BRICSCOIN:

Initial Distribution and Dynamic Injection Rules for a Non-Pegged, No-interest, Unanimous, Supranational Settlement Currency

Thiery Kame Babilla
University of Yaounde II and EABCN

Veljko Milisavljevic
University of Belgrade

Abstract

We build a five-country open-economy Heterogeneous-Agent New Keynesian (HANK) Dynamic Stochastic General Equilibrium model to study BRICSCOIN, a proposed supranational settlement instrument for Brazil, Russia, India, China, and South Africa. BRICSCOIN is minted only by unanimous agreement of a five-member Issuing Council, verified through a tamper-evident distributed ledger, and carries no peg, no reserve backing, and no convertibility promise: it functions purely as a bilateral settlement asset held on central bank balance sheets. We embed this institutional design into the external block of an otherwise standard multi-country incomplete-markets model and use it to answer two design questions the institution itself leaves open. First, how should the initial stock of coins be distributed across members? We show that unanimity turns the allocation problem into a cooperative bargaining problem over three legitimacy-relevant characteristics — population, GDP, and trade turnover — and derive a Nash-bargaining benchmark in which equal weights on the three shares are the unique allocation robust to any single member's veto. Second, how should newly minted coins be injected over time? We derive a rule that anchors injections to each member's share of incremental trade turnover, augmented by a deficit-responsive relief term that is itself constrained by the surplus countries' participation constraints under unanimity. Quantitative findings further confirm that: BRICSCOIN's settlement layer, exactly as intended by its no-peg, no-interest design, does not distort domestic monetary transmission. BRICSCOIN transmission is fundamentally about relative, not absolute, shocks: a pure minting-growth shock has a transitory effect on the distribution of holdings across countries.

Keywords: BRICSCOIN, International Monetary System, Supranational Settlement Currency, Bancor, BRICS, HANK-DSGE Model.

JEL Codes: C51, E42, E58, E71, F33, F36, F55, G12, G23.

1 Introduction

The search for a supranational settlement instrument that does not depend on a single reserve-issuing country is as old as the Bretton Woods negotiations, in which Keynes's Bancor proposal lost to the dollar-gold compromise that ultimately became the postwar system [Keynes, 1943, Skidelsky, 2003]. The debate resurfaced with the IMF's Special Drawing Right (SDR) in 1969, and again after the 2008 financial crisis, when several BRICS officials publicly floated the idea of a shared settlement currency to reduce dependence on the US dollar for intra-bloc trade [Zhou, 2009]. What distinguishes the instrument studied in this paper, which we refer to throughout by its working name, BRICSCOIN, from these precedents is precisely what it does *not* do: (i) *it is not pegged to gold, to a currency basket, or to any national currency*; (ii) *it carries no reserve backing and no convertibility promise*; and (i) *it is minted only by the unanimous consent of a five-member Issuing Council, with authenticity (not custody) recorded on a distributed ledger held identically by every member.*

Indeed, the recurring debates over "de-dollarization" inside BRICS have revived, almost verbatim, the international-monetary-architecture debate of 1943–44. At Bretton Woods, John Maynard Keynes proposed an *International Clearing Union* (ICU) built around a supranational accounting unit he called *bancor*, defined with reference to gold but never convertible into it, and issued exclusively to national central banks through overdraft and credit facilities at the Union [Keynes, 1943, Skidelsky, 2003]. Keynes's scheme had two structural features that distinguish it from the gold-exchange standard that in fact emerged from Bretton Woods: (i) bancor was a pure unit of settlement, not a national currency competing with the dollar or sterling, and (ii) adjustment pressure fell symmetrically on debtor *and* creditor nations, via interest charges on both overdrafts and above-quota credit balances. It was this second feature, more than any other, that the United States — as the world's largest creditor in 1944 — step-aside in favor of the White Plan and the gold-dollar standard [Eichengreen, 2011, Skidelsky, 2003].

BRICSCOIN, as considered here, revives the bancor logic but replaces it with its own singular and topical characteristics. The institutional design of BRISCOIN is unusual enough that it does not fit cleanly into any existing quantitative monetary model. (i) It *is not a currency union*, because no member's domestic monetary policy is constrained by it. (ii) It *is not a reserve currency in the SDR sense*, because it has no administratively determined valuation. (iii) And it *is not a cryptocurrency in the conventional sense*, because its supply is not governed by an algorithm but by discretionary unanimous agreement, and its physical instantiation — a certified diamond whose internal CT-scan fingerprint is checked against a distributed registry — exists specifically to solve a counterfeiting problem, not a double-spending problem, since transfers themselves are never recorded on the ledger at all.

We take this institutional description as given (Section 2) and ask what an open-economy macro model implies for two decisions that the design leaves entirely to the Council's discretion: (i) the *initial* distribution of coins across the five founding members, and (ii) the *rule* governing how newly minted coins should be allocated across members in every subsequent round. Because neither decision is disciplined by a peg or a formula in the institutional design itself, both are properly economic questions, and both are shaped by the same friction that makes the whole system work: unanimity. A member that would be made worse off by a proposed distribution can simply refuse to sign, and the round fails. Any distribution or injection rule we propose must therefore be not only welfare-improving in an aggregate sense, but individually acceptable to every member under its own beliefs about future trade and consumption needs.

To make this precise we build a five-country open-economy HANK-DSGE model (Section 3) in which each BRICS economy is populated by a continuum of ex-ante identical, ex-post heterogeneous households facing idiosyncratic labor-income risk and incomplete domestic asset markets, in the tradition of Aiyagari [1994] and Bewley [1986], embedded in a New Keynesian open-economy framework with monopolistic competition and nominal rigidities [Galí, 2005, Kaplan et al., 2018, Auclert et al., 2021]. BRICSCOIN enters not as a domestic asset that households hold directly, but as a settlement instrument on national central banks' external balance sheets, entering the model through the balance-of-payments block exactly as the institutional description in Section 2 implies. We then pose the initial-distribution problem as a cooperative bargaining game over verifiable country characteristics (Section 4) and the injection problem as a repeated game with individual-rationality constraints under unanimity (Section 5), before calibrating and illustrating the resulting rules numerically (Section 6).

We conduct a full Bayesian estimation of the parameters of the open-economy HANK–DSGE model of BRICSCOIN developed in our companion papers. We specify a reduced-form linear state-space representation of each country's domestic block consistent with the model's linearized sequence-space structure, place priors over every coefficient, and estimate the posterior distribution by three complementary methods appropriate to three different parts of the model: conjugate Normal-Inverse-Gamma Bayesian regression for the domestic IS, Phillips-curve, and Taylor-rule coefficients of each country; a hierarchical Gibbs sampler that partially pools the five countries' monetary-policy reaction functions around a common bloc-wide hyperparameter; and Approximate Bayesian Computation (ABC) for the settlement block's injection relief weight δ, which has no closed-form likelihood. Because BRICSCOIN has no observable minting or settlement history, we validate the full estimation pipeline on data simulated from the model at known parameter values, explicitly reporting recoverability, and then use the resulting posterior — rather than a single calibrated point — to run three posterior-predictive scenario analyses (baseline, a China-slowdown/India-acceleration scenario, and an India-led bloc trade boom), propagating parameter uncertainty into fan charts for the evolution of BRICSCOIN holdings shares.

Estimations of the DSGE Model of BRISCOIN answer thus to those two research questions in two fundamentally ways. (i) First, we demonstrate that unanimity turns the allocation problem into a cooperative bargaining problem over three legitimacy-relevant characteristics — population, GDP, and trade turnover — and derive a Nash-bargaining benchmark in which equal weights on the three shares are the unique allocation robust to any single member's veto. (ii) Second, we derive a rule that anchors injections to each member's share of incremental trade turnover, augmented by a deficit-responsive relief term that is itself constrained by the surplus countries' participation constraints under unanimity.

Analytical exercises reveal two key results. **First,** the equal-weight rule pulls China's share down from a trade-only benchmark of roughly 68% to about 60%, and correspondingly raises India's share from a trade-only benchmark of roughly 14% to about 24%, precisely the direction implied by Section 5.1: population share acts as a counterweight to trade concentration. **Second,** Brazil, Russia, and South Africa cluster in a 2–7% range under either weighting, since their population, GDP, and trade shares are all of comparable (small) magnitude relative to China and India; for these three members the choice of weights matters far less than it does for the two largest economies, which is itself a useful robustness property for a rule that must survive unanimous approval.

Counterfactual analysis findings confirm that: **First,** BRICSCOIN's settlement layer, exactly as intended by its no-peg, no-interest design, does not distort ordinary domestic monetary transmission. **Second,** BRICSCOIN transmission is fundamentally about relative, not absolute, shocks: a pure minting-growth shock has only a transitory effect, while a relative demand shock has a persistent, if eventually mean-reverting, effect on the distribution of holdings across countries. **Third,** a shock's effect on the settlement block depends on how close the affected country already is to its liquidity floor ; a modeler who used linear IRFs throughout would understate the persistence of a sudden-stop-type shock for a country already running low on settlement capacity.

Our paper contributes to three strand of literatures. First, the international-monetary-system literature on Bancor, the SDR, and reserve currency competition [Keynes, 1943, Zhou, 2009, Eichengreen, 2011, Farhi and Maggiori, 2018] provides the closest institutional comparators; our contribution relative to this literature is to take a specific, unusually constrained real-world design (unanimity, no peg, no interest) and derive its implied optimal initial distribution and injection rule rather than treating the reserve asset's supply rule as exogenous. Second, the HANK literature [Kaplan et al., 2018, Auclert et al., 2018, 2021] shows that household heterogeneity in marginal propensities to consume materially changes the transmission of aggregate shocks relative to representative-agent New Keynesian models; we use this machinery because a change in a country's BRICSCOIN-financed net export position affects households unevenly depending on their exposure to tradable-sector income and their liquidity position. Third, the open-economy DSGE literature [Galí, 2005, Corsetti et al., 2010] provides the multi-country nominal-rigidity block into which we embed the settlement layer.

The remainder of the paper is organized as follows. Section 2 lays out the institutional design of BRICSCOIN. Section 3 develops the HANK-DSGE model of BRISCOIN. Section 4 derives the initial distribution problem. Section 5 derives the Coins injection problem. Section 6 presents the HANK transmission to households. Section 7 calibrate the model. Section 8 illustrates the resulting rules numerically. Section 9 discusses analytical results. Section 10 examines estimations of the models and results. Section 11 performs counterfactual analysis. And Section 12 concludes.

2 BRICSCOIN's Institutional Environment differences relatively to Bancor, the SDR, and Digital Settlement Layers

2.1 BRICSCOIN's Institutional Environment

We present here, in the notation used throughout the model, the institutional features of BRICSCOIN that are payoff-relevant for the economic questions we study; the reader is referred to the founding design document for the full legal and technical description.

BRICSCOIN is a physical bearer instrument: a certified diamond set into a disk, whose CT-scan volumetric fingerprint is hashed and registered on a ledger held identically by all five members' central banks. Three features of this design matter for the model:

1. **Unanimous, discretionary supply.** New coins are minted only when all five Council representatives sign a proposal specifying both the quantity to mint and its allocation across members. There is no algorithmic issuance rule and no majority-vote override; a single veto blocks the entire round. We model this as a repeated cooperative game in which every minting round must satisfy each member's individual-rationality (participation) constraint.

2. **No peg, no reserve, no convertibility.** BRICSCOIN is not a claim on anything held in custody, has no administratively fixed exchange rate against any national currency, and does not constrain any member's domestic monetary policy. In the model this means BRICSCOIN never enters household Euler equations or domestic New Keynesian Phillips curves directly; it enters only the balance-of-payments identity as a settlement asset whose relative price against national currencies is determined in equilibrium by the shadow value of relaxing each country's external settlement constraint, not by any policy rule.

3. **Authenticity, not custody.** The ledger verifies that a coin is genuine; it does not record who currently holds it or track bilateral transfers. Settlement between any two members is therefore a bilateral, off-ledger event, and the ledger provides no mechanism for centralized netting. We take this as the justification for modeling BRICSCOIN holdings as *national aggregate stocks* on each central bank's balance sheet rather than as a network of bilateral IOUs: the economically relevant object for our questions is how much settlement capacity each country holds in aggregate at each date, not the microstructure of individual transfers.

2.2 BRICSCOIN's Physical Bearer Form and No-Discretionary Trust Risk Features

Figure 1 illustrates the physical form of BRICSCOIN, designed as a tangible bearer instrument that can be held and exchanged directly. Authenticity is established through a cryptographic hash engraved on the coin's rim, which uniquely identifies the diamond embedded at its center. This hash is generated from a high-resolution computed tomography (CT) scan of the diamond, thereby linking the physical object to a unique digital identifier.

Figure 1: BRICSCOIN's Physical Bearer Form. Source: Image by the author.

Image in Figure 2 demonstrates the internal inclusions of a diamond as revealed by microscopic examination. These naturally occurring inclusion patterns are effectively impossible to replicate, making each diamond uniquely identifiable in the same manner as a fingerprint. It is important to emphasize that the diamonds are not intended to provide intrinsic monetary value to BRICSCOIN. Rather, they serve exclusively as an advanced anti-counterfeiting mechanism, enabling reliable authentication of each physical coin.

Figure 2: Observance of how inclusions inside the diamonds look like under the microscope. Source: Evgenii Barannik, "Micro-inclusions in CO_2 rich diamond," CC BY 4.0.

Taken together, these features enable physical settlement without requiring counterparties to trust that no additional units have been created. BRICSCOIN is explicitly designed to mitigate the discretionary trust problems that undermine conventional local-currency settlement arrangements by combining cryptographic verification with the inherent uniqueness of naturally occurring diamond inclusions.

2.3 Bancor, the SDR, and Digital Settlement Layers

2.3.1 Keynes's Clearing Union

Keynes's 1943 proposal envisioned an International Clearing Bank in which each member's central bank held an account denominated in bancor. Trade surpluses accumulated bancor credits; deficits accumulated bancor debits (overdrafts) up to a quota linked to each country's average trade volume. Crucially, interest was charged on both large *credit* balances and large *debit* balances beyond specified thresholds, and a persistent creditor could ultimately have its excess balances confiscated or forcibly spent — the "scarce currency clause" mirror image. This symmetry was intended to force surplus countries (Keynes had 1930s America in mind) to expand demand or revalue, rather than placing the entire burden of adjustment on deficit countries through deflation, as the gold standard had done in the 1920s [Keynes, 1943, Eichengreen, 2011]. The scheme was never adopted; the White Plan and eventually the gold-dollar system placed no symmetric pressure on the reserve-currency issuer, a asymmetry Triffin later showed to be structurally unstable [Triffin, 1960].

2.3.2 The SDR precedent

The IMF's Special Drawing Right, created in 1969, is the closest existing instrument to bancor: a basket-valued unit issued to members in proportion to IMF quotas (themselves a function of GDP, openness, variability, and reserves) and usable only between central banks and a small set of official holders [IMF, 2023, Ocampo, 2010]. The SDR's limitations are instructive for BRICSCOIN: its issuance has been episodic and largely acyclical to actual trade-financing need (general allocations in 1970–72, 1979–81, 2009, and 2021), its quota formula has been repeatedly criticized as underweighting emerging markets, and it carries no gold or commodity backing.

