Cost of Holding Foreign Exchange Reserves in India

Suryashis Ghosh* Vidushi Balakrishnan†

Abstract

This paper evaluates whether India's foreign exchange reserves (including gold), now exceeding $700 billion, are economically optimal rather than merely adequate. Using quarterly data from 2000Q1 to 2025Q4, we estimate a long-run reserve demand equation via Dynamic OLS and calculate the welfare-maximizing reserve level using the Jeanne-Rancière framework, calibrated with crisis probabilities from an asymmetric complementary log-log model. We find that India's actual reserves have exceeded the optimal benchmark in every sample quarter, with excess holdings reaching $145 billion by 2025, or $191 billion once the Global Financial Safety Net is incorporated. We further show that a two-tranche reserve management framework, separating a core liquidity buffer from a dedicated sovereign wealth fund, would have generated $345 billion in cumulative wealth over 2004-2025 against $141 billion under the status quo, at an annual fiscal cost of holding foreign exchange reserves of roughly $15 billion in 2025 alone. These findings suggest that India's reserve accumulation strategy, while historically prudent, now imposes a quantifiable and avoidable fiscal drag that structural reform could substantially mitigate.

JEL Classification: F31, F32, F33, E58, G11, G23

Keywords: Foreign Exchange Reserves, Optimal Reserves, Sovereign Wealth Funds, Emerging Markets, Reserve Adequacy, Sterilization Cost, Dynamic OLS, Jeanne-Rancière Model, India, Global Financial Safety Net

*This paper is based on the author's Master's dissertation in Economics at Ashoka University. We are deeply grateful to Dr. Prachi Mishra and Dr. Adrian Peralta Alva for their invaluable guidance, insightful comments, and generous support throughout the development of this paper. We are also thankful to Mr. Rajib Das, Mr. Siddhartha Nath, and the Division of International Finance of the Department of Economic and Policy Research, Reserve Bank of India, for their valuable comments. All remaining errors are our own. The views expressed are those of the authors and do not necessarily reflect those of the Reserve Bank of India.

†Contributed to the initial idea and the literature review of this paper and provided guidance during its development.

1 Introduction

India maintains foreign exchange reserves exceeding $700 billion, attracting considerable attention from policymakers, economists, and financial analysts. The Reserve Bank of India notes that these reserves provide import cover equivalent to 11.4 months, positioning India among the world's most financially secure emerging economies by conventional adequacy standards. India's reserve levels are now comparable to or higher than many emerging economies with similar external risk profiles.

Yet beneath this appearance of financial strength lies a critical gap in India's external sector policy debate. Holding large reserves in low-yielding foreign assets implies a potential opportunity cost from alternative uses of funds. The trajectory of India's reserve accumulation reflects responses to successive crises, including the 1991 balance of payments shock, the 1997 Asian financial turbulence, and the 2008 global financial crisis. These episodes reinforced the belief in large reserve buffers as reliable insurance against sudden capital flow reversals. For nearly three decades, this precautionary approach has shaped reserve policy. Today, however, a contemporary reassessment of its economic efficiency appears warranted to balance macroeconomic stability with efficient resource allocation.

This paper addresses that reassessment through three core research questions. First, we ask: **What is the optimal reserve level for India?** To establish our empirical foundation and answer this question, we calibrate India's optimal reserve level using the welfare-based framework developed by Jeanne and Rancière (11), which balances insurance benefits against opportunity costs through explicit welfare maximisation. This calibration is supported by the reserve adequacy assessment methodology presented by Das and Nath (7), confirming that India's observed reserves exceed the levels justified by underlying macroeconomic fundamentals and external vulnerabilities. Second, we investigate: **Can India's excess reserves be deployed more productively than their current allocation in low-return foreign assets?** We simulate a two-tranche reserve management framework tailored specifically to India, separating a core liquidity buffer for the central bank from a dedicated fund for excess capital, and evaluate whether this structure can generate substantial intergenerational wealth without compromising macroeconomic stability. Third, we examine: **What is the true cost of holding India's large foreign exchange reserves?** We address this by constructing a dynamic, year-by-year decomposition of two distinct fiscal burdens: the opportunity cost, measured as the spread between the domestic 91-day Treasury bill rate and the US 3-month Treasury bill rate, and the domestic sterilisation cost, the recurring quasi-fiscal expense the central bank incurs by issuing government securities to absorb excess rupee liquidity. To quantify what India has actually lost by maintaining excess reserves, we further simulate a counterfactual scenario in which only the optimal level of reserves is held and the remaining excess is redeployed into a higher-yielding portfolio.

This paper makes four concrete contributions to the literature. First, we move beyond static adequacy benchmarks by estimating a long-run reserve demand curve for India using Dynamic OLS over 2000Q1 to 2025Q4, correcting for non-stationarity and serial correlation and delivering reliable long-run elasticities that reveal the structural drivers of central bank accumulation. Second, we formally calculate the optimal level of reserves for India using the Jeanne and Rancière (11) welfare maximisation framework, calibrated with time-varying crisis probabilities estimated through a complementary log-log model, replacing the standard Probit approach, which performs poorly on the rare, asymmetric nature of sudden stop events. Third, we simulate a two-tranche reserve management framework specifically designed for India, evaluating whether this structure is financially viable and whether it can generate substantial long-term returns without exposing India to macroeconomic risk. Fourth, we dynamically decompose the full fiscal cost of India's excess reserve accumulation over the last two decades into the opportunity cost and the domestic sterilisation cost, a dimension that adequacy-based frameworks, by design, leave entirely unaddressed.

Our paper finds that, India is persistently over-insured relative to any welfare-based benchmark, with actual reserves exceeding the dynamically calibrated optimal level in every quarter of our sample. The excess reserve position is economically large, reaching $145.02 billion by 2025Q4 under the conservative baseline and $190.7 billion once the Global Financial Safety Net is incorporated at a realistic 50 percent usability assumption. The annual opportunity cost of the total reserves reached $14.624 billion in 2025, and the foregone return from not deploying excess reserves into a diversified sovereign portfolio reached $16.83 billion in 2025 alone. A two-tranche reserve management framework is empirically viable, generating $344.58 billion in cumulative sovereign wealth over the 2004 to 2025 simulation horizon compared to $140.80 billion under the status quo.

The remainder of the paper is organised as follows. Section 2 provides institutional background and stylised facts on India's reserve accumulation. Section 3 gives the literature and Section 4 presents the econometric reserve demand estimation using Dynamic OLS. Section 5 derives the welfare-optimal reserve level using the Jeanne and Rancière (11) framework, including the cloglog crisis probability model and the GFSN-adjusted extension. Section 6 presents the two-tranche sovereign wealth framework and its simulated performance. Section 7 decomposes the fiscal cost of excess accumulation and constructs the counterfactual dividend. Section 8 discusses policy implications and implementation barriers. Section 9 concludes.

2 Institutional Background and Stylised Facts

2.1 Evolution of India's Reserve Accumulation

India's reserve accumulation strategy cannot be understood in isolation from the sequence of external crises that shaped it. The origins of modern Indian reserve policy lie in the balance of payments crisis of 1991. By mid-1991, India's usable foreign exchange reserves had fallen to less than two weeks of import cover as a combination of the Gulf War shock to oil prices, a surge in import demand, and a loss of external financing confidence produced a sudden stop. The government was compelled to pledge gold reserves to the Bank of England and the Union Bank of Switzerland to secure emergency bridge financing, and India drew on IMF resources for the first time in its post-independence history. The 1991 crisis directly preceded the landmark economic reform programme that liberalised the current account and progressively relaxed capital account restrictions, and it established reserve adequacy as a non-negotiable priority of Indian macroeconomic policy. The subsequent accumulation strategy was explicitly designed to prevent a repetition of that episode.

The Asian financial crisis of 1997 to 1998, while not directly destabilising for India owing to its partial capital account restrictions, served as a powerful lesson. Economies with apparently sound fundamentals experienced catastrophic reserve losses and currency collapses within weeks. The consensus that emerged, that reserves must be large enough to absorb sudden capital outflows rather than merely finance trade deficits, reinforced and extended the precautionary logic already established at the Reserve Bank of India. From roughly $25 billion in 2000, the reserve stock rose to over $315 billion by May 2008, driven by sustained capital inflows, active RBI intervention to manage rupee appreciation, and sterilisation through the Market Stabilisation Scheme introduced in 2004.

The global financial crisis of 2008 to 2009 produced the sharpest reserve drawdown in India's modern history. Reserves fell by approximately $70 billion between May 2008 and March 2009 as the RBI defended the rupee and supplied dollar liquidity to domestic banks facing external financing constraints. India avoided the need for an IMF programme, and the speed of the subsequent recovery reinforced institutional conviction that large pre-crisis buffers had been justified. The taper tantrum of 2013 provided the next major test. The Federal Reserve's signal of tapering in May 2013 triggered sharp capital outflows. The rupee depreciated by approximately 15 percent between May and August 2013, and the current account deficit had widened to over 5 percent of GDP in the preceding fiscal year driven by gold and oil imports. Reserves fell by approximately $15 billion over this period. The RBI's response combined reserve deployment with emergency monetary tightening and a concessional dollar swap window for banks. Each episode layered additional institutional conviction onto the precautionary strategy. Following the taper tantrum, reserves rose from approximately $290 billion in 2014

to over $640 billion by early 2021, a pace of accumulation that significantly exceeded what any standard adequacy benchmark required and that motivates the welfare-based analysis at the core of this paper.

2.2 Anatomy of India's External Vulnerability

Having traced the historical path of India's reserve accumulation, we now turn to the structural features of India's external sector that define the underlying demand for a reserve buffer. Although these historical episodes help explain the institutional foundations of India's precautionary reserve accumulation strategy, they do not fully clarify the structural factors that continue to generate demand for exceptionally large reserve holdings in the present period. Even after surpassing conventional reserve adequacy benchmarks, the Reserve Bank of India has continued to accumulate foreign exchange assets at a substantial pace. Understanding this persistence requires a closer examination of the underlying sources of India's external vulnerability and the transmission channels through which global financial shocks affect the domestic economy. India's external vulnerability profile is shaped by three structural features that together define the demand for a reserve buffer. The first and most significant is the volatility of foreign portfolio investment. Unlike FDI, which is relatively stable, portfolio flows into Indian equity and debt markets are sensitive to global risk sentiment, the interest rate differential with advanced economies, and domestic political and economic news. Net FPI has swung from large inflows to sharp outflows across multiple episodes in our sample, with peak quarterly outflows exceeding 1 percent of GDP during severe stress events. This volatility has grown as India's capital markets have deepened and become more accessible to international investors, making the capital account the dominant channel of external vulnerability.

The second structural feature is India's dependence on crude oil imports. India imports approximately 85 percent of its crude oil requirements, making the current account balance highly sensitive to global commodity prices. A $10 per barrel increase in oil prices widens the current account deficit by approximately 0.4 to 0.5 percent of GDP. Reserve drawdowns during external stress episodes are therefore often compounded by simultaneous current account deterioration, creating a dual pressure on the external position. The third feature is the composition of India's short-term external debt, which reached approximately $140 billion by 2025 on a residual maturity basis. A significant share reflects non-resident Indian deposit liabilities and external commercial borrowings, both of which are sensitive to confidence effects and interest rate differentials. The RBI operates a managed float exchange rate regime and intervenes actively to smooth excessive volatility and at times to resist appreciation that could erode export competitiveness. This intervention strategy creates a structural demand for foreign exchange and generates the sterilisation operations whose fiscal cost this paper documents. Our DOLS estimation confirms a statistically significant positive coefficient on the REER undervaluation dummy, consistent with both a precautionary and a mercantilist motive for reserve accumula-

tion.

2.3 Comparative Reserve Adequacy

Having established the institutional and structural context for India's reserve accumulation, this section benchmarks India's holdings against the four standard adequacy metrics most widely used in the international macroeconomics literature: import cover, the Greenspan-Guidotti short-term debt rule, broad money coverage, and the IMF's composite Assessing Reserve Adequacy metric. The detailed quantitative analysis is presented in Figure 1 and the discussion that follows.

The oldest and most widely used traditional indicator is the Import Cover rule, which states that a country should hold enough foreign exchange reserves to pay for at least three months of imports.[1] This rule was developed when emerging markets were primarily vulnerable to current account shocks, such as a sudden collapse in export commodity prices. It tests whether a central bank has enough foreign currency to ensure the country can continue importing essential goods while macroeconomic adjustments are made. As observed in Chart 1, India's actual reserves have consistently outpaced the three-month import requirement over the last two decades. Notably, the gap has visibly widened since 2020, showing that trade protection is no longer the binding constraint on India's reserve demand.

