

Misallocation, Dollars, and Aggregate Productivity*

Yasutaka Koike

Antonio Martner

UNC Chapel Hill

Central Bank of Chile

September 2026

Abstract

We show how nominal exchange-rate movements can affect aggregate productivity in a small open economy without changing technology. With sticky invoicing-currency prices and heterogeneous markups, exchange-rate movements reallocate resources across firms with different marginal products. A depreciation of the home currency against the dollar makes imports more expensive and shifts demand toward domestic firms, especially those with low markups and more elastic demand. It also lowers exporters' dollar costs, inducing the same pattern of reallocation when export prices adjust. Because low-markup firms are already inefficiently large relative to high-markup firms, this shift worsens resource allocation. In a quantitative model calibrated to Chile, this reallocation generates a decline in aggregate productivity. Evidence from administrative firm-level data supports the predicted reallocation and productivity decline.

JEL codes: E31, E52, F31, F41, D24, O47.

Keywords: exchange rates; dominant currency pricing; markups; misallocation; measured TFP; Solow residual; sufficient statistics.

*We thank David Baqaee, Ariel Burstein, and Oleg Itskhoki for helpful comments. The views expressed are those of the authors and do not necessarily represent the views of the Central Bank of Chile or its board members. Authors emails: yasu@unc.edu, amartner@bcentral.cl

1 Introduction

In small open economies, aggregate productivity comoves with external conditions and can fall sharply during large exchange-rate depreciations (Mendoza, 1995; Aguiar and Gopinath, 2007; Gopinath and Neiman, 2014). Yet in canonical frameworks, a currency movement does not by itself change what firms can produce, and changes in external prices have no first-order effect on aggregate productivity (Kehoe and Ruhl, 2008). This tension raises a natural question: what mechanism makes exchange-rate movements generate changes in aggregate productivity?

This paper shows that exchange-rate movements can affect aggregate productivity by reallocating resources across firms with heterogeneous markups. The mechanism combines pre-existing allocative distortions with price stickiness in the currency of invoicing, allowing nominal exchange-rate movements to alter the allocation of production without changing technology. We formalize this mechanism in a dynamic small-open-economy model, quantify it using Chilean administrative data, and test its central predictions in the same data. We find that broad dollar movements reallocate activity across firms in the predicted direction and are followed by persistent declines in aggregate productivity.

Consider a home-currency depreciation. Sticky dollar-invoiced import prices rise in home currency while sticky domestic prices remain unchanged, shifting expenditure toward domestic firms. Low-markup firms face relatively more elastic demand and expand their share of domestic production. Because these firms are already inefficiently large relative to high markup firms, this shift worsens resource allocation.

We allow the elasticity of demand to vary along the demand curve. Firms operating at different points therefore face different elasticities and charge different markups. Using a Kimball demand system to capture this variation, we derive how firms' initial markups and their responses to relative-price movements shape aggregate productivity. Under CES demand, markups are common and the first-order productivity effect from reallocation disappears.

We next introduce intermediate inputs and derive an exact decomposition of aggregate productivity into three terms. The first is the original across-firm reallocation effect: a depreciation shifts activity toward low-markup firms with lower marginal products, reducing productivity, and intermediate inputs amplify this loss because reallocation operates on gross production while productivity is measured relative to value added, so the effect is scaled by the gross-sales-to-GDP ratio. The second captures aggregate input allocation. Because domestic intermediates are underused when domestic suppliers charge markups, shifts toward domestic inputs improve allocation, whereas shifts toward

labor or imported inputs worsen it; its sign is therefore an ex-ante ambiguous equilibrium outcome. The third is an import-measurement bias: measured GDP and the cost-based productivity benchmark assign different exposures to imported inputs, so its sign follows the response of imported-input quantities and is likewise not determined ex ante.

We embed these forces in a dynamic small-open-economy model with dominant-currency pricing and discipline it using Chilean administrative and invoice data. An external borrowing-premium innovation is normalized to generate a one-percent impact depreciation. Aggregate productivity falls by 0.110 percent on impact, with across-firm reallocation contributing -0.166 percent. Intermediate inputs increase this allocation loss by about 18 percent, while the additional input margins partly offset it. The across-firm component determines the sign of aggregate productivity and remains negative throughout the first year.

We test the model's predicted differences across firms using the responses of Chilean firms to broad dollar movements. Firms are classified using their markup twelve months earlier relative to the median within their industry. Following a broad dollar appreciation, total variable cost falls more at high-markup firms than at low-markup firms, with the differential response largest after about one year. Aggregating the same panel, aggregate productivity declines at longer horizons. The aggregate response is similar when productivity is measured as output per worker.

The Chilean exercise conditions on Chile's high dollar-invoicing share. We therefore use cross-country variation to examine the second ingredient of the mechanism. Under dominant-currency pricing, common dollar movements should matter more where a larger share of trade prices is sticky in dollars. Consistent with this implication, productivity responses to broad dollar movements are more negative in economies with greater dollar-invoicing exposure. This evidence supports the invoicing channel rather than re-estimating the within-country reallocation mechanism.

Taken together, the results show how a nominal exchange-rate movement can generate aggregate productivity dynamics without any change in technology. The currency of price stickiness determines how exchange rates move relative prices; the distribution of markups determines how those price movements reallocate resources; and intermediate inputs amplify their aggregate consequences. The mechanism is therefore fundamentally distributional: aggregate exchange-rate movements matter because firms differ in their initial distortions.

The paper connects the literature on heterogeneous markups and nominal shocks to the dominant-currency paradigm. Baqaee et al. (2024), Meier and Reinelt (2024), and Fujiwara and Matsuyama (2023) study how domestic nominal disturbances interact with

markup heterogeneity. We instead study an exchange-rate movement that changes the prices at which firms compete at home and abroad. Relative to Gopinath et al. (2020), our key departure is to approximate around a heterogeneous firm distribution, allowing composition to affect productivity at first order.

The paper also relates to work on variable markups and the gains from trade (Atkeson and Burstein, 2008; Edmond et al., 2015; Arkolakis et al., 2019) and to the measurement of productivity in open economies (Burstein and Cravino, 2015; Cravino, 2017). Crisis mechanisms include imported-variety exit within firms (Gopinath and Neiman, 2014) and, using Mexican microdata, destination-specific markups and allocative inefficiency during a sudden stop (Ishide, 2026). We study business-cycle reallocation around an already distorted allocation.

2 A static small open economy

We begin with a one-period comparison between an initial allocation and the allocation after a currency movement. Firms produce with labor alone, technology is fixed, and production is linear in labor. This environment isolates the allocation of labor across firms before intermediate inputs alter the mapping from that allocation to aggregate productivity.

The home economy is small relative to an unshocked dollar block of foreign countries. Home firms sell both domestically and abroad; foreign firms sell imported varieties in the home market. Home sales are invoiced in home currency, while imports and exports are invoiced in dollars. At $t = 0$, the economy is at its initial steady state and firms inherit their prices. At $t = \frac{1}{2}$, the home currency depreciates. At $t = 1$, a fraction δ of prices remains fixed in its currency of invoicing, the remaining fraction reset its prices, and quantities are allocated at the new prices.

This section takes the realized log changes in the nominal wage and the exchange rate, (w, e) , as conditioning variables. It does not specify the primitive disturbance that produces the depreciation, an exchange-rate process, or an intertemporal asset-market closure. We first show how a depreciation affects productivity with fully fixed prices. We then allow partial price adjustment, which modifies the domestic response, introduces an export channel, and makes productivity depend on the wage. A static policy benchmark supplies w/e , allowing the resulting mapping $a(w, e)$ to be stated per unit of depreciation. The dynamic model in section 3.3 closes the intertemporal economy and generates the path of e_t from a primitive innovation.¹


¹The environment follows the small-open-economy DCP structure of Gopinath et al. (2020): the home

Notation. Uppercase letters denote levels. Let $Z_0 = \bar{Z}$ be the initial steady-state value and Z_1 the value in the post-movement allocation. Unless stated otherwise, the corresponding lowercase letter denotes the log change between the two allocations: $z \equiv \log Z_1 - \log Z_0 = \log Z_1 - \log \bar{Z}$. Time subscripts are suppressed in this section, and a tilde denotes a newly reset price. Let λ_a be firm a 's steady-state share of total sales.² Bars on cross-sectional statistics below denote steady-state weighted means. Let Y denote aggregate real output and L total labor. In the labor-only economy, aggregate productivity is output per unit of labor, and its log deviation is

$$a \equiv y - l.$$

Because firm-level technologies are fixed, a movement in a records only the effect of reallocating labor across firms. Figure 2.1 records which prices are inherited when the exchange rate moves and which firms can reset before quantities are chosen.

Figure 2.1: Timing of the conditional one-period comparison



Notes: At $t = 0$, firms inherit the initial price configuration. The currency movement e occurs at $t = \frac{1}{2}$; its source and subsequent path are not specified in the static comparison. At $t = 1$, a fraction δ of prices remains fixed in its currency of invoicing and a fraction $1 - \delta$ has reset. Firm technology is unchanged.

2.1 Kimball demand and the firm distribution

A unit continuum of home firms is indexed by $\theta \in [0, 1]$. We approximate around the heterogeneous steady-state distribution of firm sizes and markups, so changes in composition affect productivity at first order. The representative household has period

economy moves against a dollar bloc that does not move with it. Section C studies the common-dollar counterfactual used to interpret the invoicing comparison in section 4.

²For any firm-level variables z_a and v_a , define

$$\mathbb{E}_\lambda[z_a] \equiv \int_0^1 \lambda_a z_a da, \quad \text{Cov}_\lambda(z_a, v_a) \equiv \mathbb{E}_\lambda[z_a v_a] - \mathbb{E}_\lambda[z_a] \mathbb{E}_\lambda[v_a],$$

where the subscript λ indicates that both the expectation and covariance operators use the weights λ_a .

utility

$$u(C, L) = \frac{C^{1-\eta}}{1-\eta} - v_L \frac{L^{1+1/\zeta}}{1+1/\zeta},$$

where $1/\eta$ is the intertemporal elasticity of substitution, ζ is the Frisch elasticity, and logarithmic utility applies when $\eta = 1$. The household chooses labor and home and imported varieties subject to

$$\int_0^1 (P_{h\theta} C_{h\theta} + P_{f\theta} C_{f\theta}) d\theta \leq WL + \Pi + \mathcal{T},$$

where Π denotes profits and $\mathcal{T}$ net income from the rest of the world. The left-hand side equals $P^C C$ at the optimum. Home consumption C is defined by the generalized Kimball aggregator (Kimball, 1995),

$$\int_0^1 \left[(1-\gamma) \Upsilon \left(\frac{C_{h\theta}}{(1-\gamma)C} \right) + \gamma \Upsilon \left(\frac{C_{f\theta}}{\gamma C} \right) \right] d\theta = 1.$$

The function Υ is increasing and concave. Its normalizations fix the units of the aggregate and of the price aggregator; its curvature determines how the elasticity of demand varies with a variety's relative quantity. The subscript h denotes a home producer and f a foreign producer. Thus $C_{h\theta}$ is a home variety sold in the home market and $C_{f\theta}$ is an imported variety. The parameter γ is the steady-state import expenditure share. Foreign producers supplying the home market have the same type distribution; the θ on $C_{f\theta}$ indexes a foreign type, not the corresponding home firm. The function Υ is common across varieties: firms do not receive firm-specific demand curves. They face different local elasticities because their equilibrium quantities place them at different points on the common demand schedule. Section 3.5 specifies the Klenow–Willis schedule and the productivity distribution used quantitatively.

Foreign demand comes from a unit continuum of identical destination markets. The home economy has measure zero from the perspective of each destination, so its firms take destination expenditure and the destination price aggregator as given. Each destination solves the analogous expenditure-allocation problem, with home varieties entering the same Kimball schedule as local varieties. Symmetry allows us to suppress the destination index and write the common aggregate demand and price aggregator as C^* and P^* . We integrate a firm's destination-level sales and denote its resulting export quantity by $C_{h\theta}^*$. Thus the continuum keeps the foreign price aggregator exogenous while reducing the firm's foreign sales to one export branch; no destination-level heterogeneity is used below.

Each home firm θ draws fixed productivity A_θ from a cross-sectional productivity

distribution, produces one variety with $Y_\theta = A_\theta L_\theta$, and sells it at home and abroad. Define its normalized quantities in the two destinations by

$$q_{h\theta} \equiv \frac{C_{h\theta}}{(1-\gamma)C}, \quad q_{h\theta}^* \equiv \frac{C_{h\theta}^*}{\gamma C^*}.$$

Expenditure minimization gives the inverse demand schedules

$$\frac{P_{h\theta}}{P} = \Upsilon'(q_{h\theta}), \quad \frac{P_{h\theta}^*}{P^*} = \Upsilon'(q_{h\theta}^*),$$

where P and P^* are the home and foreign price aggregators entering firms' demand. The consumer price index P^C instead measures the expenditure required per unit of C . Stars denote the foreign destination, and starred prices are in foreign currency. Before integration, the second inverse-demand equation holds destination by destination; symmetry permits the destination index to be omitted. The symbol $P_{f\theta}$, used below, is the price charged in the home market by a foreign exporter.

Two steady-state restrictions align the destination problem with the firm-level aggregation:

- (A1) A home firm occupies the same normalized point q_θ on its demand curve at home and in every foreign destination. Its steady-state markup is therefore the same in all destinations.
- (A2) Every home firm divides steady-state sales between home and export markets in the common proportion $(1-\gamma) : \gamma$.³

Assumption (A1) is the normalization used in Gopinath et al. (2020). Assumption (A2) makes destination weights common across firms. Together they imply that the destination-level sales weights aggregate to λ_θ , the firm's total-sales share. Equal initial markups across destinations remove the first-order productivity effect of reallocating a given firm's labor between home and export sales. The mechanism below concerns reallocation across firms. For the exact aggregate price coefficient derived below, foreign exporters also share the home economy's Calvo parameter and type-specific desired-pass-through distribution.

