Misallocation, Dollars, and Aggregate Productivity*

Yasutaka Koike — UNC Chapel Hill

Antonio Martner — Central Bank of Chile

November 3, 2026

Abstract

We show how nominal exchange-rate movements affect aggregate productivity through resource reallocation, without changes in technology. In a model with sticky dollar-invoiced prices and heterogeneous markups, a depreciation raises import prices and shifts production and input use toward more elastic, low-markup firms, which are already inefficiently large relative to high-markup firms. This reallocation lowers aggregate productivity. We quantify the mechanism in a dynamic model calibrated to Chile and test its predictions using population-level administrative microdata on Chilean firms. We find evidence of reallocation toward low-markup firms and subsequent declines in measured aggregate productivity following broad dollar appreciations.

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 over the business cycle and can fall sharply during large exchange-rate depreciations (Mendoza, 1995; Aguiar and Gopinath, 2007; Gopinath and Neiman, 2014). Yet a nominal exchange-rate movement does not by itself change firms' technologies. In a frictionless benchmark, changes in external prices have no first-order effect on aggregate productivity (Kehoe and Ruhl, 2008). This tension raises a natural question: through what mechanism do nominal exchange-rate movements affect aggregate productivity?

This paper shows that exchange-rate movements can affect aggregate productivity by reallocating resources across firms with heterogeneous markups. High-markup firms are inefficiently small relative to low-markup firms. When prices are sticky in their currency of invoicing, nominal exchange-rate movements change relative prices and shift production across these firms, affecting aggregate productivity without changing technology. We formalize this mechanism in a dynamic small-open-economy model with dollar invoicing and apply it to Chile, where imports and exports are predominantly invoiced in dollars. Using administrative firm-level data, we quantify the mechanism and test its central predictions. Following broad dollar appreciations, we find evidence of resource reallocation from high-markup firms toward low-markup firms and a decline in aggregate productivity.

We consider a small open economy in which import and export prices are sticky in dollars, while domestic prices are sticky in home currency, as in Gopinath et al. (2020). This pricing assumption is supported by De Gregorio et al. (2024), who find that Chilean export prices respond predominantly to dollar exchange-rate movements in the short run. A depreciation of the home currency against the dollar therefore makes imports more expensive relative to domestic goods. Expenditure shifts from imported goods toward domestic firms; how much shifts toward each firm depends on its demand elasticity. We use a Kimball demand system in which these elasticities depend on firms' positions on the demand curve and determine their markups. Both low- and high-markup domestic firms expand, but low-markup firms gain a larger share of domestic production because they face relatively more elastic demand.

Low-markup firms are already inefficiently large relative to high-markup firms, so this reallocation worsens resource allocation and lowers aggregate productivity. The same reallocation toward low-markup firms occurs in export markets when some firms can adjust their dollar export prices. This first-order effect is absent under the commonly used CES demand specification, where markups are common, and in Gopinath et al.

(2020), who use Kimball demand but approximate around firms at the same point on the demand curve, with identical markups. We instead allow firms to occupy different points on the same demand curve, so their demand elasticities and initial markups differ.

We first derive a closed-form expression for the productivity response in a static model in which firms produce with labor alone. We then embed this mechanism in a dynamic small-open-economy model with intermediate inputs, calibrated using Chilean administrative firm-level data. Intermediate inputs amplify the productivity loss from reallocation through domestic production linkages. Following a depreciation of one percent on impact, aggregate productivity falls by 0.110 percent. Intermediate inputs also affect productivity through changes in firms' input mix and a measurement bias associated with imported inputs (see also Burstein and Cravino, 2015). These two effects partly offset each other, leaving the reallocation of production toward low-markup firms as the main source of the productivity decline.

We use monthly Chilean administrative firm-level data to examine whether dollar appreciations are associated with reallocation toward low-markup firms and lower aggregate productivity, as the model predicts. We compare firms' input expenditure responses across different markups and measure aggregate productivity from the output and inputs recorded in the same panel. The firm-level estimates indicate a reallocation toward low-markup firms. Aggregate productivity declines at longer horizons, and this finding also holds when productivity is measured as output per worker.

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 (2026) study how monetary policy shocks affect aggregate productivity through resource reallocation in closed economies with sticky prices and heterogeneous markups. We instead study how exchange-rate movements affect aggregate productivity by reallocating production across firms competing in domestic and foreign markets. 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). Closely related is Cravino (2017), who uses Chilean customs data to study how exchange-rate movements affect relative prices, markups, and aggregate productivity, emphasizing the role of exporters' invoicing currencies. Our mechanism instead operates even when cross-border prices are invoiced in a common currency, as heterogeneous markups generate reallocation across firms following exchange-rate movements. Related work on crisis episodes emphasizes

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). Our focus is instead on business-cycle reallocation around an already distorted allocation.

2 A static small open economy

With technology fixed, a nominal exchange-rate movement can change aggregate productivity by reallocating labor across firms that differ in their markups. To build intuition, we start with a static model that takes the exchange rate as given: with import prices sticky in dollars, a depreciation makes imports more expensive and shifts demand toward domestic firms, especially those with low markups and more elastic demand. Section 3 generalizes this static model with intermediate inputs and a dynamic general equilibrium that determines the exchange rate, and quantifies the response.

The home economy is small and takes foreign demand shifters, price indices, and marginal costs as given. Home firms sell in both domestic and export markets, while foreign firms supply the home market. Domestic sales are invoiced in home currency, whereas 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.

We take the exchange rate as given and derive how it affects aggregate productivity. We begin with fully fixed prices to show how changes in import prices reallocate production across domestic firms. We then allow some firms to reset their prices, changing how much they sell at home and abroad. Aggregate productivity then also depends on how nominal wages adjust, since wages affect the prices these firms choose. We consider two familiar cases: fixed nominal wages and monetary policy that stabilizes the CPI. In the dynamic model of section 3.3, the exchange rate and the nominal wage are jointly determined in equilibrium.