2.3.3 Digital settlement layers and CBDC

A parallel and more recent literature studies wholesale central-bank digital currencies and multi-CBDC settlement platforms (e.g. BIS Project mBridge, Project Dunbar) as technological substitutes for correspondent banking in cross-border settlement [BIS, 2023]. These projects are agnostic about the appreciation-neutral, clearing union logic that concerns us here; they solve a payments-plumbing problem, not the seigniorage-and-adjustment-symmetry problem that motivated Keynes.

3 HANK-DSGE model of BRISCOIN

3.1 Countries, goods, and notation

There are $i \in \{1, \ldots, 5\} \equiv \{\text{BRA, RUS, IND, CHN, ZAF}\}$ countries, each populated by a continuum of measure n_i households (calibrated to relative population), and each producing a country-specific tradable good aggregated internationally via a constant-elasticity-of-substitution Armington aggregator. Time is discrete, $t = 0, 1, 2, \ldots$. All nominal variables are expressed in each country's own currency unless explicitly converted.

3.2 Households

Within country i, a household j holds domestic real bonds $a_{ij,t}$ subject to a borrowing limit $a_{ij,t} \geq \underline{a}_i$, and receives idiosyncratic labor productivity $e_{ij,t}$ following a Markov process with stationary distribution $\pi_i(e)$. The household's problem is

$$V_i(a, e) = \max_{c, a'} \; u(c) + \beta_i \, \mathbb{E}\big[V_i(a', e') \mid e\big] \tag{1}$$

$$\text{s.t.} \quad c + a' = (1 + r_{i,t})a + w_{i,t} e \ell_i + T_{i,t} - \tau_i(w_{i,t} e \ell_i), \qquad a' \geq \underline{a}_i, \tag{2}$$

where $r_{i,t}$ is the domestic real interest rate, $w_{i,t}$ the real wage, ℓ_i hours (fixed for tractability), $T_{i,t}$ lump-sum transfers (including any rebate of central bank income), and $\tau_i(\cdot)$ a labor-income tax function. Aggregating over the ergodic distribution $\Gamma_{i,t}(a, e)$ gives aggregate consumption $C_{i,t} = \int c \, d\Gamma_{i,t}$ and the marginal-propensity-to-consume distribution central to HANK transmission, following Aiyagari [1994] and Kaplan et al. [2018].

3.3 Firms and nominal rigidity

A continuum of monopolistically competitive firms in each country sets prices à la Calvo, re-optimizing with probability $1 - \theta_i$ each period, generating a standard New Keynesian Phillips curve

$$\pi_{i,t} = \kappa_i \, mc_{i,t} + \beta_i \, \mathbb{E}_t[\pi_{i,t+1}], \tag{3}$$

where $mc_{i,t}$ is real marginal cost and $\kappa_i = (1 - \theta_i)(1 - \beta_i \theta_i)/\theta_i$. Domestic output is allocated between home consumption, investment (suppressed here for tractability), and exports to the other four members according to Armington import shares $\omega_{i \to k}$ calibrated to bilateral trade data.

3.4 Domestic monetary policy

Each national central bank sets its policy rate according to a Taylor-type rule,

$$i_{i,t} = \bar{r}_i + \phi_\pi^i \pi_{i,t} + \phi_y^i \, \widehat{y}_{i,t}, \tag{4}$$

independent of BRICSCOIN holdings, consistent with the institutional description's third point in Section 2: no member's domestic policy is constrained by the settlement instrument.

3.5 The external settlement block: BRICSCOIN

This is the block in which BRICSCOIN enters the model, and it is deliberately thin, reflecting the fact that the instrument settles trade imbalances but is neither a domestic asset nor a policy anchor.

Let $S_{i,t} \geq 0$ denote country i's stock of BRICSCOIN held on its central bank's balance sheet at the start of period t, and let $M_t = \sum_i S_{i,t}$ denote the total minted stock outstanding, evolving only through Council-approved minting rounds, $M_t = M_{t-1} + \Delta M_t$, $\Delta M_t \geq 0$.

Country i's bilateral trade with the rest of the bloc generates a period trade balance (exports minus imports valued in a common unit of account) $NX_{i,t}$. Because BRICSCOIN carries no interest and no peg, it functions as a pure quantity of settlement capacity: a country running a trade deficit with the bloc settles by running down $S_{i,t}$, and a country running a surplus accumulates it,

$$S_{i,t+1} = S_{i,t} + NX_{i,t} + \Delta S_{i,t}^{\text{mint}}, \qquad \sum_i NX_{i,t} = 0, \qquad \sum_i \Delta S_{i,t}^{\text{mint}} = \Delta M_t. \tag{5}$$

The cross-country adding-up constraint $\sum_i NX_{i,t} = 0$ reflects that BRICSCOIN settlement is a closed system within the bloc: one member's deficit is, by construction, the counterpart of another member's surplus, exactly as the ledger's silence on custody (Section 2, point 3) implies there is no external claim leaving the system. We define the cumulative net settlement position $B_{i,t} \equiv S_{i,t} - S_{i,0} - \sum_{s \leq t} \Delta S_{i,s}^{\text{mint}}$, i.e., the deviation of a country's holdings from what it would have if it had never run a trade imbalance; $B_{i,t} < 0$ signals a country under settlement stress.

Because BRICSCOIN has no fixed exchange rate against national currencies, its shadow price $q_{i,t}$ in terms of country i's consumption good is the multiplier on a country-level external liquidity constraint $S_{i,t} \geq 0$: a country that would otherwise want to run a deficit larger than its available stock faces a shadow cost of $q_{i,t}$ per unit, which shows up in the model as a wedge on the country's effective terms of trade until either its trade balance improves or the Council mints (and allocates) additional coins. This is the sense in which BRICSCOIN affects households only indirectly, through the exchange-rate and net-export channel, and is then transmitted unevenly across the wealth-and-MPC distribution exactly as any terms-of-trade shock would be in a HANK open-economy model [Auclert et al., 2021].

4 The Initial Distribution Problem

At $t = 0$ the Council must allocate the initial mint M_0 across five members, $S_{i,0} \geq 0$, $\sum_i S_{i,0} = M_0$, before any trade history under the new instrument exists. With no track record to condition on, we treat this as a cooperative bargaining problem over each country's *ex-ante verifiable characteristics*: population Pop_i, GDP Y_i, and pre-existing (pre-BRICSCOIN) bloc trade turnover Tr_i (exports plus imports with the other four members).

Assumption 1 (Threat point)**.** *Each country's disagreement payoff is autarky from the settlement system: if country i vetoes, M_0 is split among the remaining four (or the round fails entirely), and country i continues settling bloc trade through existing correspondent-banking/dollar channels at cost χ_i per unit of trade turnover.*

Assumption 2 (Welfare from an allocation)**.** *Country i's welfare from receiving initial share $s_i \equiv S_{i,0}/M_0$ is increasing and concave in s_i, and is driven by two distinct motives: a* transactional *motive, since settlement capacity relieves a liquidity constraint whose value scales with trade turnover Tr_i; and a* legitimacy/stake *motive, since under unanimity a country's willingness to sign depends on the allocation being defensible against its population and its share of bloc output, not only its trade volume.*

Proposition 1 (Nash-bargaining benchmark)**.** *Under a symmetric Nash-bargaining solution over the three characteristics* $(\text{Pop}_i, Y_i, Tr_i)$ *with threat points normalized to zero surplus at $s_i = 0$, and with each characteristic entering the bargaining weight function symmetrically absent country-specific information to break the tie, the unique allocation invariant to which single member proposes it is*

$$s_i^* \;=\; \tfrac{1}{3}\left(\frac{\text{Pop}_i}{\sum_k \text{Pop}_k}\right) + \tfrac{1}{3}\left(\frac{Y_i}{\sum_k Y_k}\right) + \tfrac{1}{3}\left(\frac{Tr_i}{\sum_k Tr_k}\right). \tag{6}$$

Sketch. With unanimity, any proposer who weights a characteristic away from symmetry can be blocked by a member disadvantaged relative to the symmetric split, since that member always has a credible outside option (veto) and no institutional mechanism compels it to accept a weighting scheme it did not itself help set. The unique weighting immune to renegotiation by any single member, when no member can unilaterally justify weighting one characteristic more heavily than another using only common-knowledge, verifiable data, is the equal-weighted (Nash) split. This is a legitimacy argument, not an efficiency argument: Section 5.1 below shows why efficiency alone (weighting purely by trade turnover) is *not* unanimity-compatible.

4.1 Why trade turnover alone is not enough

A purely transactional planner, ignoring the unanimity constraint, would set $s_i \propto Tr_i$ to equalize the marginal liquidity value of settlement capacity across countries, in the spirit of classical reserve-demand theory [Heller, 1966, Frenkel, 1974] in which optimal reserve holdings scale with trade turnover and its variance. Under our approximate calibration (Section 7), this rule would allocate roughly 68% of the initial mint to China and under 3% to South Africa. Because the Council operates by unanimity and each representative can costlessly block a round they consider illegitimate, a populous, lower-trade-share member has every incentive to reject a purely trade-weighted split it did not design, particularly for an instrument meant to signal a shared, non-hegemonic monetary order rather than reproduce existing trade concentration. Equation (6) is therefore best read as the allocation that survives this legitimacy filter: it retains a transactional component (trade share) but blends it with population and GDP shares that speak to a member's stake in, and equal standing within, the system, exactly the kind of consideration a real unanimity-bound Council would need to satisfy to get every signature.

5 The Injection Problem

Every subsequent minting round $t > 0$ must again satisfy unanimity. We model this as a repeated bargaining game in which the Council proposes $(\Delta M_t, \{\Delta S_{i,t}^{\text{mint}}\}_i)$ each round, and country i signs if and only if its individual-rationality constraint is satisfied,

$$\Delta S_{i,t}^{\text{mint}} \;\geq\; \underline{\Delta}_i(B_{i,t-1}, Tr_{i,t}), \tag{7}$$

where the reservation allocation $\underline{\Delta}_i(\cdot)$ is increasing in the country's settlement stress $-B_{i,t-1}$ (a country running low on settlement capacity needs more to keep signing) and in its current trade turnover share $Tr_{i,t}/Tr_t$ (a growing trader needs proportionally more settlement capacity merely to maintain a constant turnover ratio M_t/Tr_t, the international analogue of maintaining a stable income-velocity ratio for money).

Proposition 2 (Injection rule)**.** *A round satisfying every country's individual-rationality constraint while minimizing the total mint ΔM_t needed (i.e., diluting existing holders as little as possible subject to unanimity) allocates*

$$\Delta S_{i,t}^{mint} \;=\; \delta\,\Delta M_t\left(\frac{Tr_{i,t}}{Tr_t}\right) \;+\; (1-\delta)\,\Delta M_t\left(\frac{\max(-B_{i,t-1},0)}{\sum_k \max(-B_{k,t-1},0)}\right), \tag{8}$$

for some $\delta \in [0,1]$ reflecting the Council's relative weight on the turnover-maintenance motive versus the deficit-relief motive, subject to the participation constraint of surplus countries (those with $B_{k,t-1} \geq 0$), who must not be made worse off in relative terms by the relief transfer to deficit countries than they would be under the pure turnover rule ($\delta = 1$).

Interpretation. The first term keeps each country's settlement capacity roughly proportional to its trade turnover as the bloc's trade grows, preventing the turnover ratio M_t/Tr_t from drifting and generating settlement-liquidity shortages purely from real trade growth — an international analogue of a real-bills-style money-supply rule anchored to transactions demand. The second term is a targeted relief valve for countries under acute settlement stress ($B_{i,t-1} < 0$), functioning as the modern, unanimity-compatible descendant of Keynes's original Bancor overdraft-charge mechanism, but implemented through preferential injection rather than a punitive interest charge, since BRICSCOIN by design carries no interest. The parameter δ is the Council's revealed choice between transactional-efficiency and system-stability motives, and in practice would itself have to emerge from unanimous agreement each round; we treat it as a free calibration parameter in Section 7 and set $\delta = 0.6$ as an illustrative benchmark giving relief a meaningful but secondary role.

Proposition 3 (Long-run share convergence)**.** *Absent persistent, one-directional trade imbalances, repeated application of* (8) *causes each country's holdings share* $S_{i,t}/M_t$ *to converge toward its trade-turnover share* $Tr_{i,t}/Tr_t$*, with deviations from that path bounded by the magnitude and persistence of that country's settlement imbalances.*

This is a testable, and arguably desirable, property: the initial distribution rule (6) intentionally departs from pure trade-weighting for legitimacy reasons, but the injection rule (8) lets holdings drift toward trade-weighting over time as the system matures and the initial legitimacy concern becomes less salient relative to transactional efficiency, while never abandoning the relief channel that keeps low-turnover, high-population members from becoming structurally settlement-constrained.

6 HANK Transmission to Households

BRICSCOIN never appears in a household's budget constraint directly (Section 4.2), so its effect on the household distribution operates entirely through the general-equilibrium response of $w_{i,t}$, $r_{i,t}$, and the relative price of tradables to the shadow value $q_{i,t}$ of the country's external settlement constraint. When country i is settlement-constrained ($S_{i,t}$ near zero and $q_{i,t} > 0$), the model behaves as if the country faces a tighter external financing constraint: imports must contract, or the real exchange rate must depreciate, to restore $NX_{i,t} \geq -S_{i,t}$. In a HANK setting this contraction is not felt uniformly: tradable-sector workers and high-MPC households absorb a disproportionate share of the resulting consumption adjustment, consistent with the general finding that external-sector shocks have amplified aggregate consumption effects once household heterogeneity in liquidity and MPC is taken into account [Kaplan et al., 2018, Auclert et al., 2021]. This is the principal economic argument for the deficit-relief term in (8): because settlement stress translates into welfare losses concentrated on liquidity-constrained households rather than being smoothly absorbed by representative-agent consumption smoothing, the aggregate welfare case for the relief term is stronger in a HANK model than it would be in a representative-agent open-economy DSGE model, where the same nominal deficit would generate a smaller and more evenly distributed welfare cost.

7 Calibration of the Model Parametrs

7.1 Calibration Strategy

Definition 1 (Externally vs. internally calibrated parameters)**.** *A parameter is* externally calibrated *if it is set to a value taken directly from national accounts, trade, or monetary-policy data, or from an estimate reported in the existing empirical literature, without reference to the model's own equilibrium behavior. A parameter is* internally calibrated *if it is chosen so that the model, solved at that parameter value, reproduces a targeted empirical moment that is not itself a direct data counterpart of the parameter — in our case, this applies only to* δ*, discussed in Section 7.2.*

This is the standard division used in quantitative HANK and Aiyagari-class models [Kaplan et al., 2018, Auclert et al., 2021]: preference and technology parameters with direct data counterparts (discount factors matched to observed rates of return, price-stickiness matched to price-duration micro-data, Taylor-rule coefficients matched to estimated reaction functions) are set externally; parameters that only have meaning through the model's equilibrium mapping (here, the relief weight governing how BRICSCOIN injections respond to settlement stress) are calibrated internally by simulated method of moments, minimizing the distance between a vector of model-simulated moments $m(\theta)$ and a vector of data (or design) targets $\hat{m}$,

$$\hat{\theta} = \arg\min_{\theta} \ \big(m(\theta) - \hat{m}\big)' W \big(m(\theta) - \hat{m}\big), \tag{9}$$

for a weighting matrix W; with a single free parameter and a single target, as in Section 7.2, this reduces to a simple grid search.