³Allowing export intensity to vary across firms would add terms involving its covariance with markups. The common split isolates the reallocation generated by pricing and demand heterogeneity.

Elasticities, markups, and desired pass-through. At a firm's steady-state normalized quantity q_θ , define

$$\sigma_\theta \equiv \frac{\Upsilon'(q_\theta)}{-\Upsilon''(q_\theta)q_\theta}, \quad \mu_\theta \equiv \frac{\sigma_\theta}{\sigma_\theta - 1},$$

and

$$\rho_\theta \equiv \frac{\partial \log P_\theta^{\text{flex}}}{\partial \log MC} = \frac{1}{1 + q_\theta(\mu'_\theta/\mu_\theta)\sigma_\theta}.$$

Here σ_θ is the local residual demand elasticity, μ_θ the desired flexible-price markup, and $\rho_\theta \in (0, 1]$ the elasticity of the firm's desired price to its own marginal cost, holding aggregate prices fixed. The derivative μ'_θ is with respect to normalized quantity. By (A1), the same triplet $(\sigma_\theta, \mu_\theta, \rho_\theta)$ describes a firm in both destinations at the steady state.

The steady-state pricing condition makes the connection between productivity and the firm distribution explicit:

$$P\Upsilon'(q_\theta) = \mu_\theta \frac{W}{A_\theta}.$$

For a given aggregate real wage, a productivity draw therefore determines the firm's point q_θ on the common demand curve and hence its elasticity, markup, and desired pass-through. Under (A1)–(A2), its total-sales weight is

$$\lambda_\theta = \frac{\Upsilon'(q_\theta)q_\theta}{\int_0^1 \Upsilon'(q_s)q_s ds}.$$

Thus the productivity distribution and the Kimball demand system jointly induce the weighted distribution of $(q_\theta, \sigma_\theta, \mu_\theta, \rho_\theta)$. The exchange-rate experiment holds this initial cross-section fixed and changes prices and quantities around it.

Because firms occupy different points on the same residual demand curve, relative size provides an economically meaningful ordering of the cross-section. Two assumptions follow: first, desired markups rise with relative quantity (consistent with Marshall's second law of demand), or equivalently, large firms face lower residual elasticity of demand σ_θ and thus charge relatively higher markups. And second, pass-through from marginal cost to prices is inversely related to firm size, so that ρ_θ falls as relative quantity rises (consistent with Marshall's third law of demand). Therefore, larger firms are both high-markup and low-pass-through firms: after a common marginal-cost change, they change their desired prices less and absorb more of the movement into their markups. This is the joint ordering emphasized by Baqaee et al. (2024).

The elasticity-weighted sales expectation is

$$\mathbb{E}_{\lambda\sigma}[z_\theta] \equiv \frac{\mathbb{E}_\lambda[\sigma_\theta z_\theta]}{\mathbb{E}_\lambda[\sigma_\theta]}.$$

The harmonic sales-weighted markup and the two pass-through means are

$$\bar{\mu} \equiv \left(\mathbb{E}_\lambda[\mu_\theta^{-1}]\right)^{-1}, \quad \bar{\rho} \equiv \mathbb{E}_\lambda[\rho_\theta], \quad \bar{\rho}^\sigma \equiv \mathbb{E}_{\lambda\sigma}[\rho_\theta].$$

The weights $\bar{\mu}/\mu_\theta$ average to one.

2.2 Dollar pricing and the price level

Let e be the home-currency price change of one dollar, so $e > 0$ is a home depreciation. Let p be the price aggregator in the home market and p^* its foreign counterpart. Since labor is the only input and technology is fixed, a home firm's marginal-cost deviation is w . The home-currency movement leaves the foreign price aggregator and foreign dollar marginal cost unchanged, so $p^* = 0$. A home firm that resets chooses

$$\tilde{p}_{h\theta} = \rho_\theta w + (1 - \rho_\theta)p, \quad \tilde{p}_{h\theta}^* = \rho_\theta(w - e) + (1 - \rho_\theta)p^*,$$

in home-currency and dollar terms. A resetting foreign exporter has dollar marginal cost fixed; in home currency its reset price is

$$\tilde{p}_{f\theta} = \rho_\theta e + (1 - \rho_\theta)p.$$

Thus the same desired pass-through that governs a firm's response to cost also governs how far its reset price follows competitors.

The home price aggregator and the CPI are

$$p \equiv (1 - \gamma)\mathbb{E}_{\lambda\sigma}[p_{h\theta}] + \gamma\mathbb{E}_{\lambda\sigma}[p_{f\theta}], \quad p^C \equiv (1 - \gamma)\mathbb{E}_\lambda[p_{h\theta}] + \gamma\mathbb{E}_\lambda[p_{f\theta}].$$

The price aggregator uses elasticity-weighted sales shares and enters firms' demand and reset prices. The CPI uses expenditure shares. Each firm-level price averages its reset and inherited Calvo vintages.

A foreign exporter that does not reset has a fixed dollar price, so its home-currency price rises one for one with e . Aggregating the reset and inherited prices and solving the

fixed point gives (see section A.1)

$$p = \kappa_w w + \gamma e, \quad p^C = \kappa_w^C w + \gamma e,$$

where

$$\kappa_w = \frac{(1 - \gamma)(1 - \delta)\bar{p}^\sigma}{\delta + (1 - \delta)\bar{p}^\sigma}.$$

The corresponding CPI loading, $\kappa_w^C > 0$ for $\delta < 1$, describes how wage increases pass into reset prices. The coefficient on the exchange rate is exactly the import share γ .⁴ The reason is that an inherited dollar price moves one for one in home currency, while reset-price feedback is absorbed by the fixed point. This price block is the standard aggregate DCP response. We first examine its productivity effect with fully fixed prices, and then allow firms to reset.

2.3 Exchange rates and aggregate productivity

With technology fixed, aggregate productivity changes when production is reallocated across firms with different markup wedges. High-markup firms produce too little relative to low-markup firms. A shift in the composition of production toward low-markup firms therefore worsens the allocation, even if production expands at both types of firms.

A demand shifter common to all firms scales their activity together without changing this composition. Conditional on (w, e) , aggregate demand therefore determines total activity, while differences in firms' relative quantity responses determine aggregate productivity. Section A.3 derives the firm-destination aggregation underlying this distinction.

Fully fixed prices. To build intuition, begin with $\delta = 1$. Home prices and dollar export prices do not change, while the home-currency price of every imported variety rises with the exchange rate. The home price aggregator therefore rises by $p = \gamma e$.

Proposition 1 (Depreciation with fully fixed prices). *With fully fixed prices, the first-order productivity response is*


$$a = -\gamma(1 - \gamma)C_\sigma e, \quad C_\sigma \equiv \text{Cov}_\lambda\left(\frac{\bar{\mu}}{\mu_\theta}, \sigma_\theta\right).$$

Under Marshall's second law and nondegenerate markup heterogeneity, $C_\sigma > 0$. For $0 < \gamma < 1$, a depreciation therefore lowers aggregate productivity.

⁴Exactness uses the symmetry of the continuum: if foreign price-flexibility moments differ, Appendix A.1 gives the rescaled coefficient.

A depreciation makes dollar-invoiced imports more expensive relative to home varieties, shifting demand toward home producers. Low-markup firms face more elastic demand, so their quantities grow faster than those of high-markup firms. Their share of production rises, worsening an allocation in which they are already relatively too large. Panel (a) of fig. 2.2 illustrates these different quantity responses. The covariance C_σ captures how the firms that expand more are also those with lower initial markup wedges. Export quantities do not change in this case: dollar export prices, foreign expenditure, and the foreign price aggregator remain fixed.

Figure 2.2: Reallocation with non-resetting and resetting prices



Notes: Home-market Kimball demand in logs; the local slope is $-1/\sigma_\theta$. Panel (a) holds a firm's price fixed while the price aggregator rises, producing the quantity response $\sigma_\theta \Delta \log P$. Panel (b) holds marginal cost fixed and lets the firm reset and follow the price aggregator with weight $1 - \rho_\theta$, leaving the smaller response $\sigma_\theta \rho_\theta \Delta \log P$. The gray dashed segment is the fixed-price response.

Partial price adjustment. Now allow a fraction $1 - \delta$ of firms to reset their prices. At a given wage, a depreciation still shifts demand away from imports, but resetting home producers also change their own prices. Their price relative to the home price aggregator satisfies

$$p - \tilde{p}_{h\theta} = \rho_\theta(p - w).$$

Their quantity sensitivity to the gap $p - w$ is therefore $\sigma_\theta \rho_\theta$, rather than σ_θ . Under Marshall's third law, low-markup firms also have higher desired cost pass-through. They follow a rise in competitors' prices less closely and retain more of the relative-price advantage. Price adjustment thus changes the size of the response without reversing its ordering across firms, as panel (b) of fig. 2.2 illustrates.

Price adjustment also introduces an export channel. A depreciation lowers home marginal cost in dollars when the wage rises less than the exchange rate. Against the unchanged foreign price aggregator, resetting exporters choose

$$\widehat{p}_{h\theta}^* = \rho_\theta(w - e).$$

Low-markup firms pass more of this cost reduction into their dollar prices and face more elastic demand, so their exports grow faster. Exporters whose dollar prices remain fixed do not contribute to this channel. Thus the same quantity sensitivity, $\sigma_\theta \rho_\theta$, governs reallocation among resetting firms in both markets.

Combining the home and export responses gives the productivity mapping with partial price adjustment. The wage now matters because it affects reset prices in both markets and feeds back into the home price aggregator.

Proposition 2 (Exchange rates and allocative efficiency). *To first order,*

$$a = -\Psi_e e + \Psi_w w,$$

where the exchange-rate and wage loadings are

$$\begin{aligned} \Psi_e &= \underbrace{\gamma(1-\gamma) \left[\delta C_\sigma + (1-\delta) C_{\sigma\rho} \right]}_{\text{home market}} + \underbrace{\gamma(1-\delta) C_{\sigma\rho}}_{\text{export market}}, \\ \Psi_w &= (1-\delta) C_{\sigma\rho} \left[1 - (1-\gamma) \kappa_w \right] - (1-\gamma) \delta \kappa_w C_\sigma. \end{aligned}$$

The fixed-price statistic C_σ is defined in proposition 1; the corresponding statistic for resetting firms is

$$C_{\sigma\rho} \equiv \text{Cov}_\lambda \left(\frac{\bar{\mu}}{\mu_\theta}, \sigma_\theta \rho_\theta \right).$$

Under Marshall's second and third laws, nondegenerate markup heterogeneity, $0 < \gamma < 1$, and $\delta < 1$, both covariances are positive, as are Ψ_e and Ψ_w .

The home-market contribution combines the responses of non-resetting and resetting firms. The export contribution comes only from resetting exporters. Both enter the direct effect of depreciation with a negative sign: holding the wage fixed, production shifts toward low-markup firms through both markets.

At a given exchange rate, a higher wage raises marginal costs. Among resetting firms, low-markup firms raise prices more and reduce quantities more, shifting production shares toward high-markup firms and improving allocation. The rise in the home

price aggregator also increases demand for non-resetting firms, especially those with low markups, partly offsetting this gain. Under Marshall's second and third laws, the first effect dominates, so $\Psi_w > 0$. Wage increases therefore mitigate the direct allocative loss from depreciation, whereas wage declines amplify it.

A static wage benchmark. Proposition 2 maps the realized movements (w, e) into aggregate productivity. To express that result as a response per unit of depreciation, we use CPI stabilization as a static benchmark for the wage response. For $\delta < 1$, setting $p^C = 0$ in the price equation gives

$$\frac{w}{e} = -\frac{\gamma}{\kappa_w^C} < 0.$$

Corollary 1 (Exchange-rate effect under CPI stabilization). *If $\delta < 1$, a depreciation lowers aggregate productivity:*

$$\frac{a}{e} = -\left(\Psi_e + \frac{\gamma}{\kappa_w^C}\Psi_w\right) < 0.$$

Holding the wage fixed would already give $a/e = -\Psi_e < 0$. CPI stabilization requires the wage to fall, so the wage channel reinforces the direct exchange-rate effect. Section A.7 gives the expression for other static CPI responses. This substitution determines a/e conditional on a realized depreciation; it does not determine why the currency moves or its path after $t = 1$. Those objects are determined by the dynamic equilibrium in section 3.3.

Given (w, e) , household labor supply, goods-market clearing, and production determine consumption, employment, and output. These conditions change the scale of activity but not its distribution across firms: their common expenditure component drops out of the covariance in proposition 2. Appendix A.8 records the complete one-period quantity block.

The comparison above holds firm heterogeneity fixed while varying firms' ability to adjust prices. Three further benchmarks show which cross-sectional features are essential and what remains as prices become fully flexible.

Example 1 (CES demand). Under CES demand, the demand elasticity is constant. Firms may differ in productivity and size, but they all charge the same desired markup and have the same desired pass-through. The initial markup wedge is therefore constant across firms, so $C_\sigma = C_{\sigma p} = 0$. It follows that $a = 0$ for every exchange-rate movement and every monetary regime. CES removes the reallocation channel by eliminating cross-sectional variation in initial desired markups.

Example 2 (Kimball demand around a symmetric firm distribution). Kimball demand permits markups and desired pass-through to vary with a firm's position on its demand

curve. But if the model is approximated around a symmetric firm distribution, every firm occupies the same point on that curve. Markups, demand elasticities, and desired pass-through are then common at the approximation point, and both covariances again vanish. This is the firm-side symmetry in Gopinath et al. (2020). Their model retains strategic complementarity and the aggregate transmission of dollar-priced trade, but it does not generate a first-order reallocation of production across firms with different initial markup wedges. Our departure is to approximate around a heterogeneous firm distribution.