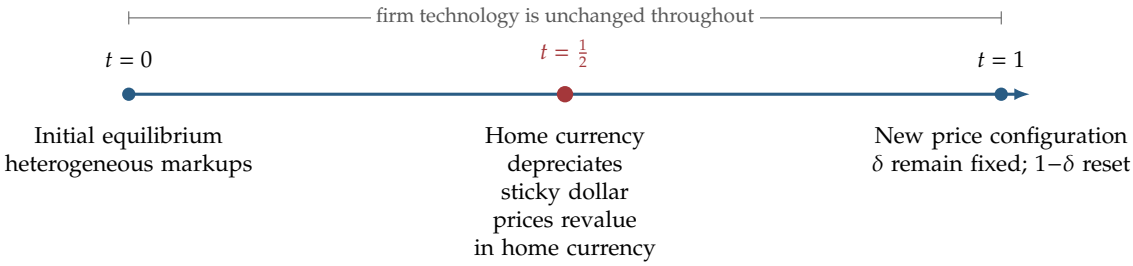
Notation. Uppercase letters denote levels. Let $Z_0 = \bar{Z}$ be the initial steady-state value and Z_1 the value in the post-movement allocation. Unless stated otherwise, the corresponding lowercase letter denotes the log change between the two allocations: $z \equiv \log Z_1 - \log Z_0 = \log Z_1 - \log \bar{Z}$. 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.[1] 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 firms' shares of production affect productivity at first order. The representative household has period utility

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

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

[1] For any firm-level variables z_a and v_a, define

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

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

imported varieties subject to

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

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 each market's demand shifter and 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 demand shifter and price aggregator as C^* and P^*. We sum a firm's sales across foreign markets and denote its total export quantity by $C^*_{h\theta}$.

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

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

Expenditure minimization gives the inverse demand schedules

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

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 expressed in dollars. 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.

We make the following assumptions about firms' domestic and export sales:

(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 $(1-\gamma) : \gamma$.[2]

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.

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

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

and

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

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

[2] Allowing 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.

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:

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

Given W/P, a productivity draw 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 $(q_\theta, \sigma_\theta, \mu_\theta, \rho_\theta)$. We study price and quantity changes around this initial steady state.

We also impose the following restrictions on the shape of demand:

(A3) Desired markups increase and desired cost pass-through decreases with normalized quantity:

$$\mu'_\theta > 0, \qquad \rho'_\theta < 0,$$

where both derivatives are with respect to q_θ.

The first inequality is Marshall's second law: larger firms face less elastic demand and charge higher markups. The second additionally makes larger firms pass through less of a cost change, absorbing more of it into their markups. Together, these restrictions give Marshall's third law in the formulation of Baqaee et al. (2024); see also Fujiwara and Matsuyama (2026). Using Belgian manufacturing data, Amiti et al. (2019) find that larger firms pass through less of their own cost changes, controlling for competitors' prices. In Chile, Albagli et al. (2026) likewise find that firms with higher market shares pass through less of their own cost changes and respond more strongly to competitors' prices.

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

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

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

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

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

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

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

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

The price aggregator uses elasticity-weighted sales shares and enters firms' demand and reset prices. The CPI uses expenditure shares. For each firm type, the average price combines newly reset prices with prices left unchanged.

A foreign exporter that does not reset has a fixed dollar price, so its home-currency price rises one for one with e. Substituting the inherited and reset prices into these aggregates and solving for p gives (see section A.1)

$$p = \kappa_w w + \gamma e, \qquad p^C = \kappa^C_w w + \gamma e,$$

where

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

The coefficient κ^C_w, positive for $\delta < 1$, measures the CPI response to a wage increase at a given exchange rate. The coefficient on the exchange rate equals the import share γ under the common price-adjustment probability and pass-through distribution assumed for home and foreign firms.[3] We next derive aggregate productivity, first with fully fixed

[3] If these differ between home and foreign firms, section A.1 gives the corresponding coefficient.

prices and then with partial price adjustment.

2.3 Exchange rates and aggregate productivity

With technology fixed, aggregate productivity changes when production is reallocated across firms with different markups. 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 changes their output proportionately, leaving production shares unchanged. Conditional on (w, e), aggregate demand therefore determines total output, 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 $\delta = 1$. A depreciation raises the home-currency price of imported varieties, while domestic prices and dollar export prices remain fixed. The increase in import prices raises the home price aggregator P, with log change $p = \gamma e > 0$ for $e > 0$.

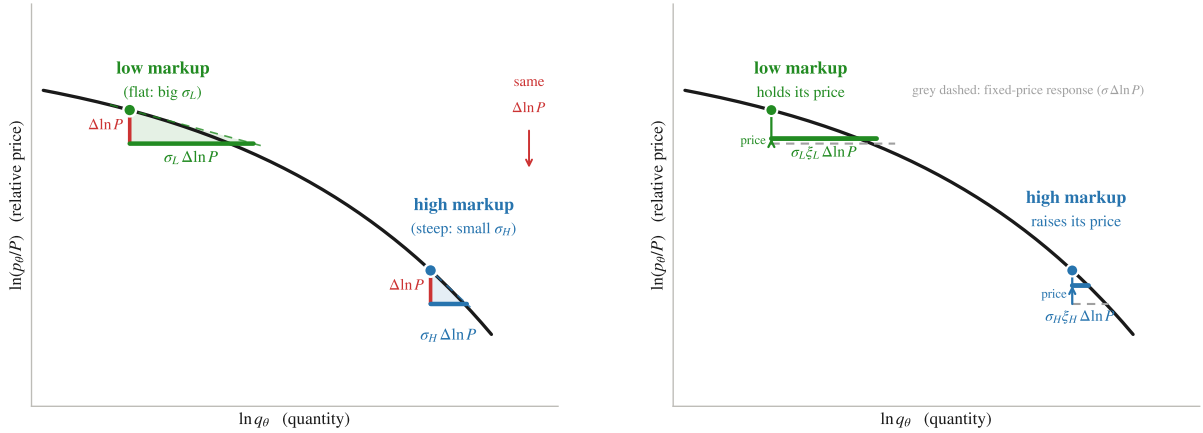
Proposition 1 (Depreciation with fully fixed prices)**.** *With fully fixed prices, the first-order productivity response is*

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

Under the second law in (A3) and nondegenerate markup heterogeneity, $C_\sigma > 0$. For $0 < \gamma < 1$, a depreciation therefore lowers aggregate productivity.

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 to the rise in P, denoted by $\Delta \ln P = p > 0$ in the figure. The covariance C_σ captures how the firms that expand more are also those with lower initial markups. Export quantities do not change in this case: dollar export prices, the foreign demand shifter C^*, and the foreign price aggregator remain fixed.

