owns the capital stock and rents it to intermediate good firms at the real rental rate r_t^k. Capital is accumulated subject to a standard convex adjustment cost, $\frac{\phi}{2}\left(K_{t+1}/K_t-1\right)^2K_t$, giving the investment Euler equation (Appendix A.3)

$$1+\phi\left(\frac{K_{t+1}}{K_t}-1\right) = \beta E_t\left[\frac{U_{1,t+1}}{U_{1t}}\left(r_{t+1}^k+1-\delta+\phi\left(\frac{K_{t+2}}{K_{t+1}}-1\right)\frac{K_{t+2}}{K_{t+1}}-\frac{\phi}{2}\left(\frac{K_{t+2}}{K_{t+1}}-1\right)^2\right)\right], \tag{10}$$

where δ is the depreciation rate and ϕ the adjustment-cost parameter. Investment is $I_t = K_{t+1}-(1-\delta)\,K_t$.

2.3 Government

Fiscal policy consists of a labor income tax τ^n, a payroll tax τ^p levied on firms' wage bill, and government purchases $G_t = \gamma_g Y_t$ of the final good, financed by a residual lump-sum transfer (so fiscal policy is Ricardian and plays no independent role in the transmission mechanism). We assume all pure rents — the profits of intermediate-good firms and of commodity producers — are fully taxed, and we choose the constant pair (τ^n, τ^p) so that the non-stochastic steady state is efficient, i.e. so that the combined effect of monopolistic

distortions in the goods and labor markets is exactly offset. As shown below, this requires

$$(1-\tau^n)\,\frac{\theta^w-1}{\theta^w}\frac{\theta-1}{\theta}=1, \tag{11}$$

where θ is the elasticity of substitution across intermediate goods, together with a level restriction on τ^p derived in Section 3.1 below. Because τ^n and τ^p are held constant, they affect only levels and drop out of the log-linearized dynamics; their sole role is to guarantee an efficient steady state.

2.4 Final-good firms

Perfectly competitive firms produce the single final good Y_t by aggregating a continuum of differentiated intermediate varieties $i \in (0,1)$,

$$Y_t=\left[\int_0^1 y_{it}^{\frac{\theta-1}{\theta}}di\right]^{\frac{\theta}{\theta-1}},$$

Cost minimization gives the usual demand for variety i and the price index

$$P_t=\left(\int_0^1 P_{it}^{1-\theta}di\right)^{\frac{1}{1-\theta}}. \tag{12}$$

There is, however, only one final good: there is no separately traded foreign final good, no Armington aggregation between home and foreign varieties, and no export demand for the final good. The economy's openness to trade operates exclusively through the commodity market, described next.

2.5 The primary commodity sector

Two commodities, x (produced at home) and m (imported), are used as inputs by intermediate-good firms. The home commodity is produced by competitive firms using labor and a fixed factor (land), normalized to unit supply, combined with a constant elasticity of substitution ω:

$$X_t=A_tB_t^{\frac{\omega}{\omega-1}},\qquad B_t\equiv(1-\rho)+\rho\left(\frac{n_t^x}{\bar{n}^x}\right)^{\frac{\omega-1}{\omega}}, \tag{13}$$

where A_t is productivity, n_t^x is labor employed in the commodity sector, $\bar{n}^x$ is a normalizing constant (chosen so that the steady state is invariant to ω), and ρ is the steady-state labor share in the sector. As $\omega\to 1$, (13) collapses to the Cobb-Douglas technology $X_t=A_t\,(n_t^x)^{\rho}$, up to a rescaling of A_t; the CES form generalizes it by allowing the labor share in commodity

production to respond to relative price movements governed by ω.

Profit maximization by commodity producers — who face payroll tax τ^p on their wage bill — gives

$$\rho P_t^{x*} A_t B_t^{\frac{1}{\omega-1}} \left(\frac{n_t^x}{\bar{n}^x}\right)^{-\frac{1}{\omega}} \frac{1}{\bar{n}^x} = \frac{w_t}{s_t}\left(1+\tau^p\right), \tag{14}$$

which is verified directly in Appendix A.4 and reduces, at $\omega = 1$, to the Cobb-Douglas commodity sector's first-order condition. Both commodities are freely traded, so the law of one price holds, $P_t^x = S_t P_t^{x*}$ and $P_t^m = S_t P_t^{m*}$.

2.6 Intermediate-good firms

Each intermediate variety i is produced by a monopolistically competitive firm using the home commodity, the imported commodity, labor, and capital, under a Cobb-Douglas technology with shares $(\eta_1, \eta_2, \eta_3, \eta_4)$ summing to one. Cost minimization (Appendix A.5) delivers input demands proportional to real marginal cost mc_t and revenue,

$$x_{it} = \eta_1 \frac{mc_t Y_t}{P_t^{x*} s_t}, \quad m_{it} = \eta_2 \frac{mc_t Y_t}{P_t^{m*} s_t}, \quad n_{it}^y = \eta_3 \frac{mc_t Y_t}{w_t\left(1+\tau^p\right)}, \quad K_{it} = \eta_4 \frac{mc_t Y_t}{r_t^k}, \tag{15}$$

and a real marginal cost function

$$mc_t = \frac{\left(P_t^{x*}\right)^{\eta_1}\left(P_t^{m*}\right)^{\eta_2}\left[w_t\left(1+\tau^p\right)\right]^{\eta_3}\left(r_t^k\right)^{\eta_4}}{Z_t} s_t^{\eta_1+\eta_2}, \tag{16}$$

where Z_t is intermediate-sector productivity: firms hire capital as well as labor and the two commodities, and the payroll tax loads on the wage term only. The marginal cost is what firms equate, at reoptimization, to the discounted expected path of future marginal costs under Calvo (1983) price rigidity, with reoptimization probability $1-\alpha^p$:

$$\Delta_t^p = \left(1-\alpha^p\right)\left[\frac{1-\alpha^p \pi_t^{\theta-1}}{1-\alpha^p}\right]^{\frac{\theta}{\theta-1}} + \pi_t^\theta \alpha^p \Delta_{t-1}^p, \tag{17}$$

$$\frac{\Psi_t^p}{\Phi_t^p} = \frac{\theta-1}{\theta}\left[\frac{1-\alpha^p \pi_t^{\theta-1}}{1-\alpha^p}\right]^{\frac{1}{1-\theta}}, \tag{18}$$

$$\Psi_t^p = U_{1t} Y_t mc_t + \alpha^p \beta E_t\left[\pi_{t+1}^\theta \Psi_{t+1}^p\right], \tag{19}$$

$$\Phi_t^p = U_{1t} Y_t + \alpha^p \beta E_t\left[\pi_{t+1}^{\theta-1} \Phi_{t+1}^p\right], \tag{20}$$

where Δ_t^p is price dispersion (aggregate output gross of which is $Y_t \Delta_t^p$), derived in Appendix A.6. The efficient-steady-state condition τ^p mentioned above is the level of the payroll-tax wedge, jointly with τ^n in (11), that sets $mc = (\theta-1)/\theta$ at the non-stochastic steady state

while simultaneously clearing the wage-setting condition (7) at $U_2/U_1 = w$: an efficient-steady-state labor subsidy operating jointly through the labor-income tax and the payroll tax.

2.7 Foreign sector, external debt, and market clearing

Net exports of the two commodities, measured in tradable-goods units, are

$$P_t^{x*}\left[X_t - x_t\right] - P_t^{m*} m_t,$$

where $x_t = \int_0^1 x_{it} di$ and $m_t = \int_0^1 m_{it} di$. The country's net foreign asset position accumulates according to

$$\frac{D_{t+1}}{1+r_t} + P_t^{x*}\left[X_t - x_t\right] - P_t^{m*} m_t = D_t, \tag{21}$$

the small open economy's analogue of the current account identity, with the interest rate given by the debt-elastic premium (3). Market clearing in the final good and the labor market is

$$\begin{aligned} (1-\gamma_g)\, Y_t &= C_t + K_{t+1} - (1-\delta)\, K_t + \frac{\phi}{2}\left(\frac{K_{t+1}}{K_t} - 1\right)^2 K_t, & (22)\\ N_t &= n_t^x + n_t^y, & (23) \end{aligned}$$

where $n_t^y = \int_0^1 n_{it}^y di$, and wage inflation is defined by $\pi_t^w = (w_t/w_{t-1})\,\pi_t$. Measured GDP is computed by double deflation at steady-state prices, $Q_t = X_t + Y_t - x_t - m_t$, and investment is $I_t = K_{t+1} - (1-\delta)\, K_t$.