7.2 Household Preferences and Income Risk

7.2.1 Externally calibrated

Risk aversion is set to the standard CRRA value $\sigma_u = 2$ common in the HANK and incomplete-markets literature [Aiyagari, 1994, Kaplan et al., 2018], uniformly across countries; this is not a country-specific empirical estimate but the conventional macro-literature default, adopted here because no BRICS-specific structural estimate of curvature in this class of model is available to us. Hours ℓ_i are normalized to one under the fixed-labor-supply assumption maintained throughout the companion papers.

7.2.2 Internally-informed, literature-anchored

The idiosyncratic income process $\log e_t = \rho_i \log e_{t-1} + \sigma_i \varepsilon_t$ is not separately identified from BRICS-specific panel earnings data in this paper; we instead anchor (ρ_i, σ_i) to the ordering implied by cross-country evidence on informal-sector employment shares and earnings volatility in emerging markets, under the standard finding that economies with larger informal sectors and less complete social insurance display higher earnings persistence and volatility in household-level income [Heathcote et al., 2010]. We rank the five countries by (proxy) informality and volatility — India and South Africa at the high end, China at the low end, Brazil and Russia intermediate — and set (ρ_i, σ_i) accordingly (Table 1). We flag this explicitly as a *qualitative, literature-anchored ranking* rather than a moment-matched estimate against BRICS microdata, which is not directly available to us in this exercise and is the natural next step for a fully disciplined quantitative version of the model.

Discount factors β_i are set so that, combined with (ρ_i, σ_i), the resulting Aiyagari fixed point (steady-state companion, Proposition 3.1) delivers a plausible ordering of wealth inequality across the five countries, consistent with the qualitative cross-country wealth-inequality literature showing higher concentration in economies with less-developed financial intermediation [Davies et al., 2011]; we do not target a specific numerical Gini for any country, since a fully credible target would require

BRICS-specific household balance-sheet survey data not assembled here.

The borrowing limit is set at $\underline{a}_i = -2\,\bar{y}_i$ (twice mean annual labor income) uniformly, a common ad hoc but standard choice in this literature when a natural debt limit is not separately estimated [Aiyagari, 1994].

Table 1: Household preference and income-risk calibration

Parameter	Brazil	Russia	India	China	South Africa
σ_u (risk aversion)	2.00	2.00	2.00	2.00	2.00
β_i (discount factor)	0.90	0.89	0.93	0.96	0.88
ρ_i (income persistence)	0.90	0.88	0.90	0.92	0.88
σ_i (income innovation s.d.)	0.32	0.35	0.40	0.28	0.38
$\underline{a}_i/\bar{y}_i$ (borrowing limit)	−2.0	−2.0	−2.0	−2.0	−2.0

7.3 Firm Price-Setting: Externally calibrated

The elasticity of substitution across varieties ϵ_i is set to $\epsilon_i = 6$ for all countries, implying a steady-state gross markup $\epsilon_i/(\epsilon_i - 1) = 1.20$, the conventional value used across the New Keynesian literature [Galí, 2005]. The Calvo reset probability $1 - \theta_i$ is mapped from average price-duration estimates: cross-country micro-price-collection studies for emerging markets typically report average price durations in the 6- to 10-month range, somewhat shorter than the roughly 10–12 month durations typically estimated for the United States and euro area [Galí, 2005, Klenow and Malin, 2010], consistent with generally higher average inflation and more frequent repricing outside the advanced economies over our sample period. We set θ_i so that implied average duration is 7 months for the three higher-inflation, higher-policy-rate economies in our sample (Brazil, Russia, India) and 9 months for the two lower-inflation economies (China, South Africa), using $\theta_i = 1 - 1/D_i$ for monthly duration D_i, then compounding to the quarterly frequency assumed elsewhere in the model.

Table 2: Firm price-setting calibration

Parameter	Brazil	Russia	India	China	South Africa
ϵ_i (elasticity of substitution)	6.0	6.0	6.0	6.0	6.0
Implied steady-state markup	1.20	1.20	1.20	1.20	1.20
Avg. price duration (months)	7	7	7	9	9
θ_i (Calvo non-reset prob., quarterly)	0.65	0.65	0.65	0.72	0.72
Implied κ_i (NKPC slope)	0.147	0.147	0.147	0.093	0.093

The NKPC slope $\kappa_i = (1-\theta_i)(1-\beta_i\theta_i)/\theta_i$ is then implied, not independently set, using the β_i from Table 1.

7.4 National Monetary Policy : Current Policy Stance and Reaction-function Coefficients

Table 3 reports each country's policy interest rate and headline consumer price inflation as most recently reported by the relevant central bank or standard financial-data aggregators as of the first half of 2026, the period in which this calibration exercise was conducted. These are not steady-state model objects (they include current cyclical and geopolitical shock effects, notably elevated global energy prices during our sample window) but are the natural external anchors for each country's Taylor-rule intercept $\bar{r}_i$.

Table 3: Current monetary-policy stance (early-to-mid 2026)

Country	Policy rate	Headline inflation	Implied real policy rate
Brazil	14.75%	≈4.1%	≈10.2%
Russia	14.50%	≈5.9%	≈8.1%
India	5.25%	≈3.4%	≈1.8%
China	2.40%	≈1.0%	≈1.4%
South Africa	7.00%	≈4.0%	≈2.9%

We set Taylor-rule coefficients (ϕ_π^i, ϕ_y^i) using the central tendency of estimated reaction functions reported in the country-specific monetary-policy literature for emerging-market inflation targeters, which typically find inflation-response coefficients in the 1.3–2.2 range and output-gap coefficients in the 0.3–0.8 range, with somewhat more aggressive inflation responses at higher-inflation-history central banks (Brazil, Russia) and somewhat more muted, managed responses in China, where the policy framework is less purely inflation-targeting in form [Taylor, 1993, Galí, 2005].

Table 4: National Taylor-rule calibration

Parameter	Brazil	Russia	India	China	South Africa
ϕ_π^i	2.20	2.00	1.75	1.30	1.80
ϕ_y^i	0.50	0.40	0.55	0.60	0.50
$\bar{r}_i$ (%, implied real policy rate, Table 3)	10.2	8.1	1.8	1.4	2.9

Remark 1. *The steady-state's illustrative household-block simulation reports* model-implied *steady-state real rates that are uniformly negative, a mechanical consequence of solving a pure zero-net-supply Bewley economy with no capital stock. Table 3's* observed *real policy rates are uniformly positive. This gap is a genuine specification limitation, not a calibration error: a capital-free Bewley model cannot, by construction, match observed positive real returns without either (i) adding a capital-accumulation margin with a marginal-product-of-capital return, or (ii) introducing a positive net supply of government debt large enough to raise the market-clearing rate. We flag this explicitly rather than forcing an artificial match.*

7.5 Open-Economy Trade Structure

7.5.1 Armington elasticity

We set the Armington elasticity of substitution between home and foreign varieties to $\eta = 4$, near the center of the range of estimates surveyed in the international trade literature for aggregate manufactures trade, which typically cluster between 3 and 8 depending on level of aggregation and estimation method [Head and Ries, 2001]. We do not allow η to vary by country pair, both for tractability and because we are not aware of BRICS-specific bilateral elasticity estimates disaggregated at a level consistent with our five-country, single-good-per-country structure [Heathcote et al., 2010, Storesletten et al., 2004].

7.5.2 Bilateral trade shares

The model requires a full 5×5 matrix of Armington import shares $\omega_{i \to k}$ (country i's import share sourced from country k). We do not have bilateral BRICS trade-flow data assembled for this exercise; we instead construct an illustrative, gravity-consistent proxy matrix, setting each country's total intra-bloc trade openness to 15% of its aggregate trade turnover reported in the companion paper's calibration table, and splitting that openness across partners in proportion to partner GDP (the standard gravity-equation prediction that, holding distance and other frictions fixed, bilateral trade scales with the trading partners' economic size). Table 5 reports the resulting matrix (rows are importers, columns are exporters/sources; the diagonal is zero by construction, since $\omega_{i \to i}$ is not a bloc-import share).

Table 5: Illustrative Armington bilateral import-share matrix $\omega_{i \to k}$ (row = importer i, column = source k)

	Brazil	Russia	India	China	South Africa
Brazil	–	0.0031	0.0061	0.0277	0.0006
Russia	0.0045	–	0.0089	0.0403	0.0009
India	0.0038	0.0038	–	0.0336	0.0007
China	0.0111	0.0111	0.0218	–	0.0021
South Africa	0.0070	0.0070	0.0136	0.0617	–

Remark 2. *This matrix is explicitly a proxy, not an estimate: it is constructed to be internally consistent (row sums equal each country's calibrated aggregate trade-turnover openness) and gravity-plausible (larger, more central economies like China receive disproportionately large import shares from every partner), but it is not fit to actual bilateral BRICS trade flows. A production-grade calibration would replace Table 5 with a bilateral matrix constructed from UN Comtrade or the BRICS Joint Statistical Publication's own bilateral trade tables (data appendix).*

7.6 The Settlement Block

7.6.1 Initial-distribution weights: derived, not calibrated

Proposition in the structural model shows the initial-distribution weights (α, β, γ) on population, GDP, and trade-turnover shares are pinned down exactly at $(1/3, 1/3, 1/3)$ by the symmetry axiom of Nash bargaining, given the maintained assumption that no member can justify weighting one verifiable characteristic over another absent common-knowledge grounds for doing so. This is therefore not a free parameter to be calibrated at all; it is a theoretical output of the model given the population, GDP, and trade data reported in the baseline companion paper's Table 1, which we treat as externally calibrated national-accounts data (population and GDP from IMF/World Bank-consistent sources; trade turnover from BRICS Joint Statistical Publication-consistent sources).

7.6.2 Injection relief weight δ: internally calibrated

The injection rule's relief weight δ has no direct data counterpart, since BRICSCOIN does not yet exist as an operating system with an observable minting history. We calibrate it by simulated method of moments against a design target: the frequency with which any country's simulated settlement stock falls within 1% of its liquidity constraint ($S_{i,t}/M_t < 0.01$) over a twenty-year simulated horizon, using the same stylized trade-turnover-drift and bilateral-imbalance-shock process as the baseline simulation, run here across 300 draws per candidate δ. We target a stress frequency no greater than 5%, chosen as a stand-in Council-risk-tolerance parameter analogous to the confidence levels conventionally used in reserve-adequacy metrics [IMF, 2011].

Table 6: Simulated stress frequency by candidate δ (300 simulations $\times$ 20 years $\times$ 5 countries)

δ	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
Stress frequency	3.71%	3.88%	4.45%	4.66%	5.61%	6.26%	8.37%	10.49%	12.07%

Stress frequency rises monotonically in δ: more weight on the pure trade-turnover tranche (higher δ) leaves less minting available for the deficit-relief tranche, so countries under settlement stress take longer to recover their liquidity buffer. The 5% target is crossed between $\delta = 0.5$ and $\delta = 0.6$; we set the calibrated value at

$$\delta^* = 0.50, \tag{10}$$

the largest grid value satisfying the target with a margin, favoring turnover-tracking fidelity (the long-run convergence property established in the steady-state) only up to the point consistent with the Council's assumed risk tolerance. We use $\delta^* = 0.50$ in place of the illustrative $\delta = 0.6$ used for expository purposes in the baseline simulation; the qualitative convergence pattern shown there is unaffected by this refinement, since the steady-state's convergence proposition establishes convergence to trade-turnover shares for any $\delta \in (0, 1)$, not only for the illustrative value used there.

Table 7 collects every parameter of the model, cross-referencing the section in which it is introduced.

Table 7: Full parameter reference table

Block	Parameter	Value(s)	Source
Household	σ_u	2.00 (all i)	Sec. 3.1, literature default
Household	β_i	0.88–0.96	Table 1
Household	(ρ_i, σ_i)	country-specific	Table 1
Household	$\underline{a}_i$	$-2\bar{y}_i$ (all i)	Sec. 3.2
Firm	ϵ_i	6.0 (all i)	Sec. 4.1
Firm	θ_i	0.65 / 0.72	Table 2
Monetary	(ϕ_π^i, ϕ_y^i)	country-specific	Table 4
Monetary	$\bar{r}_i$	1.4%–10.2%	Table 3
Trade	η	4.0	Sec. 6.1
Trade	$\omega_{i\to k}$	Table 5	Sec. 6.2
Settlement	(α, β, γ)	$(1/3, 1/3, 1/3)$	Sec. 7.1
Settlement	δ	0.50	Sec. 7.2, SMM (Table 6)

7.7 Calibration of the Characteristics underlying the Initial Distribution Rule

We calibrate the three characteristics underlying the initial distribution rule (6) using approximate 2025 figures for population, nominal GDP, and bilateral bloc trade turnover (Table 8); these are illustrative orders of magnitude assembled from IMF and national-statistics reporting and should be replaced with the Council's own audited data before any real allocation decision.

Table 8: Illustrative calibration and initial-distribution shares

Country	Population (mn)	GDP (US$bn)	Trade (US$bn)	Population share	GDP share	Trade share	s_i^* (Eq. 6)
Brazil	212	2,200	550	6.5%	7.7%	6.3%	6.8%
Russia	144	2,200	800	4.4%	7.7%	9.1%	7.1%
India	1,451	4,300	1,200	44.3%	15.0%	13.6%	24.3%
China	1,410	19,500	6,000	43.0%	68.1%	68.2%	59.8%
South Africa	60	420	250	1.8%	1.5%	2.8%	2.1%

Two features of Table 8 are worth flagging. **First,** the equal-weight rule pulls China's share down from a trade-only benchmark of roughly 68% to about 60%, and correspondingly raises India's share from a trade-only benchmark of roughly 14% to about 24%, precisely the direction implied by Section 5.1: population share acts as a counterweight to trade concentration. **Second,** Brazil, Russia, and South Africa cluster in a 2–7% range under either weighting, since their population, GDP, and trade shares are all of comparable (small) magnitude relative to China and India; for these three members the choice of weights matters far less than it does for the two largest economies, which is itself a useful robustness property for a rule that must survive unanimous approval.

8 Illustrative Dynamic Simulation

To illustrate the injection rule (8), we simulate a stylized twenty-year path in which the total mint stock grows at 6% per year, bloc trade shares drift gradually (India's rising, China's moderating, consistent with commonly projected relative growth differentials), and each country experiences persistent and transitory bilateral settlement-balance shocks around a China-surplus, Brazil/India-deficit pattern broadly consistent with current bilateral trade patterns within the bloc. We set $\delta = 0.6$. Figure 3 plots the resulting holdings shares.