Figure 2.2: Reallocation with non-resetting and resetting prices



(a) Non-resetting firms

(b) Resetting firms

Notes: $\Delta \ln P = p > 0$ is the rise in the home price aggregator caused by higher home-currency import prices. Home-market Kimball demand is shown in logs, with local slope $-1/\sigma_\theta$. Panel (a) holds a firm's price fixed while the price aggregator rises, producing the quantity response $\sigma_\theta \Delta \ln P$. 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 $\sigma_\theta \rho_\theta \Delta \ln P$. 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

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

Their quantity sensitivity to the gap $p - w$ is therefore $\sigma_\theta \rho_\theta$, rather than σ_θ. Under (A3), 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

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

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 expresses aggregate productivity as a function of the exchange rate and the wage. 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

$$\Psi_e = \underbrace{\gamma(1-\gamma)\Big[\delta C_\sigma + (1-\delta)C_{\sigma\rho}\Big]}_{\textit{home market}} + \underbrace{\gamma(1-\delta)C_{\sigma\rho}}_{\textit{export market}},$$

$$\Psi_w = (1-\delta)C_{\sigma\rho}\Big[1-(1-\gamma)\kappa_w\Big] - (1-\gamma)\delta\kappa_w C_\sigma.$$

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

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

Under (A3), 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 (A3), the first effect dominates, so $\Psi_w > 0$. Wage increases therefore mitigate the direct allocative loss from depreciation, whereas wage declines amplify it.

Wage adjustment and monetary policy. Proposition 2 expresses aggregate productivity as a function of the exchange rate and the wage. The response to a depreciation therefore depends on the conditions determining wage adjustment. Two familiar cases illustrate this dependence.

First, with fully rigid nominal wages, $w = 0$, so only the direct exchange-rate effect remains:

$$\frac{a}{e} = -\Psi_e < 0.$$

Second, suppose monetary policy stabilizes the consumer price index. Since $p^C = \kappa_w^C w + \gamma e$, CPI stabilization requires $w/e = -\gamma/\kappa_w^C < 0$ for $\delta < 1$: wages must fall to offset the increase in import prices.

Corollary 1 (Depreciation under CPI stabilization)**.** *Under the conditions of Proposition 2, CPI stabilization implies*

$$\frac{a}{e} = -\Psi_e - \frac{\gamma}{\kappa_w^C}\Psi_w < 0.$$

Thus the productivity loss is not specific to fixed wages. Under CPI stabilization, the wage decline reinforces the direct effect of depreciation. Appendix A.7 considers other CPI responses.

These comparisons remain conditional on a depreciation; the dynamic model in section 3.3 determines the exchange rate and wage jointly.[4]

The following examples clarify the roles of markup heterogeneity and price stickiness.

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. Initial markups are therefore the same across firms, so $C_\sigma = C_{\sigma\rho} = 0$. 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 exchange-rate effects on aggregate trade prices and quantities, but it does not generate a first-order reallocation of production across firms with different initial markups. Our departure is to approximate around a heterogeneous firm distribution.

[4]Conditional on (w, e), household and market-clearing conditions determine consumption, employment, and output without changing the productivity response in Proposition 2. Appendix A.8 presents these conditions.

Example 3 (Flexible prices)**.** As $\delta \to 0$, 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 heterogeneous markups, productivity depends on the change in the dollar value of the wage, $w - e$. If the nominal wage and the exchange rate rise by the same proportion, relative prices and aggregate productivity remain unchanged. Price stickiness allows a nominal depreciation to change relative prices and aggregate productivity; the currency in which prices are fixed determines which exchange-rate movements matter.

Remark (Domestic and export 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 abstract from differences between domestic and export markups to isolate reallocation across firms.

3 Intermediate inputs and quantitative dynamics

Section 2 derived how movements in the exchange rate and the wage affect aggregate productivity, taking these movements as given. We now embed this mechanism in a quantitative dynamic model that determines the exchange rate and wage in equilibrium. We first introduce intermediate inputs, which amplify the across-firm allocation loss because gross sales exceed GDP and add an aggregate input-allocation term and an import-measurement bias. We then incorporate this production structure and price setting into a monthly small-open-economy New Keynesian model in the spirit of Gopinath et al. (2020). Household optimization, goods and asset-market clearing, and monetary policy determine the joint paths of the exchange rate, prices, and quantities.

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:

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

Here $X^H_{i\theta t}$ and $X^M_{i\theta t}$ 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 $X^H_{\theta t}$ and $X^M_{\theta t}$ 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:

$$x^H_t \equiv \mathbb{E}_\lambda[x^H_{\theta t}], \qquad x^M_t \equiv \mathbb{E}_\lambda[x^M_{\theta t}].$$

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

$$x_t = (1-\gamma)x^H_t + \gamma x^M_t.$$

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,

$$mc_t = (1-\alpha)w_t + \alpha p^C_t.$$

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. Intermediate inputs change how this reallocation affects value added and introduce the aggregate input-allocation term and import-measurement bias.

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

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

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 coefficient 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 does not determine how intermediate-input quantities or employment respond to a depreciation. The dynamic model below determines these responses and hence the sign and magnitude of each term.

Two limits clarify the decomposition. If $\alpha \to 0$, then $\xi_F \to 1$ and $\xi_X, \xi_M \to 0$, 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

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

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. International borrowing takes place through the dollar bond, as in Gopinath et al. (2020). We report the aggregate dollar position as net foreign debt. The dollar borrowing rate increases with this debt to ensure stationarity (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

$$\max_{\widetilde{P}_{h\theta t}} \mathbb{E}_t \sum_{k=0}^{\infty} \delta^k \mathcal{Q}_{t,t+k} \left(\widetilde{P}_{h\theta t} - MC_{\theta,t+k} \right) C_{h\theta,t+k},$$

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

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

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 aggregator are fixed, as in the small-open-economy model of Gopinath et al. (2020). Export quantities nevertheless respond to their dollar relative price through Kimball demand. 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 price aggregator and the CPI retain their definitions from the one-period economy:

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

Define

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

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 resetters across the home and export markets. Aggregating the home-market reset conditions and inherited price vintages gives

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

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

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,*

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

Lagged and expected future reallocation enter because prices remain fixed for more than one period. 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

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

Goods, input, and labor-market clearing then give

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

The coefficients and the cost-based labor exposure are

$$\Theta_c = \frac{(1-\gamma)(\bar{\mu}-\alpha)}{\bar{\mu}-\widetilde{\alpha}}, \qquad \Theta_x = \frac{\widetilde{\alpha}}{\bar{\mu}-\widetilde{\alpha}}, \qquad \Theta_s = \frac{\bar{\mu}}{\bar{\mu}-\widetilde{\alpha}},$$

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

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 the exogenous component of its borrowing premium. Higher external debt raises the dollar borrowing rate, while UIP relates that rate to the domestic interest rate and expected depreciation. 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

Chile's imports and exports are predominantly invoiced in dollars (Boz et al., 2022). We approximate this pattern by assuming dollar invoicing for all imports and exports.