2.8 The policy rule

We abstract entirely from the question of equilibrium determinacy/implementation and simply posit that monetary policy can successfully target a combination of nominal variables, via a devaluation rule written in terms of the general marginal-cost object mc_t defined in (16),

$$\left(\frac{s_t}{s_{t-1}}\right)^{1-\nu}\left(\frac{mc_t}{mc_{t-1}}\right)^{\nu}\pi_t = \bar{\pi}, \tag{24}$$

where $\bar{\pi}$ is the (zero, in our baseline) inflation target. When $\nu = 1$, (24) stabilizes marginal cost — and hence, to a first order, domestic price inflation — exactly, letting the nominal (and real) exchange rate absorb the full brunt of commodity-price and productivity shocks: this is the flexible-exchange-rate regime, our model of Chile. When $\nu = 0$, the nominal exchange rate is instead held constant, so that all of the adjustment falls on marginal cost

and hence on domestic prices: this is a peg, our model of Ecuador, with full dollarization as the natural real-world limit. Values $\nu \in (0, 1)$ trace out the continuum of "dirty floating" regimes in between. We do not allow the payroll tax to respond to shocks as a second, wage-stabilizing instrument: τ^p is held fixed at its efficient-steady-state level throughout, so that ν is the economy's only active policy instrument. This is a deliberate simplification — it lets us isolate the trade-off between price and wage stability that a single instrument, ν, must confront, which is the object of interest in the welfare analysis below.

2.9 Equilibrium

An equilibrium is a set of stochastic processes for the exogenous variables $(A_t, Z_t, P_t^{x*}, P_t^{m*}, r_t^*)$, each following an AR(1) in logs, and allocations, prices, and policy variables satisfying equations (4), (5)–(9), (10), (14), (15)–(20), (21)–(23), and the policy rule (24), together with the definitions $s_t = S_t P_t^*/P_t$, $\pi_t = P_t/P_{t-1}$, $\pi_t^w = (w_t/w_{t-1})\,\pi_t$, and the debt-elastic premium (3). The equilibrium system features capital, investment, and the rental rate r_t^k; there is no traded-final-good block (no Armington aggregator, no foreign demand equation, and no associated tariff or export-tax instruments) and no complete state-contingent asset menu, only the single debt-elastic bond. The system is solved by casting it in state-space form around the zero-inflation, efficient, non-stochastic steady state, with the five shocks, K_t, D_t, and the lagged variables w_{t-1}, Δ_{t-1}^p, Δ_{t-1}^w, s_{t-1}, mc_{t-1} (the last two required because the policy rule (24) is forward-looking in ratios) as the predetermined block, and solving by standard linear rational-expectations methods (Klein's QZ decomposition), verifying the Blanchard-Kahn conditions at every parameterization considered.

3 Positive implications: Chile

We now calibrate the model of Section 2 to the Chilean economy and ask how well it reproduces the country's business-cycle facts, with special attention to the behavior of the real exchange rate. Throughout this section the policy rule (24) is set at $\nu = 1$ — marginal-cost, and hence to a first order price, stability — consistent with Chile's inflation-targeting regime as described in the Introduction. The calibration uses no information on exchange rates whatsoever; all targets are steady-state national-accounts ratios and supply-use shares.

3.1 Calibration

The calibration strategy is one of inversion: five observable steady-state ratios are targeted, and the technology and tax parameters are solved for in closed form from the national-

accounting identities of the non-stochastic steady state. Table 1 reports the five targets — the commodity sector's share of value added, $1 - v$; the trade-balance-to-GDP ratio, tb; the capital-income-plus-profits share, CP; the supply-use import intensity, η_1; and the government share, G/GDP — together with the labor share in commodity production, ρ, and the dynamic parameters, most of which are standard values from the small-open-economy business-cycle literature. The full closed-form derivation — five accounting identities inverted, in order, into the technology shares $(\eta_1, \eta_2, \eta_3, \eta_4)$, the labor share and the sectoral employment split, the levels of the productivity parameters, the steady-state debt stock, and finally the preference-scale parameter ς and the tax pair (τ^n, τ^p) consistent with an efficient steady state — is given in Appendix B.

Table 1: **Chile: calibration targets and dynamic parameters**

Steady-state targets

	Value
Commodity VA share, $1 - v$	0.16
Trade balance/GDP, tb	+0.8%
Capital income + profits, CP	0.35
Import intensity, η_1	0.031
Government share, G/GDP	0.20
Commodity labor share, ρ	0.13

Dynamic parameters

	Value
Risk aversion	$\gamma = 1.1$
Inverse Frisch elasticity	$\psi = 1$
Commodity labor/land elasticity	$\omega = 0.15$
Debt-elastic premium	$\chi = 0.0069$
Calvo price stickiness	$\alpha^p = 0.5$
Calvo wage stickiness	$\alpha^w = 0.82$
Investment adjustment cost	$\phi = 4$
Discount factor	$\beta = 0.99$
Price/wage markup elasticities	$\theta = \theta^w = 6$

Table 2 reports the steady-state values that follow from these targets via the identities of Appendix B. Most are genuinely implied — not chosen to match anything — but two entries reproduce targets already listed in Table 1 (marked with an asterisk) and are included only for a complete reading of the expenditure and technology shares in one place. The implied labor share, $LS = 0.511$, and the implied employment share of the commodity sector, $\bar{n}^x/N = 4.1\%$, are both reasonable given Chilean data on mining employment and compensation; the implied export share, 13.4% of GDP, is likewise in line with the data discussed in the Introduction. These steady-state accounting checks are a useful consistency test of the calibration, but they are not the paper's main test of the model: that comes from the dynamic, genuinely untargeted second moments in Table 3 below, which use none of the targets in Table 1 in any form.

The commodity-price process is estimated directly rather than calibrated by choice: an AR(1) fit to the HP-filtered real price of copper (World Bank Pink Sheet, deflated by the US

Table 2: **Chile: implied steady-state values**

Technology shares $(\eta_1^*, \eta_2, \eta_3, \eta_4)$	$(0.031, 0.153, 0.593, 0.223)$
Labor share	$LS = 0.511$
Commodity-sector employment share	$\bar{n}^x/N = 4.1\%$
Exports/GDP	13.4%
$C/I/G^*/TB^*$ (% of GDP)	$66/13.2/20/+0.8$
Capital/annual GDP	1.32

*Reproduces a target from Table 1 (η_1, G/GDP, tb); not itself an implication of the calibration.

CPI) over 2000Q1–2025Q4 delivers persistence $\rho_{px} = 0.777$ and innovation standard deviation $\sigma_{px} = 0.11$. The remaining shocks — productivity in the commodity and intermediate sectors, A_t and Z_t, and the world interest rate, r_t^* — are not estimated; we set them at conventional values ($\rho_a = \rho_z = 0.95$, $\sigma_a = \sigma_z = 0.007$ for productivity, and $\rho_{r^*} = 0.85$ with a quarterly innovation of 0.0015, about 60 basis points annualized, for the world interest rate) simply to give simulated GDP a plausible amount of non-commodity variation. All shocks are mutually independent, and the import-price shock is switched off throughout this section.

3.2 Business-cycle moments

Table 3 compares six business-cycle moments in the data with their model counterparts: the standard deviations of the real copper price, GDP, and the real exchange rate, and the three pairwise correlations among them. Unlike the steady-state values in Table 2, every single number in the model column here is genuinely untargeted: these are dynamic second moments of a stochastic simulation, and nothing about the calibration — which is entirely static, steady-state accounting plus independently estimated shock processes — has any way of influencing them except through the model's mechanism itself. Data moments are HP(1600) cycles of logged quarterly series over 2000Q1–2025Q4. Throughout the paper, wherever a real exchange rate is constructed from a US price level, we use the US core CPI (excluding food and energy) rather than the headline index: since oil and other commodities enter the headline CPI directly, using it would mechanically induce part of the very comovement between commodity prices and the real exchange rate that we are trying to explain. We do not repeat this qualification below. Model moments are population values, obtained by simulating the linearized model for 16,000 quarters (discarding the first 1,000) and applying the same filter.

It is worth recalling that none of these six numbers was targeted: the technology pa-

Table 3: **Chile: business-cycle moments, 2000Q1–2025Q4**

	Data	Model: copper shock only	Model: all shocks
sd(RER)	6.46	3.74	4.13
sd(GDP)	2.39	0.36	1.79
sd(p^{x*})	17.61	13.04	13.04
corr(RER,p^{x*})	−0.74	−0.87	−0.79
corr(RER,GDP)	−0.45	−0.99	0.07
corr(p^{x*},GDP)	0.47	0.86	0.17

Data and model are HP(1600)-filtered logs; standard deviations in percent. The "copper shock only" column shuts down every shock but the commodity price; the "all shocks" column adds sectoral productivity and the world interest rate, at the conventional values described above.

rameters come from steady-state accounting identities, and the copper-price process from an AR(1) fit to world price data alone. Seen in that light, the copper-price shock, by itself, gets the sign of all three correlations right and lands within a reasonable distance of their magnitude, along with the volatility of the real exchange rate and of the copper price itself — a purely neoclassical model with two nominal frictions was never asked to do any of this. GDP volatility from the commodity shock alone falls well short of the data (0.36% against 2.39%), which is unsurprising: a single, highly persistent relative-price shock is not, on its own, a business-cycle model of aggregate output.