Following the Asian Financial Crisis of 1997, capital account crises became the primary threat to emerging markets, leading to the creation of the Short-Term External Debt metric (also known as the Greenspan-Guidotti rule). This metric dictates that a country's reserves should equal or exceed its short-term external debt maturing within one year. The economic rationale is based on the risk of a sudden stop in capital flows; if foreign investors panic and refuse to roll over short-term loans, a country with 100% coverage can pay off all departing foreign creditors without defaulting. Chart 2 illustrates a dramatic divergence between India's reserves and its short-term external obligations. The central bank's holdings dwarf its short-term debt, indicating that the country is essentially immune to standard rollover risk.

The third traditional metric, **Broad Money**, looks at domestic vulnerabilities by suggesting that central banks hold reserves equal to 20% of broad money (M3). This captures the risk of resident-led capital flight. During a severe crisis of confidence, domestic citizens and businesses may attempt to convert their local bank deposits into foreign currency. Holding 20% of M3 in reserves provides the central bank with enough liquidity to meet sudden surges in demand for foreign currency and prevent a run on the banking system. Chart 3 shows that this comfortable margin has been closing steadily since the early 2010s. Reserves and the 20%-of-M3 benchmark run close together from around 2011 onward, and by the final years of the

[1] See Heller (1966) and IMF reserve adequacy guidelines for early discussions of the import cover benchmark as a traditional measure of external sector liquidity adequacy.

sample the two lines are nearly indistinguishable, with the broad-money benchmark briefly overtaking actual reserves toward the end of the period. This reflects the pace at which India's domestic financial system has deepened relative to the growth of the reserve stock itself: as M3 expands, the absolute rupee-equivalent buffer required to insure against resident-led capital flight expands with it, and reserve accumulation has not consistently kept pace with this domestic benchmark in the way it has with the trade- and debt-based metrics in Charts 1 and 2.

While these traditional rules are intuitive, modern financial crises often involve a combination of export collapses, sudden stops, and capital flight happening simultaneously. To address this, the IMF introduced the **Assessing Reserve Adequacy (ARA)** metric, which creates a composite risk-weighted index. For emerging markets with floating exchange rates, the baseline ARA metric is defined as:

$$ARA = 0.30 \times STED + 0.10 \times OPL + 0.05 \times M3 + 0.05 \times X \tag{1}$$

where $STED$ is short-term external debt, OPL represents other portfolio liabilities, $M3$ is broad money, and X denotes annualised export earnings. It is crucial to note that this ARA metric represents the reserve level a country should have at an absolute minimum to withstand a severe crisis; this is not the optimal level of reserves, which we will calculate later in this paper. The IMF suggests that actual reserves should sit between 100% and 150% of this calculated threshold. Chart 4 shows that India's actual reserves track this recommended band closely for most of the sample, sitting near its upper edge for extended periods and dipping toward, and briefly below, its lower boundary during the most acute stress episodes. Over the full 2000Q1-2025Q4 sample, the average ARA ratio for India is 1.48, meaning the country generally holds close to the midpoint of the IMF's recommended adequacy range rather than a large multiple of it. In only 2 of the 104 sample quarters does the ratio fall marginally below the 100% floor, both occurring during periods of acute external stress, and by 2025Q4 the ratio stands at 1.38.

The comparative evidence presents a more nuanced picture than the trade- and debt-based benchmarks alone would suggest. While India's reserves sit far above the import cover and Greenspan-Guidotti thresholds, and the broad money benchmark has tightened considerably in recent years, the composite ARA metric, the only one of the four that jointly weighs trade, debt, and capital-flow exposure, places India's holdings within, rather than substantially above, the IMF's own recommended safety band for most of the sample period. This raises the natural question of how India's position compares with that of its emerging market peers: is a reserve stock oscillating around the upper half of the recommended ARA band typical for economies of similar size and external exposure, or does it still represent a comparatively conservative stance once other large emerging markets are brought into the picture? The following section takes up this question directly, benchmarking India's ARA ratio against a panel of comparable emerging economies to place its reserve adequacy position in international context.

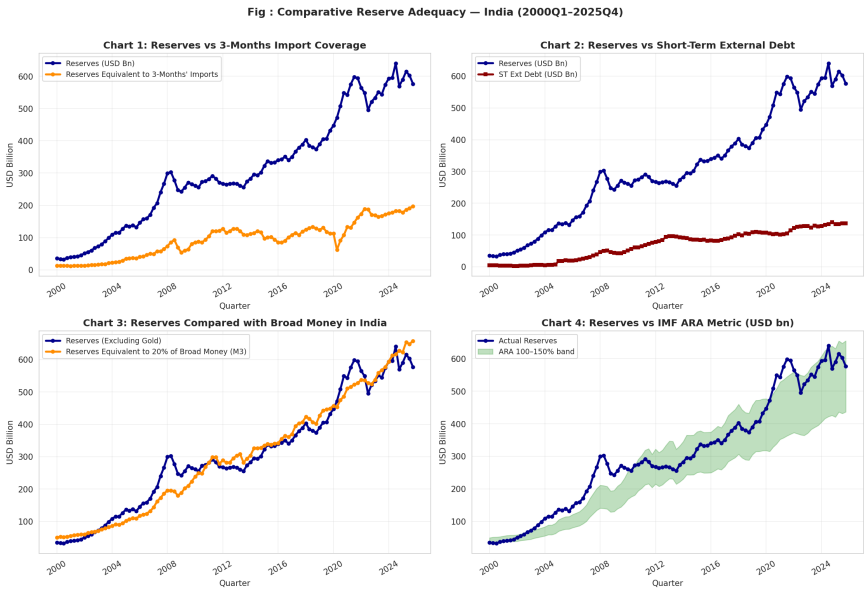


Figure 1: Comparative Reserve Adequacy - India (2000Q1-2025Q4)
Note: The figure presents four standard measures of reserve adequacy. Chart 1 compares actual reserves against the three-month import cover rule. Chart 2 compares reserves against short-term external debt (the Greenspan-Guidotti rule). Chart 3 measures reserves against 20% of broad money (M3) to gauge domestic capital flight risk. Chart 4 compares actual reserves against the IMF's composite Assessing Reserve Adequacy (ARA) metric, highlighting the recommended 100% to 150% safety band.

2.4 International Comparison of Reserve Adequacy

Having established India's own ARA trajectory, this section places it alongside a panel of comparable emerging market economies to assess whether India's position is unusual or broadly typical of countries facing similar external constraints. Peer countries are selected in three steps. First, we restrict the universe to IMF-classified Emerging Market Economies, excluding Advanced Economies. Second, we retain the top 20 EMEs by nominal GDP based on the latest IMF World Economic Outlook data. Third, we further restrict this set to countries whose ten-year average current account balance lies between -6% and $+2\%$ of GDP (approximately the 2016–2025 average, using IMF WEO and Balance of Payments data), so that the comparison group shares a broadly similar external financing profile to India rather than including large, persistent current account surplus economies for which reserve dynamics are driven by fundamentally different forces. This procedure yields nine comparator countries: Brazil, Mexico, Indonesia, the Philippines, South Africa, Turkiye, Poland, Colombia, and Chile.

Figure 2 plots India's ARA ratio against each of these nine peers individually, with the shaded band marking the IMF's recommended 100%-150% adequacy range.. India's ARA ratio has averaged 1.50 over the sample and stands at 1.38 in the latest year, placing it comfortably within the adequate band and spending zero years below the 100% floor, a stronger record than most of its peers. The Philippines and Poland are the only two countries classified as persistently *above*

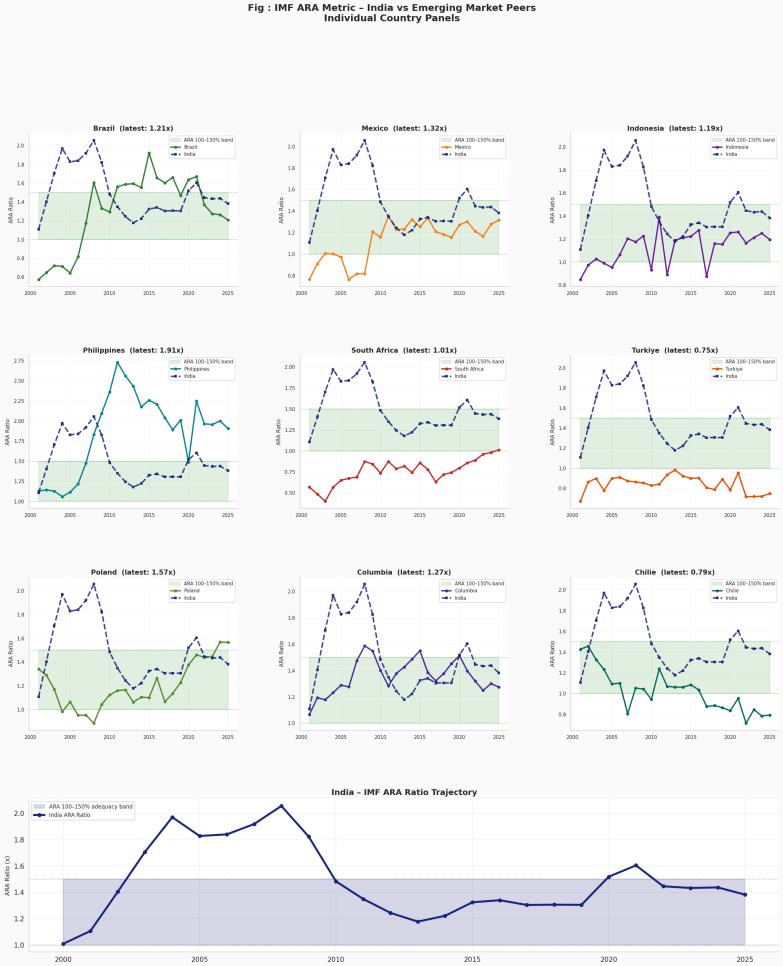


Figure 2: IMF ARA Metric – India vs Emerging Market Peers, Individual Country Panels

the adequate range, with the Philippines averaging close to 1.86 and spending 17 of 25 sample years above 150% coverage. Turkiye and Chile sit at the opposite end, classified as *below* adequate, with Turkiye spending 25 of its sample years under the 100% threshold and a latest ratio of just 0.75. South Africa, despite a low historical mean of 0.76, has converged toward the floor only in the most recent years, reaching 1.01 by the end of the sample. The remaining peers, Mexico, Colombia, Indonesia, and Brazil, cluster in the same broadly adequate range as India, though each spends more years below the 100% threshold than India does over the full sample.

Viewed in this comparative light, India's reserve position is not an outlier of over-insurance relative to peer emerging markets; if anything, it sits toward the more conservative end of a group in which several large economies have repeatedly dipped below the IMF's own adequacy floor. Adequacy-based comparisons, however, are fundamentally backward-looking and mechanical: they say nothing about whether the reserve level a country holds is the level a forward-looking, risk-averse policymaker would actually choose to hold once the insurance value of reserves is weighed against their opportunity cost. Establishing that requires moving beyond descriptive

adequacy benchmarks altogether and building a formal model of reserve demand grounded in India's own macroeconomic fundamentals.

3 Literature Review

The literature on foreign exchange reserves has evolved alongside changes in the global financial architecture and the increasing exposure of emerging market economies to capital flow volatility. Early contributions focused primarily on defining reserve adequacy and identifying the determinants of reserve demand, while later work incorporated welfare considerations, crisis insurance, and balance sheet costs. Despite these advances, the question of whether large reserve holdings are economically efficient remains unresolved, particularly for economies such as India that hold reserves far above conventional benchmarks.

The foundational literature viewed international reserves as a buffer against balance of payments pressures. **Heller (1966), "Optimal International Reserves"**, provided one of the earliest formal analyses, framing reserve demand as a trade-off between the cost of adjustment during external imbalances and the opportunity cost of holding liquid foreign assets. This approach was extended by **Frenkel and Jovanovic (1981), "Optimal International Reserves: A Stochastic Framework"**, who modeled reserves as an inventory stock depleted by stochastic external shocks. In these inventory-theoretic models, the optimal reserve level depends on the volatility of external payments, the speed of adjustment, and the interest rate differential between reserve assets and alternative uses of funds. While analytically tractable, this strand of the literature relied on reduced-form adjustment costs and lacked an explicit welfare-based objective, limiting its suitability for normative policy analysis.

Subsequent studies shifted attention toward the opportunity cost of holding reserves. **Ben-Bassat and Gottlieb (1992), "On the Effect of Opportunity Cost on International Reserve Holdings"**, emphasized that reserves represent foregone domestic investment and should therefore be evaluated relative to alternative returns. Building on this insight, **Rodrik (2006), "The Social Cost of Foreign Exchange Reserves"**, argued that reserve accumulation constitutes an expensive form of self-insurance for developing economies facing borrowing constraints. Using interest differentials as a proxy for opportunity cost, Rodrik showed that the welfare losses from holding large reserve stocks can be substantial. Although influential, this literature largely treated reserves as a static asset position and did not fully integrate crisis probabilities or output losses into a unified welfare framework.