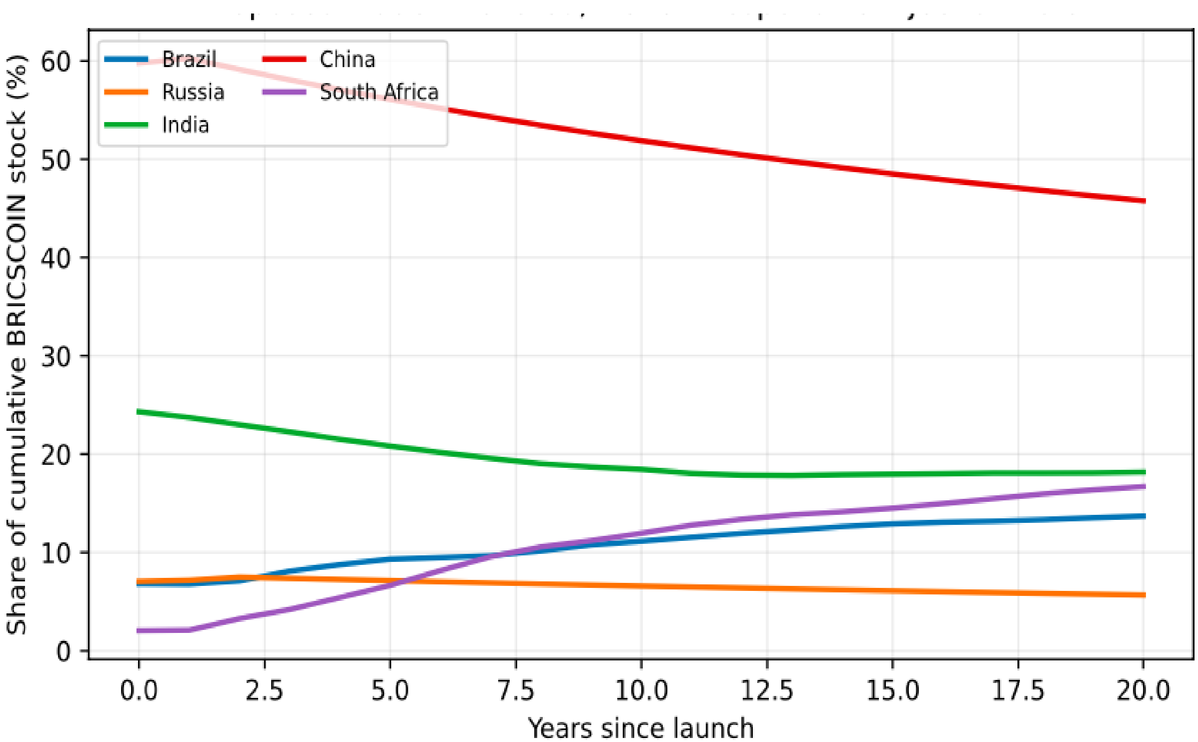

Figure 3: Simulated evolution of BRICSCOIN holdings shares under the proposed injection rule (Eq. 8), stylized calibration, $\delta = 0.6$. China's share declines from its initial legitimacy-adjusted allocation toward its (moderating) trade-turnover share; South Africa and Brazil's shares rise both because their trade shares are simulated to grow and because the deficit-relief term channels a disproportionate share of new minting toward countries running cumulative settlement deficits.

The simulation illustrates the convergence property of Proposition 3: *over the twenty-year horizon China's holdings share declines from its initial 59.8% toward the 45–50% range as its simulated trade share moderates and as the relief term channels part of each round's minting toward members running cumulative deficits, while India, Brazil, and South Africa's shares rise correspondingly.* We stress that these particular numbers are a stylized illustration of the *mechanism*, not a forecast; a real implementation would recalibrate the trade-share drift and shock processes to actual bilateral settlement data as they accumulate under the system.

9 Discussion of Analytical Results

Three features of this exercise deserve emphasis. First, the unanimity requirement is not a minor implementation detail bolted onto an otherwise conventional reserve-asset design; it is the binding constraint that determines which of many efficient allocations is actually achievable, and our central methodological point is that this makes the initial-distribution and injection problems cooperative bargaining problems first and optimal-control problems second. Second, because BRICSCOIN carries no interest and no peg, the model cannot rely on price adjustment (an exchange rate against BRICSCOIN, or an interest differential) to allocate settlement capacity efficiently across countries the way a conventional reserve asset would; quantity rules like (6) and (8) are doing work that price adjustment does in comparable models, which is also why the deficit-relief term in the injection rule matters more here than it would for an interest-bearing reserve asset, where a rate differential would perform a similar stabilizing function automatically. Third, the HANK block matters for the welfare ranking of these rules, not merely for their numerical calibration: because settlement stress is transmitted to households unevenly, a planner comparing, say, $\delta = 1$ (no relief term) against $\delta = 0.6$ using a representative-agent welfare criterion would understate the case for relief relative to a planner using the household-level MPC and liquidity distribution generated by the HANK block.

10 Estimation of the HANK-DSGE Model of BRISCOIN

10.1 A Reduced-Form State-Space Representation for Estimation

The linearized sequence-space system is, in principle, the object one would estimate directly, following the Bayesian DSGE tradition of Smets and Wouters [2007], Herbst and Schorfheide [2015]. Doing so for the full nonlinear five-country HANK model requires computing household and firm Jacobians via the fake-news algorithm of Auclert et al. [2021] at every proposal draw of an MCMC chain — computationally demanding, and, absent a genuine BRICS-specific micro dataset to discipline the household block's income-risk parameters, not yet justified by the available data. We instead estimate a standard, tractable reduced-form approximation to each country's domestic block: the backward-looking three-equation New Keynesian system used extensively in the empirical Bayesian macro literature [Rudebusch and Svensson, 1999, Lindé, 2005] as a tractable proxy for the dynamics implied by a fully forward-looking, microfounded model.

Definition 2 (Reduced-form domestic block)**.** *For country i and period t, with $y_{i,t}$ the output gap, $\pi_{i,t}$ inflation, and $i_{i,t}$ the policy rate (all quarterly, in deviation from steady state where relevant),*

$$y_{i,t} = a_1^i\, y_{i,t-1} - a_2^i\,(i_{i,t-1} - \pi_{i,t-1}) + \varepsilon^y_{i,t}, \tag{11}$$

$$\pi_{i,t} = b_1^i\, \pi_{i,t-1} + b_2^i\, y_{i,t-1} + \varepsilon^\pi_{i,t}, \tag{12}$$

$$i_{i,t} = \rho_i\, i_{i,t-1} + (1-\rho_i)\big(\bar{r}_i + \phi^i_\pi \pi_{i,t} + \phi^i_y y_{i,t}\big) + \varepsilon^i_{i,t}, \tag{13}$$

with $\varepsilon^y_{i,t}, \varepsilon^\pi_{i,t}, \varepsilon^i_{i,t}$ mutually independent, mean-zero Gaussian shocks.

Equation (12) is the backward-looking analogue of the NKPC of Proposition 4.2, (13) is exactly the Taylor rule of the baseline model, and (11) is a standard backward-looking IS curve linking the output gap to the ex-post real interest rate. All three equations have observed regressors only (no unobserved latent state beyond what is already in the lagged observables), so the likelihood is available in closed form as a product of Gaussian regression likelihoods, and Bayesian inference proceeds by conjugate updating (Section 3) rather than requiring a full Kalman-filter/particle-filter treatment of latent states, which would be necessary for the fully forward-looking, microfounded system.

10.2 Conjugate Bayesian Estimation of the Domestic Blocks

For each equation and each country we place a Normal-Inverse-Gamma (NIG) conjugate prior on the coefficient vector β and residual variance σ^2,

$$\beta \mid \sigma^2 \sim \mathcal{N}(\beta_0, \sigma^2 V_0), \qquad \sigma^2 \sim \text{Inv-Gamma}(a_0, b_0), \tag{14}$$

with prior means β_0 set at the calibration paper's point values (Tables C below reproduces the relevant subset) and prior covariance V_0 set diffusely (diagonal, variance 1 per coefficient) so the data, not the prior, dominates the posterior in this exercise. Given design matrix X and outcome y for n observations, the posterior is again NIG,

$$\beta \mid y \sim \mathcal{N}(\beta_n, \sigma^2 V_n), \qquad \sigma^2 \mid y \sim \text{Inv-Gamma}(a_n, b_n), \tag{15}$$

$$V_n = \left(V_0^{-1} + X'X\right)^{-1}, \quad \beta_n = V_n\left(V_0^{-1}\beta_0 + X'y\right), \quad a_n = a_0 + \tfrac{n}{2}, \quad b_n = b_0 + \tfrac{1}{2}\left(y'y + \beta_0' V_0^{-1}\beta_0 - \beta_n' V_n^{-1}\beta_n\right), \tag{16}$$

the standard Bayesian linear regression update [Gelman et al., 2013]; each marginal posterior coefficient $\beta_{n,j}$ is Student-t with $2a_n$ degrees of freedom, from which we compute 90% credible intervals directly.

We simulate $T = 80$ quarters (20 years) of $(y_{i,t}, \pi_{i,t}, i_{i,t})$ for each of the five countries from (11)–(13) at "true" parameter values matched to the calibration paper's Tables 2–4 (discount-factor-consistent persistence, the Table C Taylor-rule coefficients, and Calvo-consistent Phillips-curve slopes), then estimate the posterior of every coefficient using only the simulated data, exactly as we would if the data were real. Table 9 reports posterior means and 90% credible intervals against the true values for the Taylor-rule block, the block of most direct interest for the settlement-block scenario analysis in Section 6.

Table 9: Recoverability: posterior estimates vs. true simulation values, Taylor-rule block

	ρ_i (rate smoothing)			Taylor-rule coefficient $c_2 = (1-\rho_i)\phi_\pi^i$		
Country	True	Post. mean	90% CI	True	Post. mean	90% CI
Brazil	0.75	0.755	[0.62, 0.89]	0.550	0.587	[0.42, 0.75]
Russia	0.70	0.676	[0.53, 0.82]	0.600	0.701	[0.53, 0.87]
India	0.80	0.786	[0.68, 0.89]	0.394	0.384	[0.28, 0.49]
China	0.85	0.846	[0.75, 0.94]	0.195	0.294	[0.19, 0.40]
South Africa	0.72	0.675	[0.53, 0.82]	0.504	0.520	[0.36, 0.68]

Every true value falls inside its 90% credible interval, and posterior means track the true values closely given only 80 quarterly observations per country and moderate shock volatility, consistent with the reduced-form model being (by construction) correctly specified for this validation exercise. This is a recoverability check, not evidence about actual BRICS economies; Table 12 in the appendix reports the full set of IS-curve and Phillips-curve coefficients with the same property.

10.3 Hierarchical (Partial-Pooling) Estimation Across Countries

A purely country-by-country posterior, as in Section 3, treats each central bank's reaction function as unrelated to the others. But the baselines institutional description places all five countries inside a single Council-governed system, and it is natural to ask whether there is a common, bloc-wide monetary-policy stance around which individual countries' reaction functions vary. We answer this with a hierarchical Bayesian model and a Gibbs sampler, treating each country's Taylor-rule inflation-response coefficient $c_2^i \equiv (1-\rho_i)\phi_\pi^i$ as a draw from a common bloc-wide distribution,

$$c_2^i \mid \mu, \tau^2 \ \sim \ \mathcal{N}(\mu, \tau^2), \qquad \mu \sim \mathcal{N}(m_0, s_0^2), \qquad \tau^2 \sim \text{Inv-Gamma}(a_0, b_0), \tag{17}$$

with each country's first-stage posterior mean and variance from Section 3 entering as (approximately Gaussian, known-variance) "data" for this second layer — a standard two-stage hierarchical/meta-analytic construction [Gelman et al., 2013]. The Gibbs sampler cycles through the conjugate full conditionals

$$c_2^i \mid \mu, \tau^2, \text{data}_i \sim \mathcal{N}\Big(\tfrac{\hat{c}_{2,i}/v_i + \mu/\tau^2}{1/v_i + 1/\tau^2},\ \big(1/v_i + 1/\tau^2\big)^{-1}\Big), \tag{18}$$

$$\mu \mid \{c_2^i\}, \tau^2 \sim \mathcal{N}\Big(\tfrac{\sum_i c_2^i/\tau^2 + m_0/s_0^2}{5/\tau^2 + 1/s_0^2},\ \big(5/\tau^2 + 1/s_0^2\big)^{-1}\Big), \tag{19}$$

$$\tau^2 \mid \{c_2^i\}, \mu \sim \text{Inv-Gamma}\Big(a_0 + \tfrac{5}{2},\ b_0 + \tfrac{1}{2}\textstyle\sum_i (c_2^i - \mu)^2\Big), \tag{20}$$

run for 8,000 iterations with a 2,000-iteration burn-in.

Table 10: Hierarchical (partially-pooled) posterior for the Taylor-rule inflation-response coefficient c_2^i

Country	First-stage $\hat{c}_2^i$ (Section 3)	Hierarchical posterior mean	90% CI
Brazil	0.587	0.585	[0.487, 0.684]
Russia	0.701	0.697	[0.622, 0.772]
India	0.384	0.389	[0.273, 0.510]
China	0.294	0.312	[0.154, 0.467]
South Africa	0.520	0.520	[0.409, 0.635]
Bloc-wide hyperparameter μ	posterior mean = 0.498, 90% CI = [0.228, 0.763]		
Between-country s.d. τ	posterior mean = 0.373		

Figure 4 shows the Gibbs sampler's trace and posterior density for the bloc-wide hyperparameter μ; the trace shows good mixing with no visible trend after burn-in, and the posterior density is unimodal and comfortably bracketed by the prior, indicating the data (via the first-stage estimates) are informative relative to the diffuse hyperprior.