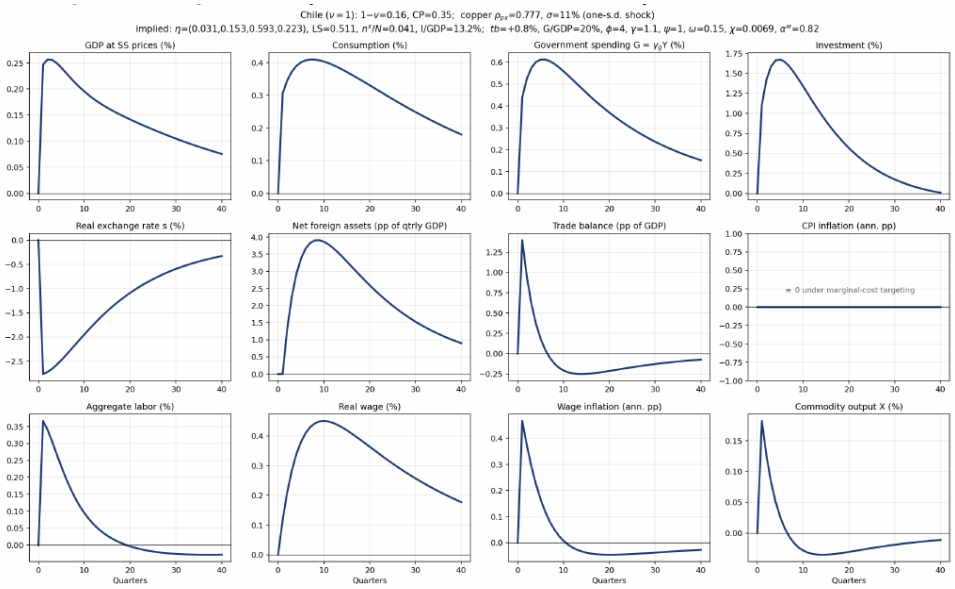
Adding the remaining shocks narrows that gap — GDP volatility rises to 1.79%, still below the data but closer — at a cost elsewhere. In particular, corr(RER,GDP) moves from −0.99, matching the data's sign, to +0.07: with productivity and world-interest-rate shocks now driving a large share of the fluctuations in both series, the model's correlation structure between the real exchange rate and output changes character, and no longer lines up with the data on this particular moment. We report both columns because we think this is the honest way to present the model's reach: the commodity-price channel we emphasize throughout the paper accounts for the exchange-rate side of the data reasonably well on its own, but a fully satisfactory account of Chilean GDP volatility evidently requires more than the three shocks considered here. We take this as a useful discipline on how much weight the rest of the paper should place on the model, rather than as a defect to be argued away.

3.3 Impulse responses

Figure 2 plots the model's impulse responses to a one-standard-deviation shock to the price of copper, at the baseline calibration. The real exchange rate appreciates on impact by close to 3% and reverts within about two years, tracking the persistence of the estimated

copper-price process; GDP, consumption, investment, and the real wage all rise and then decay over a similar horizon, and CPI inflation is identically zero by construction, since $\nu = 1$ stabilizes marginal cost — and hence domestic prices — exactly to a first order.

Figure 2: **Chile: model impulse responses to a one-s.d. copper-price shock**

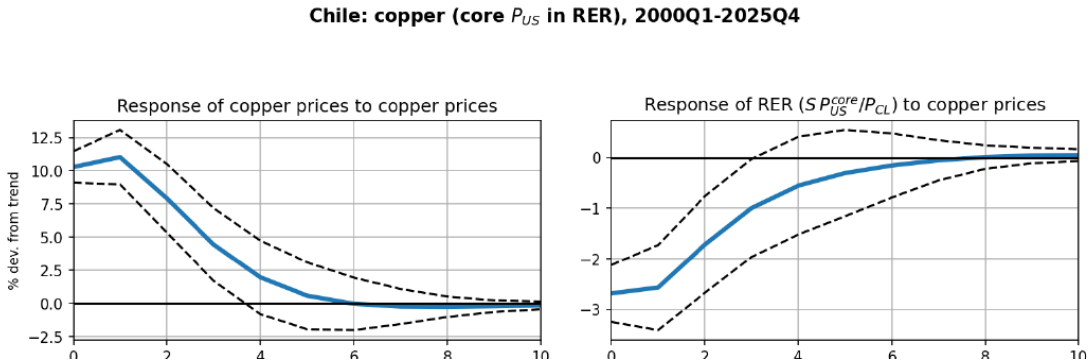


Note: $\nu = 1$; copper price $\rho_{px} = 0.777$, $\sigma_{px} = 11\%$ (one-s.d. shock); baseline calibration of Tables 1–2.

For comparison, Figures 3 and 4 report the corresponding empirical impulse responses from a bivariate VAR estimated on Chilean data over the same sample, with the copper price ordered first. The shape of the exchange-rate response is very similar across model and data — a sharp appreciation on impact, bottoming out within one or two quarters, and reverting to trend by roughly two years — though the empirical trough, at around 3%, is close to the model's. The GDP response is where the shortfall identified in Table 3 is visible directly: the empirical response peaks at close to 0.8% around the second and third quarter after the shock, several times the model's response to the copper shock alone, though the timing and hump shape of the two responses are qualitatively alike.

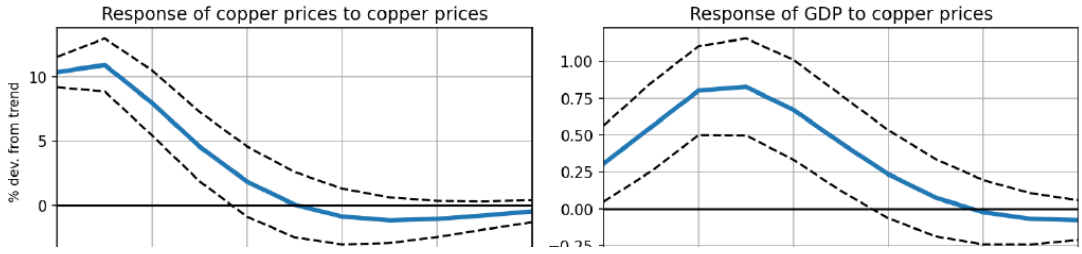
Taken together, the moments and the impulse responses tell a consistent story: a purely neoclassical model with two nominal frictions, driven by commodity-price shocks and calibrated without any reference to exchange-rate data, accounts for the exchange-rate side of the Chilean business cycle reasonably well, and captures a meaningful — if partial — share of the fluctuations in real activity. We take this as evidence that the mechanism is worth taking seriously, and as the starting point for the rest of the paper, rather than as a finished quantitative account of the Chilean business cycle.

Figure 3: **Chile: empirical impulse response of the real exchange rate to a copper-price shock**



Note: Bivariate VAR, 2000Q1–2025Q4, copper price ordered first. Dashed lines are 90% bootstrapped confidence bands.

Figure 4: **Chile: empirical impulse response of GDP to a copper-price shock**



Note: Bivariate VAR, 2000Q1–2025Q4, copper price ordered first. Dashed lines are 90% bootstrapped confidence bands.