The recurrence of financial crises in emerging markets during the 1990s and early 2000s motivated a new generation of models that explicitly linked reserves to insurance against sudden stops in capital flows. A central contribution in this literature is **Jeanne and Rancière (2011), "The Optimal Level of International Reserves for Emerging Market Countries"**,

which develops a welfare-based insurance model in which reserves smooth consumption during episodes of external financial distress. The model delivers a closed-form expression for optimal reserve holdings as a function of crisis probability, output loss, risk aversion, and the opportunity cost of reserves. Calibrations suggest that observed reserve levels in many emerging markets can be justified by plausible crisis risks, though explaining very high reserve holdings requires assuming large output losses or high risk aversion. This framework represents a major advance by grounding reserve accumulation in explicit welfare maximization rather than heuristic adequacy rules.

Parallel to the theoretical literature, policy-oriented research developed practical reserve adequacy benchmarks. IMF contributions, including the Assessing Reserve Adequacy framework, introduced composite indicators that account for trade exposure, short-term debt, and capital flow risks. These metrics have been widely adopted by policymakers and central banks. In the Indian context, **Das and Nath (2015) applied a similar adequacy-based approach**, constructing a composite reserve adequacy index using multiple conventional metrics to assess whether India's reserves were sufficient relative to its external vulnerabilities. Their analysis confirmed that India's reserves were broadly adequate by standard benchmarks. However, such frameworks including that of Das and Nath (2015) are explicitly precautionary in nature and do not address the economic cost of holding reserves beyond recommended thresholds. They are better suited for identifying whether a country holds enough reserves than for evaluating whether it holds too many. This paper departs from that approach in three ways. First, rather than asking whether India's reserves are adequate, we ask whether they are optimal using the **Jeanne and Rancière (2011) welfare-maximization framework to derive a level of reserves that explicitly balances the insurance benefit against the opportunity cost**. Second, while Das and Nath (2015) estimate a reserve demand equation using standard OLS, this approach produces unreliable estimates when the underlying variables are non stationary and serially correlated, both of which our diagnostics confirm for India's macroeconomic data. We therefore estimate the long-run reserve demand curve using Dynamic OLS, which corrects for these problems and delivers unbiased long-run elasticities. Additionally, for estimating crisis probabilities, we replace the standard Probit model with an asymmetric complementary log-log specification, which is better suited to the rare-event, heavily skewed nature of sudden stop episodes. Third, we go beyond identifying excess reserves to quantifying what that excess has actually cost India fiscally, decomposing both the opportunity cost and the sterilization burden over the last two decades, a dimension that adequacy-based frameworks, by design, leave entirely unaddressed.

In the Indian context, much of the literature has focused on explaining the surge in reserve accumulation rather than evaluating its optimality. **Ramachandran (2006), "On the Upsurge of Foreign Exchange Reserves in India"**, documented the structural and policy drivers of reserve growth, emphasizing capital inflows, exchange rate management, and precautionary motives following the balance of payments crisis of the early 1990s. Subsequent studies exam-

ined sterilization costs, monetary policy implications, and exchange rate dynamics associated with large reserve holdings. However, relatively little work has quantified the economic cost of excess reserves in India or compared observed reserve levels with welfare-based optimal benchmarks. Even when opportunity costs are discussed, they are rarely embedded within a formal insurance framework that accounts for crisis risk and output losses.

A related literature examines the management of large external asset positions through sovereign wealth funds. Reviews such as **"Sovereign Wealth Funds: A Literature Review"** highlight how countries with persistent external surpluses or large reserve buffers have separated short-term liquidity needs from long-term investment objectives. While this literature offers valuable institutional insights, it is often disconnected from the reserve optimality debate and provides limited guidance for economies like India that do not consistently run current account surpluses.

This paper makes four concrete contributions to the literature on foreign exchange reserve management in India. First, we move beyond static adequacy benchmarks by estimating a long-run reserve demand curve for India using Dynamic OLS over 2000Q1–2025Q4. Unlike standard OLS used in prior work, our specification corrects for non-stationarity and serial correlation in India's macroeconomic data, delivering reliable long-run elasticities that reveal the true structural drivers of central bank accumulation financial deepening, exchange rate undervaluation, and sterilization costs. Second, we formally calculate the optimal level of reserves for India using the Jeanne and Rancière (2011) welfare maximization framework, calibrated with time-varying crisis probabilities estimated through a complementary log-log model, replacing the standard Probit approach, which performs poorly on the rare, asymmetric nature of sudden stop events. This allows us to quantify, for the first time, the precise magnitude of India's excess reserves in each quarter from 2000 to 2025. Third, we simulate a two-tranche reserve management framework specifically designed for India, in which the calibrated optimal reserve is deployed as per the central bank's current deployment and the identified excess reserve is channeled into a dedicated sovereign wealth fund. We evaluate whether this structure is financially viable and whether it can generate substantial long-term returns without exposing India to macroeconomic risk. Finally, we dynamically decompose the full fiscal cost of India's excess reserve accumulation over the last two decades into two distinct components: the opportunity cost, measured as the interest rate spread between domestic Treasury bills and US Treasury bills, and the domestic sterilization cost, which the central bank incurs by issuing securities to absorb rupee liquidity each time it purchases foreign exchange. We further simulate a counterfactual in which only the optimal level of reserves is held and the excess is redeployed, showing what India has foregone by maintaining the status quo.

4 Econometric Specification: Reserve Demand via Dynamic OLS (DOLS)

As established previously, simple adequacy rules fail to capture the true motives behind India's massive accumulation of foreign exchange reserves. If the central bank were solely concerned with covering import bills or short-term debt, accumulation would have plateaued long ago. The persistence of excess reserves implies that the RBI is optimizing against a broader set of macroeconomic variables, balancing the precautionary need for liquidity against the financial cost of holding these assets. To uncover these structural drivers, we must move beyond static thresholds and estimate a formal, long-run reserve demand equation.

4.1 Data Construction and Variables

To empirically investigate the structural drivers of reserve demand, we utilize quarterly macroeconomic data spanning 2000Q1 to 2025Q4, defining our dependent variable as the natural logarithm of the ratio of foreign exchange reserves to annualized gross domestic product (GDP). Our explanatory framework incorporates a suite of log-transformed variables to facilitate their direct interpretation as long-run elasticities. We capture traditional trade vulnerabilities and domestic capital flight risks via the imports-to-GDP and broad money (M3)-to-GDP ratios, respectively, alongside the Real Effective Exchange Rate (REER) index to account for broad currency movements. To explicitly model the central bank's precautionary response to external instability, we construct rolling 24-quarter coefficients of variation (the standard deviation divided by the mean) for both current account receipts and foreign portfolio investment (FPI) scaled by GDP. Finally, the quasi-fiscal burden of sterilization is quantified by the opportunity cost, calculated as the yield spread between the domestic 91-day Treasury bill and the US 3-month Treasury bill, and we augment the specification with the CBOE Volatility Index (VIX) as a measure of global risk appetite and a push factor for capital flows into and out of emerging markets.

4.2 Econometric Specification

To estimate these long-run relationships more formally, we next develop an econometric specification linking reserve accumulation to its underlying macroeconomic determinants. Our baseline framework for modeling reserve demand builds on the core selection of variables established in the emerging market literature.[2] Initially, we estimated the long-run demand curve via standard Ordinary Least Squares (OLS) approach. However, preliminary diagnostics reveal

[2]The selection of fundamental variables builds upon the framework explored by Das and Nath (2015), who assessed reserve adequacy and demand drivers in India. Das, & Nath, S. (2015). Assessing Reserve Adequacy in India. Reserve Bank of India Occasional Papers, 35 & 36(1 & 2), ISSN 0972-7493.

significant issues. Augmented Dickey-Fuller (ADF) tests indicate that most of our core variables including reserves, import cover, REER, and opportunity cost are non-stationary, specifically integrated of order one, I(1). Furthermore, Breusch-Godfrey tests on a baseline OLS specification reject the null hypothesis of no serial correlation ($p < 0.001$), pointing to severe endogeneity and autocorrelation. Estimating standard OLS on I(1) variables risks producing a spurious regression.

To resolve these econometric challenges, we employ the Dynamic Ordinary Least Squares (DOLS) estimator developed by Stock and Watson (1993). The DOLS approach provides a statistically robust framework for estimating long-run cointegrating vectors in small samples while mathematically correcting for endogeneity and serial correlation.

The DOLS model is formally specified as:

$$y_t = \alpha + \beta' X_t + \sum_{j=-q}^{p} \gamma_j \Delta X_{t-j} + \epsilon_t \tag{2}$$

In this equation, y_t represents the dependent variable, the natural logarithm of the reserves-to-GDP ratio at time t. The term X_t is a $k \times 1$ vector containing our structural macroeconomic fundamentals in levels (imports/GDP, M3/GDP, REER, the volatility metrics, opportunity cost, and the VIX). The vector β represents the long-run cointegrating elasticities that we seek to estimate.

To correct for endogenous feedback effects from the dependent variable to the regressors, the Stock-Watson method augments the equation with ΔX_{t-j}, which represents the first differences of the independent variables. Finally, α is the intercept and ϵ_t is the stationary stochastic error term. We estimate this model utilizing Heteroskedasticity and Autocorrelation Consistent (HAC) standard errors, specifically the Newey-West estimator, to address any residual serial correlation.

We perform an Engle-Granger test on the DOLS residuals to check for cointegration. The Augmented Dickey-Fuller statistic on the residuals is -4.6658, yielding a p-value of 0.0001. This confirms that the residuals are strictly stationary, proving that a stable, long-run cointegrating relationship exists between India's foreign exchange reserves and our selected macroeconomic fundamentals. The model also demonstrates a strong overall fit, explaining approximately 86 percent of the variance in the reserves-to-GDP ratio over the two-and-a-half-decade sample.

4.3 Results and Interpretation

Analyzing the specific coefficients reveals several crucial insights into the behavior of the central bank. Broad money (M3-to-GDP) emerges as the most powerful and highly significant driver of reserve demand, yielding an elasticity of 1.45 ($p < 0.001$). This indicates that for ev-

Table 1: Dynamic Ordinary Least Squares (DOLS) Structural Elasticities

Dependent Variable: ln(Reserves / GDP)	Coefficient	P-Value	Significance
Constant	3.1010	0.0004	***
ln(Imports / GDP)	-0.1085	0.2867	
ln(Volatility CA)	-0.0193	0.5917	
ln(Volatility FPI)	-0.1523	0.0006	***
ln(M3 / GDP)	1.4530	0.0000	***
Opportunity Cost	-0.0362	0.0000	***
ln(REER)	-1.0683	0.0000	***
VIX	0.0073	0.0000	***
Observations: Model R-squared: 0.8579			

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Differenced lead and lag terms are included in the estimation to correct for endogeneity but are suppressed from this table for brevity.

ery one percent increase in the size of the domestic financial system relative to GDP, the central bank increases its reserve holdings by 1.45 percent. This strongly supports the precautionary capital flight hypothesis: as India's financial sector deepens, the central bank aggressively scales up its reserves to insure against the risk of resident deposit flight.

The opportunity cost of holding reserves is also highly significant and carries the expected negative sign (coefficient of -0.0362, $p < 0.001$). This suggests that higher interest rate differentials between domestic Treasury bills and US Treasury bills are associated with lower reserve accumulation. Intuitively, as the domestic cost of sterilizing foreign exchange interventions rises relative to returns earned on reserve assets, the incentive to continue accumulating reserves weakens.

The real effective exchange rate enters with a large, negative, and highly significant coefficient (-1.0683, $p < 0.001$), indicating that reserve accumulation is systematically higher when the rupee is comparatively weak. This is consistent with the central bank leaning against currency depreciation and accumulating reserves during periods of REER weakness rather than strength, reinforcing the precautionary character of India's reserve policy documented throughout this paper. The VIX also enters positively and significantly (coefficient of 0.0073, $p < 0.001$), confirming that elevated global risk aversion is associated with higher, rather than lower, reserve holdings; this is consistent with the central bank building additional buffers precisely when global financial conditions tighten, rather than being forced into passive drawdowns. Foreign Portfolio Investment volatility is also significant, showing a negative elasticity, which suggests that high financial market turbulence often forces the central bank to deploy reserves, thereby reducing reserve holdings.