Example 3 (Flexible prices). As $\delta \rightarrow 0$, the fixed-price channel vanishes and $\Psi_e = \Psi_w = \gamma(2 - \gamma)C_{\sigma\rho}$. Hence

$$a = \gamma(2 - \gamma)C_{\sigma\rho}(w - e).$$

With heterogeneity retained, only the real gap between home costs and foreign competitors remains. Thus flexible prices do not eliminate real reallocation, but they remove the nominal traction: if a nominal depreciation is matched by the wage, nothing moves. Price stickiness gives nominal exchange-rate movements traction, and the invoice currency determines which exchange-rate movements transmit (section C).

Remark (Destination-specific markups). Under assumption (A1), a firm charges the same initial markup on domestic and export sales. This removes the first-order productivity effect of reallocating the firm's labor between home and export sales. Evidence from Chile and Mexico shows that, within firms, export markups are on average lower than domestic markups (Blum et al., 2024; Ishide, 2026). With such differences, a depreciation that shifts production toward export sales within firms would add an allocative loss, reinforcing the across-firm channel. We set aside this previously studied destination margin to isolate reallocation across firms.

3 Intermediate inputs and quantitative dynamics

Section 2 treated the realized wage and exchange-rate movements as arguments of a conditional productivity mapping. We now construct a quantitative model that determines their paths. First, we add intermediate inputs and derive their mapping into aggregate productivity. They amplify the across-firm allocation loss because gross sales exceed GDP, and they introduce an aggregate input-allocation term and an import-measurement bias. Second, we embed the resulting production and pricing blocks in a monthly small-open-economy New Keynesian model in the spirit of Gopinath et al. (2020). The household, asset-market, market-clearing, and policy conditions close the dynamic general equilibrium and determine the exchange rate jointly with quantities and prices.

3.1 Production with intermediate inputs

To distinguish a firm in its role as an input buyer from the varieties it purchases, index the buyer by i and the supplier by θ . Firm i produces with labor and a composite intermediate input,

$$Y_{it} = A_i L_{it}^{1-\alpha} X_{it}^\alpha,$$

where $0 \leq \alpha < 1$ is the elasticity of gross output with respect to the intermediate composite. The composite is defined by the same Kimball aggregator as consumption:

$$\int_0^1 \left[(1-\gamma) \Upsilon \left(\frac{X_{i\theta t}^H}{(1-\gamma)X_{it}} \right) + \gamma \Upsilon \left(\frac{X_{i\theta t}^M}{\gamma X_{it}} \right) \right] d\theta = 1.$$

Here $X_{i\theta t}^H$ and $X_{i\theta t}^M$ are purchases from home and imported suppliers of type θ . At the reference allocation, paired home and imported supplier types occupy the same point q_θ on the Kimball schedule. This normalization makes γ the steady-state import expenditure share of the input basket.

The aggregator is homothetic, and all buyers face the same input prices. Their input vectors are therefore proportional, so they can be aggregated across buyers. Let $X_t \equiv \int X_{it} di$, and let $X_{\theta t}^H$ and $X_{\theta t}^M$ denote aggregate purchases from each supplier type. Define the source-specific Divisia quantity changes using the steady-state expenditure weights λ_θ introduced in section 2.1:

$$x_t^H \equiv \mathbb{E}_\lambda[x_{\theta t}^H], \quad x_t^M \equiv \mathbb{E}_\lambda[x_{\theta t}^M].$$

The common Kimball schedule and matched reference quantities make these the input-expenditure weights within both sources. Differentiating the aggregator then gives

$$x_t = (1-\gamma)x_t^H + \gamma x_t^M.$$

Appendix A.9 derives this relation from the Kimball aggregator.

Because firms purchase the same basket of varieties at the same prices as households, its unit-expenditure index is the CPI. Cost minimization therefore gives a marginal cost common across firms,

$$mc_t = (1-\alpha)w_t + \alpha p_t^C.$$

The static pricing equations continue to hold with mc in place of w . The cross-sectional covariances are unchanged: they describe how a common marginal-cost and exchange-rate movement redistributes demand across the same firm types. What changes is the

value-added consequence of that redistribution and the set of aggregate input margins measured by the Solow residual.

3.2 Aggregate productivity

Aggregate productivity is

$$a_t \equiv y_t - \widetilde{\Lambda}_L l_t,$$

where y_t is Divisia real GDP and $\widetilde{\Lambda}_L$ is the cost-based labor exposure of GDP. This is the residual obtained by valuing the production structure at marginal cost. In the labor-only economy, $\widetilde{\Lambda}_L = 1$ and the object reduces to $a = y - l$ in section 2. With intermediate inputs it has the following quantity representation.

The firm-level reallocation mechanism from section 2 is unchanged. With marginal cost replacing the wage, its static mapping is

$$a_F = -\Psi_e e + \Psi_w mc.$$

We use $a_{F,t}$ for the corresponding across-firm component in the dynamic economy. Inherited price vintages determine its path through proposition 4.

Proposition 3 (Aggregate productivity). *To first order,*

$$a_t = \underbrace{\xi_F a_{F,t}}_{\text{across-firm reallocation}} + \underbrace{\xi_X \left[(1 - \alpha)(x_t - l_t) + \gamma(x_t^H - x_t^M) \right]}_{\text{aggregate input allocation}} + \underbrace{\xi_M x_t^M}_{\text{import-measurement bias}},$$

where

$$\xi_F = \frac{\bar{\mu}}{\bar{\mu} - \alpha}, \quad \xi_X = \frac{\widetilde{\alpha}(\bar{\mu} - 1)}{(\bar{\mu} - \alpha)(1 - \widetilde{\alpha})}, \quad \xi_M = \frac{\alpha\gamma(\bar{\mu} - 1)}{(\bar{\mu} - \alpha)(1 - \widetilde{\alpha})}, \quad \widetilde{\alpha} \equiv \alpha(1 - \gamma).$$

The first term is the allocation loss of proposition 2, evaluated at marginal cost and multiplied by ξ_F . The multiplier is the ratio of gross sales to GDP for this reallocation direction. It is greater than one whenever firms use intermediate inputs. Including intermediate inputs, therefore amplifies the same across-firm loss; it does not create a different reallocation mechanism.

The second term compares the domestically produced intermediate with its two alternatives, labor and imported intermediates. Its bracket can equivalently be written as

$$(1 - \widetilde{\alpha})x_t^H - (1 - \alpha)l_t - \alpha\gamma x_t^M.$$

A domestic supplier charges a markup over marginal cost, so its customer faces a private input price above the input's domestic resource cost. Home intermediates are consequently underused, while labor and imported intermediates are overused in their place. A shift from the home intermediate toward either alternative deepens this input distortion and lowers aggregate productivity. A shift toward the home intermediate closes the wedge and raises aggregate productivity. This is why x_t^H enters positively, while l_t and x_t^M enter negatively, and why the loading ξ_X is positive when $\bar{\mu} > 1$. At unit markups, the private and social input costs coincide and $\xi_X = 0$.

The final term is a measurement bias, not a productive contribution of imports. Real GDP already subtracts imported intermediates. The bias arises because the national accounts make that subtraction using the observed import-expenditure share, whereas the cost-based productivity benchmark assigns imports a larger exposure when domestic production carries markups. Their difference is $\xi_M > 0$, so an imported-input quantity change leaves $\xi_M x_t^M$ in the measured residual even though technology is unchanged. A fall in imported inputs therefore contributes negatively to aggregate productivity through this bias. At unit markups the two exposures coincide, $\xi_M = 0$, and the bias vanishes. Appendix A.10 derives the exposure difference from firm sales, variable costs, and the national accounts and records an equivalent price-and-expenditure representation.

The decomposition is an accounting identity and does not determine how the two input quantities respond to a depreciation. Their signs are equilibrium outcomes of the dynamic model below; the quantitative response then maps those movements back into the economic margins defined here.

Two limits clarify the decomposition. If $\alpha \rightarrow 0$, then $\xi_F \rightarrow 1$ and $\xi_X, \xi_M \rightarrow 0$, returning to proposition 2. Under CES demand the firm covariances, and hence a_F , are zero, although input allocation and the import-measurement bias can remain when the average markup exceeds one. With unit markups and flexible prices, all three terms vanish, nesting the frictionless neutrality benchmark of Kehoe and Ruhl (2008).

3.3 Dynamic environment

We now embed the production and demand blocks above in a monthly small-open- economy equilibrium. Unlike the static comparison, the exchange rate is an endogenous equilibrium variable. Firm productivity and Kimball demand type remain fixed; there is no entry, capital accumulation, or technology shock. The equilibrium combines four familiar blocks: household optimization, Calvo price setting, goods and input market clearing, and an asset-market and policy closure. We state the household and firm problems here

and collect their first-order equilibrium conditions in the next subsection.

Here and below, $\mathbb{E}_t$ denotes expectation conditional on information available at date t . The representative household has preferences

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\eta}}{1-\eta} - v_L \frac{L_t^{1+1/\zeta}}{1+1/\zeta} \right],$$

where $1/\eta$ is the intertemporal elasticity of substitution and ζ is the Frisch elasticity; logarithmic utility is understood when $\eta = 1$. Let B_{t+1} and $B_{t+1}^\$$ denote one-period claims that pay one unit of home currency and one dollar, respectively, at date $t + 1$. If R_t and $R_t^\$$ are their gross returns, the household's nominal flow budget is

$$P_t^C C_t + \frac{B_{t+1}}{R_t} + E_t \frac{B_{t+1}^\$}{R_t^\$} = W_t L_t + \Pi_t + B_t + E_t B_t^\$.$$

The home-currency bond is in zero net supply. The aggregate dollar position is reported below as net foreign debt, and a debt-elastic borrowing premium stationarizes it as in Schmitt-Grohe and Uribe (2003). The household's first-order conditions give the labor-supply, consumption-Euler, and uncovered-interest-parity equations reported below. Nominal wages are flexible.

Prices remain fixed in their currency of invoicing with probability δ each month. Home sales are priced in home currency, and imports and exports are priced in dollars. Consider first a home firm's price in the home market. Conditional on a reset opportunity, it solves

$$\max_{\bar{p}_{h\theta t}} \mathbb{E}_t \sum_{k=0}^{\infty} \delta^k Q_{t,t+k} (\bar{p}_{h\theta t} - MC_{\theta,t+k}) C_{h\theta,t+k},$$

subject to its Kimball demand schedule at every date for which the price remains in effect. Here $Q_{t,t+k}$ is the household's nominal stochastic discount factor. Linearizing its first-order condition gives

$$\bar{p}_{h\theta,t} = (1 - \beta\delta) [\rho_\theta mc_t + (1 - \rho_\theta)p_t] + \beta\delta \mathbb{E}_t \bar{p}_{h\theta,t+1}.$$

The expression in brackets is the one-period reset price in section 2.2; the dynamic condition averages the same target over the expected duration of the price spell. The export and import problems are analogous in dollars. Their corresponding targets are $\rho_\theta(mc_t - e_t) + (1 - \rho_\theta)p_t^*$ for a home export and $(1 - \rho_\theta)(p_t - e_t)$ for a foreign supplier's dollar price. The latter equals $\rho_\theta e_t + (1 - \rho_\theta)p_t$ in home currency on the reset date. Appendix B.1 records the complete price-state recursions.

The remaining conditions clear labor, home goods, each intermediate input, and the two asset markets. Home-currency bonds are in zero net supply, while the current account changes the dollar position. The domestic interest rate follows an inertial CPI-inflation rule; uncovered interest parity links it to the dollar borrowing rate and the exchange rate. Foreign expenditure, foreign marginal cost, and the foreign dollar price benchmark are fixed. Export quantities nevertheless respond to their dollar relative price through Kimball demand. This is the small-economy asymmetry in Gopinath et al. (2020): the home currency moves against an unmoved dollar bloc. The primitive disturbance in the quantitative experiment is an external borrowing-premium innovation. The full equilibrium determines the exchange-rate path; after solving it, we rescale all responses so that the impact depreciation is one percent. This normalization puts the dynamic results in the per-unit-depreciation units used by the static mapping, but it does not impose the path of e_t . The domestic monetary-policy innovation is zero.

3.4 Dynamic equilibrium and reallocation

Log-linearizing the household and firm first-order conditions and the market-clearing conditions gives the equilibrium system. We present its blocks in the order in which they enter the reallocation mechanism: price setting first, household choices and market clearing next, and the asset-market and policy closure last. At each date, the strategic price index and the CPI retain their definitions from the one-period economy:

$$p_t \equiv (1 - \gamma)\mathbb{E}_{\lambda\sigma}[p_{h\theta,t}] + \gamma\mathbb{E}_{\lambda\sigma}[p_{f\theta,t}], \quad p_t^C \equiv (1 - \gamma)\mathbb{E}_\lambda[p_{h\theta,t}] + \gamma\mathbb{E}_\lambda[p_{f\theta,t}].$$

Define

$$\widehat{mc}_t \equiv (1 - \gamma)mc_t + \gamma e_t - p_t, \quad \psi \equiv \frac{(1 - \delta)(1 - \beta\delta)}{\delta}.$$

The variable $\widehat{mc}_t$ is the trade-weighted home-cost gap relative to the strategic price index. It is the common cost pressure faced by price reseters across the home and export markets. Aggregating the home-market reset conditions and inherited price vintages gives

$$\Delta(p_t - \gamma e_t) = \beta \mathbb{E}_t \Delta(p_{t+1} - \gamma e_{t+1}) + \psi \bar{\rho}^\sigma \widehat{mc}_t,$$

$$\Delta(p_t^C - \gamma e_t) = \beta \mathbb{E}_t \Delta(p_{t+1}^C - \gamma e_{t+1}) + \psi [\bar{\rho} \widehat{mc}_t - (p_t^C - p_t)].$$

Writing the equations in terms of exchange-rate-adjusted price indices makes the common revaluation of inherited dollar prices explicit. The second equation is the CPI Phillips curve. The first is also required away from CES because Kimball reset prices depend on

the elasticity-weighted price index rather than the expenditure-weighted CPI. Under CES, $p_t^C = p_t$ and $\bar{\rho} = \bar{\rho}^\sigma = 1$, so the two equations coincide.