4 Policy analysis: a Chilean counterfactual

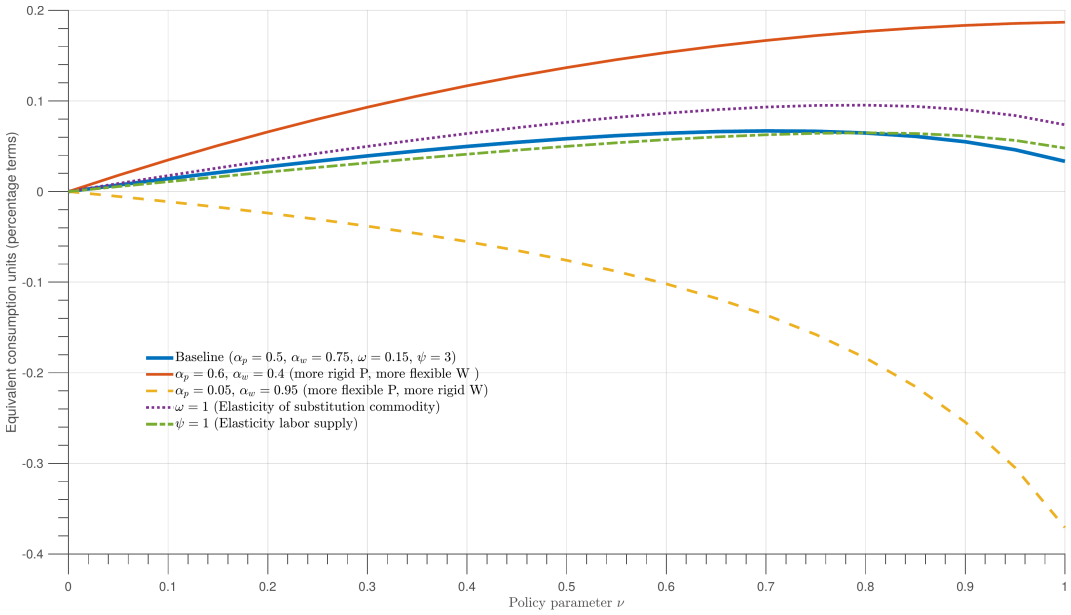
We now turn to the model's normative implications. We stress at the outset that what follows is an illustration, not a full welfare-optimal policy exercise for either country: the calibration underlying this section is Chile's, so the experiment is a Chilean counterfactual — what would happen, in a model calibrated to Chile, if Chile had instead chosen to peg. The point we want to make is a simple one. With two independent nominal distortions — sticky prices and sticky wages — and a single policy instrument, the degree of exchange-rate flexibility ν, the planner cannot in general stabilize both at once, so a genuine trade-off emerges. That trade-off is what disciplines the optimal choice of ν, and it is the sense in which some degree of "fear of floating" can be optimal even absent any of the financial frictions emphasized elsewhere in the literature.

Figure 5 plots the welfare effect, in units of lifetime consumption, of alternative values of $\nu \in [0, 1]$, relative to the hard peg $\nu = 0$, for the baseline parameterization and four variations on it. Two features stand out. First, at the baseline parameterization, the optimal policy is interior: welfare rises with ν up to a value of roughly 0.7–0.75, past which further exchange-rate flexibility starts to reduce welfare, because it comes increasingly at the cost of wage stability. Some managed floating, in other words, dominates both a hard peg and a pure float. Second, the magnitude of the gain from choosing the optimal ν over a hard peg is modest — on the order of a tenth of a percentage point of lifetime consumption at this parameterization — and the gain from choosing the optimum over pure price stability ($\nu = 1$) is smaller still.

The other curves in Figure 5 show how this trade-off depends on the relative degree of price and wage rigidity. When prices are more rigid than wages (the curve with more rigid P, more flexible W), the welfare gain from a more flexible exchange rate is large and monotonic over the range considered: with wages comparatively free to adjust, there is little cost, and a large benefit, to letting the nominal exchange rate absorb the shock. The opposite case — flexible prices and rigid wages — reverses the conclusion, and does so starkly: a pure float is now *worse* than a hard peg, by an amount approaching half a percentage point of lifetime consumption at $\nu = 1$. When wages are the binding constraint, stabilizing prices through a floating exchange rate does nothing to address the distortion that matters, while it actively destabilizes the marginal cost that wages are, in effect, chasing. This is the cleanest illustration in the paper of the point that motivates the whole exercise: a single instrument cannot serve two distinct nominal targets, and which target it should prioritize depends on which friction binds harder.

We regard this exercise as exactly what its title promises: an illustration, using the

Figure 5: **Welfare gains relative to a hard peg, as a function of the policy parameter ν**



Note: Equivalent variation in units of lifetime consumption, relative to $\nu = 0$. Baseline and four one-parameter variations, as labeled.

Chilean calibration of Section 3, of the trade-off that a single exchange-rate instrument faces once both price and wage rigidity are present. It is not, at this stage, a welfare-optimal policy analysis conducted with the capital-augmented model of Section 2; extending it to that model, and eventually to Ecuador's calibration as well, is a natural next step.

5 Ecuador

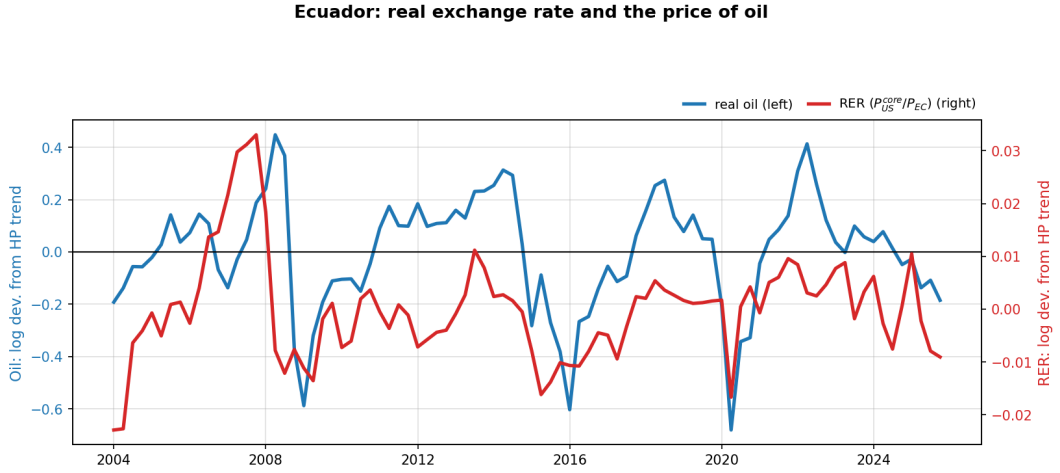
We now confront the same model, at the same calibration, with the Ecuadorian experience. Ecuador dollarized in 2000, giving us a peg-side counterpart to the Chilean float studied in Section 3. We set $\nu = 0$ in the policy rule (24), change nothing else in the calibration of Tables 1–2, and replace the copper-price process with an AR(1) fit to the real price of oil ($\rho_{px} = 0.744$, $\sigma_{px} = 0.14$). Any difference between the model's Chilean and Ecuadorian implications is therefore driven entirely by the change in policy regime, not by a different economy.

5.1 An empirical complication, and how we address it

Ecuador's first years under dollarization were not a normal business cycle: the transition itself — the withdrawal of the sucre, the associated one-time price-level adjustment — dominates the data through roughly 2003. Because that transition is a one-time event rather than the recurring commodity-price mechanism this paper is about, we exclude it: all Ecuadorian data moments and impulse responses reported below use the sample 2004Q1–2025Q4, rather than the full post-dollarization sample used for Chile.

Dropping these years matters a great deal, and not in the direction we expected. The volatility of the real exchange rate falls sharply, from 3.94% (full sample) to 0.99% (2004Q1–2025Q4) — confirming that the transition years accounted for a large share of the variance in the longer sample. But the correlation between the real exchange rate and the oil price, if anything, becomes *more* strongly positive: 0.40 in the restricted sample against 0.27 in the full one. Excluding the transition years sharpens the puzzle rather than resolving it. Figure 6 plots the two HP-filtered series directly, in the same format as Figure 1 for Chile and Figure 10 for Saudi Arabia below: the positive comovement is visually clear throughout the sample, not driven by any particular episode.

Figure 6: **Ecuador: real exchange rate and the price of oil**

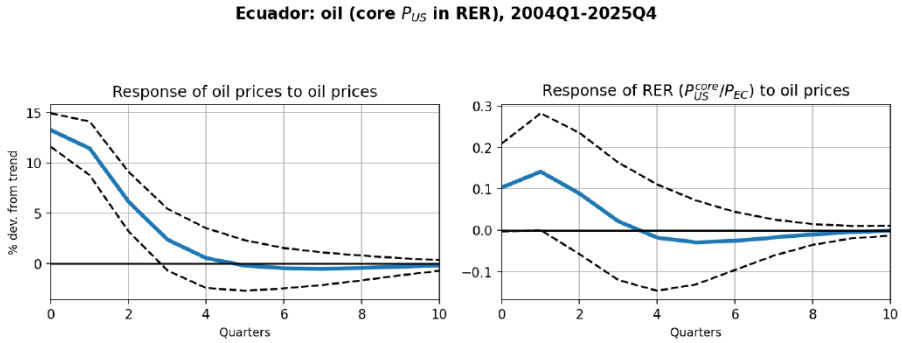


Note: HP-filtered real oil price and real exchange rate, 2004Q1–2025Q4.

It is also worth being precise about what does, and does not, look different in the restricted sample. The bivariate VAR impulse response of the real exchange rate to an oil-price shock (Figure 7) is small and, judging from the confidence bands, not clearly distinguishable from zero at any horizon — a similar picture to what we found before restricting the sample. The static correlation and the dynamic impulse response are not contradictory — a VAR-

identified shock and an unconditional correlation are different objects — but together they mean we cannot claim the "wrong sign" is a dynamic artifact that disappears once the right price index and sample are used: the contemporaneous comovement in the data is real, and, on this measure, slightly stronger than it looked before.