Interestingly, traditional trade metrics continue to fail to explain long-run reserve demand. Both the imports-to-GDP ratio and current account volatility remain statistically insignificant in the DOLS framework, corroborating our findings from the adequacy metrics section: trade shocks and import cover rules no longer dictate India's foreign exchange policy. Instead, modern

reserve accumulation in India is primarily driven by the need to secure a growing domestic financial system, manage exchange rate weakness, respond to global risk conditions, and balance the fiscal costs of sterilization.

Relative to Das and Nath (2015), whose adequacy-based approach relies on a standard OLS specification, our DOLS estimation resolves the non-stationarity and serial correlation present in the underlying macroeconomic series and delivers a substantially better overall fit, with the model explaining close to 86 percent of the variation in the reserves-to-GDP ratio against the considerably weaker explanatory power typically reported in adequacy-based OLS specifications. The core precautionary finding, however, is broadly consistent across both approaches: financial deepening, proxied by broad money, remains the single most powerful and robust determinant of India's reserve accumulation.

4.4 Fitted Reserve Demand and Structural Shocks

To visualize the explanatory power of our estimated long-run demand curve, Figure 3 plots the observed reserves-to-GDP ratio against the fitted values generated by the DOLS model. The solid black line represents India's actual reserve holdings scaled by GDP, while the dashed blue line represents the fundamental demand predicted by our structural variables. The shaded blue region illustrates the ± 3 standard deviation confidence band, providing a statistical boundary for normal reserve fluctuations.

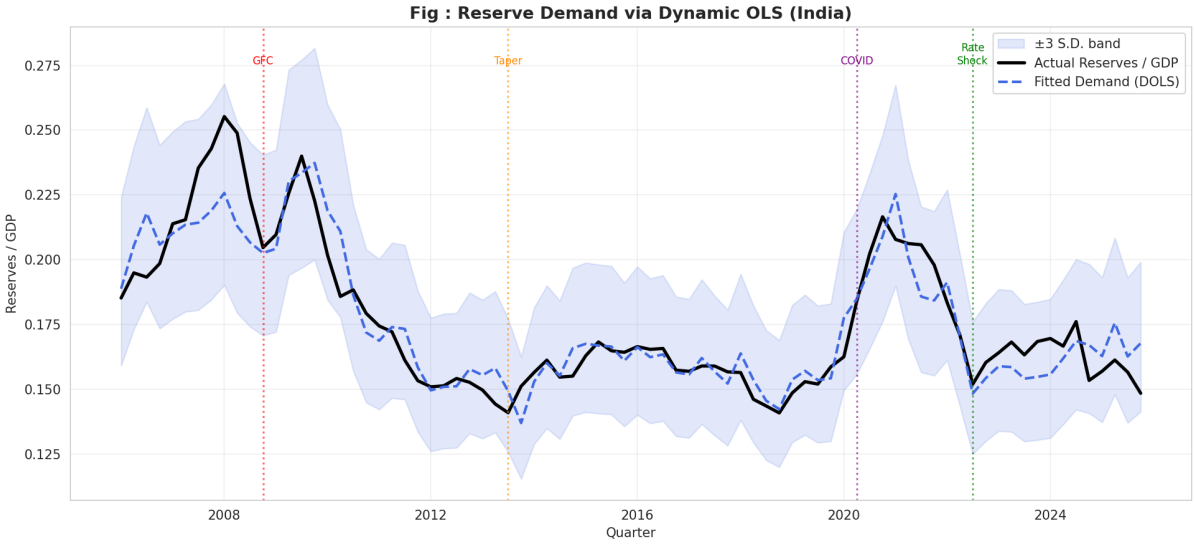


Figure 3: Reserve Demand via Dynamic OLS (India)

Note: The chart displays actual Reserves/GDP against the fitted demand from the DOLS model. Vertical dotted lines mark major structural shocks: the Global Financial Crisis (2008Q4), the Taper Tantrum (2013Q3), the COVID-19 pandemic (2020Q2), and the global Rate Shock (2022Q3). The shaded area represents the ± 3 standard deviation band.

The visual evidence confirms that the model tracks the historical data exceptionally well.

Rather than accumulating reserves randomly, the central bank's behavior is highly systematic and responsive to the macroeconomic fundamentals we identified. The graph clearly captures the central bank's countercyclical interventions during major structural shocks. During the 2008 Global Financial Crisis (GFC) and the 2013 Taper Tantrum, the actual reserves line dips sharply, reflecting the central bank's active deployment of foreign currency to supply market liquidity and defend the rupee against sudden capital outflows. Conversely, during the onset of the COVID-19 pandemic in 2020, the ratio spikes dramatically. This reflects both a mechanical increase due to the temporary contraction in GDP and a deliberate precautionary hoarding of dollars amid severe global uncertainty. During the aggressive global monetary tightening cycle in 2022 (Rate Shock), reserves appropriately contract again as the RBI intervenes to smooth exchange rate volatility.

However, proving that reserve accumulation is predictable is not the same as proving it is economically optimal. Our DOLS estimation successfully explains *why* the central bank holds such a massive reserve specifically, to insure a growing domestic financial system and manage exchange rate competitiveness despite rising holding costs. However, a reserve demand equation explains the determinants of reserve accumulation rather than assessing whether the observed level of reserves is optimal relative to underlying external risks. Given that holding excess foreign assets carries a substantial opportunity cost, we must move beyond descriptive regressions. To determine whether India is currently holding too much or too little insurance, we must now build a formal optimization framework that explicitly equates the marginal benefit of crisis prevention with the marginal cost of holding reserves.

5 Optimal Reserve Calculation

To evaluate whether India's massive reserve accumulation is economically justified, we must transition from a descriptive approach to a normative one by calculating the strictly optimal level of reserves. To achieve this, we calibrate the inter-temporal welfare maximization model developed by Jeanne and Rancière (2011) to the Indian economy spanning 2000Q1 to 2025Q4. This model views reserve accumulation through the lens of a risk-averse policymaker managing a small open economy. The primary objective is to smooth domestic consumption against a "sudden stop", a severe crisis where the country abruptly loses access to external credit. By accumulating reserves during normal periods and paying an opportunity cost, the central bank secures a vital insurance policy to cushion the economy, sustain imports, and mitigate output collapse when a crisis strikes.

5.1 Defining External Crisis Episodes

To estimate the probability of a sudden stop, we first construct a time-series of historical crisis quarters for India spanning 2000Q1 to 2025Q4. Rather than relying on a single threshold or an

externally imposed crisis chronology, we build a multi-dimensional Exchange Market Pressure (EMP) composite index tailored to India's external sector structure.

- **Currency misalignment** (C1): The real effective exchange rate (REER) deviates more than one standard deviation below its Hodrick-Prescott trend, indicating a structurally undervalued currency under external pressure.
- **Reserve drawdown** (C2): The quarterly change in foreign exchange reserves falls more than 1.5 standard deviations below its sample mean, capturing active central bank intervention to defend the exchange rate.
- **Portfolio capital flight** (C3): Net foreign portfolio investment scaled by GDP falls more than 0.75 standard deviations below its mean, reflecting a broad-based reversal of foreign investor sentiment.
- **Sharp depreciation** (C4): The REER depreciates by more than 2 percent quarter-on-quarter, capturing acute and rapid currency stress that cannot be attributed to gradual trend movement.
- **Rollover pressure** (C5): The ratio of short-term external debt to reserves rises more than one standard deviation above its mean, proxying for the Greenspan-Guidotti rollover risk that precedes sudden stop episodes.
- **Portfolio liability outflow** (C6): The quarterly change in total outstanding portfolio liabilities falls more than one standard deviation below its mean, capturing large-scale foreign exit from domestic asset markets.

We define a crisis quarter as any period where at least two of these six criteria are simultaneously breached. To align with standard empirical practices in the sudden stop and exchange market pressure literature, consecutive crisis quarters are clustered into distinct episodes, separated by a minimum of two tranquil quarters. This baseline specification identifies 27 crisis quarters (26.0% of the sample) spanning 11 distinct stress episodes. We specifically favor this threshold over a stricter criteria (which identifies only 7 quarters, or 6.7%). A stricter threshold structurally fails to identify the onset of the 2020 COVID-19 shock, as modern external shocks often manifest as sudden, symmetric capital freezes rather than immediate, severe fundamental misalignments. Furthermore, our algorithm correctly abstains from flagging the 2000:Q4 Dot-com bust as a crisis. Because India's capital account was relatively insulated in 2000, trade and currency shocks were muted compared to global peers; thus, the algorithm's omission of this period validates its fidelity to India's actual macroeconomic reality rather than mere global timelines.

5.2 Estimating Crisis Probability: A Cloglog Approach

To operationalize the Jeanne and Rancìere (2011) framework, we require a time-varying estimate of the quarterly probability of a sudden stop, π_t. Following the established literature, including Das and Nath (2015) and Jeanne and Rancìere (2011), we begin by estimating a standard Probit model. The binary crisis indicator is regressed on three lagged structural vulnerabilities: real effective exchange rate overvaluation, financial openness measured by net foreign portfolio investment flows, and the ratio of short-term external liabilities to broad money. To match the annual identification horizon used in the prior literature, each predictor is averaged across its fourth and eighth quarterly lags. The resulting model yields a McFadden Pseudo R^2 of only 0.089, with two of the three regressors statistically insignificant at conventional levels. This poor fit is a direct consequence of model misspecification rather than weak fundamentals: the standard Probit estimator relies on a symmetric cumulative normal distribution, which assigns equal probability mass to each tail. External financial crises, however, are inherently rare and right-skewed events, with tranquil periods vastly outnumbering crisis quarters in the sample. Imposing symmetry on such an asymmetric data-generating process structurally suppresses the model's discriminatory power.

To overcome this limitation, we replace the Probit link function with an asymmetric complementary log-log (cloglog) specification and augment the predictor set with the lagged US Federal Funds rate as a global push factor. This is the **methodological contribution** of this paper relative to the prior literature. Unlike Probit frameworks, the cloglog link function is naturally asymmetric. It approaches an expected probability of zero relatively slowly but approaches one sharply, making it mathematically superior for predicting rare, extreme events like sudden stops. The conditional probability of a crisis occurring, $\pi_t = P(Y_t = 1|X_t)$, is specified as:

$$\pi_t = 1 - \exp\left(-\exp(\alpha + \beta' X_{t-k})\right) \tag{3}$$

which can be transformed into the linear predictor:

$$\log(-\log(1 - \pi_t)) = \alpha + \beta' X_{t-k} \tag{4}$$

In this specification, X_{t-k} represents a vector of lagged structural vulnerabilities, including real effective exchange rate (REER) overvaluation, financial openness (measured by net foreign portfolio investment flows), and the ratio of short-term external liabilities to broad money. Crucially, we augment these domestic predictors with a global push factor: the lagged US Federal Funds rate. Incorporating the US policy rate captures the external financial cycle, as aggressive tightening by the Federal Reserve historically triggers capital flight from emerging markets. Transitioning to the asymmetric cloglog framework and including this global push

factor drastically improves the model's forecasting ability. The McFadden Pseudo R^2 increases to 0.13860, and the model achieves a robust Area Under the Receiver Operating Characteristic Curve (AUC-ROC) score of 0.73698, indicating strong discriminatory power between crisis and tranquil periods. Full regression results and diagnostic metrics are provided in the Appendix.

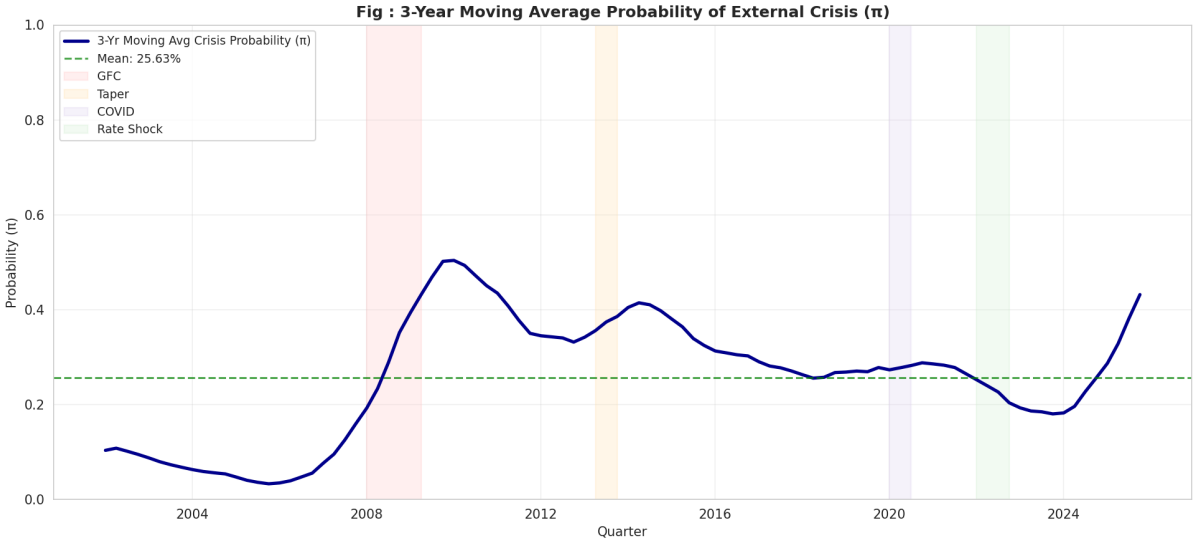


Figure 4: 3-Year Moving Average Probability of External Crisis (π_t)
Note: The solid line represents the smoothed (3-year moving average) predicted probability of a sudden stop derived from the cloglog regression model. The horizontal dashed line marks the sample mean of 25.63 percent. Shaded vertical regions indicate actual identified crisis episodes based on the Exchange Market Pressure (EMP) composite index.