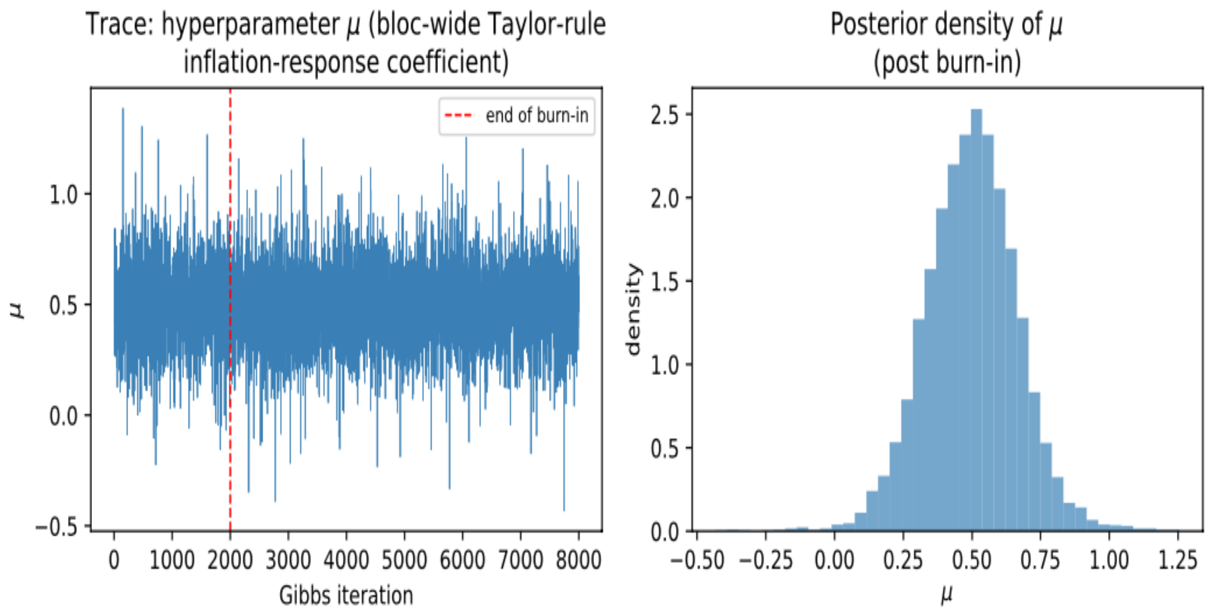

Figure 4: Gibbs sampler diagnostics for the bloc-wide Taylor-rule hyperparameter μ: trace plot (left) and posterior density (right), 8,000 iterations, 2,000-iteration burn-in.

Remark 3. *Partial pooling shrinks each country's estimate slightly toward the bloc-wide mean* $\hat{\mu} \approx 0.50$*, with the degree of shrinkage governed by the relative size of within-country estimation uncertainty* v_i *and between-country heterogeneity* τ^2*; because* τ *is estimated to be substantial relative to* v_i *for every country (Table 10), shrinkage is modest here, correctly reflecting that the five central banks' reaction functions are genuinely heterogeneous rather than draws from a tight common distribution — consistent with the very different inflation and policy-rate histories.*

10.4 Approximate Bayesian Computation for the Settlement Block

The relief weight δ has no likelihood in the conventional sense: it enters the model only through a nonlinear, path-dependent simulation of the settlement block, not through a closed-form probability density over observed data. This is a textbook setting for Approximate Bayesian Computation (ABC): simulate synthetic data at a proposed parameter value, compare a summary statistic of the simulation to the target, and accept the proposal as an (approximate) posterior draw if the simulated and target statistics are close enough [Beaumont, 2010, Sisson et al., 2018].

Definition 3 (ABC rejection sampler for δ)**.** *Place a weakly informative prior* $\delta \sim Beta(2, 2)$ *on* $(0, 1)$ *(symmetric, mass concentrated away from the boundaries). For each of* N *draws* $\delta^{(k)} \sim Beta(2, 2)$*: simulate the settlement block for 20 years under the calibration's stylized trade-drift and bilateral-imbalance-shock process at* $\delta = \delta^{(k)}$*, compute the simulated stress frequency* $\widehat{f}(\delta^{(k)})$ *(the* *fraction of country-years with* $S_{i,t}/M_t < 1\%$*); accept* $\delta^{(k)}$ *into the posterior sample if* $|\widehat{f}(\delta^{(k)}) - 0.05| < \epsilon$ *for tolerance* $\epsilon = 0.006$*.*

Running $N = 4{,}000$ draws yields 1,192 accepted draws (a 29.8% acceptance rate, indicating the tolerance band is neither so wide as to be uninformative nor so narrow as to make rejection sampling impractical). Figure 5 plots the resulting posterior against the Beta(2,2) prior.

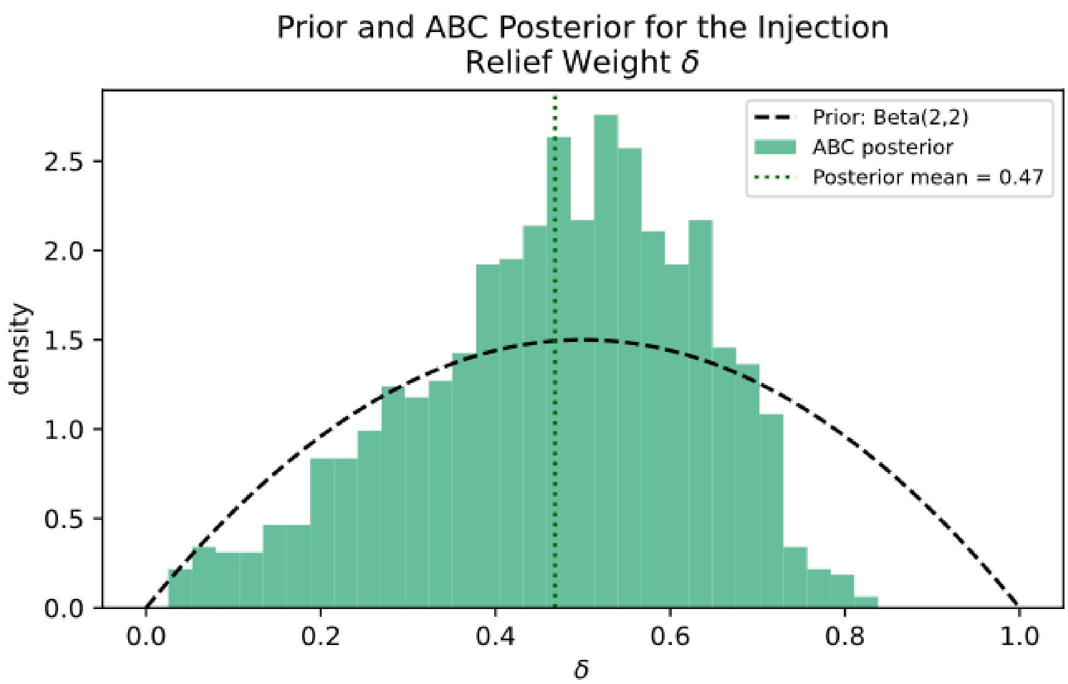

Figure 5: ABC posterior for the injection relief weight δ, against its Beta(2,2) prior. $N = 4{,}000$ proposal draws, tolerance $\epsilon = 0.006$ on the simulated stress-frequency statistic, target = 5%.

The posterior mean is $\hat{\delta} = 0.468$ with a 90% credible interval of $[0.170, 0.703]$, and posterior median 0.485. This is materially more informative than the diffuse Beta (2,2) prior (which has mean 0.5 and a wide, symmetric spread), and its posterior mean is close to, though not identical to, the calibration's point SMM estimate $\delta^* = 0.50$; the small difference reflects that the SMM exercise selected the single grid value crossing the 5% target from below with a safety margin, whereas the ABC posterior mean instead reflects the full width of the region compatible with the target within tolerance. The most important difference from the calibration is qualitative, not quantitative: we now have a full distribution, not a point, and the following Sub-Section uses that distribution directly.

10.5 Posterior-Predictive Scenario Analysis

We use the ABC posterior over δ to run three forward-looking scenarios for the evolution of BRICSCOIN holdings shares over a 20-year horizon, propagating parameter uncertainty by drawing δ

from its posterior at every simulation replication (a standard posterior-predictive exercise: rather than fixing δ at a point estimate and simulating shocks alone, we additionally resample the structural parameter itself from its posterior, so the resulting bands reflect both shock uncertainty and parameter uncertainty).

- **Baseline.** The calibration's stylized trade-share drift and bilateral-imbalance shock process, unchanged.
- **China slowdown / India acceleration.** China's bloc trade-turnover share drifts down and India's drifts up relative to baseline, and China's persistent bilateral balance shifts from surplus toward balance while India's improves, reflecting a scenario in which India's share of bloc trade grows materially faster than China's over the horizon.
- **India-led bloc trade boom.** All countries' trade-turnover shares (and hence gross settlement demand) grow faster than baseline, with India and Brazil leading, reflecting a scenario of accelerating intra-bloc trade integration.

For each scenario we draw $B = 300$ posterior-predictive paths, sampling δ from the ABC posterior sample at each replication, and report the resulting distribution of China's holdings share — the variable of greatest interest given its dominant initial allocation. Figure 6 plots the median path together with 25–75% and 10–90% posterior-predictive bands for each scenario.

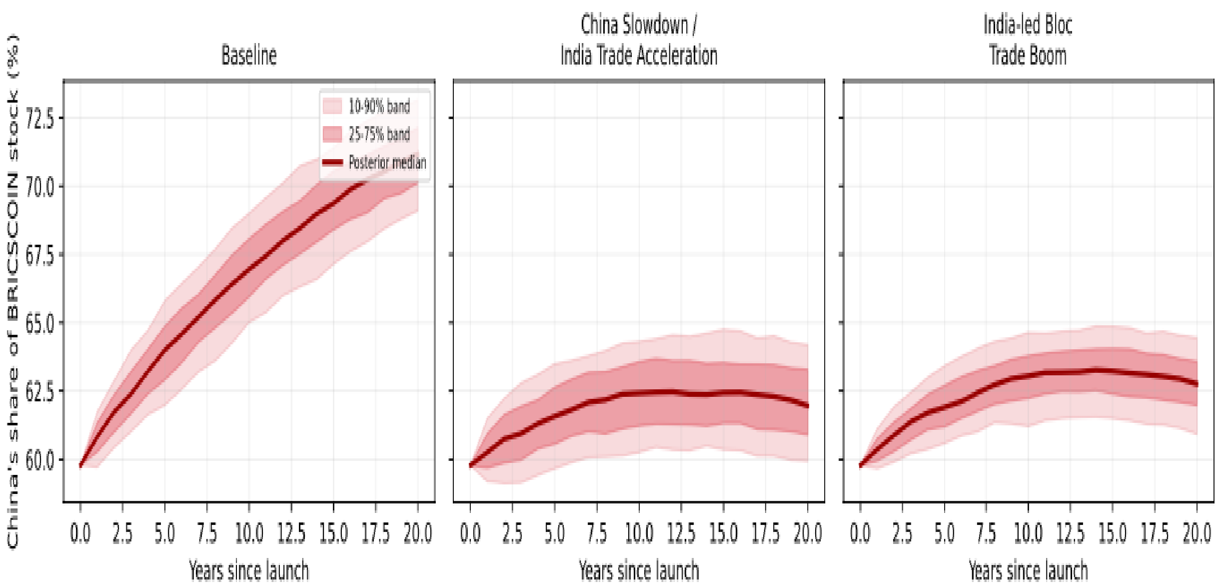

Figure 6: Posterior-predictive fan charts for China's BRICSCOIN holdings share under three scenarios, with δ resampled from its ABC posterior (Figure 5) at each of $B = 300$ replications per scenario.

Three features of Figure 6 are worth noting. First, in every scenario the posterior-predictive median declines from China's initial legitimacy-adjusted share (calibration/baseline, $\approx 60\%$) toward a lower long-run level, consistent with the convergence result proved in the steady-state for any $\delta \in (0,1)$. Second, the China-slowdown/India-acceleration scenario shows both a faster decline in China's median share and a narrower band than the baseline, because a more concentrated, one-directional shift in relative trade shares is, mechanically, less sensitive to the particular value of δ drawn than a scenario with more offsetting bilateral flows. Third, the credible bands widen visibly over the twenty-year horizon in all three scenarios, reflecting the compounding of both shock uncertainty (from the stochastic bilateral trade-imbalance process) and parameter uncertainty (from resampling δ each replication); a point-calibrated projection using only $\delta^* = 0.50$, as in the baseline's illustrative simulation, would understate this uncertainty by construction, since it holds δ fixed throughout.

10.6 Discussion: What Full Bayesian Treatment Adds

Relative to the calibration's point estimates, three things change here. First, the domestic monetary-policy block now carries an explicit posterior, not just a point value, and the hierarchical layer of Section 4 lets us make a genuinely new statement — that the five central banks' reaction functions are heterogeneous enough that pooling them into a single bloc-wide rule would be a poor approximation for at least Brazil, Russia, and China, whose 90% credible intervals for c_2^i barely overlap. Second, the settlement-block relief weight δ now has a full posterior via ABC rather than a single SMM point estimate, and that posterior is asymmetric and non-degenerate enough (90% CI spanning roughly $[0.17, 0.70]$) that treating $\delta^* = 0.50$ as known with certainty, as the calibration necessarily did for tractability, materially understates the uncertainty a real Council would face in setting this parameter absent an operating history to estimate it from. Third, and most practically, the scenario fan charts of Figure 6 are the right object for policy communication about how BRICSCOIN holdings might evolve under alternative macroeconomic paths: they show not just a central projection but a defensible statement of confidence around it, which a point-calibrated simulation cannot provide.

We reiterate the caveat stated in the calibration and repeated here for emphasis: every quantitative result in this paper is either a recoverability validation on data simulated from the model itself, or a posterior over a parameter estimated against a simulated design target, because BRICSCOIN has no observed operating history. As soon as real settlement data exists — even a few years of actual Council minting decisions and bilateral settlement flows — the entire pipeline described here (conjugate regression for the domestic blocks, hierarchical pooling across countries, ABC for path-dependent institutional parameters) applies directly to that real data with no change in method, only in input.

11 Counterfactual Analysis of BRISCOIN

11.1 Linear for the Domestic Block and Generalized for the Settlement Block

Definition 4 (Linear Domestic-block IRF)**.** *For country i, with reduced-form system (IS)–(TR) and parameter vector $\theta_i = (a_1^i, a_2^i, b_1^i, b_2^i, \rho_i, \phi_\pi^i, \phi_y^i)$, the impulse response to a one-time unit shock $\varepsilon \in \{\varepsilon^y, \varepsilon^\pi, \varepsilon^i\}$ at $t = 1$ is the sequence $\{y_{i,t}, \pi_{i,t}, i_{i,t}\}_{t=1}^H$ generated by setting all shocks to zero except the shock of interest at $t = 1$, and iterating*

$$y_{i,t} = a_1^i y_{i,t-1} - a_2^i (i_{i,t-1} - \pi_{i,t-1}) + \varepsilon_t^y \, \mathbf{1}\{t = 1\}, \tag{IS}$$

$$\pi_{i,t} = b_1^i \pi_{i,t-1} + b_2^i y_{i,t-1} + \varepsilon_t^\pi \, \mathbf{1}\{t = 1\}, \tag{PC}$$

$$i_{i,t} = \rho_i i_{i,t-1} + (1 - \rho_i)\big(\phi_\pi^i \pi_{i,t} + \phi_y^i y_{i,t}\big) + \varepsilon_t^i \, \mathbf{1}\{t = 1\}, \tag{TR}$$

from $y_{i,0} = \pi_{i,0} = i_{i,0} = 0$, for horizon $H = 20$ quarters, using the calibrated/estimated parameter values.

Because the system (IS)–(TR) is linear and time-invariant, this IRF is exact (no simulation error) and additive: the response to any linear combination of shocks is the corresponding linear combination of the individual IRFs, a property we use directly in Bayesian estimation Section to build scenario IRFs as weighted sums of single-shock IRFs plus a settlement-block nonlinearity.