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 estimate markups 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 C.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,

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

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

[6] This 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.

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 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 $q_\theta \Upsilon'(q_\theta)$ 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 C.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. The ratio of the own-cost coefficient to the sum of the own-cost and competitor-price coefficients, $0.24/(0.24 + 0.38)$, implies $\epsilon_K = 7.92$. We match this ratio rather than equating either estimated dynamic response with static desired pass-through.

Firm productivity follows a Pareto upper tail,

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

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

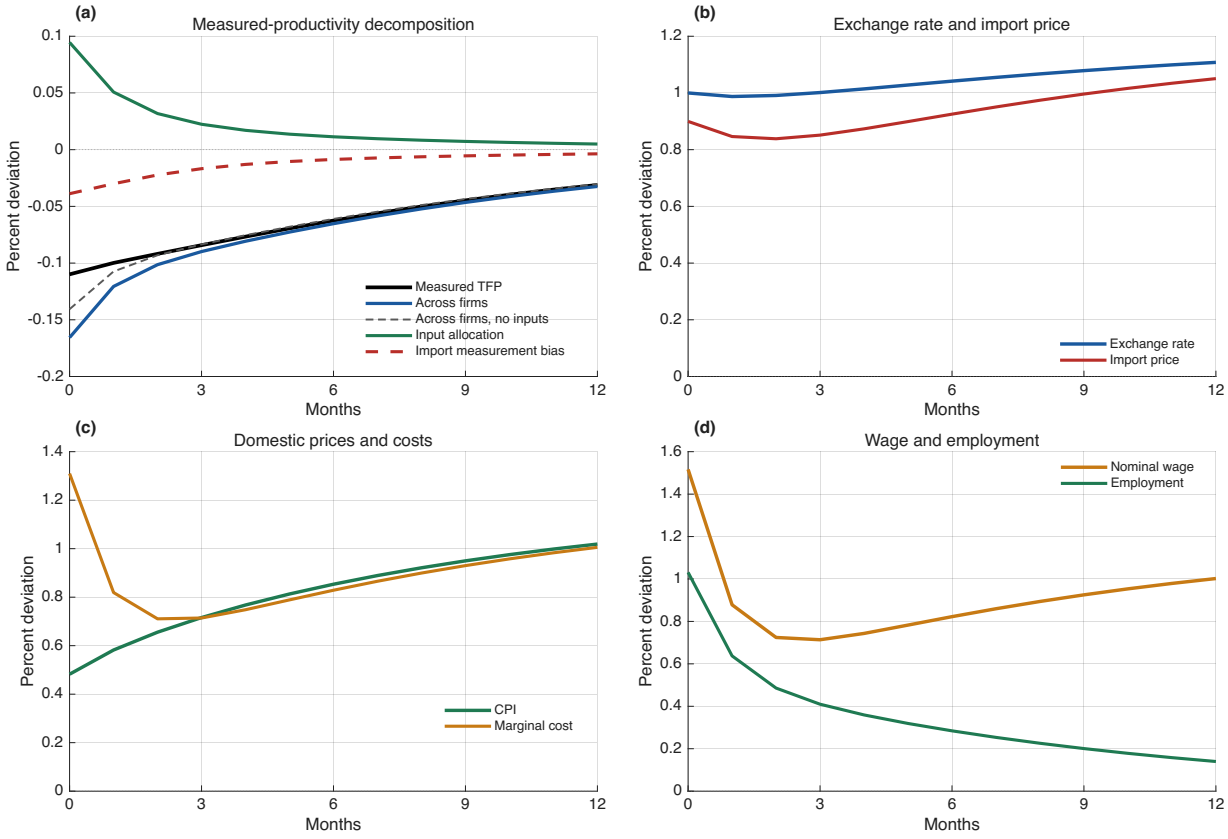
Parameter or moment	Value	Calibration
Panel A. Kimball demand and firm distribution		
Elasticity level, $\bar{\sigma}_K$	6.00	Reference-point normalization
Superelasticity, ϵ_K	7.92	Chilean invoice data
Pareto tail, κ_A	3.93	Chilean firm distribution
Harmonic sales-weighted markup, $\bar{\mu}$	1.20	Production-function estimates
Panel B. Aggregate and dynamic parameters		
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 across-firm productivity response in an otherwise identical economy without intermediate inputs, showing how intermediate inputs amplify the reallocation loss. 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 aggregate input-allocation term and import-measurement bias therefore partly offset one another, with a net contribution of 0.056 percent. Across-firm reallocation is about three times that

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.

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 to -0.166 percent, an amplification of about 18 percent. The total productivity decline is smaller because the two additional terms 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 input-allocation term and import-measurement bias alter the total response without reversing its sign.

4 Reduced-form evidence

The empirical analysis examines two implications of the model using Chilean administrative data. First, we study whether dollar appreciations reallocate activity across firms along the markup dimension predicted by the model. Second, we examine the accompanying response of aggregate measured productivity. Because Chilean trade is predominantly invoiced in dollars, these exercises hold the invoicing environment fixed and exploit time-series variation in the value of the U.S. dollar. Details on variable construction, local-projection implementation, and inference are reported in Appendix C.

For the Chilean exercises, exchange-rate variation comes from the Federal Reserve broad nominal dollar index. We use the broad dollar rather than the bilateral peso–dollar exchange rate because the latter responds directly to Chile-specific monetary, demand, and financial conditions, whereas broad dollar movements are unlikely to be driven by developments in Chile. Let $s_t = \Delta \log E_t^{\$}$, where $E_t^{\$}$ denotes the broad dollar index, so that $s_t > 0$ corresponds to an appreciation of the U.S. dollar. Because monthly changes in the broad dollar are serially correlated, our baseline shock z_t is the residual from a regression of s_t on twelve of its own lags and calendar-month fixed effects. This isolates the component of the current dollar movement that is unexpected given its recent history and provides a cleaner timing experiment for comparison with the model. We refer to z_t as the broad-dollar innovation.