Figure 4 plots the three-year moving average of the estimated sudden stop probability (π_t) generated by our model. Over the 2000–2025 sample, the smoothed crisis probability averages 25.63 percent. The trajectory is highly intuitive and successfully tracks the cyclical buildup and dissipation of macroeconomic vulnerabilities. We observe distinct probability spikes leading into and during major global stress episodes, most notably the 2008 Global Financial Crisis and the 2013 Taper Tantrum.

Furthermore, the graph captures a pronounced upward trend in crisis probability beginning in late 2022 and extending through 2025. This recent elevation is mechanically driven by the aggressive tightening cycle of the US Federal Funds rate and heightened global financial volatility, which collectively increase the likelihood of sudden capital reversals for emerging markets. By utilizing this dynamic, time-varying probability rather than a static historical average, our subsequent optimal reserve calculations can accurately adapt to the shifting global risk environment, ensuring that the prescribed reserve buffer appropriately reflects real-time external threats.

5.3 Calibration of Macroeconomic Stress Scenarios

Having established a reliable, time-varying estimate of crisis probability π_t through the cloglog model, which yields a sample-mean sudden stop probability of 25.63 percent and an AUC-ROC of 0.734, we now turn to the second key input of the Jeanne and Rancìere (2011) framework: the magnitude of the shock itself. Knowing how *likely* a crisis is tells us only half the story; we must also quantify how*severe* it would be in order to determine how large a reserve buffer is actually worth holding. Before formally optimizing the reserve model, it is necessary to calibrate the magnitude of the exogenous shocks that trigger the precautionary motive. To isolate structural fundamentals from high-frequency noise, we apply a Hodrick-Prescott (HP) filter ($\lambda = 1600$) to our key macroeconomic variables. We define a severe stress event as a 1.5 standard deviation (σ) negative deviation from this long-run trend.

This specific shock calibration is rooted in the empirical foundations of global liquidity frameworks. The International Monetary Fund's (2011) *Assessing Reserve Adequacy* metric is explicitly designed to cover tail events, which the IMF mathematically calibrates to the 10th percentile of historical outflows observed during exchange market pressure episodes. Under a Gaussian distribution, the 10th percentile corresponds exactly to a 1.28 standard deviation shock. By calibrating our model to a 1.5 standard deviation drawdown threshold ($Trend - 1.5\sigma$), we structurally harmonize our intertemporal optimization parameters with the IMF's tail-risk methodology while embedding a slightly more conservative buffer. This conservatism is empirically warranted given the elevated volatility of emerging market capital flows.

Figure 5 illustrates the application of our Hodrick-Prescott stress calibration across four critical macroeconomic dimensions: the propensity to export, portfolio liabilities, real GDP growth, and the REER. The graphical evidence confirms the empirical validity of the 1.5σ threshold (green dash-dot lines), which actual realizations (solid blue lines) violently breach during known sudden stops like the 2008 Global Financial Crisis and the 2020 COVID-19 shock. Beyond validating the statistical thresholds, these dynamics reveal the evolving anatomy of external crises in emerging markets. Notably, we observe a "twin drought" of dollar liquidity, where severe capital flight coincides with a collapse in export revenues, perfectly capturing the dual nature of the modern financing gap (λ). Furthermore, the unprecedented output collapse in 2020 highlights that macroeconomic shocks are developing fatter tails with larger potential output losses (γ), while sharp real depreciations during events like the 2013 Taper Tantrum act as a severe penalty that abruptly inflates foreign-denominated debt costs (ΔQ). By establishing these mathematically rigorous and economically intuitive shock parameters, we can seamlessly transition into calibrating the intertemporal utility-maximizing optimal reserve framework.

With both inputs to the optimization, a dynamic quarter-by-quarter crisis probability π_t from Section 5.2 and empirically grounded shock magnitudes λ, γ, and ΔQ from the stress cali-

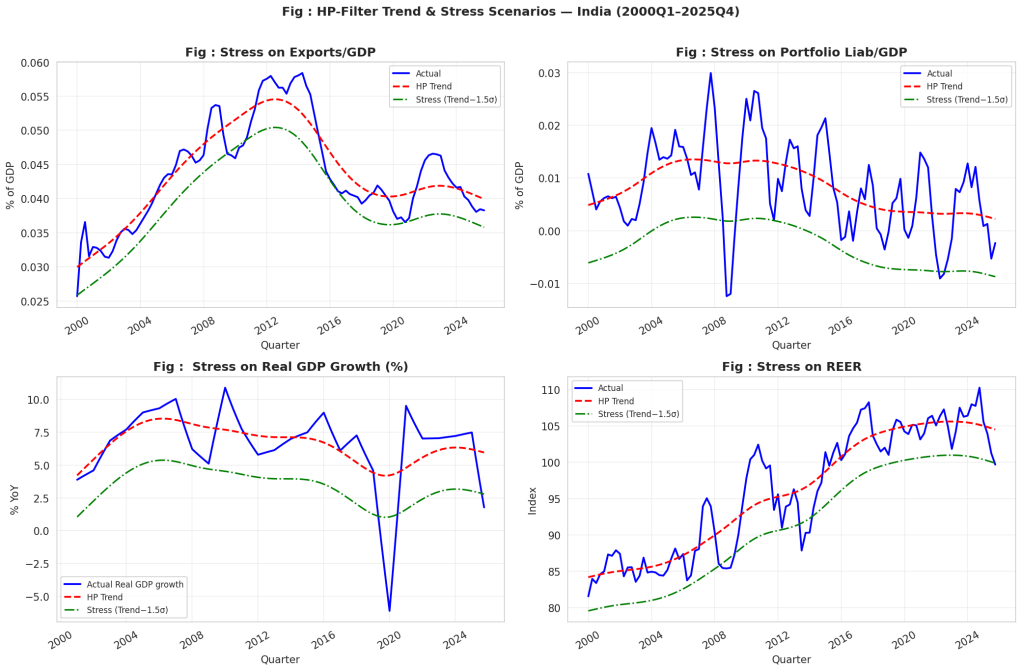


Figure 5: HP-Filter Trend & Stress Scenarios-India (2000:Q1 - 2025:Q4)

Notes: The figure plots the actual series (solid blue), the structural HP trend (dashed red), and the calibrated severe stress threshold corresponding to a 1.5 standard deviation negative shock (dash-dot green). Breaches of the green line historically align with severe real-world sudden stops.

bration above, we have everything required to solve the Jeanne and Rancìere (2011) welfare-maximization problem. Section 5.4 plugs these parameters directly into the closed-form optimal reserve formula and derives the precise level of reserves India should have held in each quarter from 2000 to 2025, against which we then benchmark the actual observed holdings.

5.4 Optimal Reserve Calibration: An Intertemporal Utility Framework

To compute the strictly optimal level of foreign exchange reserves, we apply the seminal intertemporal utility-maximizing framework developed by Jeanne and Rancière,2011. The model postulates a small open economy that accumulates reserves in tranquil times to smooth domestic consumption during a sudden stop in capital flows. The optimal reserve-to-GDP ratio, ρ_t^*, is derived by equating the marginal cost of holding reserves (the term premium) with the marginal benefit of mitigating an output collapse. The structural formula is given by:

$$\rho_t^* = \frac{\lambda + \gamma - \left[1 - \frac{r_t - g_t}{1+g_t}\lambda\right]\left[1 - p_t^{\frac{1}{\sigma}}\right] + \frac{1+r_t}{1+g_t}\lambda\Delta Q}{1 - (\pi_t + \delta)\left[1 - p_t^{\frac{1}{\sigma}}\right] + (1 - \pi_t - \delta)\Delta Q} \tag{5}$$

where the marginal rate of substitution between consumption in the crisis state and the tranquil

state, p_t, is defined as:

$$p_t = \frac{1 - \pi_t - \delta}{\pi_t + \delta} \cdot \frac{\pi_t}{1 - \pi_t} \cdot (1 + \Delta Q) \tag{6}$$

The model requires rigorous calibration of both state-dependent and time-varying parameters. The coefficient of relative risk aversion (σ) is set to 2.0, adhering to the standard baseline utilized in the macroeconomic literature for intertemporal consumption smoothing, including Jeanne and Rancière 2011 and Das and Nath 2015. The term premium (δ), is derived empirically from the US sovereign yield curve. Using historical data from the Federal Reserve Economic Data (FRED) for our 2000–2025 sample, we calculate the average spread between the US 10-Year Treasury Constant Maturity Rate (mean: 3.32%) and the Effective Federal Funds Rate (mean: 2.01%), yielding a structural term premium of 1.31% ($\delta = 0.0131$).

The time-varying parameters include the risk-free rate (r_t), proxied by the 3-month US Treasury yield; the domestic potential growth rate (g_t), extracted via an HP filter of YoY real GDP growth; and the time-varying probability of a sudden stop (π_t), supplied by our smoothed Cloglog estimates. Finally, the expected crisis shock magnitudes are mapped directly from our 1.5 standard deviation threshold calibrations. This yields an expected sudden stop size (λ) of 3.70% of GDP (comprising combined export and portfolio liability collapses), an expected output loss (γ) of 7.75% of GDP, and an expected real exchange rate depreciation (ΔQ) of 12.86%.

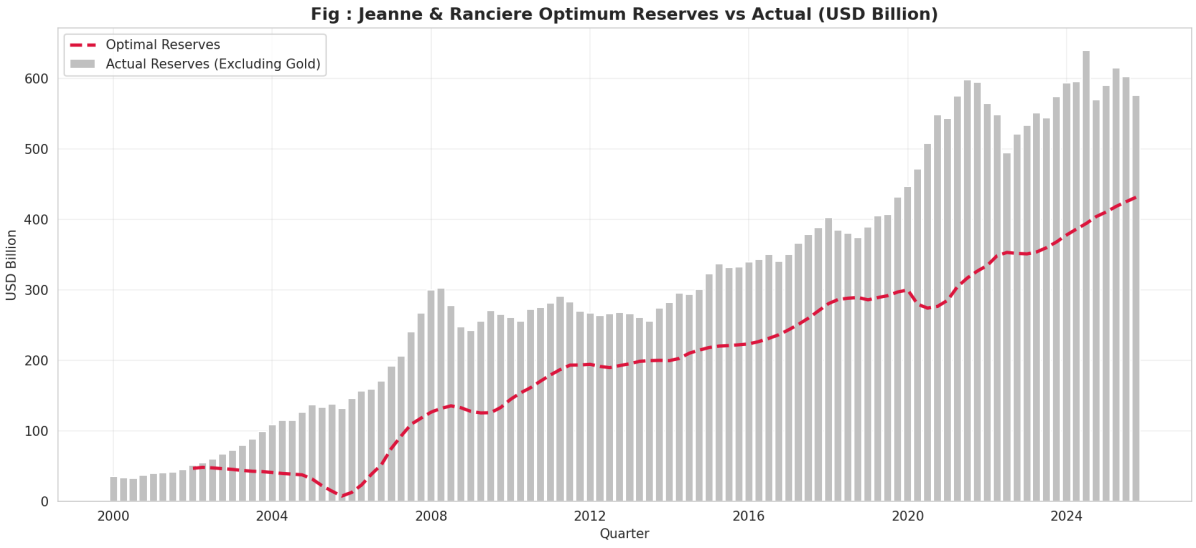


Figure 6: Jeanne & Rancière Optimum Reserves vs Actual (USD Billion).
Notes: The figure plots the actual reserves excluding gold (solid grey bars) against the dynamically calculated optimal reserve level (dashed crimson line). The gap between the bars and the line quantifies the magnitude of excess reserves.