The dynamic counterpart of proposition 2 is as follows.

Proposition 4 (Dynamic across-firm reallocation). *Given aggregate prices and marginal cost,*

$$\begin{aligned} (1 + \beta + \psi)a_{F,t} &= a_{F,t-1} + \beta\mathbb{E}_t a_{F,t+1} \\ &\quad + \psi \left[C_{\sigma\rho} \left[mc_t - (1 - \gamma)p_t - \gamma e_t \right] - (1 - \gamma)\bar{\rho}^\sigma C_\sigma \widehat{mc}_t \right] \\ &\quad - (1 - \gamma)\gamma C_\sigma (\Delta e_t - \beta\mathbb{E}_t \Delta e_{t+1}). \end{aligned}$$

The lag and lead carry the distributional consequences of past price vintages. The two terms in square brackets are the cost gap faced by firms that reset and the relative-price gap that moves quantities at firms whose prices remain fixed. The last line is the revaluation of inherited dollar prices. Both cross-sectional channels vanish under CES.

Using the definition of $a_{F,t}$, the relative-price response entering aggregate labor demand separates into a sales-weighted mean and the across-firm allocation component. Common expenditure movements scale every firm and drop out of both. Household optimality and the common marginal-cost condition are

$$\begin{aligned} w_t - p_t^C &= \zeta^{-1}l_t + \eta c_t, \\ c_t &= \mathbb{E}_t c_{t+1} - \eta^{-1}(i_t - \mathbb{E}_t \Delta p_{t+1}^C), \\ mc_t &= (1 - \alpha)w_t + \alpha p_t^C. \end{aligned}$$

Goods, input, and labor-market clearing then give

$$\begin{aligned} l_t &= \Theta_c c_t + \Theta_x (mc_t - p_t^C) \\ &\quad - \Theta_s \left(\mathbb{E}_\lambda \left[\sigma_\theta \left((1 - \gamma)(p_{h\theta,t} - p_t) + \gamma p_{h\theta,t}^* \right) \right] + a_{F,t} \right) + mc_t - w_t, \\ y_t &= a_t + \widetilde{\Lambda}_L l_t. \end{aligned}$$

The coefficients and the cost-based labor exposure are

$$\begin{aligned} \Theta_c &= \frac{(1 - \gamma)(\bar{\mu} - \alpha)}{\bar{\mu} - \widetilde{\alpha}}, \quad \Theta_x = \frac{\widetilde{\alpha}}{\bar{\mu} - \widetilde{\alpha}}, \quad \Theta_s = \frac{\bar{\mu}}{\bar{\mu} - \widetilde{\alpha}}, \\ \widetilde{\Lambda}_L &= \frac{(1 - \alpha)(\bar{\mu} - \widetilde{\alpha})}{(\bar{\mu} - \alpha)(1 - \widetilde{\alpha})}. \end{aligned}$$

The expectation in labor demand is the average relative-price response of home firms across their home and export markets, using the same sales weights as the rest of the firm

aggregation. The term $a_{F,t}$ adds the change in firm composition. Appendix B.1 proves the equivalence to variable-cost aggregation and records the law of motion.

The source-specific input quantities in proposition 3 are the Divisia quantities defined in section 3.1. Cost minimization and Kimball input demand determine their equilibrium movements. Appendix B.1 records their log-linear mapping from input prices. Given these quantities and $a_{F,t}$, proposition 3 determines aggregate productivity a_t , and $y_t = a_t + \widetilde{\Lambda}_L l_t$ determines real GDP.

The final block closes the asset markets and determines the exchange rate jointly with quantities. Let p_t^M and p_t^E denote the home-currency import and export price indices. The Divisia GDP deflator is

$$p_t^{GDP} = p_t^C + \gamma \xi_F (p_t^E - p_t^M).$$

The dollar borrowing rate, UIP, the domestic policy rule, and the current account obey

$$\begin{aligned} i_t^\$ &= \phi_b b_t + u_t, & u_t &= \rho_u u_{t-1} + \varepsilon_t^u, \\ i_t &= i_t^\$ + \mathbb{E}_t e_{t+1} - e_t, & i_t &= \rho_i i_{t-1} + (1 - \rho_i) \phi_\pi \Delta p_t^C, \\ nx_t &= y_t + p_t^{GDP} - c_t - p_t^C, & b_t &= \beta^{-1} b_{t-1} - nx_t. \end{aligned}$$

The first row prices dollar borrowing and specifies its innovation. The second requires the domestic interest rate to satisfy both UIP and the policy reaction. The final row links the current-account surplus to the dollar debt position. Here b_t and nx_t are level deviations scaled by monthly steady-state GDP, ϕ_b is the corresponding monthly debt-premium coefficient, $i_t^\$$ is the home economy's dollar borrowing rate, and u_t is its external borrowing premium. The debt-elastic premium and UIP link the external position to the exchange rate and the domestic interest rate. The policy rule responds to CPI inflation and contains no output term or domestic policy innovation in the reported experiment. Appendix B.1 derives the cross-border price targets and records the unit and shock-normalization details.

Together with the cross-border price and input-demand equations in Appendix B.1, the inherited states, lagged debt and interest rate, and the innovation ε_t^u jointly determine all prices, quantities, asset positions, and aggregate productivity.

3.5 Calibration

Data and firm-level markups. The calibration and the reduced-form analysis use the same monthly panel of Chilean firms. We combine VAT filings on sales, material expenditure, and investment with wage and employment records, annual income-tax information

on capital, and the universe of electronic tax invoices. The invoices report product-level prices and quantities from 2014 onward and are used to construct firm-level Tornqvist indices for output and material inputs. The administrative shock window is 2005–2023; objects that require invoice-based prices or quantities use the 2014–2023 sample. The unit of observation is a firm-month, and each firm is treated as producing one composite good.⁵ The underlying administrative records and their coverage are also described by Giarda and Martner (2026).⁶

We estimate monthly firm-level markups using quantity-based production functions. We estimate a Cobb–Douglas production function by six-digit industry and year, recover the output elasticity of materials, and divide it by the firm’s monthly material-expenditure share, following De Loecker and Warzynski (2012). The invoice-based quantity indices keep output-price variation out of the production-function dependent variable. Because product-level prices are observed, the production function is estimated in physical output rather than revenue. This avoids the revenue-based non-identification problem emphasized by Bond et al. (2021) when firm-level output prices are unobserved. The resulting markup distribution supplies the harmonic sales-weighted markup used below. Construction of the quantity indices, production-function estimation, and the fixed firm-level sorting measure used in the reduced-form exercise are detailed in section D.2.

Kimball demand and the firm distribution. For the quantitative exercise we specialize the generalized demand system in section 2.1 to a common Klenow–Willis schedule,

$$\Upsilon'(q) = \frac{\bar{\sigma}_K - 1}{\bar{\sigma}_K} \exp\left(\frac{1 - q^{\epsilon_K/\bar{\sigma}_K}}{\epsilon_K}\right), \quad \sigma(q) = \bar{\sigma}_K q^{-\epsilon_K/\bar{\sigma}_K}.$$

Here q is a firm’s quantity relative to the destination aggregate, $\bar{\sigma}_K$ is the demand elasticity at $q = 1$, and ϵ_K governs how rapidly elasticity falls as relative quantity rises: $-d \log \sigma(q)/d \log q = \epsilon_K/\bar{\sigma}_K$. Thus the schedule maps a firm’s equilibrium relative quantity

⁵Koike and Martner (2026) use the same records while retaining the buyer-product dimension to study production-network distortions with multiproduct firms.

⁶This study was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in economic and financial affairs of its competence. The CBC has access to anonymized information from various public and private entities, by virtue of collaboration agreements signed with these institutions. To secure the privacy of workers and firms, the CBC mandates that the development, extraction and publication of the results should not allow the identification, directly or indirectly, of natural or legal persons. Officials of the Central Bank of Chile processed the disaggregated data. All the analysis was implemented by the authors and did not involve nor compromise the Servicios de Impuestos Internos de Chile. The information contained in the databases of the Chilean IRS is of a tax nature originating in self-declarations of taxpayers presented to the Service; therefore, the veracity of the data is not the responsibility of the Service.

into its elasticity, markup, and desired pass-through.

The firm distribution is obtained by combining this schedule with the cross-sectional productivity distribution. For each draw A_θ , the steady-state pricing condition in section 2.1 is solved for q_θ . Revenue $q_\theta Y'(q_\theta)$ then determines the sales weight λ_θ , and evaluating the local slope and curvature of the KW schedule gives σ_θ , μ_θ , and ρ_θ . The weighted means and covariances in proposition 2 are computed from this induced distribution; they are not calibrated independently.

We set the reference elasticity to $\bar{\sigma}_K = 6.0$ and choose the common real-cost normalization so that the induced harmonic sales-weighted markup is 1.20, the value estimated from the firm panel after rounding. The scale normalization and the cross-sectional markup match are distinct; their exact mapping is given in section D.3. Conditional on that scale, Albagli et al. (2026) discipline the curvature of demand using own-cost and competitor-price responses estimated from the same Chilean electronic-invoice records. Their relative own-cost loading, $0.24/(0.24 + 0.38)$, implies $\epsilon_K = 7.92$. We use this relative loading rather than equating either dynamic response coefficient by itself with a static desired pass-through.

Firm productivity follows a Pareto upper tail,

$$\Pr(A_\theta \geq a) = \left(\frac{A_{\min}}{a} \right)^{\kappa_A}, \quad a \geq A_{\min},$$

with $\kappa_A = 3.93$, estimated from the Chilean firm distribution reported by Giarda and Martner (2026). Together with the normalization described above, the productivity draws and the common KW schedule determine the cross-section. Desired pass-through and the two reallocation covariances are therefore implications of the calibrated cross-section rather than separate targets, and no aggregate impulse response is used in the calibration. Table 3.1 summarizes the resulting parameters and moments.

Aggregate and dynamic parameters. The model is monthly, matching the firm panel and invoice-price data. We set the production elasticity of intermediate inputs to $\alpha = 0.20$, the import share to $\gamma = 0.30$, the Frisch elasticity to $\zeta = 0.50$, and the intertemporal elasticity to one. The monthly frequency of regular price changes is 0.405, implying $\delta = 0.595$ (Albagli et al., 2026). We set $\beta = 0.99^{1/3}$, $\rho_i = 0.50^{1/3}$, and $\phi_\pi = 1.5$. The annual debt-premium elasticity is $\phi_b^A = 0.000742$, following Schmitt-Grohe and Uribe (2003), and the external-premium innovation has monthly persistence $\rho_u = 0.90$. Wages are flexible.

Table 3.1: Calibration

Object	Value	Discipline
<i>Panel A. Kimball demand and firm distribution</i>		
Elasticity level, $\bar{\sigma}_K$	6.00	Reference-point normalization
Superelasticity, ϵ_K	7.92	Chilean invoice data
Pareto tail, κ_A	3.93	Chilean firm distribution
Harmonic sales-weighted markup, $\bar{\mu}$	1.20	Production-function estimates
<i>Panel B. Aggregate and dynamic parameters</i>		
Intermediate-input elasticity, α	0.20	Production
Import share, γ	0.30	Expenditure
Monthly price non-reset probability, δ	0.595	Invoice prices
Frisch elasticity, ζ	0.50	Household labor supply
External-premium persistence, ρ_u	0.90	Monthly process


3.6 Quantitative response

Figure 3.1 reports the first twelve months after the external-premium innovation, normalized to a one-percent impact depreciation. The figure follows the variables entering the mechanism derived above: the exchange rate and import price, the CPI and marginal cost, the wage and employment, and the three components of aggregate productivity. Its first panel also reports the otherwise identical no-input across-firm response, making the intermediate inputs amplification directly visible. The currency movement raises import prices and employment. Marginal cost and the wage rise, but the movement in home costs remains smaller than the competitive effect of the depreciation. Across-firm reallocation is therefore negative from impact through month 12.

Aggregate productivity falls by 0.110 percent on impact. Across-firm reallocation contributes -0.166 percent and remains the largest component. Aggregate input allocation contributes 0.095 percent: total intermediate use rises relative to labor, and firms substitute home inputs for imports. The import-measurement bias contributes -0.039 percent because the rise in dollar import prices reduces imported-input quantities. The two input margins therefore partly offset, with a net contribution of 0.056 percent. Across-firm reallocation is about three times that net input contribution in absolute value and determines the sign of the total response.

In an otherwise identical economy without intermediate inputs, the impact across-firm component is -0.141 percent. Introducing the calibrated input share raises its magnitude

Figure 3.1: Response to a one-percent depreciation



Notes: Monthly percent log-level deviations from steady state. The exchange-rate path is an equilibrium outcome; after solving the model, the response to the external-premium innovation is rescaled to a one-percent impact depreciation. The monetary-policy innovation is zero. Wages are flexible. All panels show levels, not monthly growth rates or cumulative responses. In Panel (a), the thin gray dashed line is the across-firm response in the otherwise identical economy without intermediate inputs.

to -0.166 percent, an amplification of about 18 percent. The total response falls by less because the two additional input margins have a positive net contribution on impact.

The same ordering persists after impact. Both aggregate productivity and the across-firm component remain negative throughout the first year, while the two input terms shrink and continue to offset one another. Intermediate inputs thus matter principally because they magnify the allocation loss across firms; the additional input margins modify, but do not overturn, that mechanism.