Figure 7: **Ecuador: empirical impulse response of the real exchange rate to an oil-price shock**



Note: Bivariate VAR, 2004Q1–2025Q4, oil price ordered first. Dashed lines are 90% bootstrapped confidence bands.

5.2 The model under a peg

As with the Chilean moments in Table 3, every number in the model columns below is a genuinely untargeted implication of the calibration — the same static, steady-state calibration used for Chile, with only the policy rule and the commodity-price process changed. Table 4 compares the data with the model at $\nu = 0$, oil-producer prices set freely at the world price ("laissez-faire").

Two things stand out, and neither is an improvement on what a naive extrapolation from the Chilean results would have predicted. First, the model's correlations between the real exchange rate and both the oil price and GDP have the *wrong sign*: strongly negative in the model, positive in the data, regardless of whether we restrict attention to the oil shock alone or add the other two shocks. Second, on volatility the ranking has flipped relative to Chile: where the model *understated* Chilean GDP volatility, it now *overstates* Ecuadorian real exchange rate volatility by a factor of two to three, a direct consequence of the data's RER volatility having collapsed once the transition years are excluded. Simply changing the policy rule from $\nu = 1$ to $\nu = 0$, holding the rest of the economy fixed, does not deliver

Table 4: **Ecuador: business-cycle moments, 2004Q1–2025Q4, laissez-faire**

	Data	Model: oil shock only	Model: all shocks
sd(p^{x*})	21.56	16.27	16.27
sd(GDP)	3.02	3.67	4.01
sd(RER)	0.99	2.49	2.66
corr(p^{x*},GDP)	0.62	0.99	0.91
corr(p^{x*},RER)	0.40	−0.57	−0.53
corr(GDP,RER)	0.23	−0.53	−0.44

Data and model are HP(1600)-filtered logs; standard deviations in percent. Same calibration as Chile (Tables 1–2), with $\nu = 0$ and the oil-price process described in the text. Data column uses the 2004Q1–2025Q4 sample throughout.

anything close to the data.

5.3 Does export-price stabilization help?

A natural next step is to ask whether accounting for actual Ecuadorian policy narrows the gap. Many oil-producing countries, Ecuador among them, insulate the domestic economy from swings in the world oil price through some form of price stabilization — typically an export levy whose proceeds (or shortfalls) are absorbed by the government rather than passed through to domestic producers and users. We introduce this as a Ecuador-specific extension, layered on top of the peg, rather than as a feature of the general model of Section 2: a wedge τ_t^x such that the price relevant for domestic commodity producers' and intermediate-good firms' decisions — equations (14) and (16) — is $(1 - \tau_t^x)\, S_t P_t^{x*}$ rather than $S_t P_t^{x*}$, with τ_t^x chosen each period so that this domestic price is held exactly at its steady-state value, $(1 - \tau_t^x)\, S_t P_t^{x*} = \bar{P}^x$. The current account, and hence net exports in equation (21), continue to value the commodity at the true world price P_t^{x*}, since that is what actually crosses the border; the difference between the world-price value and the stabilized-price value of domestic commodity output is a fiscal transfer, rebated to households lump sum. In short, the policy shuts down the domestic commodity-price channel while leaving the country's foreign income unaffected.

Table 5 repeats the exercise under this regime.

Stabilizing the domestic price of oil moves the model in the right direction on two of the six moments — RER volatility falls from 2.49% to 1.85% (versus a data target of 0.99%, so the gap narrows but does not close), and GDP volatility ticks down very slightly closer to the data as well. It does essentially nothing for the moments that matter most for the puzzle we set out to address: both correlations involving the real exchange rate remain

Table 5: **Ecuador: business-cycle moments, 2004Q1–2025Q4, oil-price stabilization**

	Data	Model: oil shock only	Model: all shocks
sd(p^{x*})	21.56	16.27	16.27
sd(GDP)	3.02	3.01	3.41
sd(RER)	0.99	1.85	2.08
corr(p^{x*},GDP)	0.62	1.00	0.88
corr(p^{x*},RER)	0.40	−0.51	−0.46
corr(GDP,RER)	0.23	−0.50	−0.37

Construction as in Table 4. Under stabilization the domestic producer price of oil is held at its steady-state value by an export levy whose proceeds are rebated lump sum.

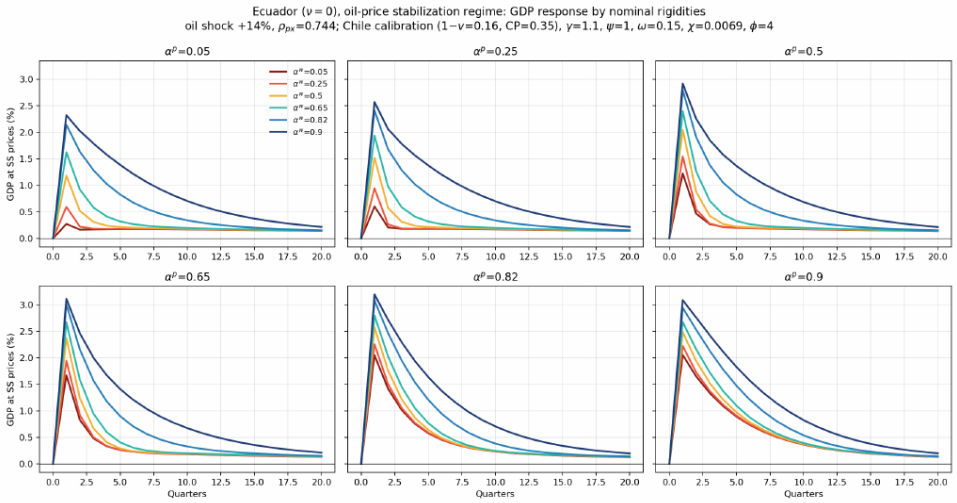
strongly negative, barely moved from the laissez-faire economy. Shutting down the domestic commodity-price channel was, if anything, the natural candidate mechanism for turning the sign of corr(p^{x*},RER) around, since it is precisely that channel that generates the negative correlation in the model. That it does not do so is informative: the wrong sign in the model is not primarily coming from the domestic pass-through of the oil price into costs and wages, which this policy successfully mutes, but from something else in the transmission mechanism that this exercise has not identified.

5.4 Is the sign robust to the degree of nominal rigidity?

A natural question is whether the wrong sign in Table 5 is an artifact of the particular Calvo parameters $(\alpha^p, \alpha^w) = (0.5, 0.82)$ used throughout, rather than a robust feature of the mechanism. Figures 8 and 9 answer this by tracing out the impulse response of GDP and the real exchange rate, under the oil-price stabilization regime, to a 14% oil-price shock, across the full grid $\alpha^p, \alpha^w \in \{0.05, 0.25, 0.5, 0.65, 0.82, 0.9\}$.

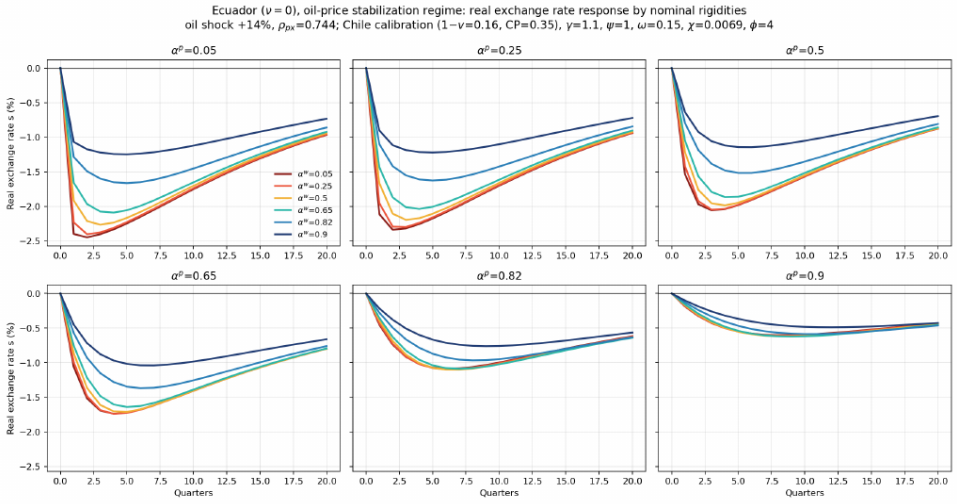
The degree of nominal rigidity matters a great deal for magnitude and shape, but not at all for sign. GDP responses range from barely perceptible (very flexible prices and wages) to a peak of over 3% (very sticky prices, sticky wages), and the real exchange rate response ranges from a shallow, brief dip of about half a percentage point (very sticky prices, $\alpha^p = 0.9$, where the response is nearly flat regardless of α^w) to a deep trough approaching 2.5 percentage points (very sticky prices, flexible wages). Wage stickiness in particular flattens the real exchange rate response: at any given α^p, the most wage-sticky curve ($\alpha^w = 0.9$) is consistently the shallowest. But in every one of the thirty-six parameter combinations plotted, the real exchange rate response to a positive oil-price shock is negative, exactly as in the baseline. The wrong sign relative to the data is therefore not a knife-edge feature of

Figure 8: **Ecuador: GDP response to an oil-price shock, across degrees of price and wage stickiness**



Note: Oil-price stabilization regime, $\nu = 0$; oil shock +14%, $\rho_{px} = 0.744$; Chile calibration. Each panel fixes α^p; curves within a panel vary α^w.