Figure 6 contrasts[3] the time-varying J&R optimal reserve baseline against India's actual hold-

[3]As shown in Figure 6, the optimal reserve series begins later due to lag requirements in the Cloglog model

ings, exposing a persistent and structurally widening regime of over-insurance. Beyond visual trends, translating these gaps into tangible policy metrics reveals a meaningful prudential overhang. By the end of the sample in 2025Q4, the model estimates that a 1.5σ sudden stop would trigger a total macroeconomic resource shortfall of $444.36 billion, comprising a $143.64 billion external financing gap (λ) compounded by a $300.72 billion collapse in output (γ). To optimally smooth domestic consumption against this severe tail-risk event, the Reserve Bank of India strictly requires a baseline holding of $431.45 billion. However, with actual reserves stockpiled at $576.47 billion, India is maintaining $145.02 billion in excess, low-yielding liquidity. Intuitively, this means the central bank holds enough "spare" reserves to entirely absorb an additional 33 percent of a full-scale systemic crisis without ever breaching its fundamental macroeconomic safety floor.

5.5 The Global Financial Safety Net and Usability Frictions

While the baseline Jeanne and Rancìere (2011) framework assumes the sovereign relies exclusively on self-insurance, the modern international monetary system provides India with supplementary external liquidity through what the IMF terms the Global Financial Safety Net (GFSN). For India, this net has expanded considerably over the sample period, growing from approximately USD 6.0 billion prior to 2007 to USD 86.4 billion by 2020, and comprises four distinct unconditional or partially unconditional facilities: the IMF reserve tranche (USD 4.0 billion), the unconditional tranche of the BRICS Contingent Reserve Arrangement (USD 5.4 billion, representing 30 percent of India's USD 18 billion quota that does not require a concurrent IMF programme), the SAARC currency swap framework (USD 2.0 billion), and the bilateral swap arrangement with the Bank of Japan (USD 75.0 billion, confirmed and renewed in 2023).[4]

However, nominal access does not equate to immediately deployable liquidity during a sudden stop. The IMF (2016) formally recognises a set of *usability frictions* that constrain the effective drawdown capacity of these facilities in real time. Three frictions are particularly binding for India. First, activating multilateral facilities such as the BRICS CRA carries severe market stigma, as the act of drawing is itself interpreted as a signal of distress, potentially accelerating the capital outflow it is intended to arrest. Second, 70 percent of India's BRICS CRA entitlement is contractually contingent on an active IMF programme, which India has historically avoided on political economy grounds, rendering this tranche effectively unavailable. Third, bilateral swap arrangements such as the Japan BSA, while large in nominal terms, require contemporaneous diplomatic negotiation to activate and are therefore subject to operational delays

and subsequent smoothing, which consume initial observations before estimation is feasible.

[4]India has never applied for an IMF Flexible Credit Line or Precautionary and Liquidity Line. Given the well-documented stigma associated with Fund precautionary instruments among systemically important emerging markets (IMF, 2016), these facilities are treated as unavailable and excluded from the GFSN calculation throughout the sample.

that are incompatible with the speed of a modern sudden stop episode.

To formally incorporate these frictions into the Jeanne and Rancìere (2011) optimisation, we follow the scenario-based discounting methodology of IMF (2016) and introduce a usability rate $u \in (0, 1]$ that parameterises the fraction of nominal GFSN access that is effectively deployable during a crisis. Let $GFSN_t$ denote the nominal USD value of the safety net at time t. The usability-adjusted financing gap, $\tilde{\lambda}_t(u)$, is defined by reducing the baseline sudden stop shock λ by the usable fraction of the GFSN scaled to GDP:

$$\tilde{\lambda}_t(u) \;=\; \max\!\left(\lambda \;-\; u \cdot \frac{GFSN_t}{GDP_t},\; \varepsilon\right) \tag{7}$$

where $\varepsilon > 0$ is an arbitrarily small lower bound imposed to ensure mathematical tractability and to prevent the financing gap from collapsing to zero under optimistic assumptions. The adjusted gap $\tilde{\lambda}_t(u)$ is then substituted for λ in Equation (5) to yield the GFSN-adjusted optimal reserve ratio $\tilde{\rho}^*_t(u)$. To reflect genuine uncertainty around real-world deployment capacity, we simulate three usability scenarios following the haircut taxonomy of IMF (2016): a conservative scenario of $u = 0.25$, reflecting high stigma and severe activation delays; a realistic baseline of $u = 0.50$, capturing partial unconditional drawing rights; and an optimistic scenario of $u = 0.75$, assuming strong bilateral cooperation and minimal friction.

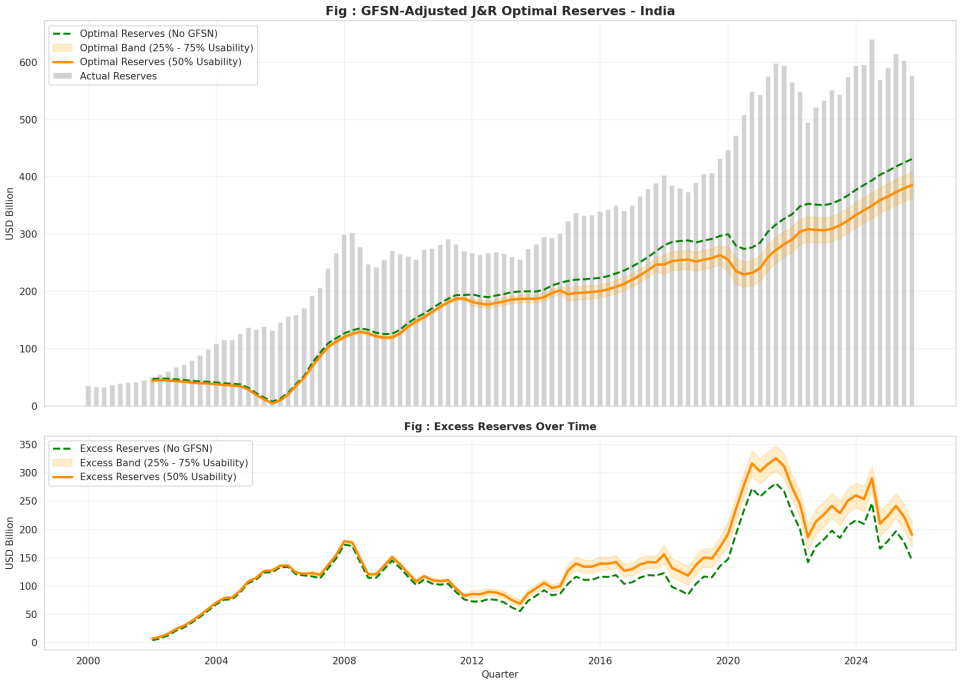


Figure 7: GFSN-Adjusted Optimal Reserves and Excess Bands.
Notes: The top panel displays actual reserves against the baseline optimal requirement and the lower GFSN-adjusted optimal bands. The bottom panel isolates the magnitude of excess reserves, demonstrating that factoring in external safety nets significantly increases the volume of surplus reserves.

As illustrated in Figure 7, integrating the Global Financial Safety Net (GFSN) mechanically

depresses the optimal self-insurance requirement, thereby exposing a significantly larger magnitude of true excess reserves. Translating these adjustments into tangible policy metrics for 2025Q4 reveals a meaningful prudential overhang. While a 1.5σ sudden stop dictates a total macroeconomic resource shortfall of $444.36 billion, assuming a realistic 50 percent usability of India's $86.40 billion nominal GFSN access reduces the optimal self-insurance burden to just $385.77 billion. Consequently, the true excess reserve buffer rises to $190.70 billion. Intuitively, this means that when accounting for external liquidity backstops, the central bank holds enough "spare" reserves to entirely absorb an additional 49 percent of a full-scale systemic crisis without ever breaching its fundamental macroeconomic safety floor. Nevertheless, from a methodological standpoint, our subsequent analysis of the opportunity cost of reserve holding relies exclusively on the unadjusted baseline (0% usability, yielding the more conservative $145.02 billion excess). This deliberate choice establishes a strict lower bound.

6 Excess Reserve Allocation: A Two-Tranche Sovereign Wealth Framework

Now that we have established the size of India's excess reserves, the policy focus shifts to how these funds should be optimally allocated. While global liquidity backstops are available, the International Monetary Fund (2016) notes that "usability frictions" can severely limit a central bank's ability to quickly access these external credit lines during a fast-moving crisis. To account for these practical frictions, we take a highly conservative approach and base our proposed allocation strategy solely on the unadjusted excess reserves, assuming zero reliance on the Global Financial Safety Net. To reduce the high opportunity cost of holding low-yielding assets without compromising macroeconomic stability, we propose transitioning India's foreign exchange management to a Two-Tranche Sovereign Wealth Framework. Under this model, the strictly optimal level of reserves (Tranche 1) remains under the control of the Reserve Bank of India. These core funds will continue to be invested in traditional, highly liquid safe assets to ensure they are immediately available for crisis management. Meanwhile, the structural excess reserves (Tranche 2) are separated and transferred to a newly established Sovereign Wealth Fund (SWF).

6.1 Formulating the SWF Target Return Mandate

Establishing a mathematically rigorous target return is a fundamental prerequisite for sovereign asset allocation. According to Das, Mazarei, and van der Hoorn (2010) in their IMF framework for sovereign wealth economics, a well-defined investment mandate must avoid arbitrary return goals; rather, it should explicitly link the fund's objectives to the sovereign's broader macroeconomic vulnerabilities and aggregate liability profile. To operationalize this within a Sovereign Asset and Liability Management (SALM) framework, we define the SWF's target nominal

return, denoted as R^*, using a "Max Rule." This rule dictates that the required portfolio return must clear the maximum of two distinct systemic benchmarks: a macroeconomic hurdle (R_{macro}) and a financial hurdle (R_{fin}). Formally, the target mandate is expressed as:

$$R^* = \max\{R_{macro}, R_{fin}\} \tag{8}$$

The macroeconomic hurdle (R_{macro}) ensures that the nominal size of the excess reserve buffer grows at a pace sufficient to match the structural expansion of the country's external vulnerabilities. If the SWF fails to meet this threshold, the real economic value of the buffer will mechanically erode relative to the size of the economy. We quantify this hurdle using the compound annual growth rate (CAGR) of India's IMF Assessing Reserve Adequacy (ARA) metric over a trailing 10-year window ($n = 10$):

$$R_{macro} = \left(\frac{ARA_t}{ARA_{t-n}} \right)^{\frac{1}{n}} - 1 \tag{9}$$

Over the trailing decade of our sample, the empirical CAGR of the Indian ARA metric yields $R_{macro} = 5.77$ percent, establishing the minimum growth pace required to simply maintain the country's relative external defense shield.

Conversely, the financial hurdle (R_{fin}) represents the sovereign cost of capital, the absolute minimum return required to economically compensate the state for external debt servicing, global equity risk, and operational frictions. We derive this hurdle empirically using financial market data spanning 2001 to 2025. The sovereign cost of capital is formally specified as:

$$R_{fin} = r_f + s_{sov} + \lambda_{erp} + \tau \tag{10}$$

where r_f is the global risk-free rate, proxied by the sample average yield on the US 10-Year Treasury note (3.20 percent). The term s_{sov} captures India's sovereign credit spread to reflect actual borrowing costs; utilizing the ICE BofA BBB US Emerging Markets Option-Adjusted Spread, this premium averages 1.19 percent. Furthermore, because deploying capital into an SWF carries an opportunity cost compared to simply retiring outstanding sovereign debt, the mandate must require compensation for taking on market risk. We capture this via a global equity risk premium, λ_{erp}, calculated as the S&P 500 CAGR minus the risk-free rate over the sample period (3.76 percent). Finally, τ represents standard institutional management frictions (0.10 percent). Summing these empirical components yields a comprehensive financial hurdle rate of $R_{fin} = 8.25$ percent.

Applying the SALM policy rule, we evaluate the binding constraint for the portfolio mandate:

$$R^* = \max\{5.77\%, 8.25\%\} = 8.25\% \tag{11}$$

Because the fundamental cost of sovereign capital strictly exceeds the growth rate of India's external vulnerabilities, the financial hurdle constitutes the binding constraint. Consequently, the target return mandate for India's proposed Tranche 2 Sovereign Wealth Fund is mathematically anchored at 8.25 percent nominal.

6.2 Asset Proxies and Mean-Variance Optimization

To empirically simulate the performance of the proposed Tranche 2 Sovereign Wealth Fund, we utilize highly liquid, low-cost index funds as proxies for the three primary asset classes. We select the **Vanguard Total Stock Market Index (VTSMX)** to represent global equities and the **Vanguard Total Bond Market Index (VBMFX)** for global fixed income. While direct sovereign investments in real estate and private markets involve complex, illiquid cash flow valuations (Gupta and Van Nieuwerburgh, 2021), we utilize the **Vanguard Real Estate Index (VGSLX)** as a standard, highly liquid empirical proxy for real assets. To construct the optimal strategic asset allocation, we employ a dynamic **Markowitz mean-variance optimization** framework. Crucially, to eliminate look-ahead bias, we utilize an expanding-window approach; at each quarter t, the optimization relies strictly on historical data available up to $t-1$.