4 Reduced-form evidence

The empirical analysis separates two implications of the model. In the Chilean firm-month panel introduced in section 3.5, the firm comparison asks whether activity moves along the markup dimension predicted by the model, while an aggregate TFP projection records the companion aggregate response at the same horizons. Because Chilean trade is predominantly invoiced in dollars, these exercises hold the invoice currency fixed. An annual country panel uses cross-country variation in dollar invoicing to examine the complementary exposure margin: whether the productivity response to a broad dollar movement is concentrated in economies whose trade prices are sticky in dollars. Details on variable construction, local-projection implementation, and inference are reported in Appendix D.

For the Chilean exercises, exchange-rate variation comes from changes in the Federal Reserve broad nominal dollar index described in D. We use the broad dollar index rather than the bilateral peso-dollar exchange rate because the latter also responds directly to Chile-specific monetary, demand, and financial shocks, whereas broad dollar movements provide variation that is unlikely to be driven by domestic Chilean conditions. We define $z_t = \Delta \log E_t^\$$, where $E_t^\$$ denotes the broad dollar index, so that a positive shock corresponds to a broad appreciation of the U.S. dollar. Because both the shock and the outcomes are expressed in logs, the estimated coefficients in this section are elasticities: they give the log-point response of the outcome to a one-log-point change in the broad dollar index.

4.1 Firm-level reallocation

The theory predicts that the reallocation response varies monotonically along the markup distribution. For the empirical analysis, we use a representation of the reallocation term in proposition 3 that remains valid with intermediate inputs and under weaker assumptions than the closed-form response in proposition 2.

Corollary 2 (Variable-cost representation of reallocation). *Let tvc_θ denote the log deviation of firm θ 's nominal total variable cost and y_θ its gross-output log deviation. Under constant returns and a marginal-cost change mc common across firms,*

$$\text{tvc}_\theta = mc + y_\theta, \quad a_F = -\text{Cov}_\lambda \left[\frac{\bar{\mu}}{\mu_\theta}, \text{tvc}_\theta \right].$$

Thus the reallocation component in proposition 3 can be written in terms of firm-level markups, nominal total variable costs, and sales weights.

Because the common marginal-cost term drops out of the covariance, this representation remains valid with intermediate inputs and provides the firm-level outcome used below. It identifies the across-firm term in proposition 3; aggregate input allocation and the import-measurement bias require aggregate input data. The derivation is in section A.10.

We implement this implication with a predetermined high-low split rather than estimating the covariance directly. Let $\mu_{i,t-12}$ denote firm i 's markup twelve months earlier, constructed from the monthly estimates described in section D.2, and define

$$D_{it}^H = \mathbf{1}\{\mu_{i,t-12} > \text{median}_{j \in s(i)} (\mu_{j,t-12})\},$$

The high- and low-markup classification is based on markups observed twelve months earlier and is therefore predetermined relative to the current shock. Define $\Delta_h x_{t+h} \equiv x_{t+h} - x_{t-1}$. For horizon h , we estimate

$$\Delta_h \log TVC_{i,t+h} = \alpha_{i,m}^h + \tau_{s(i),t}^h + \delta_h D_{it}^H + \beta_h^{H-L} (D_{it}^H \times z_t) + u_{i,t+h},$$

where TVC is total variable cost⁷, the wage bill plus material expenditure; z_t is the broad-dollar movement; and $\alpha_{i,m}^h$ and $\tau_{s(i),t}^h$ denote firm-by-month-of-year and industry-by-calendar-month effects. Baseline standard errors are clustered by industry; Appendix D reports robustness to two-way clustering by industry and calendar month.

Industry-by-calendar-month effects absorb movements common to firms within an industry-month cell, so identification comes from the differential response of high- and low-markup firms within those cells. The coefficient β_h^{H-L} is the high-minus-low-markup response. A depreciation shifts activity toward the more elastic, low-markup firms, so the prediction is $\beta_h^{H-L} < 0$: high-markup firms' variable costs rise less, or fall more, than those of low-markup firms.

⁷Total variable cost is the outcome because it appears directly in the covariance representation of corollary 2.

4.2 Aggregate productivity

The firm panel is aggregated to the same monthly frequency. The baseline productivity measure is a capital and labor TFP residual,

$$a_t^{KL} = \log Y_t^{VA} - s_t^L \log L_t - s_t^K \log K_t,$$

and the local projections use its twelve-month moving average to reduce high-frequency noise. For comparison, we also construct labor productivity,

$$a_t^L = \log Y_t^{VA} - \log L_t,$$

which does not require a monthly capital stock. For either measure, the aggregate local projection uses the same broad-dollar movement as the firm regression,

$$\Delta_h a_{t+h}^j = \beta_h^j z_t + \gamma_m^h + v_{t+h}^j, \quad j \in \{KL, L\},$$

where γ_m^h denotes month-of-year effects.

Figure 4.1 places the firm-level reallocation response and the aggregate TFP response side by side. Panel A asks whether activity moves along the markup dimension, and Panel B reports the baseline capital-and-labor TFP response to the same shock and at the same horizons. Neither aggregate productivity measure isolates the across-firm component on its own; that role belongs to the firm-level variable-cost comparison. The corresponding labor-productivity response is reported in Appendix D.

The responses are consistent with the model's reallocation mechanism. Following a broad dollar appreciation, total variable cost declines more at high-markup firms than at low-markup firms, with the differential response reaching its largest magnitude after about one year and subsequently dissipating. Aggregate TFP displays little persistent response over the first year, but declines thereafter, with the fall becoming more pronounced at horizons of two to three years. Thus the cross-firm reallocation predicted by the model precedes a deterioration in aggregate measured productivity. The results are robust to two-way clustering by industry and calendar month, which leaves the point estimates unchanged but yields wider confidence intervals, and to using labor productivity in place of capital-and-labor TFP; both robustness exercises are reported in Appendix D.


4.3 Invoice-currency exposure across countries

The Chilean exercises examine reallocation conditional on dollar invoicing; they do not compare alternative invoice currencies. The country panel addresses that second ingre-

Figure 4.1: Exchange-rate responses in the Chilean firm panel

A. Reallocation across markup groups

B. Aggregate measured TFP



TBD

Notes: Local-projection responses to changes in the Federal Reserve broad nominal dollar index; shaded areas are 95% confidence intervals. Panel A reports the high-minus-low-markup response of total variable cost, where markup groups are defined using markups lagged twelve months. Panel B reports the response of aggregate capital-and-labor TFP using its twelve-month moving average.

dient. Proposition 5 implies that a pure common-dollar movement has no direct effect on cross-border prices under producer- or local-currency pricing, whereas sticky dollar prices transmit it into local currency. The corresponding reduced-form implication is that productivity responses to the broad dollar factor should be concentrated in economies with greater dollar-invoicing exposure.

Annual labor productivity is real GDP per hour from the Penn World Table, the dollar factor is the Federal Reserve broad nominal dollar index, and dollar-invoicing exposure is the dollar share of imports in Boz et al. (2022).

Table 4.1 reports the response for the high-dollar-invoicing group, the high-minus-low difference, and the continuous invoicing gradient. The first specification uses country effects without year effects; the two relative specifications use both country and year effects. Exact estimating equations, controls, and sample construction are in section D.

Table 4.1: The dollar factor and annual productivity across countries

	(1) High-DCP response	(2) difference	(3) gradient
$h = 0$	−0.09***	−0.09**	−0.21***
$h = 1$	−0.13**	−0.12*	−0.33**
$h = 2$	−0.15**	−0.18**	−0.50***
$h = 3$	−0.16*	−0.24**	−0.66***

Notes: Local projections of cumulative log labor productivity (real GDP per hour from the Penn World Table) on the broad dollar factor. Column (1) reports the dollar response for economies whose dollar-invoiced import share is at least 65 percent and includes country effects but no year effects. Column (2) reports the difference between those economies and the rest; column (3) reports the coefficient on the dollar factor interacted with the continuous dollar-invoicing share. Horizons are in years. The sample contains 16 high-dollar-invoicing economies and 18 others and excludes the euro area. Columns (2) and (3) include country and year fixed effects. Standard errors are clustered by country. One, two, and three stars indicate statistical significance at the 10%, 5%, and 1% level.

The estimates have the predicted sign and become larger with the horizon. Relative to the other economies, measured labor productivity in the high-dollar-invoicing group is 0.09 percent lower on impact and 0.24 percent lower after three years following a one-percent broad dollar appreciation. The continuous exposure gradient steepens from −0.21 on impact to −0.66 at three years. This concentration supplies evidence for the invoicing assumption that underlies the Chilean exercise. It does not identify the cross-firm reallocation term or distinguish that mechanism from other effects of dollar-priced trade.

5 Conclusion

Exchange rates are nominal variables, yet aggregate productivity moves with them. This paper shows how that comovement can arise without any change in technology. When cross-border prices are sticky in their currency of invoicing, exchange-rate movements change relative prices; when firms have heterogeneous markups, the resulting expenditure shifts reallocate resources across firms with different marginal products. A depreciation shifts activity toward more elastic, low-markup firms and thereby worsens the initial allocation. Intermediate inputs amplify this across-firm loss through the production network and introduce additional input-allocation and measurement margins. In the quantitative model, these additional margins partly offset one another, while the amplified across-firm component determines the sign of measured productivity.

The empirical evidence is consistent with this mechanism. In the Chilean firm panel,

broad dollar movements reallocate variable costs across high- and low-markup firms in the direction predicted by the model, and aggregate measured productivity subsequently declines. Across countries, productivity responses to broad dollar movements are stronger where dollar invoicing is more prevalent, supporting the pricing condition that gives the mechanism traction. The broader implication is that aggregate productivity dynamics need not reflect changes in technology: nominal shocks can move measured productivity by changing where production takes place across an initially distorted distribution of firms.

References

- Akerberg, Daniel A., Kevin Caves, and Garth Frazer, “Identification Properties of Recent Production Function Estimators,” *Econometrica*, 2015, 83 (6), 2411–2451.
- Aguiar, Mark and Gita Gopinath, “Emerging Market Business Cycles: The Cycle Is the Trend,” *Journal of Political Economy*, 2007, 115 (1), 69–102.
- Albagli, Elias, Francesco Grigoli, Emanuele Luttini, Felipe Quevedo, and Matias Rojas, “Beyond Costs: The Dominant Role of Strategic Complementarities in Pricing,” 2026. Central Bank of Chile, working paper.
- Arkolakis, Costas, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare, “The Elusive Pro-Competitive Effects of Trade,” *Review of Economic Studies*, 2019, 86 (1), 46–80.
- Atkeson, Andrew and Ariel Burstein, “Pricing-to-Market, Trade Costs, and International Relative Prices,” *American Economic Review*, 2008, 98 (5), 1998–2031.
- Baqae, David Rezza and Emmanuel Farhi, “Productivity and Misallocation in General Equilibrium,” *Quarterly Journal of Economics*, 2020, 135 (1), 105–163.
- , —, and Kunal Sangani, “The Supply-Side Effects of Monetary Policy,” *Journal of Political Economy*, 2024, 132 (4).
- Blum, Bernardo S., Sebastian Claro, Ignatius Horstmann, and David A. Rivers, “The ABCs of Firm Heterogeneity When Firms Sort into Markets: The Case of Exporters,” *Journal of Political Economy*, 2024, 132 (4), 1162–1208.
- Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch, “Some Unpleasant Markup Arithmetic: Production Function Elasticities and Their Estimation from Production Data,” *Journal of Monetary Economics*, 2021, 121, 1–14.
- Boz, Emine, Camila Casas, Georgios Georgiadis, Gita Gopinath, Helena Le Mezo, Arnaud Mehl, and Tra Nguyen, “Patterns of Invoicing Currency in Global Trade: New Evidence,” *Journal of International Economics*, 2022, 136, 103604.

- Burstein, Ariel and Javier Cravino**, “Measured Aggregate Gains from International Trade,” *American Economic Journal: Macroeconomics*, 2015, 7 (2), 181–218.
- Cravino, Javier**, “Exchange Rates, Aggregate Productivity, and the Currency of Invoicing of International Trade,” 2017. Working Paper.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu**, “Competition, Markups, and the Gains from International Trade,” *American Economic Review*, 2015, 105 (10), 3183–3221.
- Fujiwara, Ippei and Kiminori Matsuyama**, “Competition and the Phillips Curve,” 2023. Working paper.
- Giarda, Mario and Antonio Martner**, “Markup Distribution and Inflation Dynamics,” 2026. Working Paper.
- Gopinath, Gita and Brent Neiman**, “Trade Adjustment and Productivity in Large Crises,” *American Economic Review*, 2014, 104 (3), 793–831.
- , **Emine Boz, Camila Casas, Federico J. Díez, Pierre-Olivier Gourinchas, and Mikkel Plagborg-Møller**, “Dominant Currency Paradigm,” *American Economic Review*, 2020, 110 (3), 677–719.
- Ishide, Akira**, “Allocative Inefficiency during a Sudden Stop,” 2026. Working Paper.
- Kehoe, Timothy J. and Kim J. Ruhl**, “Are Shocks to the Terms of Trade Shocks to Productivity?,” *Review of Economic Dynamics*, 2008, 11 (4), 804–819.
- Kimball, Miles S.**, “The Quantitative Analytics of the Basic Neomonetarist Model,” *Journal of Money, Credit and Banking*, 1995, 27 (4), 1241–1277.
- Koike, Yasutaka and Antonio Martner**, “Aggregating Distortions in Networks with Multi-Product Firms,” 2026. Working Paper.
- Loecker, Jan De and Frederic Warzynski**, “Markups and Firm-Level Export Status,” *American Economic Review*, 2012, 102 (6), 2437–2471.
- Meier, Matthias and Timo Reinelt**, “Monetary Policy, Markup Dispersion, and Aggregate TFP,” *Review of Economics and Statistics*, 2024, 106.
- Mendoza, Enrique G.**, “The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations,” *International Economic Review*, 1995, 36 (1), 101–137.
- Schmitt-Grohe, Stephanie and Martin Uribe**, “Closing Small Open Economy Models,” *Journal of International Economics*, 2003, 61 (1), 163–185.