4.1 Firm-level reallocation

The theory predicts that the reallocation induced by a depreciation is ordered by firms' markups: low-markup firms expand relatively more than high-markup firms. For the em-

pirical analysis, we exploit the representation of the across-firm component a_F in proposition 3. This representation remains valid in the presence of intermediate inputs and under weaker assumptions than the closed-form firm-level response characterized in proposition 2, making it the natural object to take to the data.

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,*

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

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

Under the assumptions of corollary 2, the common marginal-cost term drops out of the covariance, allowing us to use nominal total variable costs to examine the reallocation across firms. We measure total variable cost as the sum of the wage bill and material expenditures. This approach provides an empirical counterpart to the across-firm component a_F in proposition 3, but does not separately identify the aggregate input-allocation or import-measurement terms. The derivation is provided in section A.10.

We measure firm-level markups using the labor first-order condition of the gross-output production function, assuming that labor is a flexible input. Under this assumption, $\mu^L_{it} = \sigma^L_{s(i)}/(WB_{it}/R_{it})$, where $\sigma^L_{s(i)}$ is the industry-specific output elasticity of labor, WB_{it} is the firm's wage bill, and R_{it} is nominal sales.

We implement the model's prediction using a high–low markup classification rather than estimating the covariance directly. To ensure that the classification is predetermined relative to the current exchange-rate shock, we use each firm's markup observed twelve months earlier. Specifically, we define

$$D^H_{it} = \mathbf{1}\left\{\mu^L_{i,t-12} > \text{median}_{j\in s(i),t}\left(\mu^L_{j,t-12}\right)\right\},$$

where the median is computed within each industry among firms observed in month t. Thus, D^H_{it} identifies firms whose lagged markups exceed the industry-month median.

To express the response in the same units as the model, we estimate local projections using the bilateral peso–dollar exchange rate, instrumented by the broad-dollar innovation. Let $\Delta_h x_{i,t+h} \equiv x_{i,t+h} - x_{i,t-1}$ denote the cumulative change through horizon h. We

estimate

$$\Delta_h \log TVC_{i,t+h} = \alpha^h_{i,m} + \tau^h_{s(i),t} + \delta_h D^H_{it} + \beta^{IV}_h \left(D^H_{it} \times \Delta \log E^{CLP/USD}_t \right) + u_{i,t+h},$$

where $E^{CLP/USD}_t$ is the peso price of one U.S. dollar, and we instrument $D^H_{it} \times \Delta \log E^{CLP/USD}_t$ with $D^H_{it} \times z_t$. The fixed effects $\alpha^h_{i,m}$ and $\tau^h_{s(i),t}$ denote firm-by-month-of-year and industry-by-calendar-month effects, respectively. The coefficient β^{IV}_h measures the differential cumulative response of high- relative to low-markup firms to a peso depreciation induced by an unexpected broad-dollar appreciation.

The broad-dollar innovation predicts movements in the bilateral exchange rate. In a monthly regression, a 1 percent unexpected broad-dollar appreciation is associated with a 1.35 percent depreciation of the Chilean peso.[7] Industry-by-calendar-month fixed effects absorb movements common to firms within each industry and month, so identification comes from the differential response of high- and low-markup firms within these cells. Baseline standard errors are clustered by industry, and Appendix C reports robustness to two-way clustering by industry and calendar month. Because broad-dollar innovations may also affect firms through channels other than the bilateral exchange rate, we interpret the IV estimates as differential responses scaled by the induced peso depreciation, rather than as isolating the effects of the bilateral exchange rate alone.

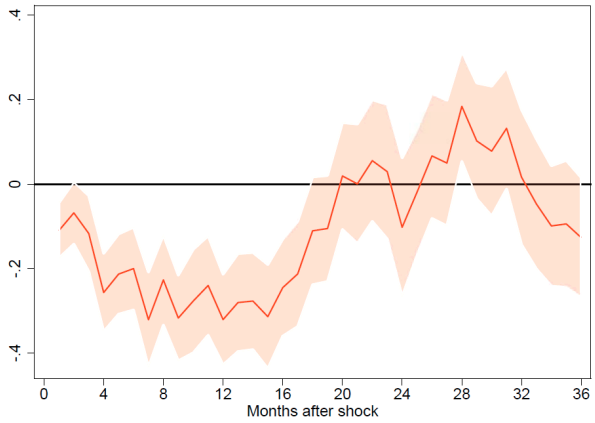
Panel A of fig. 4.1 supports the model's reallocation mechanism. Following a broad-dollar-induced peso depreciation, total variable cost grows less at high-markup firms than at low-markup firms. The differential response is negative at short horizons, increases in magnitude over the first year, and subsequently dissipates. At twelve months, a 1 percent peso depreciation is associated with a 0.32 percent lower cumulative response of total variable cost at high-markup firms relative to low-markup firms. This pattern is consistent with the model's prediction that a depreciation shifts production toward low-markup firms, worsening the allocation of resources across firms. Although the empirical response builds more gradually than in the calibrated model, the direction and eventual dissipation of the reallocation are consistent with its predictions.

Appendix C reports two robustness exercises. First, we classify firms by their average markup over the sample, net of sector-level seasonal and annual variation, rather than by a single lagged value, and obtain a similar pattern of reallocation. Second, we allow for two-way clustering by industry and calendar month, which leaves the point estimates unchanged but produces wider confidence intervals. The main pattern of reallocation is preserved in both specifications.

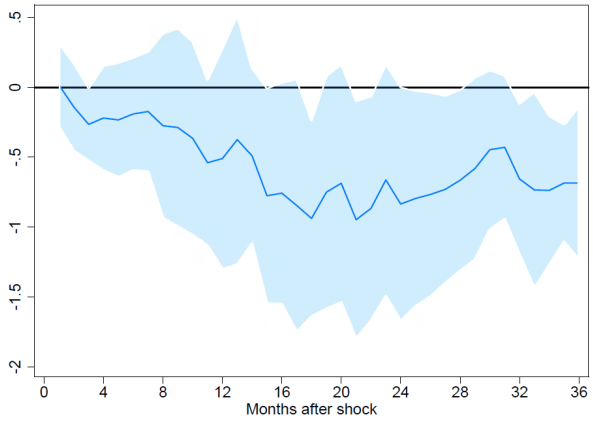
[7] Newey–West standard error of 0.14, based on 203 monthly observations.