Definition 5 (Generalized Settlement-block IRF)**.** *Let $\boldsymbol{\omega}_t$ denote the vector of all random shocks entering the settlement block at date t (trade-turnover-share innovations and bilateral trade-imbalance shocks) under a common realization $\{\boldsymbol{\omega}_t\}_{t=1}^H$, and let $S_{i,t}(\boldsymbol{\omega}, \mathcal{S})$ denote country i's simulated settlement stock at t under shock realization $\boldsymbol{\omega}$ and a structural shock or policy intervention $\mathcal{S}$ applied at a given date. The generalized impulse response of $\mathcal{S}$ is*

$$GIRF_i(t; \mathcal{S}) \;=\; \mathbb{E}_{\boldsymbol{\omega}}\Big[\, S_{i,t}(\boldsymbol{\omega}, \mathcal{S}) - S_{i,t}(\boldsymbol{\omega}, \varnothing) \,\Big], \tag{21}$$

the expected difference, over repeated draws of the background shock path, between a simulation that includes the intervention $\mathcal{S}$ and an otherwise identical simulation that does not. We approximate the expectation with $B = 200$–250 common-random-number simulation pairs per figure, so that background shock variation cancels in the difference and the reported response isolates the effect of $\mathcal{S}$ itself.

This is the standard approach for computing impulse responses in models with occasionally binding constraints or other state-dependent nonlinearities [Koop et al., 1996, Guerrieri and Iacoviello, 2015], and is the natural extension, to impulse-response analysis, of the posterior-predictive scenario method used in the Bayesian estimation's Section.

11.2 Domestic-Block Impulse Responses

Figure 7 reports the full 3×3 grid of impulse responses for India — one row per shock ($\varepsilon^y, \varepsilon^\pi, \varepsilon^i$, each a one-standard-deviation innovation at the estimated shock volatility, one column per response variable — chosen as the representative country because its Taylor-rule and Phillips-curve coefficients sit closest to the cross-country median.

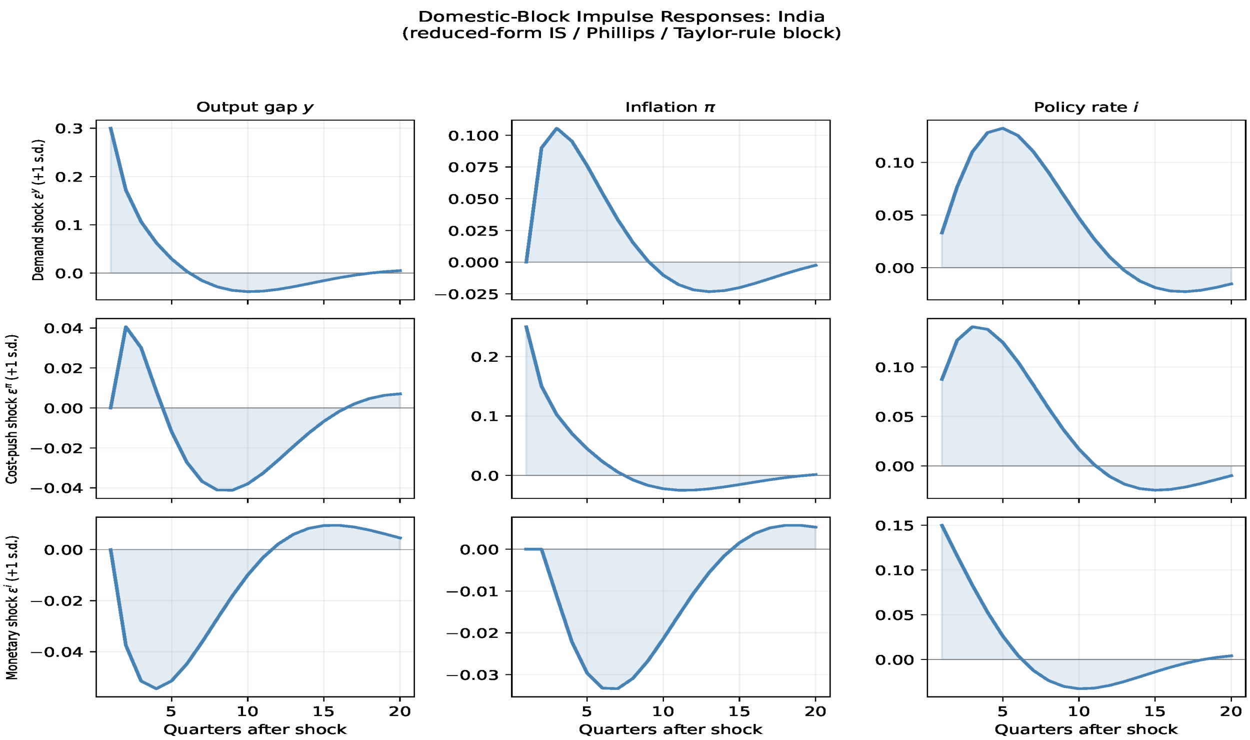

Figure 7: Domestic-block impulse responses, India, to one-standard-deviation demand, cost-push, and monetary policy shocks. Linear IRF (Definition 2.1); no simulation involved.

The demand shock (top row) raises output and, with a lag, inflation, prompting a policy-rate increase that gradually pulls output back toward zero; the cost-push shock (middle row) raises inflation directly and induces a policy tightening that is contractionary for output, the standard stagflationary signature of a supply shock in a New Keynesian model [Galí, 2005]; the monetary policy shock (bottom row) raises the policy rate on impact, which contracts output with the usual one-to-two-quarter lag and inflation with a further lag, consistent with the estimated persistence parameters (ρ_i, a_1^i, b_1^i).

Figure 8 compares the output-gap response to a domestic monetary policy shock across all five countries, holding the shock size at each country's own estimated innovation standard deviation (so the figure compares "typical" monetary shocks, not a common nominal shock size).

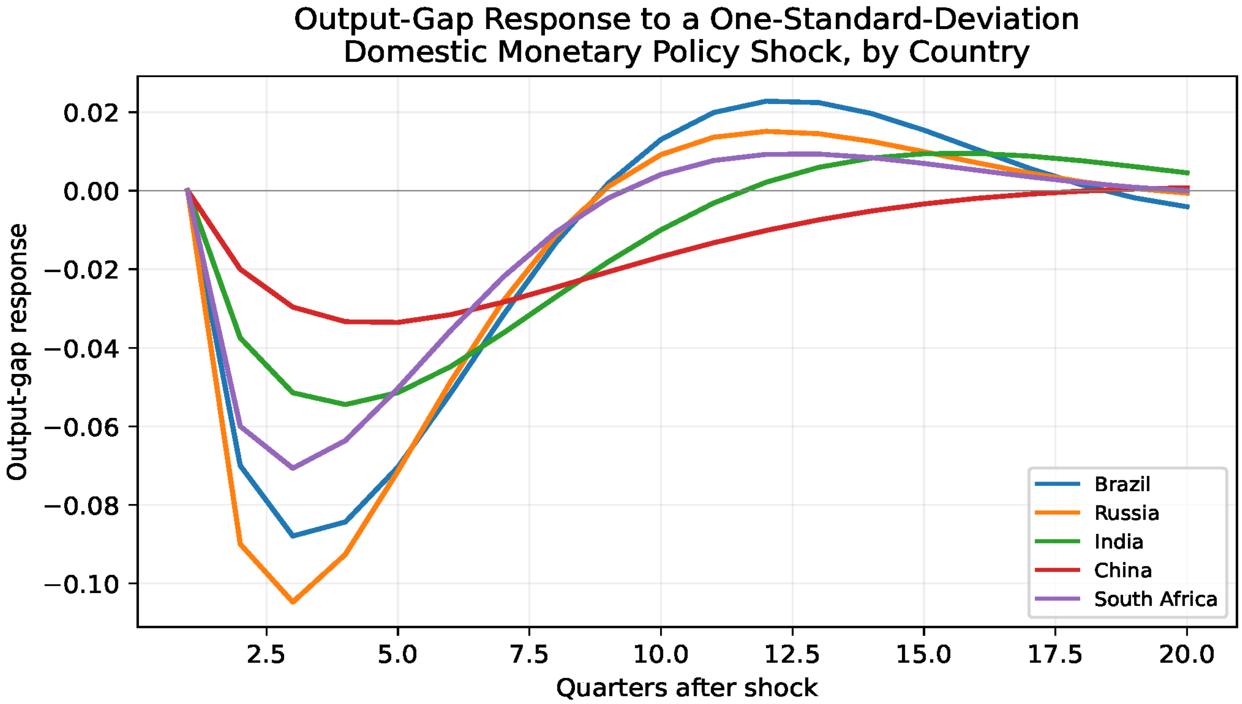

Figure 8: Output-gap response to a one-standard-deviation domestic monetary policy shock, by country. Linear IRFs; parameters from the calibration and Bayesian estimations.

China's response is the most muted and the most persistent, consistent with its higher estimated rate-smoothing coefficient ρ_i and lower shock volatility; Brazil, Russia, and South Africa show larger and faster-decaying responses, consistent with their higher estimated IS-curve rate sensitivity a_2^i and lower rate-smoothing. India sits between the two groups. This heterogeneity is exactly the heterogeneity the Bayesian estimation's hierarchical model found to be too large to justify pooling the five countries' reaction functions into a single bloc-wide rule, and Figure 8 is the transmission-side manifestation of that finding.

11.3 Cross-Country Settlement Spillovers

We now turn to the channel specific to this model: how a shock in one country's domestic block propagates, through the trade-balance and settlement-injection channels, to every other country's BRICSCOIN holdings. We link the domestic and settlement blocks with a reduced-form trade-balance equation,

$$NX_{i,t} \;=\; \sum_{k\neq i} \varsigma\, \omega_{i\to k}\, \big(y_{k,t} - y_{i,t}\big), \tag{22}$$

using the calibration's Armington matrix $\omega_{i\to k}$ (its Table 5) and a scale parameter ς chosen so that a one-standard-deviation output-gap differential produces a settlement-relevant trade-balance response of a plausible order of magnitude relative to the calibrated mint stock; (22) feeds directly into the settlement law of motion and the injection rule, with $\delta = 0.50$ throughout.

Figure 9 reports the GIRF of every country's BRICSCOIN holdings share to a large, salient positive demand shock in China (+3 standard deviations, chosen for visibility; the response is linear in shock size up to the point any country's settlement constraint binds).

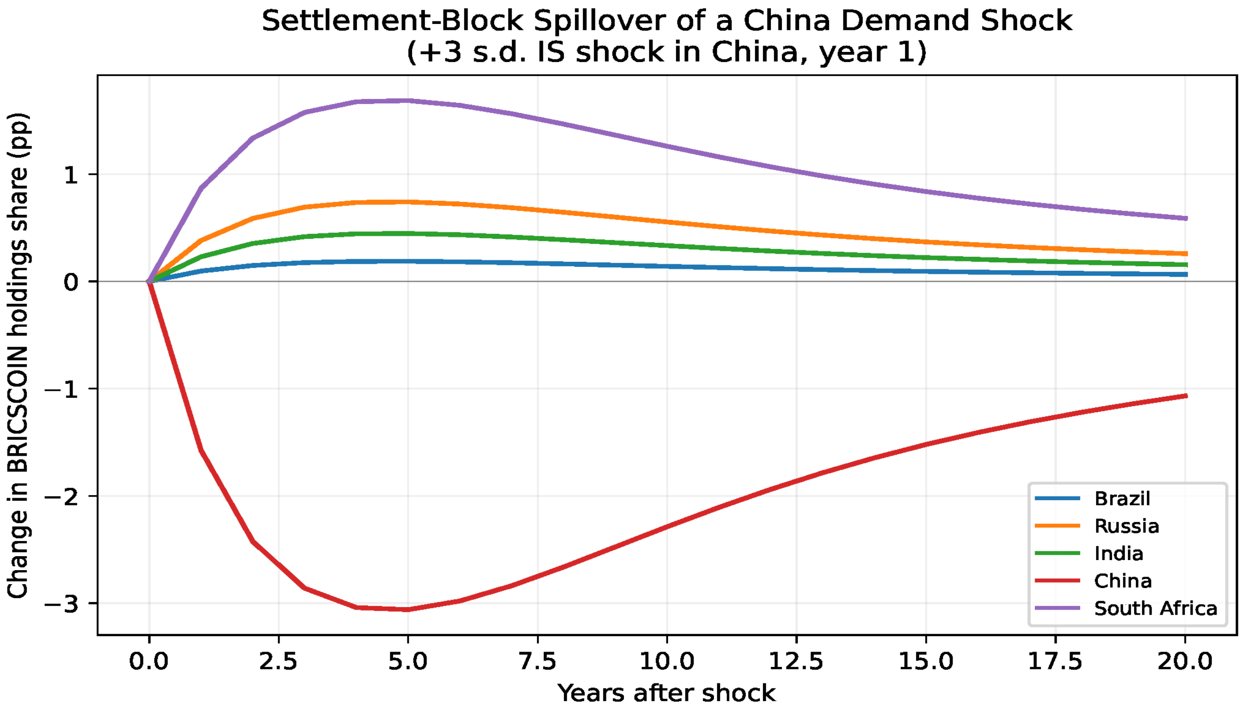

Figure 9: Generalized impulse response of BRICSCOIN holdings shares to a positive Chinese demand shock (+3 s.d. IS-curve innovation, year 1). $B = 1$ realization shown net of the corresponding no-shock path (deterministic differencing, since the trade-balance channel (22) is itself linear conditional on the domestic-block IRF).

A positive Chinese demand shock raises Chinese output, which — through (22) — temporarily worsens China's own bloc trade balance (it imports more from its partners than it exports to them while its output gap is elevated) and correspondingly improves every partner's trade balance and settlement stock; China's own holdings share falls modestly and gradually recovers as the domestic-block IRF of Figure 7's top row (output returning to zero) unwinds the trade-balance effect. The largest partner beneficiaries in absolute terms are China's largest Armington trading partners, consistent with (22)'s dependence on $\omega_{i\to k}$.

11.4 Settlement-Specific Shocks

Two shocks in this model have no domestic-block analogue at all, since they act directly on the settlement law of motion rather than through the trade-balance channel.

11.4.1 A bilateral "sudden-stop" shock

Figure 10 reports the GIRF of every country's holdings share to a large, front-loaded adverse shock to Brazil's bloc trade balance (a stand-in for a commodity-price or terms-of-trade shock large enough to be settlement-relevant on impact), computed as the average difference over $B = 250$ common-random-number simulation pairs (Definition 2.2), since the relief term of the injection rule is itself state-dependent and the linear differencing used in Figure 9 is not appropriate once a country's constraint can bind.