Appendix

A Static and measurement derivations

A.1 The price aggregator

This subsection derives the static fixed point for the price aggregator p used in section 2.2. For a domestic firm that resets, the home-market price is

$$\tilde{p}_{h\theta} = \rho_\theta w + (1 - \rho_\theta)p.$$

For a foreign firm selling to home under DCP, the reset price in home currency is

$$\tilde{p}_{f\theta} = \rho_\theta e + (1 - \rho_\theta)p,$$

and a non-resetting foreign price is $p_{f\theta} = e$. Under the Kimball aggregator, a resetter's first-order condition weights competitors' prices by their sales and demand elasticities. Averaging reset and inherited prices under the common Calvo parameter and common pass-through moments gives

$$\begin{aligned} p &= (1 - \gamma)(1 - \delta) [\bar{\rho}^\sigma w + (1 - \bar{\rho}^\sigma)p] \\ &\quad + \gamma \{ (1 - \delta)(1 - \bar{\rho}^\sigma)p + [\delta + (1 - \delta)\bar{\rho}^\sigma] e \}. \end{aligned}$$

Collecting terms in p yields

$$[\delta + (1 - \delta)\bar{\rho}^\sigma] p = (1 - \gamma)(1 - \delta)\bar{\rho}^\sigma w + \gamma [\delta + (1 - \delta)\bar{\rho}^\sigma] e,$$

and hence

$$p = \frac{(1 - \gamma)(1 - \delta)\bar{\rho}^\sigma}{\delta + (1 - \delta)\bar{\rho}^\sigma} w + \gamma e \equiv \kappa_w w + \gamma e.$$

If foreign exporters instead have Calvo parameter δ_f and elasticity-weighted pass-through $\bar{\rho}_f^\sigma$, the exchange-rate coefficient becomes

$$\frac{\gamma [\delta_f + (1 - \delta_f)\bar{\rho}_f^\sigma]}{(1 - \gamma) [\delta + (1 - \delta)\bar{\rho}^\sigma] + \gamma [\delta_f + (1 - \delta_f)\bar{\rho}_f^\sigma]}.$$

It equals the import share γ under the symmetric pricing moments maintained in the main text.

A.2 The price block: pass-through into the CPI

The CPI uses sales weights rather than elasticity weights. Direct aggregation gives

$$p^C = (1 - \gamma)(1 - \delta) [\bar{\rho}w + (1 - \bar{\rho})p] + \gamma \{ (1 - \delta)(1 - \bar{\rho})p + [\delta + (1 - \delta)\bar{\rho}]e \}.$$

Using the fixed point of section A.1,

$$p^C = [(1 - \gamma)(1 - \delta)\bar{\rho} + (1 - \delta)(1 - \bar{\rho})\kappa_w]w + \gamma [(1 - \delta)(1 - \bar{\rho}) + \delta + (1 - \delta)\bar{\rho}]e.$$

The exchange-rate coefficient is

$$\gamma [(1 - \delta)(1 - \bar{\rho}) + \delta + (1 - \delta)\bar{\rho}] = \gamma.$$

The wage coefficient is

$$\kappa_w^C = (1 - \gamma)(1 - \delta)\bar{\rho} + (1 - \delta)(1 - \bar{\rho})\kappa_w = \frac{(1 - \delta)(1 - \gamma) [\delta\bar{\rho} + (1 - \delta)\bar{\rho}^\sigma]}{\delta + (1 - \delta)\bar{\rho}^\sigma}.$$

A.3 Firm-destination aggregation

Lemma 1 (Open-economy reallocation statistic). *Let $\ell_{h\theta}$ and ℓ_θ^* denote the log deviations of firm θ 's labor used for home and export sales. Under (A1)–(A2), the first-order change in aggregate productivity is*

$$a = \text{Cov}_\lambda \left[-\frac{\bar{\mu}}{\mu_\theta}, (1 - \gamma)\ell_{h\theta} + \gamma\ell_\theta^* \right].$$

Proof. Apply the first-order productivity accounting of Baqaee and Farhi (2020) to firm–destination cells (θ, m) , where $m \in \{h, x\}$ denotes home and export sales. With labor as the only input and fixed firm technologies,

$$a = -\text{Cov}_{\lambda(\theta, m)} \left[\frac{\bar{\mu}}{\mu_{\theta m}}, \ell_{\theta m} \right].$$

Under (A1), $\mu_{\theta h} = \mu_{\theta x} = \mu_\theta$, so the inverse-markup weight is constant within a firm. Under (A2), the firm–destination sales shares satisfy

$$\lambda_{\theta h} = (1 - \gamma)\lambda_\theta, \quad \lambda_{\theta x} = \gamma\lambda_\theta, \quad \lambda_\theta \equiv \lambda_{\theta h} + \lambda_{\theta x}.$$

Thus the cell weights aggregate to the firm's total-sales share, and the corresponding harmonic markup is $\bar{\mu} = [\mathbb{E}_\lambda(\mu_\theta^{-1})]^{-1}$. Therefore

$$\mathbb{E}_{\lambda(\theta, m)} \left[\frac{\bar{\mu}}{\mu_{\theta m}} \ell_{\theta m} \right] = \mathbb{E}_{\lambda} \left[\frac{\bar{\mu}}{\mu_{\theta}} \left((1 - \gamma) \ell_{h\theta} + \gamma \ell_{\theta}^* \right) \right].$$

The covariance form follows because $\mathbb{E}_{\lambda}[\bar{\mu}/\mu_{\theta}] = 1$ and the aggregate component of employment growth drops out of the covariance:

$$a = \text{Cov}_{\lambda} \left[-\frac{\bar{\mu}}{\mu_{\theta}}, (1 - \gamma) \ell_{h\theta} + \gamma \ell_{\theta}^* \right].$$

At the initial allocation, branch revenue equals $\mu_{\theta} \bar{W}$ times branch labor, so within-firm sales shares also equal labor shares. Holding each firm's total labor fixed therefore implies $(1 - \gamma) \ell_{h\theta} + \gamma \ell_{\theta}^* = 0$ for every θ . Such a reallocation between destinations has no first-order productivity effect; changes in firms' total labor allocations can still affect a . ■

A.4 Proof of Proposition 1

With $\delta = 1$, home prices and dollar export prices are unchanged, and $p = \gamma e$. Since technology and foreign expenditure are fixed, the branch labor responses are

$$\ell_{h\theta} = c + \sigma_{\theta} \gamma e, \quad \ell_{\theta}^* = 0.$$

Substitution into lemma 1 gives

$$a = -\text{Cov}_{\lambda} \left[\frac{\bar{\mu}}{\mu_{\theta}}, (1 - \gamma)c + \gamma(1 - \gamma)\sigma_{\theta} e \right] = -\gamma(1 - \gamma)C_{\sigma} e.$$

The common expenditure component drops out. Because $\bar{\mu}/\mu_{\theta} = \bar{\mu}(1 - 1/\sigma_{\theta})$ is strictly increasing in σ_{θ} , the covariance is positive for a nondegenerate markup distribution.

A.5 The covariance identities

Propositions 1 and 2 define the two reallocation covariances C_{σ} and $C_{\sigma\rho}$. This appendix records their closed forms in calibrated moments and decomposes the resetting covariance into the fixed-price covariance and the sorting of elasticity with desired pass-through.

In the zero-inflation steady state every firm charges its desired markup, so $\mu_{\theta} = \sigma_{\theta}/(\sigma_{\theta} - 1)$ and $\bar{\mu}/\mu_{\theta} = \bar{\mu}(1 - 1/\sigma_{\theta})$. Write $\bar{\sigma}_{\lambda} \equiv \mathbb{E}_{\lambda}[\sigma_{\theta}]$ for the sales-weighted mean elasticity. This is not the unit-relative-quantity parameter $\bar{\sigma}_K$ used in the KW calibration of section 3.5. Then

$$C_\sigma = \mathbb{E}_\lambda [\bar{\mu}(\sigma_\theta - 1)] - \bar{\sigma}_\lambda = (\bar{\mu} - 1) \bar{\sigma}_\lambda - \bar{\mu} \geq 0,$$

nonnegative because $\bar{\mu}/\mu_\theta$ is increasing in σ_θ , and, by the definition of $\bar{\rho}^\sigma$,

$$\text{Cov}_\lambda[\sigma_\theta, \rho_\theta] = \bar{\sigma}_\lambda (\bar{\rho}^\sigma - \bar{\rho}).$$

The identity used in proposition 2 follows from aggregation. A resetting firm's quantity moves with $\sigma_\theta \rho_\theta$ per unit of the common cost gap, so $C_{\sigma\rho}$ summarizes the response of firms that reset their prices. Using $\mathbb{E}_\lambda[\bar{\mu}/\mu_\theta] = 1$ and $\mathbb{E}_\lambda[\sigma_\theta \rho_\theta] = \bar{\sigma}_\lambda \bar{\rho}^\sigma$,

$$C_{\sigma\rho} = \mathbb{E}_\lambda [\bar{\mu}(\sigma_\theta - 1)\rho_\theta] - \bar{\sigma}_\lambda \bar{\rho}^\sigma = (\bar{\mu} - 1) \bar{\sigma}_\lambda \bar{\rho}^\sigma - \bar{\mu} \bar{\rho} = \bar{\rho} C_\sigma + (\bar{\mu} - 1) \text{Cov}_\lambda[\sigma_\theta, \rho_\theta].$$

These identities relate the fixed-price and resetting statistics used in propositions 1 and 2.

A.6 Proof of Proposition 2

With foreign expenditure fixed, demand changes are

$$\ell_{h\theta} = c - \sigma_\theta(p_{h\theta} - p), \quad \ell_\theta^* = -\sigma_\theta(p_{h\theta}^* - p^*), \quad p^* = 0.$$

For domestic sales,

$$p_{h\theta} - p = \begin{cases} \rho_\theta(w - p), & \text{if the firm resets,} \\ -p, & \text{if the price is sticky.} \end{cases}$$

For export sales under DCP,

$$p_{h\theta}^* = \begin{cases} \rho_\theta(w - e), & \text{if the firm resets,} \\ 0, & \text{if the price is sticky.} \end{cases}$$

Taking the type average over the Calvo states gives

$$(1 - \gamma)\ell_{h\theta} + \gamma\ell_\theta^* = (1 - \gamma)c - (1 - \delta)\sigma_\theta\rho_\theta[(1 - \gamma)(w - p) + \gamma(w - e)] + \delta(1 - \gamma)\sigma_\theta p.$$

Substituting this expression into lemma 1 eliminates the common expenditure component and gives

$$a = (1 - \delta)C_{\sigma\rho}[(1 - \gamma)(w - p) + \gamma(w - e)] - (1 - \gamma)\delta C_{\sigma}p.$$

By the fixed point of section A.1,

$$(1 - \gamma)(w - p) + \gamma(w - e) = [1 - (1 - \gamma)\kappa_w]w - \gamma(2 - \gamma)e.$$

Substitution yields

$$a = -\gamma[(1 - \delta)C_{\sigma\rho}(2 - \gamma) + (1 - \gamma)\delta C_{\sigma}]e + [(1 - \delta)C_{\sigma\rho}[1 - (1 - \gamma)\kappa_w] - (1 - \gamma)\delta C_{\sigma}\kappa_w]w.$$

This is $a = -\Psi_e e + \Psi_w w$. The market-specific expression in proposition 2 follows by writing $2 - \gamma = (1 - \gamma) + 1$: the first term is the home-market contribution and the second is the export contribution.

To sign the wage loading, use the covariance identity in section A.5:

$$C_{\sigma\rho} - \bar{\rho}^\sigma C_{\sigma} = \bar{\mu}(\bar{\rho}^\sigma - \bar{\rho}).$$

Substituting this identity and $\kappa_w = (1 - \gamma)(1 - \delta)\bar{\rho}^\sigma / [\delta + (1 - \delta)\bar{\rho}^\sigma]$ into the coefficient on w gives

$$\Psi_w = (1 - \delta)[\gamma(2 - \gamma)\bar{\rho}^\sigma C_{\sigma} + \{1 - (1 - \gamma)\kappa_w\}\bar{\mu}(\bar{\rho}^\sigma - \bar{\rho})].$$

Marshall's second law implies $C_{\sigma} \geq 0$, and the third implies $\bar{\rho}^\sigma - \bar{\rho} \geq 0$. Since $1 - (1 - \gamma)\kappa_w > 0$, $\Psi_w \geq 0$. The inequality is strict when $\delta < 1$ and the firm distribution is nondegenerate.

A.7 Static monetary regimes

The main text uses strict CPI stabilization. More generally, let the static CPI response be $p^C = -\chi e$. Combining the price block $p^C = \kappa_w^C w + \gamma e$ with proposition 2 gives

$$\frac{w}{e} = -\frac{\gamma + \chi}{\kappa_w^C}, \quad \frac{a}{e} = -\Psi_e - \frac{\Psi_w}{\kappa_w^C}(\gamma + \chi).$$

When $\delta < 1$ and firm heterogeneity is nondegenerate, $\Psi_w > 0$, so a larger χ increases the allocative loss. Strict CPI stabilization sets $\chi = 0$; a fixed nominal wage sets $\chi = -\gamma$ and gives $a/e = -\Psi_e$. A hard peg sets $e = 0$ and lies outside the exchange-rate experiment in proposition 2. These formulas map the wage response under each regime into allocative efficiency.