The objective is to minimize portfolio volatility while strictly satisfying the target return threshold (R^*) derived in the previous section. Let $\mathbf{w}_t = [w_{eq,t}, w_{bond,t}, w_{ra,t}]^T$ represent the portfolio weights at time t. The formal out-of-sample optimization problem is defined as:

$$\min_{\mathbf{w}_t} \quad \mathbf{w}_t^T \boldsymbol{\Sigma}_{t-1} \mathbf{w}_t \tag{12}$$

$$\text{subject to:} \quad \mathbf{w}_t^T \boldsymbol{\mu}_{t-1} \geq R^* \tag{13}$$

$$\mathbf{1}^T \mathbf{w}_t = 1 \tag{14}$$

$$w_{ra,t} = 0.05 \tag{15}$$

$$\mathbf{w}_t \geq 0 \tag{16}$$

where $\boldsymbol{\Sigma}_{t-1}$ and $\boldsymbol{\mu}_{t-1}$ are the annualized covariance matrix and mean return vector estimated over the expanding historical window. We explicitly constrain real asset exposure ($w_{ra,t}$) to exactly 5.0 percent. While modern SWFs exhibit heterogeneous and evolving mandates that increasingly favor alternative investments to generate alpha (Aguilera, Capapé, and Santiso, 2023), a 5.0 percent cap reflects the practical institutional deployment constraints and illiquidity profiles of unlisted real assets. This securely anchors our model to the operational realities of mature funds, matching the exact real estate target allocation utilized by Norway's Government Pension Fund Global (NBIM, 2024).

Finally, we impose a long-only constraint ($\mathbf{w}_t \geq 0$) to align with standard sovereign mandates. Running this expanding-window optimization through the end of our sample in 2025Q4 yields an optimal strategic allocation of 60.1 percent to global equities, 34.9 percent to global bonds,

and 5.0 percent to real assets. This mathematically demonstrates that successfully clearing R^* strictly necessitates a substantial equity risk premium. Consequently, our empirical model robustly validates the broader structural shift toward 60/40 asset allocation frameworks currently observed among global sovereign wealth managers seeking to maximize intertemporal utility .

6.3 Empirical Performance of the SWF Tranche (2004-2025)

To evaluate the economic viability of the proposed Sovereign Wealth Fund, we simulate the historical out-of-sample performance of Tranche 2 from 2004:Q1 through 2025:Q4. In each period t, the dynamically optimized portfolio weights ($\mathbf{w}_t$) are applied exclusively to the strictly excess reserve balance identified in $t-1$. This strict sequencing ensures that the absolute 1.5σ optimal reserve threshold (Tranche 1) is entirely immunized and never compromised for the sake of yield. Figure 8 visualizes the financial trajectory of this strategic allocation. The top panel displays the absolute quarterly dollar returns, exhibiting a highly resilient 72.4 percent "win rate" (positive returns in 63 out of 87 quarters) and an average quarterly generation of $2.96 billion. The bottom panel tracks the compounding effect of these returns, revealing a steady structural ascent that culminates in a staggering $257.86 billion of net new sovereign wealth by the end of 2025.

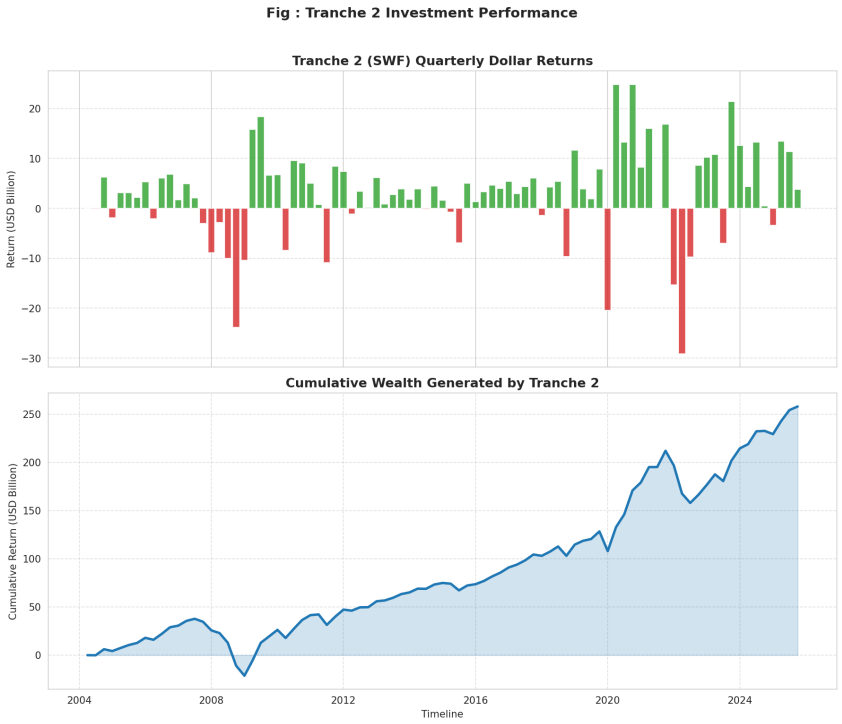


Figure 8: Tranche 2 (SWF) Financial Performance (2004:Q1 - 2025:Q4).
Notes: The top panel displays the absolute quarterly dollar returns (USD Billion) generated by the optimized excess reserve portfolio. The bottom panel traces the cumulative sovereign wealth generated over the sample horizon.

Beyond the headline figures, the trend reveals a profound policy insight regarding sovereign

risk tolerance. As expected, the equity-heavy Tranche 2 portfolio experiences severe mark-to-market drawdowns during global systemic shocks, violently contracting during the 2008 Global Financial Crisis (losing $23.82 billion in a single quarter) and the 2022 aggressive monetary tightening cycle (a $29.08 billion drawdown). However, under a two-tranche framework, this financial volatility is fundamentally benign. Because the central bank's acute crisis-fighting liquidity is already mathematically guaranteed by Tranche 1, the SWF is never forced to liquidate assets at distressed prices. This structural patience allows the SWF to effortlessly absorb short-term macroeconomic turbulence and remain invested to capture the aggressive, asymmetric market recoveries that historically follow crises. Over the two-decade simulation, the compounding global equity risk premium overwhelmingly dominates these intermittent shock losses. Ultimately, this demonstrates that bifurcating the reserve mandate transforms deadweight opportunity costs into a quarter-trillion-dollar fiscal windfall, significantly expanding the Indian state's intergenerational financial capacity.

6.3.1 Tranche 1: The Liquidity Portfolio and Traditional Deployment

While Tranche 2 optimizes the excess reserves for long-term growth, the dynamically calibrated Jeanne and Rancière optimal reserves constitute **Tranche 1 (The Liquidity Portfolio)**. Because these funds must be instantaneously deployable during a sudden stop, they are strictly constrained to highly liquid, risk-free assets. To calculate the baseline return of this tranche, we empirically replicate the RBI's traditional reserve deployment strategy.

Based on the RBI's reported reserve composition and the September Foreign Exchange Reserves deployment framework, India's reserves (excluding gold) consist predominantly of Foreign Currency Assets (FCA), alongside Special Drawing Rights (SDR) and the IMF Reserve Tranche Position. The FCA itself is deployed across three sub-categories: sovereign securities ($w_{sec} = 84.52\%$), deposits with other central banks and the BIS ($w_{cb} = 7.96\%$), and deposits with commercial banks ($w_{comm} = 7.52\%$)[5].

To simulate the return profile of this strictly constrained liquidity portfolio, we apply standard empirical proxies widely utilized in central bank reserve accounting literature. The securities allocation (R_{SHY}) is proxied by the iShares 1-3 Year Treasury Bond ETF (SHY), which accurately captures the capital gains and yield dynamics of short-duration US sovereign debt, the primary instrument for global central bank liquidity. For central bank deposits and IMF/SDR holdings (R_{TB3}), we utilize the US 3-Month Treasury Bill yield to represent the baseline global risk-free rate. Finally, commercial bank deposits are modeled using the 3-Month Treasury Bill yield plus a standard 25 basis point spread ($R_{TB3} + 0.0025/4$) to appropriately compensate for commercial counterparty credit risk.

[5]RBI has published the FCA deployment data for 2025; since earlier data were unavailable, we assumed the FCA was invested in that ratio for all periods.

The quarterly return on the FCA ($R_{FCA,t}$) and the total blended return for Tranche 1 ($R_{T1,t}$) are mathematically formalized as:

$$R_{FCA,t} = 0.8452 \cdot R_{SHY,t} + 0.0796 \cdot R_{TB3,t} + 0.0752 \cdot (R_{TB3,t} + 0.000625) \tag{17}$$

$$R_{T1,t} = w_{FCA,t} R_{FCA,t} + w_{SDR,t} R_{TB3,t} + w_{IMF,t} R_{TB3,t} \tag{18}$$

where $w_{FCA,t}$, $w_{SDR,t}$, and $w_{IMF,t}$ are the dynamically extracted quarterly weights of each component from the RBI's historical data.

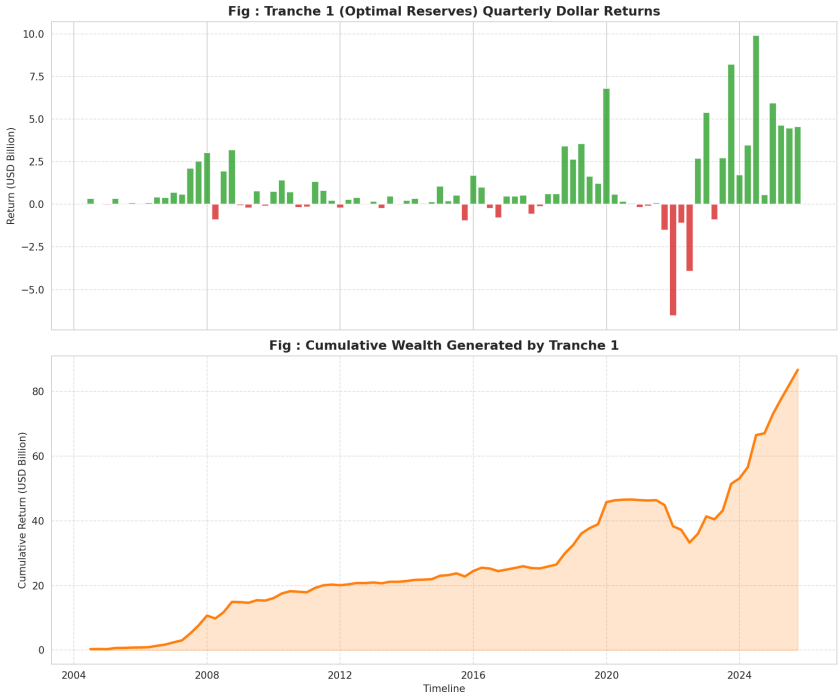


Figure 9: Tranche 1 (Optimal Reserves) Financial Performance.
Notes: Displays the quarterly dollar returns and cumulative wealth generated by investing strictly the Jeanne and Rancière optimal reserve baseline into the replicated RBI liquidity portfolio.

Applying this empirical rate of return strictly to the optimal reserve levels calculated in Section 5.4 yields the performance of Tranche 1 (Figure 9). Over the 2004–2025 horizon, the ultra-conservative Liquidity Portfolio generated an annualized return (CAGR) of 1.96% with a negligible annualized volatility of 1.63%, accumulating USD 86.70 billion in cumulative wealth.

6.4 Aggregate Mechanism Performance: Status Quo vs. Two-Tranche Framework

Having modeled the Tranche 2 Sovereign Wealth Fund (SWF) by itself, we now look at the overall performance of the entire foreign exchange portfolio. By combining the highly safe optimal reserves (Tranche 1) with the dynamically invested excess reserves (Tranche 2), the total portfolio achieves a blended annualized return of 4.68 percent. The overall volatility stays highly manageable at 5.12 percent.

Figure 10 breaks down exactly how this combined portfolio performs over time. The top panel shows the total quarterly dollar returns. While there are noticeable red bars representing drawdowns during major global shocks (such as the 2008 financial crisis and the 2022 rate hikes), the vast majority of the bars are green. In fact, the portfolio delivers positive returns in 75.6 percent of the quarters. The bottom panel explains why this works so well by tracking the cumulative wealth. The blue dashed line (Tranche 1) acts as a massive stabilizing anchor, providing a slow, steady stream of safe interest income without any dips. Riding on top of that safe foundation is the orange dashed line (Tranche 2), which acts as the primary growth engine. Together, they drive the total combined wealth (the purple shaded area) to a staggering $344.58 billion by 2025.