Figure 9: **Ecuador: real exchange rate response to an oil-price shock, across degrees of price and wage stickiness**



Note: Construction as in Figure 8.

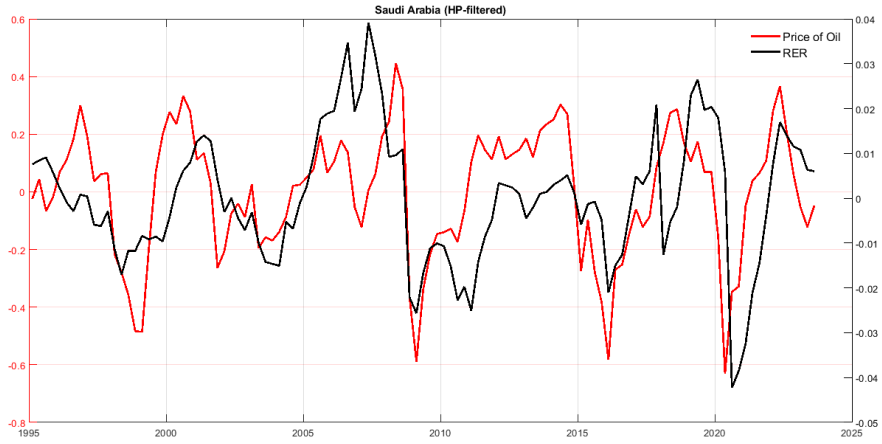
one calibration choice; it survives across the entire plausible range of Calvo price and wage stickiness.

5.5 Is this an Ecuador-specific feature?

A natural concern with a single-country negative result is that it might reflect something idiosyncratic to Ecuador — its particular oil dependence, its particular fiscal institutions, or simply noise in a relatively short time series — rather than a general feature of fixed exchange rate regimes in commodity-exporting economies. We can partly address this by looking at Saudi Arabia, another oil exporter with a long-standing, essentially fixed exchange rate regime (the riyal has been pegged to the US dollar since 1986), using the same construction as for Ecuador and Chile.

Figure 10 plots the HP-filtered bilateral real exchange rate against the price of oil for Saudi Arabia, 1995–2023.

Figure 10: **Saudi Arabia: real exchange rate and the price of oil**

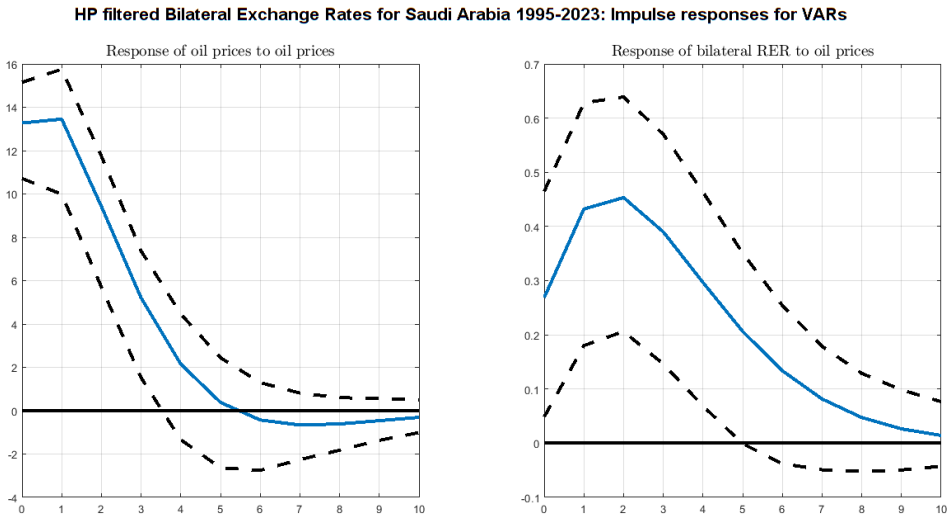


Note: HP-filtered bilateral real exchange rate and real oil price, 1995–2023.

As in Ecuador, the correlation is positive: the real exchange rate depreciates, rather than appreciates, when the price of oil rises. Figure 11 confirms this dynamically with a bivariate VAR, identified in the same way as our Chilean and Ecuadorian systems, with the oil price ordered first.

The response of the bilateral real exchange rate to a positive oil-price shock is positive and, for the first several quarters, clearly outside the confidence bands — the same sign as in Ecuador, and the same wrong sign relative to what a Dutch-disease mechanism under a peg would predict. This is, admittedly, evidence from a single additional country, and

Figure 11: **Saudi Arabia: impulse response of the bilateral real exchange rate to an oil-price shock**



Note: Bivariate VAR, HP-filtered, 1995–2023, oil price ordered first. Dashed lines are 90% bootstrapped confidence bands.

Saudi Arabia differs from Ecuador along many dimensions — the oil share of the economy, the presence of a sovereign wealth fund, the absence of dollarization *per se* (the peg is a conventional one, not a currency union) — so we do not read this as a formal test. But it is at least suggestive that the puzzle we document for Ecuador is not an idiosyncrasy of that country: a similarly signed, similarly positive comovement between the real exchange rate and the oil price shows up in another fixed-exchange-rate oil exporter, which is consistent with the model being wrong about a general feature of pegged commodity exporters, rather than about something particular to Ecuador.

5.6 Taking stock

None of the ingredients we have varied for Ecuador — the exchange rate regime, the presence and degree of price and wage rigidity across the entire plausible Calvo grid, and a realistic fiscal instrument that stabilizes the domestic oil price — overturns the sign of the correlation between the real exchange rate and the price of the exportable commodity that characterizes the Ecuadorian data, a correlation that, if anything, is more pronounced once the dollarization-transition years are excluded. Nor is the puzzle Ecuador-specific: the same sign shows up in Saudi Arabia, another fixed exchange rate commodity exporter. We view this as a genuine and useful negative result rather than a loose end to be explained

away: it tells us that whatever mechanism accounts for the experience of fixed exchange rate commodity exporters like Ecuador and Saudi Arabia is not simply the combination of ingredients that worked reasonably well for Chile. Identifying that mechanism — candidates include a different role for the fixed factor in commodity production, balance-sheet or financial channels of the kind emphasized in the capital-flows literature, or a more structural treatment of the dollarization transition itself — is a central item on our agenda for future work.

References

Ariel Burstein and Gita Gopinath. Chapter 7 - international prices and exchange rates. In Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, editors, *Handbook of International Economics*, volume 4 of *Handbook of International Economics*, pages 391–451. Elsevier, 2014. doi: https://doi.org/10.1016/B978-0-444-54314-1.00007-0. URL https://www.sciencedirect.com/science/article/pii/B9780444543141000070.

Guillermo A. Calvo and Carmen M. Reinhart. Fear of floating. *The Quarterly Journal of Economics*, 117(2):379–408, 2002. ISSN 00335533, 15314650. URL http://www.jstor.org/stable/2696430.

Isabel Correia, Juan Pablo Nicolini, and Pedro Teles. Optimal fiscal and monetary policy: Equivalence results. *Journal of Political Economy*, 116(1):141–170, 2008. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/10.1086/533504.

Jordi Galí and Tommaso Monacelli. Monetary Policy and Exchange Rate Volatility in a Small Open Economy. *The Review of Economic Studies*, 72(3):707–734, 07 2005. ISSN 0034-6527. doi: 10.1111/j.1467-937X.2005.00349.x. URL https://doi.org/10.1111/j.1467-937X.2005.00349.x.

Jordi Galí and Tommaso Monacelli. Understanding the gains from wage flexibility: The exchange rate connection. *American Economic Review*, 106(12):3829–68, December 2016. doi: 10.1257/aer.20131658. URL https://www.aeaweb.org/articles?id=10.1257/aer.20131658.

Constantino Hevia and Juan Pablo Nicolini. Optimal Devaluations. *IMF Economic Review*, 61(1):22–51, April 2013. URL https://ideas.repec.org/a/pal/imfecr/v61y2013i1p22-51.html.

Constantino Hevia and Juan Pablo Nicolini. Monetary Policy and Dutch Disease: The Case of Price and Wage Rigidity. In Rodrigo Caputo and Roberto Chang, editors, *Commodity Prices and Macroeconomic Policy*, volume 22 of *Central Banking, Analysis, and Economic Policies Book Series*, chapter 3, pages 051–089. Central Bank of Chile, March 2015. URL https://ideas.repec.org/h/chb/bcchsb/v22c03pp051-089.html.