A.8 Conditional one-period quantities

For any realized pair (w, e) , household optimality, goods-market clearing, and production imply

$$\begin{aligned} w - p^C &= \eta c + \frac{l}{\zeta}, \\ y &= (1 - \gamma)c - \mathbb{E}_\lambda \left[\sigma_\theta \left\{ (1 - \gamma)(p_{h\theta} - p) + \gamma p_{h\theta}^* \right\} \right], \\ y &= a + l. \end{aligned}$$

The expectation is the sales-weighted demand displacement generated by firms' relative prices; when positive, it lowers aggregate sales at given expenditure. These equations determine (c, l, y) conditional on (w, e) . No equation in the one-period comparison selects the source or subsequent path of e . Their common expenditure component scales firms proportionally and therefore drops out of the covariance in section A.3. The implied external balance is residual in the one-period exercise; section 3.3 supplies the intertemporal UIP and current-account closure.

A.9 Aggregation of the Kimball input basket

Index an input-buying firm by i and a supplying variety by θ . The intermediate composite used by buyer i is the unique X_{it} satisfying

$$\int_0^1 \left[(1 - \gamma) \Upsilon \left(\frac{X_{i\theta t}^H}{(1 - \gamma) X_{it}} \right) + \gamma \Upsilon \left(\frac{X_{i\theta t}^M}{\gamma X_{it}} \right) \right] d\theta = 1. \quad (\text{A.1})$$

The ratios inside Υ make this aggregator homogeneous of degree one. At the reference allocation, let the normalized quantity supplied by either source be q_θ , and define

$$D_X \equiv \int_0^1 \Upsilon'(q_\theta) q_\theta d\theta, \quad \lambda_\theta^X \equiv \frac{\Upsilon'(q_\theta) q_\theta}{D_X}. \quad (\text{A.2})$$

Log differentiating (A.1) around that allocation gives

$$\begin{aligned} 0 &= (1 - \gamma) \int_0^1 \Upsilon'(q_\theta) q_\theta (x_{i\theta t}^H - x_{it}) d\theta \\ &\quad + \gamma \int_0^1 \Upsilon'(q_\theta) q_\theta (x_{i\theta t}^M - x_{it}) d\theta. \end{aligned} \quad (\text{A.3})$$

Dividing by D_X therefore yields

$$x_{it} = (1 - \gamma)\mathbb{E}_{\lambda^X}[x_{it}^H] + \gamma\mathbb{E}_{\lambda^X}[x_{it}^M]. \quad (\text{A.4})$$

All buyers face the same supplier prices. Homotheticity then implies $X_{it}^s = d_{\theta t}^s X_{it}$ for $s \in \{H, M\}$, where $d_{\theta t}^s$ is common across buyers. Hence, with

$$X_t \equiv \int X_{it} di, \quad X_{\theta t}^s \equiv \int X_{it}^s di,$$

we have $X_{\theta t}^s = d_{\theta t}^s X_t$. The aggregate input vector is therefore collinear with every buyer's input vector and satisfies the same Kimball constraint with composite X_t . Applying (A.3) to that aggregate vector and defining

$$x_t^H \equiv \mathbb{E}_{\lambda^X}[x_{\theta t}^H], \quad x_t^M \equiv \mathbb{E}_{\lambda^X}[x_{\theta t}^M]$$

therefore gives

$$x_t = (1 - \gamma)x_t^H + \gamma x_t^M. \quad (\text{A.5})$$

The same weights arise from expenditure minimization. If P_t^X denotes the price aggregator associated with (A.1), its first-order conditions are $P_{\theta t}^s/P_t^X = \Upsilon'(q_{\theta t}^s)$. Thus λ_{θ}^X is the within-source expenditure share at the reference allocation. Because the two source distributions are matched, their total expenditure shares are $1 - \gamma$ and γ . The exact unit-expenditure index consequently satisfies

$$\begin{aligned} p_t^C &= (1 - \gamma)p_t^H + \gamma p_t^M, \\ p_t^s &\equiv \mathbb{E}_{\lambda^X}[p_{\theta t}^s], \quad s \in \{H, M\}. \end{aligned} \quad (\text{A.6})$$

Under the common-basket normalization of the main text, $\lambda_{\theta}^X = \lambda_{\theta}$, and the unit-expenditure index is the CPI. Buyer-specific input prices or input baskets would break the proportionality used above and introduce additional aggregation terms.

A.10 Proof of Proposition 3

The result can be obtained directly from firm sales, variable costs, and the national accounts. Index a firm and destination by (θ, m) , with $m \in \{h, x\}$ for home and export sales. Let $R_{\theta m}$ be steady-state revenue and let $V_{\theta m}$ be the variable cost assigned to that branch. Constant returns, no fixed costs, and equal steady-state markups across a firm's

destinations imply

$$V_{\theta m} = \frac{R_{\theta m}}{\mu_{\theta}}.$$

Writing $S \equiv \sum_{\theta, m} R_{\theta m}$ and $V \equiv \sum_{\theta, m} V_{\theta m}$, define branch sales and variable-cost weights by

$$\lambda_{\theta m} \equiv \frac{R_{\theta m}}{S}, \quad \omega_{\theta m} \equiv \frac{V_{\theta m}}{V} = \lambda_{\theta m} \frac{\bar{\mu}}{\mu_{\theta}}, \quad V = \frac{S}{\bar{\mu}}. \quad (\text{A.7})$$

Under (A2), $\lambda_{\theta h} = (1 - \gamma)\lambda_{\theta}$ and $\lambda_{\theta x} = \gamma\lambda_{\theta}$. Define $\omega_{\theta} \equiv \sum_m \omega_{\theta m}$, so that $\omega_{\theta} = \lambda_{\theta} \bar{\mu} / \mu_{\theta}$. If $y_{\theta m}$ is the branch quantity deviation and $y_{\theta} = (1 - \gamma)y_{\theta h} + \gamma y_{\theta x}$, the difference between sales-weighted and variable-cost-weighted gross output is

$$\begin{aligned} a_F &\equiv \mathbb{E}_{\lambda}[y_{\theta}] - \mathbb{E}_{\omega}[y_{\theta}] \\ &= -\text{Cov}_{\lambda}\left[\frac{\bar{\mu}}{\mu_{\theta}}, y_{\theta}\right]. \end{aligned} \quad (\text{A.8})$$

This is the covariance in lemma 1. It follows here from the different weights used to aggregate the same firm quantities, rather than from an aggregate production function.

Let s and v denote the log deviations of S and V , and let $\mu_{h\theta}$ and $\mu_{x\theta}$ denote log deviations of realized markups in the two destinations. Since the marginal-cost deviation is common across firms,

$$\begin{aligned} s &= mc + \mathbb{E}_{\lambda}[y_{\theta}] + (1 - \gamma)\mathbb{E}_{\lambda}[\mu_{h\theta}] + \gamma\mathbb{E}_{\lambda}[\mu_{x\theta}], \\ v &= mc + \mathbb{E}_{\omega}[y_{\theta}] = w + l. \end{aligned} \quad (\text{A.9})$$

The last equality uses $WL = (1 - \alpha)V$: the labor bill and total variable cost have the same log change because their ratio is fixed.

Steady-state nominal GDP is domestic gross sales less intermediate expenditure,

$$\text{nominal GDP} = S - \alpha V = S \frac{\bar{\mu} - \alpha}{\bar{\mu}}.$$

Hence gross sales and intermediate expenditure have steady-state GDP exposures ξ_F and $\xi_F - 1$, respectively. Differentiating this level identity and using (A.8)–(A.9) gives

$$y + p^{\text{GDP}} = w + l + \xi_F a_F + \xi_F [(1 - \gamma)\mathbb{E}_{\lambda}[\mu_{h\theta}] + \gamma\mathbb{E}_{\lambda}[\mu_{x\theta}]], \quad \xi_F = \frac{\bar{\mu}}{\bar{\mu} - \alpha}. \quad (\text{A.10})$$

Thus the multiplier on a_F comes directly from converting gross sales into GDP.

We next subtract the GDP deflator. In the balanced-trade steady state, consumption

equals GDP. Final imports therefore enter consumption and imports with opposite signs and cancel from the deflator. Exports have GDP exposure $\gamma\xi_F$, while imported intermediates have GDP exposure $\alpha\gamma/(\bar{\mu} - \alpha)$, and domestic final expenditure has share $1 - \gamma$. The Divisia GDP deflator is consequently

$$p^{GDP} = (1 - \gamma)p^H + \gamma\xi_F p^E - \frac{\alpha\gamma}{\bar{\mu} - \alpha} p^M. \quad (\text{A.11})$$

To first order, $p^H = mc + \mathbb{E}_\lambda[\mu_{h\theta}]$ and $p^E = mc + \mathbb{E}_\lambda[\mu_{x\theta}]$. Substituting these expressions and (A.10) into $a = y - \widetilde{\Lambda}_L l$ cancels the export-markup term and gives

$$\begin{aligned} a = & \xi_F a_F + w + (1 - \widetilde{\Lambda}_L)l + (1 - \gamma)(\xi_F - 1)\mathbb{E}_\lambda[\mu_{h\theta}] \\ & - (1 - \gamma + \gamma\xi_F)mc + \frac{\alpha\gamma}{\bar{\mu} - \alpha} p^M. \end{aligned} \quad (\text{A.12})$$

It remains to collect the common-input terms. Let $\widetilde{\alpha} \equiv \alpha(1 - \gamma)$. The external-demand exposure to domestic gross output is $1 - \gamma + \gamma\xi_F = (\bar{\mu} - \widetilde{\alpha})/(\bar{\mu} - \alpha)$. Repeated use of domestic intermediates multiplies that exposure by $1 + \widetilde{\alpha} + \widetilde{\alpha}^2 + \dots$, and labor has share $1 - \alpha$ in each round. Hence

$$\widetilde{\Lambda}_L = (1 - \alpha) \frac{\bar{\mu} - \widetilde{\alpha}}{(\bar{\mu} - \alpha)(1 - \widetilde{\alpha})}.$$

The intermediate-input basket has the CPI cost index, $p^C = (1 - \gamma)p^H + \gamma p^M$, and $mc = (1 - \alpha)w + \alpha p^C$. Using $p^H = mc + \mathbb{E}_\lambda[\mu_{h\theta}]$ therefore yields

$$mc = \frac{(1 - \alpha)w + \widetilde{\alpha}\mathbb{E}_\lambda[\mu_{h\theta}] + \alpha\gamma p^M}{1 - \widetilde{\alpha}}. \quad (\text{A.13})$$

Substituting (A.13) into (A.12) and collecting coefficients gives the equivalent price-and-expenditure representation

$$a = \xi_F a_F - \xi_X \mathbb{E}_\lambda[\mu_{h\theta}] + \xi_M (w + l - p^M), \quad (\text{A.14})$$

where

$$\xi_X = \frac{\widetilde{\alpha}(\bar{\mu} - 1)}{(\bar{\mu} - \alpha)(1 - \widetilde{\alpha})}, \quad \xi_M = \frac{\alpha\gamma(\bar{\mu} - 1)}{(\bar{\mu} - \alpha)(1 - \widetilde{\alpha})}.$$

In particular, $1 - \widetilde{\Lambda}_L = \xi_M$ and $\xi_M/\xi_X = \gamma/(1 - \gamma)$.

Finally, section A.9 and cost minimization give

$$p^C = (1 - \gamma)p^H + \gamma p^M, \quad x = (1 - \gamma)x^H + \gamma x^M, \quad w + l = p^C + x.$$

Using these identities, $mc = (1 - \alpha)w + \alpha p^C$, and the coefficient ratio above, the last two terms in (A.14) satisfy

$$\begin{aligned}
& -\xi_X(p^H - mc) + \xi_M(w + l - p^M) \\
& = \xi_X(1 - \alpha)(x - l) + \xi_M x \\
& = \xi_X \left[(1 - \alpha)(x - l) + \gamma(x^H - x^M) \right] + \xi_M x^M.
\end{aligned} \tag{A.15}$$

The final term in the last line is the import-measurement bias. In the notation of the open-economy accounting decomposition, ξ_M is the gap between the cost-based exposure of GDP to imported inputs and the import-expenditure share used in measured real GDP. Thus $\xi_M x^M$ is not a technology contribution of imported services. It is the part of an imported-input quantity change left in the measured residual because the two exposure weights differ. The preceding bracket retains the associated input-allocation movements. Combining (A.14)–(A.15) proves proposition 3.

The proof uses the common input elasticity α , common import share γ , and common input basket imposed in the model. Firm-specific input shares or firm-specific input prices would add further covariance terms.

A.11 Variable-cost representation used in the empirical analysis

For each firm, nominal total variable cost satisfies

$$\text{tvc}_\theta = mc + y_\theta.$$

The change in marginal cost is common across firms under the common input basket. It therefore drops out of (A.8), yielding

$$a_F = -\text{Cov}_\lambda \left[\frac{\bar{\mu}}{\mu_\theta}, \text{tvc}_\theta \right].$$

This proves corollary 2. The demand system and price-setting assumptions determine how an exchange-rate movement changes the covariance; the representation itself requires neither a parametric demand curve nor a particular asset-market closure.

B Dynamic derivations and unit conventions

B.1 Cross-border price and quantity aggregation

Section 3.4 states the complete dynamic equilibrium. This appendix derives the cross-border reset targets and the quantity aggregation used in labor demand, then records the unit conventions used in the external block. Unless otherwise stated, variables are log deviations from the zero-inflation steady state. Net foreign debt and the current-account surplus are level deviations scaled by steady-state GDP.