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

A. Reallocation across markup groups



B. Aggregate measured TFP



Notes: Both panels report instrumental-variable local-projection responses to a one-percent peso depreciation induced by an unexpected broad-dollar appreciation. Shaded areas denote 95% confidence intervals. Panel A reports the cumulative high-minus-low-markup response of total variable cost. Markup groups are defined using labor-based markups lagged twelve months. The interaction between the bilateral exchange-rate change and the high-markup indicator is instrumented by the broad-dollar innovation interacted with the same indicator. Regressions include firm-by-month-of-year and industry-by-calendar-month fixed effects, with standard errors clustered by industry. Panel B reports the cumulative response of aggregate capital-and-labor TFP using the 2007–2019 sample. The bilateral exchange-rate change is instrumented by the broad-dollar innovation, and standard errors are heteroskedasticity- and autocorrelation-consistent.

4.2 Aggregate productivity

We next examine the aggregate productivity response to dollar movements. Using the Chilean administrative data, we construct monthly real value added, aggregate labor, and capital. Our baseline productivity measure is a capital-and-labor TFP residual, calculated as

$$\Delta a_t^{KL} = \Delta \log Y_t^{VA} - s_{t-1}^{L} \Delta \log L_t - s_{t-1}^{K} \Delta \log K_t,$$

where Y_t^{VA} is real value added, L_t and K_t are aggregate labor and capital, and s_{t-1}^{L} and s_{t-1}^{K} are their lagged factor shares. We cumulate these monthly changes to obtain an index of aggregate log TFP. The baseline analysis uses the 2007–2019 period, excluding the pandemic and subsequent period of exceptional economic disruptions.

For comparison, we also construct labor productivity ($a_t^{L} = \log Y_t^{VA} - \log L_t$), which does not depend on the capital-stock measure. For either productivity measure, we estimate the cumulative response using the same instrumental-variables approach as in the firm-level exercise:

$$a_{t+h}^{j} - a_{t-1}^{j} = \beta_h^{j} \Delta \log E_t^{CLP/USD} + \gamma_m^{h} + v_{t+h}^{j}, \qquad j \in \{KL, L\},$$

where the bilateral peso–dollar depreciation is instrumented by the broad-dollar innovation z_t, and γ_m^h denotes month-of-year effects. We estimate the responses over horizons of up to 36 months, using heteroskedasticity- and autocorrelation-consistent standard errors with horizon-dependent bandwidths.

Panel B of fig. 4.1 reports the response of aggregate measured TFP to a broad-dollar-induced peso depreciation. TFP declines at short horizons, with the negative response becoming more pronounced after approximately one year and remaining persistent at longer horizons. The response is consistent with the model's prediction that exchange-rate-induced reallocation can reduce aggregate productivity without changes in technology. The corresponding labor-productivity estimates display a similar pattern and are reported in Appendix C. Neither aggregate measure isolates the across-firm reallocation component, but together with the firm-level evidence, these results support the aggregate implications of the mechanism.

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. In the model, a depreciation shifts production 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. 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

Ackerberg, 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, Elías, Francesco Grigoli, Emiliano Luttini, Dagoberto Quevedo, and Marco Rojas, "Beyond costs: The dominant role of strategic complementarities in pricing," *Journal of Monetary Economics*, 2026, p. 103966.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings, "International Shocks, Variable Markups, and Domestic Prices," *Review of Economic Studies*, 2019, *86* (6), 2356–2402.

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.

Baqaee, 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," *Journal*

of Monetary Economics, 2026, *159*, 103920.

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.

Gregorio, José De, Pablo García, Emiliano Luttini, and Marco Rojas, "From Dominant to Producer Currency Pricing: Dynamics of Chilean Exports," *Journal of International Economics*, 2024, *149*, 103934.

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

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

Collecting terms in p yields

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

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

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

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

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

Using the fixed point of section A.1,

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

The exchange-rate coefficient is

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

The wage coefficient is

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

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*

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

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,

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

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

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

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:

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

At the initial allocation, revenue from each market equals $\mu_{\theta}\bar{W}$ times the labor used to produce for that market. Within each firm, domestic and export sales shares therefore equal the corresponding 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 responses of labor used for domestic and export sales are

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

Substitution into lemma 1 gives

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

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

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

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

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

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 $\mathbb{E}_\lambda[\sigma_\theta\rho_\theta] = \bar{\sigma}_\lambda\bar{\rho}^\sigma$,

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

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

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

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

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

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

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

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

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

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 show that $\Psi_w > 0$, 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

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

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

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

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

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

A positive value of the expectation indicates that relative-price changes reduce aggregate demand for domestic output 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

$$\int_0^1 \left[(1-\gamma)\Upsilon\left(\frac{X^H_{i\theta t}}{(1-\gamma)X_{it}}\right) + \gamma\Upsilon\left(\frac{X^M_{i\theta t}}{\gamma X_{it}}\right)\right] d\theta = 1. \tag{A.1}$$

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

$$D_X \equiv \int_0^1 \Upsilon'(q_\theta) q_\theta \, d\theta, \qquad \lambda^X_\theta \equiv \frac{\Upsilon'(q_\theta) q_\theta}{D_X}. \tag{A.2}$$

Log differentiating (A.1) around that allocation gives

$$\begin{aligned} 0 = (1-\gamma) &\int_0^1 \Upsilon'(q_\theta) q_\theta (x^H_{i\theta t} - x_{it}) \, d\theta \\ &+ \gamma \int_0^1 \Upsilon'(q_\theta) q_\theta (x^M_{i\theta t} - x_{it}) \, d\theta. \end{aligned} \tag{A.3}$$

Dividing by D_X therefore yields

$$x_{it} = (1-\gamma)\mathbb{E}_{\lambda^X}[x^H_{i\theta t}] + \gamma\mathbb{E}_{\lambda^X}[x^M_{i\theta t}]. \tag{A.4}$$