Stephanie Schmitt-Grohe and Martín Uribe. Pegs and Pain. CEPR Discussion Papers 8275, February 2011. URL https://ideas.repec.org/p/cpr/ceprdp/8275.html.

Stephanie Schmitt-Grohe and Martín Uribe. Downward nominal wage rigidity, currency pegs, and involuntary unemployment. *Journal of Political Economy*, 124(5):1466 – 1514, 2016. URL https://EconPapers.repec.org/RePEc:ucp:jpolec:doi:10.1086/688175.

A Derivations for Section 2

A.1 The household's problem and the consumption Euler equation

The household's period budget constraint, in nominal domestic-currency terms, is

$$P_t C_t + S_t \left(1 + r_{t-1}\right) D_{t-1} = W_t \left(1 - \tau^n\right) N_t + \Pi_t + S_t D_t,$$

where D_t is the stock of (net) external debt issued at t and repaid at $t+1$ with interest r_t, denominated in tradable-goods (foreign currency) units, and Π_t collects the (fully taxed, hence redistributed lump sum) profits of intermediate-good and commodity firms. Aggregating this across the continuum of household members — who differ only in the wage they currently charge, and who fully pool their labor income and consumption, as is standard under the Erceg-Henderson-Levin assumption of perfect within-household risk sharing — delivers the representative-household budget constraint used in the text.

The household chooses C_t and D_t to maximize (1) subject to this constraint. The first-order condition for D_t is

$$\frac{S_t}{P_t} U_{1t} = \beta \left(1 + r_t\right) E_t \left[\frac{S_{t+1}}{P_{t+1}} U_{1,t+1}\right].$$

Define the real exchange rate $s_t \equiv S_t P_t^* / P_t$, where P_t^* is a foreign price level normalized, without loss of generality for the log-linear system, to grow at the (zero) rate of foreign inflation; then $S_t / P_t = s_t / P_t^*$, and dividing both sides of the first-order condition by the (common, non-stochastic) P_t^* gives exactly

$$s_t U_{1t} = \beta \left(1 + r_t\right) E_t \left[s_{t+1} U_{1,t+1}\right],$$

which, upon defining $\lambda_t \equiv U_{1t} s_t$, is equations (4) in the text. The economics: since the bond pays off in tradable-goods units while the household cares about domestic consumption units, s_t is exactly the exchange rate needed to convert one type of marginal utility into the other; the resulting Euler equation is a completely standard one once expressed in a common numeraire.

A.2 Wage setting

A household member h who is free to reset her wage at t chooses $w_{ht} = w_t$ (by symmetry, all resetters choose the same wage) to maximize the expected discounted utility flow from supplying labor at that wage for as long as it remains in effect, i.e. with probability $(\alpha^w)^j$ at horizon j. Labor demand at h for a given aggregate N_t is the standard Dixit-Stiglitz demand $n_{ht} = \left(w_{ht} / W_t\right)^{-\theta^w} N_t$ (in real terms, $\left(w_{ht} / w_t\right)^{-\theta^w}$ relative to the aggregate real wage w_t).

The member's problem is

$$\max_{w_t} E_t \sum_{j=0}^{\infty} (\alpha^w \beta)^j \left[-\varsigma \frac{\left(\Delta^w_{t+j} N_{t+j}\right)^{1+\psi}}{1+\psi} \left(\frac{w_t}{w_{t+j}}\right)^{-\theta^w(1+\psi)} \left(\frac{w_{t+j}}{w_t}\right)^{\theta^w \psi} + U_{1,t+j} \left(1-\tau^n\right) w_t \left(\frac{w_t}{w_{t+j}}\right)^{-\theta^w} N_{t+j} \right],$$

using (5) to substitute for individual labor supply in terms of $\Delta^w_{t+j} N_{t+j}$. Differentiating with respect to w_t and rearranging (a standard but lengthy manipulation, identical in structure to the price-setting problem below) yields the optimal reset wage as a ratio of two recursively defined objects, Ψ^w_t and Φ^w_t, satisfying (8)–(9), with the optimal markup condition (7). Substituting the reset wage back into the definition of the aggregate wage index $\left[\int_0^1 w_{ht}^{1-\theta^w} dh\right]^{1/(1-\theta^w)} = w_t$ and using the Calvo aggregation across resetters and non-resetters delivers the wage-dispersion recursion (6). The payroll tax does not appear in this block, since τ^p is a tax on firms' wage bill, not on the wage received by workers.

A.3 Investment

The household chooses $\{K_{t+1}\}$ to maximize (1), given the budget constraint augmented with capital income $r^k_t K_t$ and investment expenditure $I_t = K_{t+1} - (1-\delta) K_t$, subject to the adjustment-cost technology

$$K_{t+1} = (1-\delta) K_t + I_t - \frac{\phi}{2} \left(\frac{I_t}{K_t} - \delta\right)^2 K_t,$$

or, in the equivalent and more convenient form used in the text, with the adjustment cost expressed directly on the growth rate of capital. The first-order condition for K_{t+1}, using the envelope condition for the shadow value of capital (equated, in this frictionless-Tobin's-q formulation, to the marginal utility of consumption since there is no separate financial friction on the household's capital position), is

$$U_{1t}\left[1 + \phi\left(\frac{K_{t+1}}{K_t} - 1\right)\right] = \beta E_t \left\{ U_{1,t+1} \left[r^k_{t+1} + (1-\delta) + \phi\left(\frac{K_{t+2}}{K_{t+1}} - 1\right)\frac{K_{t+2}}{K_{t+1}} - \frac{\phi}{2}\left(\frac{K_{t+2}}{K_{t+1}} - 1\right)^2 \right] \right\},$$

which is (10) after dividing through by U_{1t}.

A.4 The commodity sector's first-order condition

Commodity producers choose n^x_t to maximize profit, measured in tradable-goods units, $P^{x*}_t X_t - (w_t/s_t)(1+\tau^p) n^x_t$, where w_t/s_t is the tradable-goods value of the real wage. Differentiating $X_t = A_t B_t^{\omega/(\omega-1)}$ with respect to n^x_t,

$$\frac{\partial X_t}{\partial n^x_t} = A_t \frac{\omega}{\omega-1} B_t^{\frac{1}{\omega-1}} \cdot \rho \frac{\omega-1}{\omega} \left(\frac{n^x_t}{\bar{n}^x}\right)^{-\frac{1}{\omega}} \frac{1}{\bar{n}^x} = \rho A_t B_t^{\frac{1}{\omega-1}} \left(\frac{n^x_t}{\bar{n}^x}\right)^{-\frac{1}{\omega}} \frac{1}{\bar{n}^x},$$

and setting $P_t^{x*}\partial X_t/\partial n_t^x = (w_t/s_t)(1+\tau^p)$ gives exactly (14). We verified this derivative and first-order condition by computer algebra; the same computation confirms that at $\omega = 1$, $X_t = A_t\,(n_t^x/\bar{n}^x)^\rho = A_t\,(\bar{n}^x)^{-\rho}\,(n_t^x)^\rho$, a Cobb-Douglas technology up to a rescaling of A_t, and (14) collapses to $\rho P_t^{x*} A_t\,(n_t^x)^{\rho-1} = (w_t/s_t)(1+\tau^p)$.

A.5 Intermediate-good cost minimization

Each intermediate-good firm solves, for a given required output level $\bar{y}$,

$$\min_{x,m,n^y,K} P_t^{x*} s_t x + P_t^{m*} s_t m + w_t(1+\tau^p)\,n^y + r_t^k K \quad \text{s.t.} \quad \bar{\eta} Z_t x^{\eta_1} m^{\eta_2} (n^y)^{\eta_3} K^{\eta_4} = \bar{y},$$

where $\bar{\eta} = \eta_1^{-\eta_1}\eta_2^{-\eta_2}\eta_3^{-\eta_3}\eta_4^{-\eta_4}$ is the constant that makes the unit cost function take the clean Cobb-Douglas form below. Forming the Lagrangian and solving the four first-order conditions jointly with the constraint — a mechanical but lengthy elimination of the multiplier, which we carried out and checked by computer algebra, both symbolically and against a numerical example — gives the conditional factor demands (15) and the real marginal cost

$$mc_t = \frac{(P_t^{x*} s_t)^{\eta_1}\,(P_t^{m*} s_t)^{\eta_2}\,[w_t(1+\tau^p)]^{\eta_3}\,\left(r_t^k\right)^{\eta_4}}{Z_t},$$

which is (16) once the two powers of s_t are collected, $s_t^{\eta_1+\eta_2}$. As a check, note that setting $\eta_4 = 0$ (no capital) collapses this expression, and the accompanying input demands, to the analogous three-input marginal cost and factor demands, with the payroll tax entering only through the labor term $w_t(1+\tau^p)$.