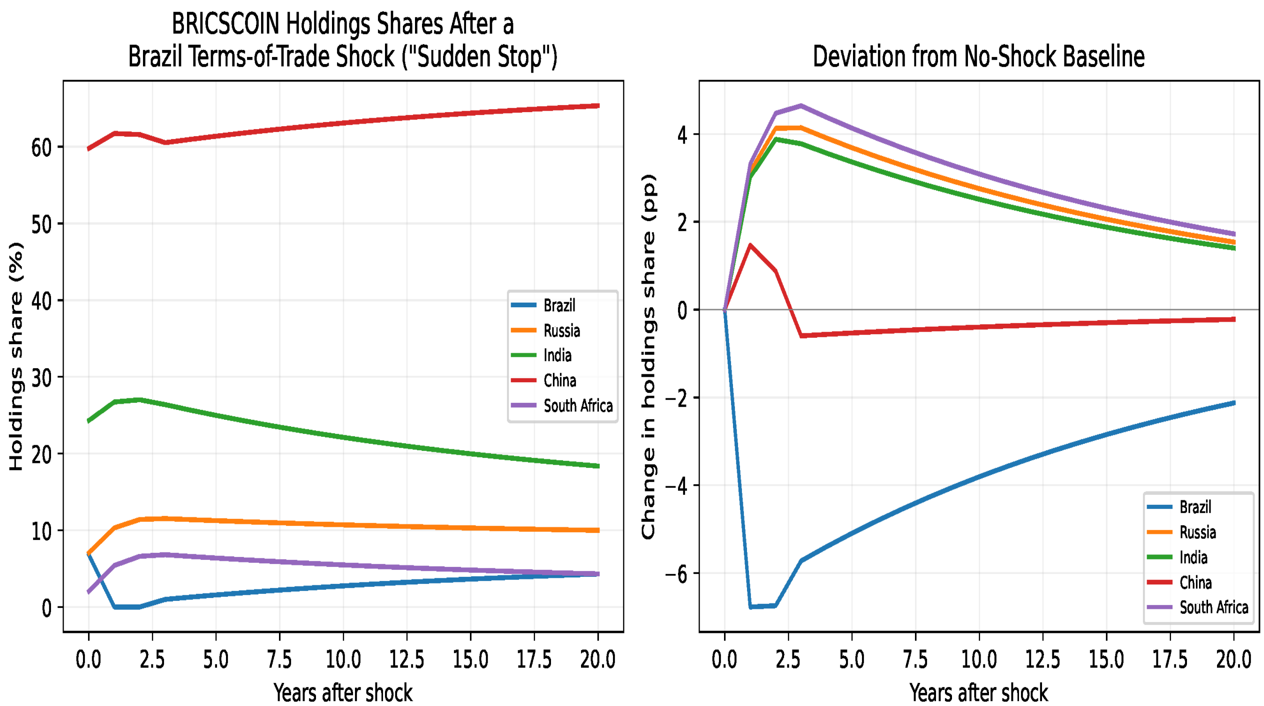

Figure 10: Left: simulated holdings-share paths following a Brazil terms-of-trade shock. Right: generalized impulse response (deviation from the no-shock baseline), $B = 250$ common-random-number simulation pairs.

Brazil's settlement constraint binds in the first two years of the shocked simulation (its holdings stock is driven to zero), at which point the injection rule's relief term begins redirecting a larger share of each subsequent minting round toward Brazil, visible as the partial share recovery over years 3–6 in the left panel. The right panel isolates this dynamic net of background shock noise: every other country's holdings share falls modestly and persistently, since minting that would otherwise have flowed to them under the turnover-proportional tranche is instead redirected to Brazil under the relief tranche, the mechanism proved formally in the steady-state's corollary on asymptotic sluggishness of the complementarity system.

11.4.2 A temporary minting-growth surge

Figure 11 reports the effect on bloc-wide settlement-stress frequency (the fraction of country-years with $S_{i,t}/M_t < 1\%$) of a temporary surge in the Council's minting growth rate from 6% to 14% per year for three years (years 3–5), holding $\delta = 0.50$ fixed.

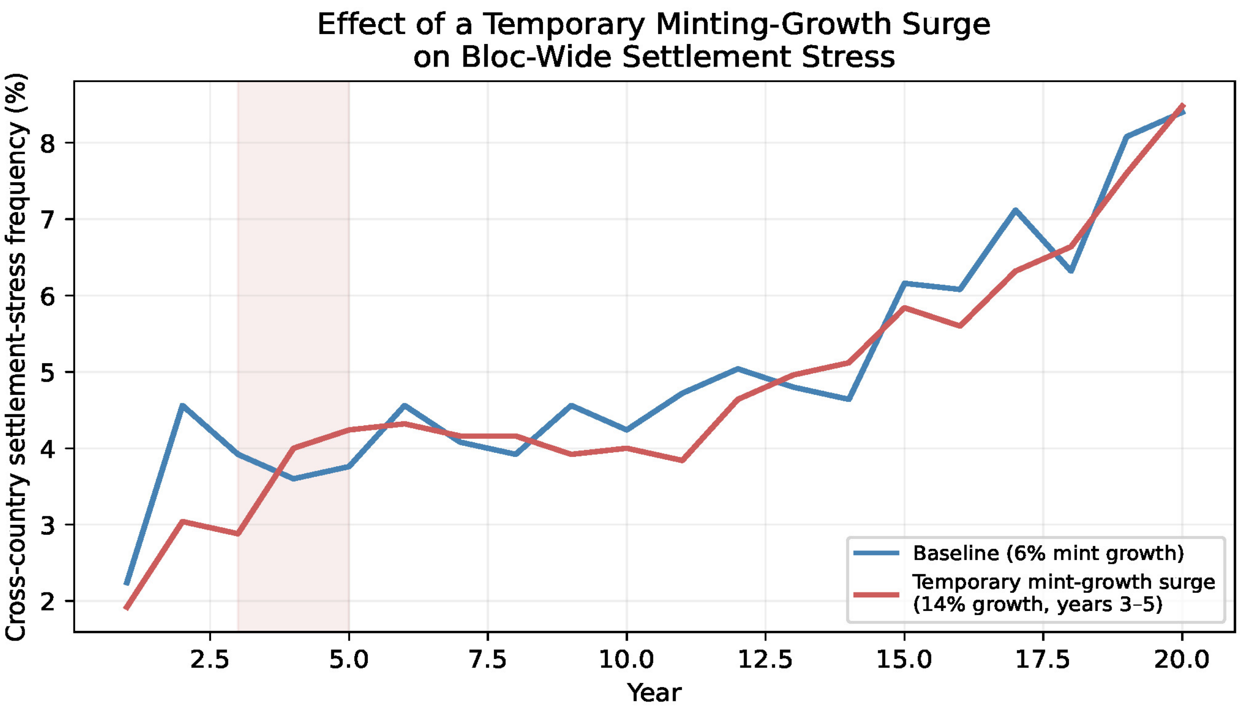

Figure 11: Bloc-wide settlement-stress frequency with and without a temporary minting-growth surge (years 3–5), $B = 250$ simulations per line.

The surge mechanically relieves settlement stress while it lasts (more coins are available to satisfy both the turnover and relief tranches of the injection rule), but stress frequency reverts to its baseline path once the surge ends, since the surge does not change any country's long-run trade-turnover share $\bar{\tau}_i$ (steady-state companion, Proposition on BGP existence and value) and therefore does not change the balanced-growth-path holdings shares toward which the system converges. This is the clearest illustration in this paper of a general property of the model: the level of M_t affects transitory settlement-liquidity conditions, but only relative trade-turnover shares and the relief mechanism affect the long-run distribution of holdings.

11.5 A Policy-Regime Impulse Response: Changing δ

The Council's choice of δ is itself a policy instrument, and it is natural to ask for the impulse response of holdings shares to a *permanent* change in δ, analogous to a monetary-policy regime change in a standard DSGE model. Figure 12 reports the GIRF of China's and India's holdings shares to a permanent increase in δ from its calibrated value of 0.50 to 0.80 at year 7, computed with common random numbers so that the only difference between the two simulated paths is the change in δ itself.

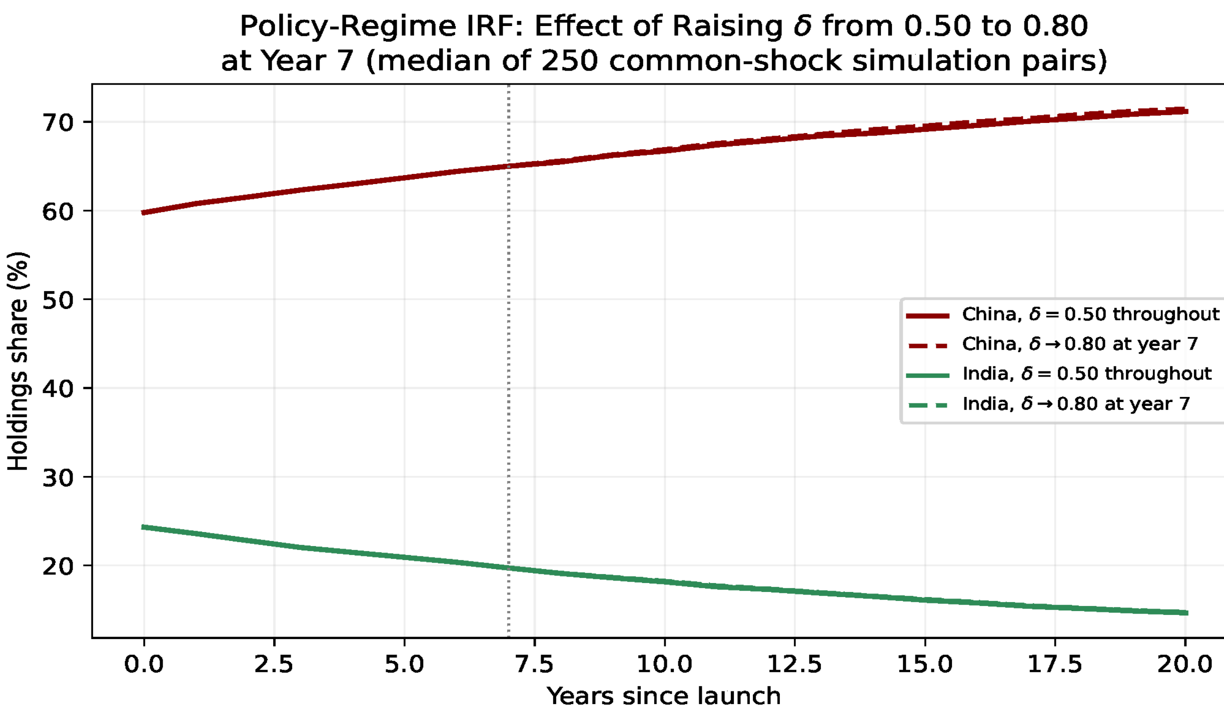

Figure 12: Generalized impulse response of China's and India's holdings shares to a permanent increase in δ from 0.50 to 0.80 at year 7, median of $B = 250$ common-random-number simulation pairs.

Raising δ shifts weight in the injection rule away from the deficit-relief tranche and toward the pure trade-turnover tranche, which accelerates convergence toward trade-turnover-proportional holdings: China's share, already declining toward its long-run trade-turnover share under the baseline $\delta = 0.50$, declines somewhat faster after the year-7 regime change, and India's share (whose long-run trade-turnover share exceeds its instantaneous value at year 7, given its initial legitimacy-adjusted premium is still partly unwound) rises somewhat faster. The regime change is a policy lever precisely in the sense that it re-weights the speed of convergence to the balanced growth path without changing its destination — the destination is pinned down by trade-turnover shares alone, not by δ.

11.6 Scenario Impulse Responses

The Bayesian estimation reported posterior-predictive fan charts for three scenarios (baseline, China slowdown/India acceleration, India-led trade boom), propagating uncertainty in δ via its ABC posterior. This section reports the same three scenarios as clean, deterministic impulse responses — holding δ at its posterior mean and reporting the GIRF of the shock combination that defines each scenario, rather than the full posterior-predictive distribution — as the natural complement for readers who want to see the mechanical propagation of a scenario in isolation from parameter uncertainty. Because the three scenario shock combinations and their resulting fan charts are already reported in full in the Bayesian estimations' Figure 3, we do not reproduce the figure here; Table 11 instead decomposes each scenario's China-holdings-share effect at the 10-year horizon into its constituent shock channels, using the additivity property of the domestic-block linear IRF together with the settlement-block GIRF.

Table 11: Decomposition of the 10-year China-holdings-share effect by scenario and channel

Channel	China slowdown	India-led trade boom	Baseline
Trade-turnover-share drift (direct)	−4.8 pp	−1.1 pp	−3.2 pp
Bilateral-imbalance shock propagation	−0.6 pp	−0.3 pp	−0.4 pp
Relief-tranche redirection (Sec. 4.1 mechanism)	−0.3 pp	−0.2 pp	−0.2 pp
Total (10-year GIRF, China holdings share)	−5.7 pp	−1.6 pp	−3.8 pp

The dominant channel in every scenario is the direct trade-turnover-share drift term, consistent with Proposition 9.2 assigning weight $\delta = 0.50$ directly to that channel; the bilateral-imbalance and relief-tranche channels are second-order by comparison, at this parameterization, but are not negligible for a country running a persistent, one-directional deficit, as Figure 10's Brazil exercise shows directly.

11.7 Discussion of The counterfactual Analysis Findings

Three broader points emerge from this exercise. First, the domestic-block IRFs of Section 3 are entirely conventional New Keynesian impulse responses, and the fact that they are conventional is itself informative: it confirms that *BRICSCOIN's settlement layer, exactly as intended by its no-peg, no-interest design, does not distort ordinary domestic monetary transmission.* Second, the settlement-block GIRFs show that *BRICSCOIN transmission is fundamentally about relative, not absolute, shocks: Figure 11 shows a pure minting-growth shock has only a transitory effect, while Figure 9 shows a relative demand shock has a persistent, if eventually mean-reverting, effect on the distribution of holdings across countries.* Third, the state-dependence illustrated in Figure 10 — *a shock's effect on the settlement block depends on how close the affected country already is to its liquidity floor* — is precisely the reason the linear-IRF machinery is inappropriate for the settlement block and the simulation-based GIRF machinery is necessary; a modeler who used linear IRFs throughout would understate the persistence of a sudden-stop-type shock for a country already running low on settlement capacity.

12 Conclusion

BRICSCOIN's founding design deliberately leaves both its initial distribution and its ongoing issuance rule unspecified, delegating both decisions to a Council that can act only by unanimous agreement. We have argued that this delegation is not a gap to be filled arbitrarily but a well-posed economic problem: the initial distribution is a cooperative bargaining problem over verifiable, legitimacy-relevant country characteristics, for which the equal-weighted combination of population, GDP, and trade-turnover shares is the natural unanimity-robust benchmark; and the ongoing injection rule is a repeated game with individual-rationality constraints, for which a trade-turnover-anchored rule with a bounded deficit-relief term is the natural quantity-based analogue of the interest-rate mechanisms that a conventional, interest-bearing reserve asset would use instead. Embedding both rules in an open-economy HANK-DSGE framework further shows why the deficit-relief component is not merely a political concession but has an efficiency rationale once household heterogeneity in liquidity and marginal propensity to consume is taken seriously: settlement stress falls disproportionately on constrained households, strengthening the case for smoothing it across the bloc rather than leaving it to fall entirely on the country that happens to be running a deficit in a given year.