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

$$X_t \equiv \int X_{it}\, di, \qquad X^s_{\theta t} \equiv \int X^s_{i\theta t}\, di,$$

we have $X^s_{\theta t} = d^s_{\theta t} X_t$. 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

$$x^H_t \equiv \mathbb{E}_{\lambda^X}[x^H_{\theta t}], \qquad x^M_t \equiv \mathbb{E}_{\lambda^X}[x^M_{\theta t}]$$

therefore gives

$$x_t = (1-\gamma)x^H_t + \gamma x^M_t. \tag{A.5}$$

The same weights arise from expenditure minimization. If P^X_t denotes the price aggregator associated with (A.1), its first-order conditions are $P^s_{\theta t}/P^X_t = \Upsilon'(q^s_{\theta t})$. Thus λ^X_θ 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

$$\begin{aligned} p^C_t &= (1-\gamma)p^H_t + \gamma p^M_t, \\ p^s_t &\equiv \mathbb{E}_{\lambda^X}[p^s_{\theta t}], \qquad s \in \{H, M\}. \end{aligned} \tag{A.6}$$

Under the common-basket normalization of the main text, $\lambda^X_\theta = \lambda_\theta$, 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

$$\lambda_{\theta m} \equiv \frac{R_{\theta m}}{S}, \qquad \omega_{\theta m} \equiv \frac{V_{\theta m}}{V} = \lambda_{\theta m}\frac{\bar{\mu}}{\mu_\theta}, \qquad V = \frac{S}{\bar{\mu}}. \tag{A.7}$$

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

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

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,

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

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,

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

The steady-state ratios of gross sales and intermediate expenditure to GDP are therefore ξ_F and $\xi_F - 1$, respectively. Differentiating this level identity and using (A.8)–(A.9) gives

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

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

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

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

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

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 $1+\widetilde{\alpha}+\widetilde{\alpha}^2+\cdots$, 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

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

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

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

where

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

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

Finally, section A.9 and cost minimization give

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

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

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

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

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

The three targets used in section 3.4 are

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

With its dollar marginal cost fixed, a foreign supplier's reset price responds to the home price aggregator expressed in dollars, $p_t - e_t$. A domestic exporter's reset price responds to its dollar marginal cost, $mc_t - e_t$, while foreign expenditure and the foreign dollar price aggregator remain 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,

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

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 aggregate demand across the domestic and export markets, define each firm's relative price averaged across these markets 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

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

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

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

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

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

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

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

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 Data and estimation details

C.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 indices. Giarda and Martner (2026) use the same administrative sources and provide further details on their coverage.

The administrative records cover 2005–2023. Electronic invoice data used to construct firm-specific price indices are available from 2014 onward and are used to estimate industry-level production elasticities. These elasticities are held fixed over time and combined with monthly labor expenditure shares to construct firm-level markups. The baseline aggregate productivity analysis uses 2007–2019, while specifications including the pandemic and subsequent years are reported as robustness exercises. The baseline observation is a firm-month, and each firm is treated as producing one composite good.

C.2 Firm-level markup estimation

We estimate markups using a quantity-based production-function approach. Electronic invoices identify product-level prices and quantities for goods sold and intermediate

inputs purchased by formal firms. We construct firm-level Törnqvist price indices for output and materials using products observed in consecutive periods. We then use these indices to deflate nominal sales and material expenditures, respectively, obtaining real output and material-input indices that account for firm-specific price movements.

For firm i in industry s and period t, we assume a Cobb–Douglas gross-output production function in labor, capital, and materials:

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

where lower-case variables denote log real input and output indices, ω_{it} is productivity observed by the firm when choosing inputs, and ε_{it} is an unanticipated disturbance. We estimate the production function using the control-function approach of Ackerberg et al. (2015), allowing production elasticities to differ across industries while holding them fixed over time.

For the firm-level reallocation exercise, we recover markups from the labor first-order condition, following De Loecker and Warzynski (2012). Assuming that labor is a flexible input, let

$$s_{it}^l \equiv \frac{WB_{it}}{R_{it}}$$

denote the labor expenditure share in revenue. The labor-based markup is

$$\widehat{\mu}_{it}^L = \frac{\widehat{\vartheta}_s^l}{s_{it}^l}.$$

Because the production elasticity is constant within each industry, within-industry variation in estimated markups reflects differences in firms' labor expenditure shares. We use this measure in the reallocation exercise because materials constitute a large share of total variable cost. Constructing the sorting variable from the materials expenditure share would directly link it to a major component of the outcome. The labor-based measure avoids this direct overlap, although it still requires the assumption that labor satisfies the variable-input first-order condition.

C.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 $\epsilon_K = 7.92$.

C.4 Chilean local-projection implementation

The firm-level regressions use monthly total variable cost, defined as the sum of the wage bill and material expenditures. Firms are classified using labor-based markups from the gross-output production function, as described in section C.2. The baseline classification compares each firm's markup twelve months before the shock with the median lagged markup among firms in the same industry and month. The local projections estimate cumulative responses relative to $t - 1$ through 36 months, absorbing firm-by-month-of-year and industry-by-calendar-month fixed effects. Baseline standard errors are clustered by industry.

The baseline firm-level specification instruments the interaction between the high-markup indicator and the monthly change in the bilateral peso–dollar exchange rate with the interaction between the same indicator and the broad-dollar innovation. To assess the relevance of the instrument, we estimate the monthly regression