A.6 Calvo price setting

A reoptimizing intermediate-good firm at t chooses its reset price p_t to maximize expected discounted profits over the (random) interval during which the price remains in effect, taking as given the demand curve $y_{it} = Y_t\,(P_{it}/P_t)^{-\theta}$ implied by the final-good aggregator (12) and the real marginal cost mc_t from (16), which is common across firms. This is the standard Calvo (1983) pricing problem; the algebra — differentiating the discounted profit stream with respect to p_t, substituting the demand curve, and using the household's stochastic discount factor $\beta U_{1,t+j}/U_{1t}$ — is identical in structure to the wage-setting problem above, and delivers (17)–(20). Aggregating the price index (12) across resetting and non-resetting firms, using the fraction $(1-\alpha^p)\,(\alpha^p)^j$ of firms last resetting j periods ago, gives the price-dispersion recursion (17). We omit the remaining lines of algebra, since the price-setting problem depends on marginal cost only through the single object mc_t, whatever its microfoundation.

B Calibration: closed-form derivation

This appendix gives the closed-form inversion underlying Table 1–2. Throughout, we work in units of steady-state GDP and normalize $s = P^{x*} = P^{m*} = 1$ and $N = 1$; a bar denotes a steady-state value, and mc is the steady-state real marginal cost.

B.1 Stage 1: the nominal block and the tax wedge

At the baseline zero-inflation target $\bar{\pi} = 1$, reset prices and wages equal their respective averages, $\Delta^p = \Delta^w = 1$, and the pricing recursions (17)–(20) collapse to $mc = (\theta - 1)/\theta$. The static wage condition similarly requires the combined tax-and-markup wedge to equal one,

$$(1 - \tau^n)\,\frac{\theta^w - 1}{\theta^w}\,\frac{\theta - 1}{\theta} = 1,$$

which is equation (11) in the text and pins down τ^n given $\theta = \theta^w$. The payroll tax τ^p is pinned down jointly, once the technology shares are known (Stage 3 below), by the requirement that the commodity-sector first-order condition (14) and the marginal-cost condition (16) both hold at the efficient allocation; in practice this amounts to a second, analogous wedge restriction on the labor cost entering (16).

B.2 Stage 2: returns

The bond Euler equation (4) at a deterministic steady state gives $1 = \beta(1 + \bar{r})$, so $\beta = 1/(1 + \bar{r}^*)$ and $\bar{r} = \bar{r}^*$; the debt-elastic premium (3) then holds exactly at $D = \bar{D}$, leaving χ free to be disciplined by dynamics alone. The investment Euler equation (10) at $K_{t+1} = K_t$ gives $r^k = \beta^{-1} - (1 - \delta)$; every adjustment-cost term vanishes at the steady state, so ϕ, like χ, is a pure dynamic parameter.

B.3 Stage 3: technology shares from three accounting identities

Two accounting facts convert the targets into share equations. First, double deflation nets out intermediate usage, so measured GDP equals $X_t + Y_t - x_t - m_t = Y_t + TB_t$, i.e. $\bar{Y} = (1 - tb)\,GDP$. Second, cost minimization — equations (15) together with (16) — makes each input's bill an η-share of variable cost $mc\bar{Y}$,

$$\bar{x} = \eta_1 mc\bar{Y}, \qquad \bar{m} = \eta_2 mc\bar{Y}, \qquad w\bar{n}^y = \eta_3 mc\bar{Y}, \qquad r^k\bar{K} = \eta_4 mc\bar{Y},$$

with monopoly profits $\bar{Y}(1 - mc) = \bar{Y}/\theta$.

Trade identity. Net exports are $TB = \bar{X} - \bar{x} - \bar{m}$ with $\bar{X} = (1 - v)\,GDP$, so

$$tb = (1 - v) - (\eta_1 + \eta_2)\,mc\,(1 - tb) \quad \Longrightarrow \quad \eta_1 + \eta_2 = \frac{(1 - v) - tb}{mc\,(1 - tb)}. \tag{B.1}$$

Only the sum $\eta_1 + \eta_2$ is identified by the trade balance — purchasing the home commodity domestically displaces imports one for one — and the split is pinned down separately by the supply-use target η_1, with η_2 residual.

Capital-income identity. The target CP is rental income plus profits, $CP \cdot GDP = r^k \bar{K} + \bar{Y}/\theta$, so

$$\eta_4 = \frac{CP - (1 - tb)/\theta}{mc\,(1 - tb)}. \tag{B.2}$$

The remaining share is $\eta_3 = 1 - \eta_1 - \eta_2 - \eta_4$.

B.4 Stage 4: factor incomes and the labor allocation

The commodity-sector first-order condition (14), at $\omega = 1$ for this purpose, gives $w\bar{n}^x = \rho\bar{X}$; adding $w\bar{n}^y = \eta_3 mc\bar{Y}$,

$$LS \equiv \frac{wN}{GDP} = \rho\,(1 - v) + \eta_3 mc\,(1 - tb)\,, \qquad \bar{n}^x = \frac{\rho\,(1 - v)}{LS}, \qquad \bar{n}^y = 1 - \bar{n}^x. \tag{B.3}$$

Substituting (B.1) eliminates the tradable shares from the labor share entirely,

$$LS = mc\,(1 - tb)\,(1 - \eta_4) - (1 - \rho)\,(1 - v) + tb, \tag{B.4}$$

so that, given mc and the balance of payments, the labor share depends only on the capital share and the sector size, not on the η_1/η_2 split.

B.5 Stage 5: levels and the debt stock

Normalizing A so that $\bar{X} = (\bar{n}^x)^\rho$ (the Cobb-Douglas limit, also the CES normalization point, so the choice is ω-free), the commodity first-order condition gives $\bar{w} = \rho\,(\bar{n}^x)^{\rho-1}$. Levels follow recursively,

$$\bar{Y} = \frac{\bar{w}\bar{n}^y}{\eta_3 mc}, \qquad \bar{K} = \frac{\eta_4}{\eta_3}\frac{\bar{w}}{r^k}\bar{n}^y, \qquad \bar{I} = \delta\bar{K}, \qquad \bar{x} = \frac{\eta_1}{\eta_3}\bar{w}\bar{n}^y, \qquad \bar{m} = \frac{\eta_2}{\eta_3}\bar{w}\bar{n}^y,$$

and $GDP = \bar{X} + \bar{Y} - \bar{x} - \bar{m}$, with Z recovered from (16). The debt stock comes from the steady-state current account (21), $TB = (1 - \beta)\,\bar{D}$,

$$\bar{D} = \frac{tb \cdot GDP}{1 - \beta}, \qquad \frac{debt}{\text{annual } GDP} = \frac{tb}{4\,(1 - \beta)}, \tag{B.5}$$

so the trade balance and the debt-to-GDP ratio are a single target, not two.

B.6 Stage 6: fiscal share and absorption

With $G_t = \gamma_g Y_t$,

$$\gamma_g = \frac{G/GDP}{1 - tb}, \qquad \bar{C} = (1 - \gamma_g)\,\bar{Y} - \delta\bar{K}, \tag{B.6}$$

from the final-good market-clearing condition (22) at $K_{t+1} = K_t$. We verify $\bar{C} > 0$ and that the expenditure shares $C/I/G/TB$ sum to one.

B.7 Stage 7: preference scale and verification

Finally, ς is whatever value makes the static wage condition hold at the constructed allocation, with $N = 1$ and $\Delta^w = 1$:

$$\varsigma = (1 - \tau^n) \frac{\theta^w - 1}{\theta^w} \bar{w} \bar{C}^{-\gamma}. \tag{B.7}$$

Every parameterization is verified numerically to satisfy (i) steady-state residuals below 10^{-10} on all twenty-four equilibrium conditions plus the policy rule, and (ii) exactly the right number of stable generalized eigenvalues for the Blanchard-Kahn conditions to be satisfied, using Klein's QZ decomposition. The steady state of Stages 1–7 is invariant to ω, ψ, χ, ϕ, α^p, and α^w at $\bar{\pi} = 1$; these affect the dynamics only, and γ enters the steady state solely through (B.7).

We verified the values reported in Tables 1–2 against equations (B.1)–(B.4) directly: with $1 - v = 0.16$, $tb = 0.008$, $CP = 0.35$, $\eta_1 = 0.031$, $\rho = 0.13$, and $mc = 5/6$, (B.1)–(B.2) give $\eta_1 + \eta_2 = 0.184$ and $\eta_4 = 0.223$, so $\eta_2 = 0.153$ and $\eta_3 = 0.593$; (B.4) then gives $LS = 0.511$, and $\bar{n}^x = \rho\,(1 - v)\,/LS = 0.041$ — matching Table 2 exactly.