For any Calvo price index p_t^j with reset target z_t^j , the reset price and the inherited-price index obey

$$\tilde{p}_t^j = (1 - \beta\delta)z_t^j + \beta\delta\mathbb{E}_t\tilde{p}_{t+1}^j, \quad p_t^j = \delta p_{t-1}^j + (1 - \delta)\tilde{p}_t^j. \quad (\text{B.1})$$

The three targets used in section 3.4 are

$$z_t^{M,\$} = (1 - \bar{\rho})(p_t - e_t), \quad z_t^{E,\$} = \bar{\rho}(mc_t - e_t), \quad z_t^{M,\$, \sigma} = (1 - \bar{\rho}^\sigma)(p_t - e_t). \quad (\text{B.2})$$

An exporter to the home market has fixed foreign marginal cost and resets against the home competitive benchmark in dollars, $p_t - e_t$. A home exporter instead resets against its own dollar marginal cost, $mc_t - e_t$, because foreign expenditure and the foreign dollar benchmark are fixed. Replacing expenditure weights by elasticity weights gives the third target. Eliminating the reset price from (B.1) gives

$$(1 + \beta\delta^2)p_t^j = \delta p_{t-1}^j + \beta\delta\mathbb{E}_t p_{t+1}^j + (1 - \delta)(1 - \beta\delta)z_t^j.$$

Substituting the three targets in (B.2) produces the price-state equations used in the quantitative equilibrium. In home currency,

$$p_t^M = p_t^{M,\$} + e_t, \quad p_t^E = p_t^{E,\$} + e_t, \quad p_t^{M,\sigma} = p_t^{M,\$, \sigma} + e_t.$$

The first two indices value imports and exports in the national accounts. The third is the elasticity-weighted import-price index entering Kimball input demand.

To interpret the quantity block, define a firm's destination-weighted relative price by

$$\varphi_{\theta t} \equiv (1 - \gamma)(p_{h\theta,t} - p_t) + \gamma p_{h\theta,t}^*.$$

After collecting expenditure movements common to all firms, Kimball demand gives the firm quantity response $-\sigma_\theta \varphi_{\theta t}$. Let $\omega_\theta = \lambda_\theta \bar{\mu} / \mu_\theta$ denote the firm's steady-state variable-

cost weight. From $a_{F,t} = \mathbb{E}_\lambda[y_{\theta t}] - \mathbb{E}_\omega[y_{\theta t}]$,

$$a_{F,t} = \mathbb{E}_\omega[\sigma_\theta \varphi_{\theta t}] - \mathbb{E}_\lambda[\sigma_\theta \varphi_{\theta t}].$$

Hence the relative-price term in main-text labor demand is $\mathbb{E}_\lambda[\sigma_\theta \varphi_{\theta t}] + a_{F,t}$, which equals the variable-cost-weighted response $\mathbb{E}_\omega[\sigma_\theta \varphi_{\theta t}]$. Write $\bar{\sigma}_\lambda \equiv \mathbb{E}_\lambda[\sigma_\theta]$ for the sales-weighted mean demand elasticity. Calvo-vintage aggregation then gives

$$\begin{aligned} (1 + \beta + \psi)\mathbb{E}_\omega[\sigma_\theta \varphi_{\theta t}] &= \mathbb{E}_\omega[\sigma_\theta \varphi_{\theta, t-1}] + \beta \mathbb{E}_t \mathbb{E}_\omega[\sigma_\theta \varphi_{\theta, t+1}] \\ &\quad + \psi \left[\gamma \bar{\sigma}_\lambda \bar{\rho}^\sigma (2 - \gamma)(mc_t - e_t) + C_{\sigma\rho} \left[mc_t - (1 - \gamma)p_t - \gamma e_t \right] \right. \\ &\quad \left. - (1 - \gamma) \bar{\rho}^\sigma C_\sigma \widehat{mc}_t \right] \\ &\quad - (1 - \gamma) \gamma (\bar{\sigma}_\lambda + C_\sigma) (\Delta e_t - \beta \mathbb{E}_t \Delta e_{t+1}). \end{aligned} \tag{B.3}$$

Together with the law of motion for $a_{F,t}$ in proposition 4, this is equivalent to aggregating firm demand under sales and variable-cost weights. Common expenditure growth scales every firm and drops out of both cross-sectional objects.

B.2 Aggregation and monthly units

Let total variable cost be V_t . Cost minimization gives $W_t L_t = (1 - \alpha)V_t$ and $P_t^C X_t = \alpha V_t$. Since both expenditure shares are constant, their log changes imply

$$x_t = w_t + l_t - p_t^C. \tag{B.4}$$

Kimball input demand and the Divisia identity for the common input basket give

$$x_t^M = x_t - \bar{\sigma}_\lambda (p_t^{M,\sigma} - p_t), \quad x_t^H = \frac{x_t - \gamma x_t^M}{1 - \gamma}. \tag{B.5}$$

These equations determine the source-specific quantities used in proposition 3.

The main text scales both net foreign debt b_t and the monthly current- account surplus nx_t by monthly steady-state GDP. This normalization gives

$$i_t^\$ = \phi_b b_t + u_t, \quad b_t = \beta^{-1} b_{t-1} - nx_t. \tag{B.6}$$

To connect this notation to an annual convention, let $b_t^A = b_t/12$ denote debt relative to

annual steady-state GDP and let ϕ_b^A be the annual debt-premium elasticity. Then

$$\phi_b = \frac{\phi_b^A}{144}, \quad i_t^\$ = \frac{\phi_b^A}{12} b_t^A + u_t, \quad b_t^A = \beta^{-1} b_{t-1}^A - \frac{1}{12} n x_t. \quad (\text{B.7})$$

Thus (B.6) is exactly the monthly-unit representation of the annual calibration used in the quantitative implementation.

B.3 Shock normalization

The equilibrium is solved at first order at monthly frequency, with ε_t^u as the primitive innovation and e_t as an endogenous variable. We first compute the response to a unit innovation in ε_t^u , then divide every response by the impact response of e_t . The displayed paths therefore set $e_0 = 1$ and report percent log-level deviations per one-percent impact depreciation. This normalization changes units but neither imposes the exchange-rate path nor changes the equilibrium decomposition.

C Invoice currency and common-dollar exposure

The body studies a small economy moving against an unmoved dollar bloc. The annual evidence in table 4.1 asks a different but narrower question: whether a broad dollar movement has a larger effect where more trade is invoiced in dollars. The connection is a necessary condition involving invoicing exposure. Markup heterogeneity and sticky prices determine the allocative consequence once cross-border prices move; the invoice currency determines whether the common dollar factor moves those prices in the first place.

Proposition 5 (Common-dollar exposure). *Consider a pure common-dollar movement in which all non-dollar currencies depreciate by e against the dollar while bilateral exchange rates among them remain unchanged. Suppose a fraction ϕ of cross-border prices is sticky in dollars and the remainder is sticky in producer or destination currency. The direct effect of the common dollar movement on cross-border price indices is proportional to ϕ and vanishes at $\phi = 0$. Under pure producer- or local-currency pricing, the movement is neutral under the conditions stated below.*

To see the result, consider first a non-resetting import price. If it is fixed in dollars, its local-currency value is $p^\$ + e$, so the common dollar movement passes through one for one. If it is fixed in the producer or destination currency, the relevant bilateral exchange rate between two non-dollar economies has not moved, and the local-currency price is

unchanged. With a non-reset probability δ , the direct contribution to the import-price index is therefore $\phi\delta e$; expenditure shares add the corresponding trade weights to the CPI and strategic price indices.

Reset prices do not add a separate term in the common dollar movement. Expressed in a common local currency, the seller's marginal cost and competitive price index move together. For example, under dollar invoicing a foreign exporter's reset price converts from

$$\rho_\theta(mc - e) + (1 - \rho_\theta)(p - e)$$

in dollars to

$$\rho_\theta mc + (1 - \rho_\theta)p$$

in local currency. The common dollar component cancels. Hence only the sticky dollar-priced cells transmit the pure common factor directly, and their mass is proportional to ϕ .

At $\phi = 0$, prices are sticky only in producer or destination currencies. Since the relevant bilateral rates among non-dollar economies are unchanged, both fixed and reset prices are unchanged. With no other disturbance, the coefficient on the common dollar factor is zero, and the equilibrium response is zero. This neutrality assumes that direct trade with the dollar bloc has zero measure, that the movement is a pure common factor rather than a collection of bilateral shocks, and that wages and domestic prices are set in local currency.

The proposition supplies a necessary condition, not a quantitative relationship between an invoicing share and a country's full equilibrium response. It is the restriction taken to the annual panel: common-dollar effects should be concentrated where dollar invoicing is prevalent. The country regressions do not re-estimate the within-country reallocation mechanism or a symmetric global equilibrium.

D Data and estimation details

D.1 Data sources and sample

We assemble the firm-month panel from Chilean VAT filings for sales, materials, and investment; wage and employment records; annual tax filings for capital; and electronic invoice records used to construct firm-level output and input price and quantity indices. Giarda and Martner (2026) use the same administrative sources and provide further details on their coverage.

The administrative shock window is 2005 to 2023. This matches the main shock window in Giarda and Martner (2026). Electronic invoice price and quantity indices are available from 2014 onward. The markup split and aggregate productivity measures that require those indices therefore use 2014–2023; specifications using only nominal administrative records can use the full window. We do not impute pre-2014 quantity series.

The baseline observation is a firm-month, and each firm is treated as producing one composite good. The analysis therefore does not use buyer–product cells or transaction-level invoicing observations.

D.2 Firm-level markup estimation

We estimate markups using a quantity-based production-function approach. Giarda and Martner (2026) provide a parallel implementation using the same administrative sources. Electronic invoices identify product-level prices and quantities for goods sold and intermediate inputs purchased by formal firms. We first construct firm-level Tornqvist indices for output and materials from seller-product observations that appear in consecutive periods. These indices remove firm-specific output and input price movements before the production function is estimated.

For firm i in six-digit industry s and year t , the production function is Cobb–Douglas in labor, capital, and materials:

$$q_{it} = \vartheta_{st}^l l_{it} + \vartheta_{st}^k k_{it} + \vartheta_{st}^m m_{it} + \omega_{it} + \varepsilon_{it},$$

where lower-case variables are logs of physical quantities, ω_{it} is productivity observed by the firm when it chooses inputs, and ε_{it} is an unanticipated disturbance. We estimate the production function by industry-year using the control-function method of Akerberg et al. (2015). Materials are the flexible input. Let

$$s_{it}^m \equiv \frac{P_{it}^m M_{it}}{P_{it}^y Y_{it}}$$

denote the firm’s material-expenditure share in revenue. The materials first-order condition then gives the monthly markup estimate

$$\widehat{\mu}_{it} = \frac{\widehat{\vartheta}_{st}^m}{s_{it}^m},$$

as in De Loecker and Warzynski (2012). Because q_{it} and m_{it} are quantity indices, variation in a firm's own output price does not enter the dependent variable mechanically. The maintained Cobb–Douglas specification makes the estimated material-output elasticity common within an industry-year, while the monthly material share generates within-year markup variation.

The calibration uses the harmonic sales-weighted mean,

$$\bar{\mu} = \left(\mathbb{E}_\lambda \left[\widehat{\mu}_{it}^{-1} \right] \right)^{-1}.$$

For the reduced-form high–low comparison, we classify firms using their estimated markup twelve months before the current shock. At each date, firm i is assigned to the high-markup group if $\widehat{\mu}_{i,t-12}$ exceeds the median lagged markup among firms in the same industry. Thus the sorting variable varies over time but is predetermined relative to the contemporaneous exchange-rate movement.

D.3 Demand calibration details

Under the KW schedule in section 3.5, $\sigma(q) = \bar{\sigma}_K q^{-\epsilon_K/\bar{\sigma}_K}$. The Lerner relation therefore implies

$$\bar{\mu}^{-1} = 1 - \frac{\mathbb{E}_\lambda[q_\theta^{\epsilon_K/\bar{\sigma}_K}]}{\bar{\sigma}_K}.$$

The reference elasticity $\bar{\sigma}_K = 6.0$ fixes the local markup at $q = 1$; by itself it does not fix the harmonic markup of a heterogeneous firm distribution. After solving the steady-state quantity and sales-weight distribution, we choose the common real-cost normalization so that $\mathbb{E}_\lambda[q_\theta^{\epsilon_K/\bar{\sigma}_K}] = 1$. The induced harmonic sales-weighted markup is then 1.20, matching the rounded firm-panel estimate.

At the reference quantity, the KW desired own-cost response is

$$\rho(1) = \frac{\bar{\sigma}_K - 1}{\bar{\sigma}_K - 1 + \epsilon_K}.$$

Albagli et al. (2026) estimate price responses to own costs and competitors' prices of 0.24 and 0.38 using the same Chilean electronic-invoice records. We match their relative own-cost loading, $0.24/(0.24 + 0.38)$, rather than either dynamic coefficient in isolation. This gives $\epsilon_K = 7.92$.


D.4 Chilean local-projection implementation

Total variable cost is the wage bill plus material expenditure. At each date, firms are classified as high markup when their markup twelve months earlier is above the median lagged markup within their industry. The firm-level local projection absorbs firm-by-month-of-year and industry-by-calendar-month effects and includes the high-markup indicator and its interaction with the broad-dollar movement. Responses are cumulative relative to $t-1$ and are estimated through 36 months. Baseline standard errors are clustered by industry.

Aggregate value added is sales minus material expenditure. Employment and the nominal capital stock are aggregated across firms each month, and nominal value added is deflated using the GDP deflator. The capital-and-labor TFP residual is constructed using monthly labor and capital cost shares, normalized to sum to one. Both TFP and labor productivity are expressed as twelve-month moving averages. Their local projections use the same broad-dollar movement and include month-of-year effects.

Figure D.1 reports the robustness exercises referenced in the main text. Panel A uses two-way clustering by industry and calendar month for the firm-level reallocation regression; this leaves the point estimates unchanged but widens the confidence intervals. Panel B replaces capital-and-labor TFP with labor productivity and yields a similar aggregate response.

Figure D.1: Robustness of the Chilean local projections



Notes: Panel A reproduces the firm-level response in fig. 4.1 with standard errors two-way clustered by industry and calendar month. Panel B reports the aggregate local-projection response using labor productivity rather than capital-and-labor TFP. Shaded areas are 95% confidence intervals.