$$\Delta \log E_t^{CLP/USD} = \pi z_t + \gamma_m + \nu_t,$$

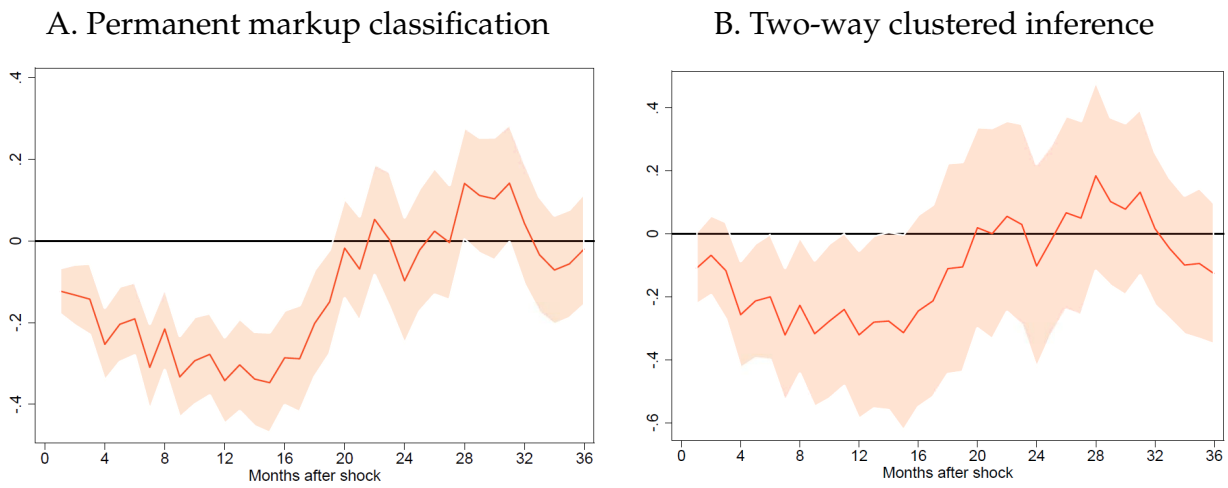
where γ_m denotes month-of-year effects. The estimated coefficient is $\widehat{\pi} = 1.35$, with a Newey–West standard error of 0.14 using twelve lags and 203 monthly observations. The monthly regression provides an independent time-series assessment of the relationship between the broad-dollar innovation and the bilateral exchange rate.

As a robustness exercise, we construct a permanent markup classification to capture persistent differences across firms. Specifically, we regress the logarithm of monthly

labor-based markups on firm fixed effects, sector-by-month-of-year fixed effects, and sector-by-year fixed effects. The estimated firm fixed effects capture persistent markup differences after removing sector-specific seasonal and annual variation. We classify firms as high markup when their estimated firm fixed effect exceeds a fixed median calculated across firms within the same industry. We then reestimate the baseline LP-IV specification using this classification, keeping the outcome, instrument, fixed effects, and clustering unchanged.

Figure C.1 reports two robustness exercises. Panel A uses the permanent markup classification and yields a reallocation response similar to the baseline, indicating that the result is not driven by transitory variation in measured markups. Panel B retains the baseline classification but allows for two-way clustering by industry and calendar month. This leaves the point estimates unchanged but widens the confidence intervals, while preserving the main pattern of reallocation.

Figure C.1: Firm-level reallocation: markup classification and inference



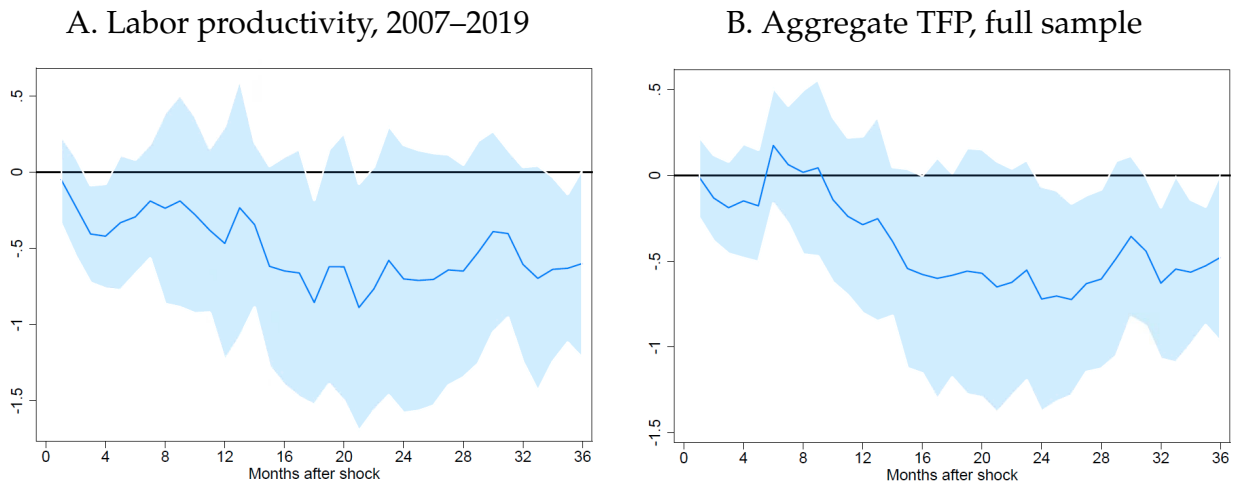
Notes: Both panels report instrumental-variable local-projection estimates of the cumulative high-minus-low-markup response of total variable cost to a one-percent broad-dollar-induced peso depreciation. Panel A replaces the baseline twelve-month-lagged classification with a permanent markup classification and clusters standard errors by industry. Panel B retains the baseline classification and clusters standard errors by industry and calendar month. Both specifications include firm-by-month-of-year and industry-by-calendar-month fixed effects. Shaded areas denote 95% confidence intervals.

For aggregate productivity, nominal value added is constructed from sales net of material expenditures and deflated using the GDP deflator. Aggregate labor and capital are obtained by summing the corresponding firm-level measures. The capital-and-labor TFP index cumulates monthly Solow residuals calculated using lagged labor and capital cost shares, normalized to sum to one. Labor productivity is measured as real value added per worker. Both aggregate measures are constructed at the monthly frequency.

The baseline aggregate analysis covers 2007–2019, preceding the exceptional disruptions associated with the COVID-19 pandemic and subsequent inflationary period. Aggregate local projections use the bilateral peso–dollar depreciation instrumented by the broad-dollar innovation, include month-of-year effects, and employ heteroskedasticity- and autocorrelation-consistent standard errors with horizon-dependent bandwidths. The baseline uses unsmoothed productivity series.

Figure C.2 reports two complementary exercises. Panel A replaces capital-and-labor TFP with labor productivity over the baseline sample, obtaining a similar negative response. Panel B extends the TFP analysis to the full available sample, including the pandemic and subsequent years. The medium-run negative response remains visible, although the short-run estimates are considerably more irregular.

Figure C.2: Aggregate productivity: alternative measures and sample



Notes: Instrumental-variable local-projection responses to a one-percent broad-dollar-induced peso depreciation. Panel A reports the cumulative response of aggregate labor productivity for 2007–2019. Panel B reports the cumulative response of aggregate capital-and-labor TFP using the full available sample. Both regressions include month-of-year effects. Shaded areas denote 95% confidence intervals based on heteroskedasticity- and autocorrelation-consistent standard errors.