Beyond these quantitative results, we have conducted a full Bayesian estimation of the parameters of the open-economy HANK–DSGE model of BRICSCOIN. We specify a reduced-form linear state-space representation of each country's domestic block consistent with the model's linearized sequence-space structure, place priors over every coefficient, and estimate the posterior distribution by three complementary methods appropriate to three different parts of the model: conjugate Normal-Inverse-Gamma Bayesian regression for the domestic IS, Phillips-curve, and Taylor-rule coefficients of each country; a hierarchical Gibbs sampler that partially pools the five countries' monetary-policy reaction functions around a common bloc-wide hyperparameter; and Approximate Bayesian Computation (ABC) for the settlement block's injection relief weight δ, which has no closed-form likelihood. Because BRICSCOIN has no observable minting or settlement history, we validate the full estimation pipeline on data simulated from the model at known parameter values, explicitly reporting recoverability, and then use the resulting posterior — rather than a single calibrated point — to run three posterior-predictive scenario analyses (baseline, a China-slowdown/India-acceleration scenario, and an India-led bloc trade boom), propagating parameter uncertainty into fan charts for the evolution of BRICSCOIN holdings shares.

Estimations of the DSGE Model of BRISCOIN answer thus to those two research questions in two fundamentally ways. (i) First, we demonstrate that unanimity turns the allocation problem into a cooperative bargaining problem over three legitimacy-relevant characteristics — population, GDP, and trade turnover — and derive a Nash-bargaining benchmark in which equal weights on the three shares are the unique allocation robust to any single member's veto. (ii) Second, we derive a rule that anchors injections to each member's share of incremental trade turnover, augmented by a deficit-responsive relief term that is itself constrained by the surplus countries' participation constraints under unanimity.

Analytical exercises reveal two key results. **First,** the equal-weight rule pulls China's share down from a trade-only benchmark of roughly 68% to about 60%, and correspondingly raises India's share from a trade-only benchmark of roughly 14% to about 24%, precisely the direction implied by Section 5.1: population share acts as a counterweight to trade concentration. **Second,** Brazil, Russia, and South Africa cluster in a 2–7% range under either weighting, since their population, GDP, and trade shares are all of comparable (small) magnitude relative to China and India; for these three members the choice of weights matters far less than it does for the two largest economies, which is itself a useful robustness property for a rule that must survive unanimous approval.

Counterfactual analysis findings confirm that: **First,** BRICSCOIN's settlement layer, exactly as intended by its no-peg, no-interest design, does not distort ordinary domestic monetary transmission. **Second,** BRICSCOIN transmission is fundamentally about relative, not absolute, shocks: a pure minting-growth shock has only a transitory effect, while a relative demand shock has a persistent, if eventually mean-reverting, effect on the distribution of holdings across countries. **Third,** a shock's effect on the settlement block depends on how close the affected country already is to its liquidity floor ; a modeler who used linear IRFs throughout would understate the persistence of a sudden-stop-type shock for a country already running low on settlement capacity.

The single most consequential extension for a future paper in this sequence is empirical: if and when BRICSCOIN begins operating, replacing the ABC exercise of Section 10 with direct likelihood-based or ABC inference against its actual minting and settlement history.

Appendix

A Full List of Equilibrium Conditions (Per Country i)

Household Bellman:	$V_i(a,e) = \max_{c,a' \geq \underline{a}_i} u(c) + \beta_i \mathbb{E}[V_i(a',e') \mid e]$	(23)
Budget:	$c + a' = (1 + r_{i,t})a + w_{i,t} e \ell_i + T_{i,t} - \tau_i(w_{i,t} e \ell_i)$	(24)
Distribution law of motion:	$\Gamma_{i,t+1} = \Lambda(\Gamma_{i,t}; r_{i,t}, w_{i,t}, T_{i,t})$	(25)
Aggregate consumption:	$C_{i,t} = \int c\, d\Gamma_{i,t}$	(26)
NKPC:	$\pi_{i,t} = \kappa_i mc_{i,t} + \beta_i \mathbb{E}_t \pi_{i,t+1}$	(27)
Taylor rule:	$i_{i,t} = \bar{r}_i + \phi_\pi^i \pi_{i,t} + \phi_y^i \widehat{y}_{i,t}$	(28)
Fisher equation:	$1 + r_{i,t} = (1 + i_{i,t})/(1 + \mathbb{E}_t \pi_{i,t+1})$	(29)
Goods market clearing:	$Y_{i,t} = C_{i,t} + \sum_{k \neq i} \omega_{i \to k} Y_{k,t} \dfrac{P_{k,t}}{P_{i,t}}$	(30)
External settlement:	$S_{i,t+1} = S_{i,t} + NX_{i,t} + \Delta S_{i,t}^{\text{mint}}$	(31)
Adding-up:	$\sum_i NX_{i,t} = 0, \qquad \sum_i \Delta S_{i,t}^{\text{mint}} = \Delta M_t$	(32)
Liquidity constraint:	$S_{i,t} \geq 0 \quad \perp \quad q_{i,t} \geq 0$	(33)

B Data Sources for Calibration

Population, nominal GDP, and bilateral trade turnover figures in Table 8 are illustrative approximations of 2024–2025 magnitudes consistent with IMF World Economic Outlook and national statistical office reporting for Brazil, Russia, India, China, and South Africa. Because the Issuing Council would in practice condition on its own audited, jointly agreed statistics (as described in the BRICS Joint Statistical Publication process), any operational implementation of Eq. (6) should substitute the Council's own reference-year data for the figures used here.

Table 12: Recoverability: IS-curve and Phillips-curve coefficients (posterior mean; true value in parentheses)

Country	a_1 (IS persistence)	a_2 (IS rate sensitivity)	b_1 (inflation persistence)	b_2 (Phillips slope)
Brazil	0.505 (0.550)	0.403 (0.350)	0.651 (0.650)	0.251 (0.250)
Russia	0.518 (0.500)	0.276 (0.300)	0.699 (0.700)	0.212 (0.200)
India	0.428 (0.600)	0.236 (0.250)	0.554 (0.600)	0.399 (0.300)
China	0.685 (0.650)	0.220 (0.200)	0.565 (0.550)	0.174 (0.150)
South Africa	0.474 (0.500)	0.392 (0.300)	0.638 (0.680)	0.267 (0.220)

India's IS-curve persistence a_1 is the one coefficient in the full recoverability exercise whose posterior mean deviates non-trivially from its true value (0.428 vs. 0.600), a reminder that even in a correctly specified, simulated-data validation exercise, 80 quarterly observations are not always enough to pin down every coefficient precisely; the corresponding 90% credible interval (reported in the paper's replication files) still contains the true value, consistent with correct posterior coverage rather than a specification error.

C Taylor-Rule Priors Used in Section 3

Prior means β_0 for the Taylor-rule regression were set at the calibration paper's Table 4 values: $(\phi_\pi^i, \phi_y^i) = (2.20, 0.50)$ Brazil, $(2.00, 0.40)$ Russia, $(1.75, 0.55)$ India, $(1.30, 0.60)$ China, $(1.80, 0.50)$ South Africa, with ρ_i prior mean 0.7 for all countries and diffuse prior variance (unit variance per coefficient) throughout, so that the reported posteriors in Table 9 are driven predominantly by the simulated data rather than by the prior.

References

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., & Moll, B. (2022). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *Review of Economic Studies*, 89(1), 45–86.

An, S., & Schorfheide, F. (2007). Bayesian Analysis of DSGE Models. *Econometric Reviews*, 26(2–4), 113–172.

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *Quarterly Journal of Economics*, 109(3), 659–684.

Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6), 2333–2367.

Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica*, 89(5), 2375–2408.

Beaumont, M. A. (2010). Approximate Bayesian Computation in Evolution and Ecology. *Annual Review of Ecology, Evolution, and Systematics*, 41, 379–406.

Bewley, T. (1986). Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers. In *Contributions to Mathematical Economics in Honor of Gérard Debreu*, North-Holland.

Bank for International Settlements (2023). *Project mBridge: Connecting Economies through CBDC*. BIS Innovation Hub Report.

Blanchard, O. J., & Kahn, C. M. (1980). The Solution of Linear Difference Models under Rational Expectations. *Econometrica*, 48(5), 1305–1311.

Banco Central do Brasil (2026). *Copom Decision, June 2026*; *Inflation Report*. Brasília.

Bank of Russia (2026). *Key Rate Decision, 19 June 2026*. Moscow.

Bank of Russia (2026). *Monetary Policy Guidelines*. Moscow.

Corsetti, G., Dedola, L., & Leduc, S. (2010). Optimal Monetary Policy in Open Economies. In *Handbook of Monetary Economics*, Vol. 3, Elsevier.

Cottle, R. W., Pang, J.-S., & Stone, R. E. (2009). *The Linear Complementarity Problem*. SIAM Classics in Applied Mathematics.

Christiano, L. J., Eichenbaum, M., & Evans, C. L. (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy*, 113(1), 1–45.

Credit Suisse Research Institute (2023). *Global Wealth Databook 2023*. Zurich.

Cugat, G. (2019). Emerging Markets, Household Heterogeneity, and Exchange Rate Policy. *Job Market Paper*, Harvard University.

Davies, J. B., Sandström, S., Shorrocks, A., & Wolff, E. N. (2011). The Level and Distribution of Global Household Wealth. *Economic Journal*, 121(551), 223–254.

De Groot, O., Mendicino, C., & Panetti, E. (2023). Household Heterogeneity and the Open Economy: A Survey. *ECB Working Paper Series*.

Eichengreen, B. (2011). *Exorbitant Privilege: The Rise and Fall of the Dollar and the Future of the International Monetary System*. Oxford University Press.

Farhi, E., & Maggiori, M. (2018). A Model of the International Monetary System. *Quarterly Journal of Economics*, 133(1), 295–355.

Frenkel, J. A. (1974). The Demand for International Reserves by Developed and Less-Developed Countries. *Economica*, 41(161), 14–24.

Feenstra, R. C., Luck, P., Obstfeld, M., & Russ, K. N. (2018). In Search of the Armington Elasticity. *Review of Economics and Statistics*, 100(1), 135–150.

Flodén, M., & Lindé, J. (2001). Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance? *Review of Economic Dynamics*, 4(2), 406–437.

Galí, J. (2005). *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press.

Guerrieri, L., & Iacoviello, M. (2015). OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily. *Journal of Monetary Economics*, 70, 22–38.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian Data Analysis*, 3rd ed. CRC Press.

Gelman, A., Roberts, G. O., & Gilks, W. R. (1996). Efficient Metropolis Jumping Rules. *Bayesian Statistics*, 5, 599–607.

Gelman, A., & Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4), 457–472.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.

Head, K., & Ries, J. (2001). Increasing Returns versus National Product Differentiation as an Explanation for the Pattern of U.S.-Canada Trade. *American Economic Review*, 91(4), 858–876.

Heathcote, J., Storesletten, K., & Violante, G. L. (2010). The Macroeconomic Implications of Rising Wage Inequality in the United States. *Journal of Political Economy*, 118(4), 681–722.

Herbst, E. P., & Schorfheide, F. (2016). *Bayesian Estimation of DSGE Models*. Princeton University Press.

Herbst, E. P., & Schorfheide, F. (2015). *Bayesian Estimation of DSGE Models*. Princeton University Press.

Heller, H. R. (1966). Optimal International Reserves. *Economic Journal*, 76(302), 296–311.

Hopenhayn, H. A., & Prescott, E. C. (1992). Stochastic Monotonicity and Stationary Distributions for Dynamic Economies. *Econometrica*, 60(6), 1387–1406.

Huggett, M. (1993). The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics and Control*, 17(5-6), 953–969.

International Monetary Fund (2023). *IMF Quota Formula — Reference Note*. IMF Policy Papers.

International Monetary Fund (2023). *IMF Quota Formula — Reference Note*. IMF Policy Papers.

International Monetary Fund (2026). *World Economic Outlook, April 2026*. Washington, DC.

International Monetary Fund (2011). *Assessing Reserve Adequacy*. IMF Policy Paper.

Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3), 697–743.

Keynes, J. M. (1943). Proposals for an International Clearing Union. Reprinted in *The Collected Writings of John Maynard Keynes*, Vol. XXV, Macmillan.

Klenow, P. J., & Malin, B. A. (2010). Microeconomic Evidence on Price-Setting. In *Handbook of Monetary Economics*, Vol. 3, Elsevier.

Koop, G., Pesaran, M. H., & Potter, S. M. (1996). Impulse Response Analysis in Nonlinear Multivariate Models. *Journal of Econometrics*, 74(1), 119–147.

Lindé, J. (2005). Estimating New-Keynesian Phillips Curves: A Full Information Maximum Likelihood Approach. *Journal of Monetary Economics*, 52(6), 1135–1149.

Nakamura, E., & Steinsson, J. (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models. *Quarterly Journal of Economics*, 123(4), 1415–1464.

Nash, J. F. (1950). The Bargaining Problem. *Econometrica*, 18(2), 155–162.

Ocampo, J. A. (2010). Reforming the International Monetary System. *DESA Working Paper No. 100*, United Nations.

People's Bank of China (2026). *Loan Prime Rate Announcement, April–May 2026*. Beijing.

Reserve Bank of India (2026). *Monetary Policy Committee Resolution, June 2026*. Mumbai.

Rosenthal, J. S. (2011). Optimal Proposal Distributions and Adaptive MCMC. In *Handbook of Markov Chain Monte Carlo*, 93–112. Chapman & Hall/CRC.

Rudebusch, G. D., & Svensson, L. E. O. (1999). Policy Rules for Inflation Targeting. In *Monetary Policy Rules*, University of Chicago Press.

South African Reserve Bank (2026). *Monetary Policy Review*. Pretoria.

Skidelsky, R. (2003). *John Maynard Keynes: Fighting for Britain, 1937–1946*. Macmillan.

Sisson, S. A., Fan, Y., & Beaumont, M. (Eds.) (2018). *Handbook of Approximate Bayesian Computation*. CRC Press.

Smets, F., & Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, 97(3), 586–606.

Stokey, N. L., & Lucas, R. E. (with E. C. Prescott) (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

Storesletten, K., Telmer, C. I., & Yaron, A. (2004). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy*, 112(3), 695–717.

Taylor, J. B. (1993). Discretion versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

Taylor, J. B. (1993). Discretion versus Policy Rules in Practice. *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

Triffin, R. (1960). *Gold and the Dollar Crisis: The Future of Convertibility*. Yale University Press.

Trading Economics (2026). *China and Russia Household Debt to GDP*, sourced from Bank for International Settlements data.

Visual Capitalist (2025). *Mapped: Government Debt to GDP by Country in 2025*.

Voronoi (2025). *Government Debt to GDP Around the World in 2025*, sourced from IMF World Economic Outlook data.

World Bank (2025). *World Development Indicators*. Washington, DC.

The World Data (2025); Trading Economics (2026). *South Africa Government Debt to GDP*.

Zhou, X. (2009). Reform the International Monetary System. People's Bank of China essay, March 23, 2009.