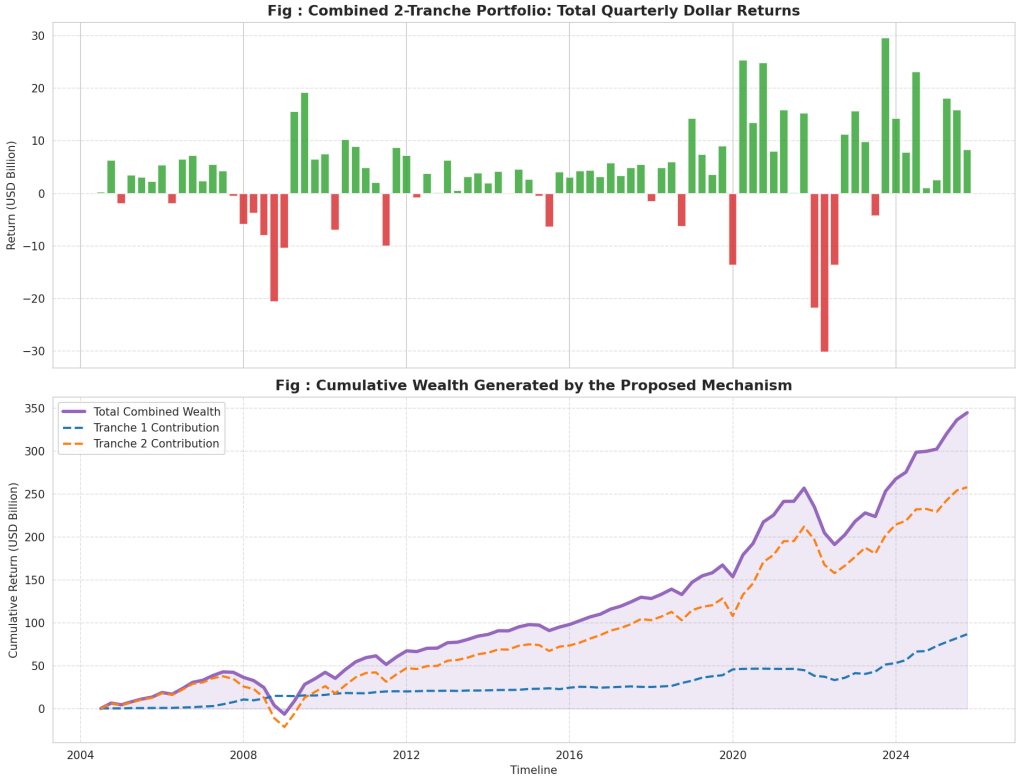


Figure 10: Combined 2-Tranche Portfolio Performance (Optimal + Excess)

To clearly see the financial benefit of this new framework, we must compare it against a simu-

lated “Status Quo” baseline. Under the Status Quo, we assume the central bank simply continues its historical practice of keeping 100 percent of its total reserves in traditional, low-yielding safe assets.

Figure 11 illustrates this stark contrast. In the top panel, the gray dashed line represents the Status Quo strategy, which grows very slowly over the two decades. The purple solid line represents our proposed Two-Tranche model. The expanding green shaded area between the two lines perfectly captures the extra sovereign wealth generated by the new strategy. Notice how this green gap widens significantly after 2010, as the compounding effects of the SWF’s equity investments take over. The bottom panel confirms this success by plotting the quarterly outperformance; the proposed mechanism beats the traditional Status Quo in almost every single quarter, easily making up for the rare periods of market underperformance.

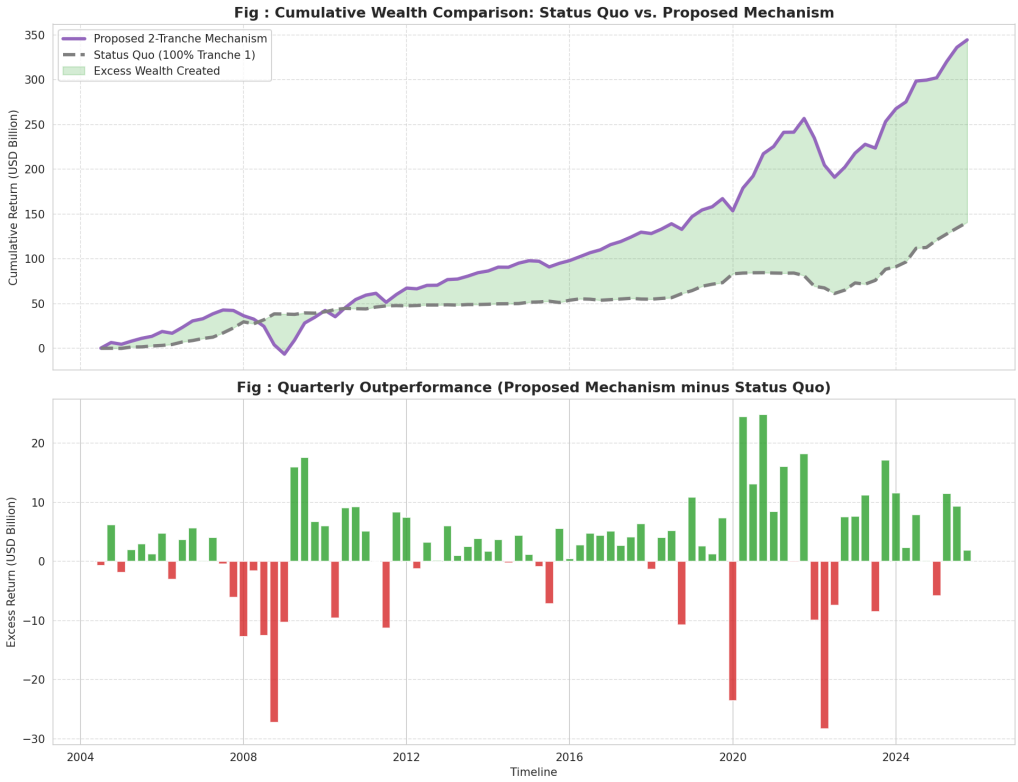


Figure 11: Cumulative Wealth Comparison: Status Quo vs. Proposed Mechanism

The final numbers emphasize the visual trend. Over the twenty-year simulation, the Status Quo strategy yields an annualized return of just 1.96 percent, generating $140.80 billion in cumulative wealth. In stark contrast, the proposed Two-Tranche mechanism generates $344.58 billion over the exact same period. By simply separating the excess reserves into a dedicated wealth fund, the country gains exactly $203.78 billion in extra financial capacity, all while keeping its core crisis insurance fully intact and ready for sudden stops.

However, showing that a diversified wealth fund generates higher returns than safe bonds is only half of the equation. To truly understand why accumulating and holding massive amounts

of traditional reserves is problematic for an emerging market, we must look at the liabilities side of the central bank's balance sheet. When a central bank buys foreign currency, it must issue domestic debt to sterilize those purchases and control local inflation. This creates a direct financial burden on the domestic economy, which leads us to formally calculate the true, carrying cost of holding these massive reserves.

7 The Fiscal Cost of Accumulation and the Counterfactual Dividend

While accumulating foreign exchange reserves helps protect emerging market economies from sudden stops, this insurance comes at a price. When a central bank buys foreign currency, it usually issues domestic liabilities to sterilize the monetary impact and control local inflation. This dynamic means the government effectively borrows at higher domestic interest rates to invest in lower yielding foreign safe assets. Rodrik (2006) highlights this negative spread as the social cost of reserve accumulation. To measure this macroeconomic inefficiency, we break down the total realized fiscal cost (FC_t) into two parts: the Opportunity Cost (OC_t) of holding the existing stock of reserves, and the Sterilization Cost (SC_t) of actively absorbing domestic liquidity.

The Opportunity Cost is defined as the spread between the long term domestic sovereign borrowing rate and the blended yield on the reserve portfolio, applied to the total reserve stock. Formally:

$$OC_t = \left(i_{dom,t}^{10Y} - r_t^{res}\right) \times R_{t-1} \tag{19}$$

where $i_{dom,t}^{10Y}$ is the yield on India's 10-Year government bond, r_t^{res} is the empirical return on the traditional foreign exchange portfolio, and R_{t-1} is the stock of reserves.

The Sterilization Cost captures the marginal cost of reserve accumulation. We measure this using the spread between the short term domestic policy proxy (the 91-day Treasury Bill) and the reserve return, applied to the change in reserves:

$$SC_t = \left(i_{dom,t}^{91D} - r_t^{res}\right) \times \Delta R_t \tag{20}$$

The aggregate fiscal cost is the sum of these two frictions:

$$FC_t = OC_t + SC_t \tag{21}$$

Applying this methodology to India from 2000 to 2025 shows a clear upward trend in the cost of holding reserves. Figure 12 illustrates this burden.

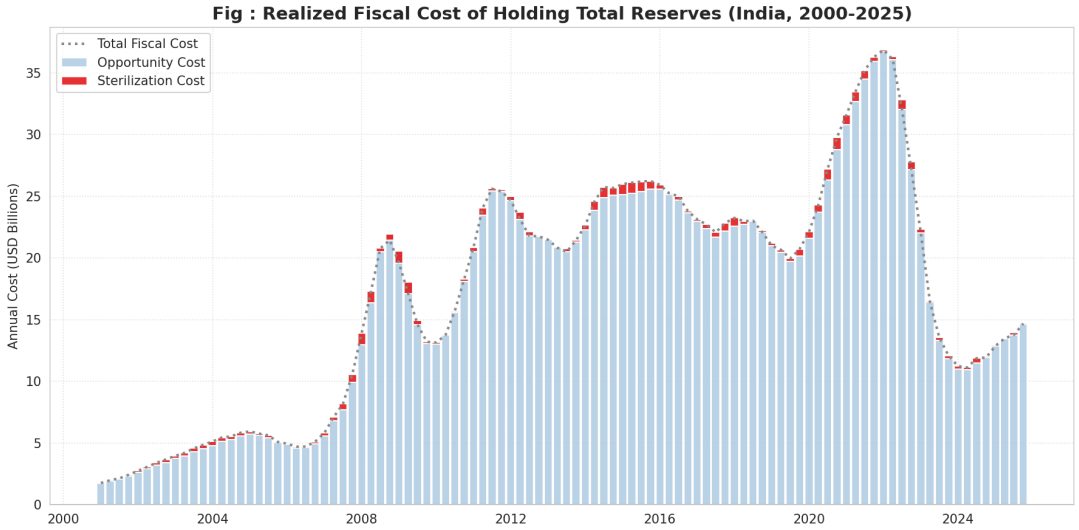


Figure 12: Realized Fiscal Cost of Holding Total Reserves (India, 2000 to 2025).
Notes: The light blue bars represent the opportunity cost of the reserve stock, while the red bars denote the marginal sterilization cost.

In the early 2000s, the total fiscal cost was manageable at around $5 billion annually. However, it grew significantly over time, reaching $35.904 billion in 2021, which represented about 1.209 percent of India's GDP for that year. The opportunity cost (light blue bars) is the primary driver of this expense. The sterilization cost (red bars) fluctuates more and sometimes turns negative when reserves decrease.

Beyond the headline trend, the data reveal an important structural feature of reserve holding costs in emerging market economies. The fiscal burden associated with maintaining large reserve stocks tends to widen precisely during periods of global financial stress. This pattern arises from the asymmetric response of interest rates across markets: during episodes of global monetary easing, yields on the safe foreign assets in which reserves are invested compress toward zero, while domestic sovereign borrowing rates in emerging markets remain comparatively elevated, reflecting domestic inflation dynamics and risk premia that do not fully co-move with global cycles. The result is a cyclical widening of the negative interest rate spread that mechanically inflates the opportunity cost of reserve holding at exactly the moment when external buffers are most actively needed. This is not a policy failure; it is an inherent structural feature of the asymmetric integration of emerging markets into the global financial system, and it underscores why a static, one-size reserve stock is a suboptimal response to a time-varying external risk environment.

To understand what is lost by keeping excess reserves in these safe assets, we look at a counterfactual scenario. We calculate the foregone alpha by taking the difference between the return of the proposed Sovereign Wealth Fund (r_t^{SWF}) and the traditional Tranche 1 return (r_t^{T1}), applied

only to the excess reserve balance:

$$OC_t^{SWF} = \left(r_t^{SWF} - r_t^{T1}\right) \times Excess_R_t \tag{22}$$

Figure 13 visualizes this counterfactual opportunity cost.

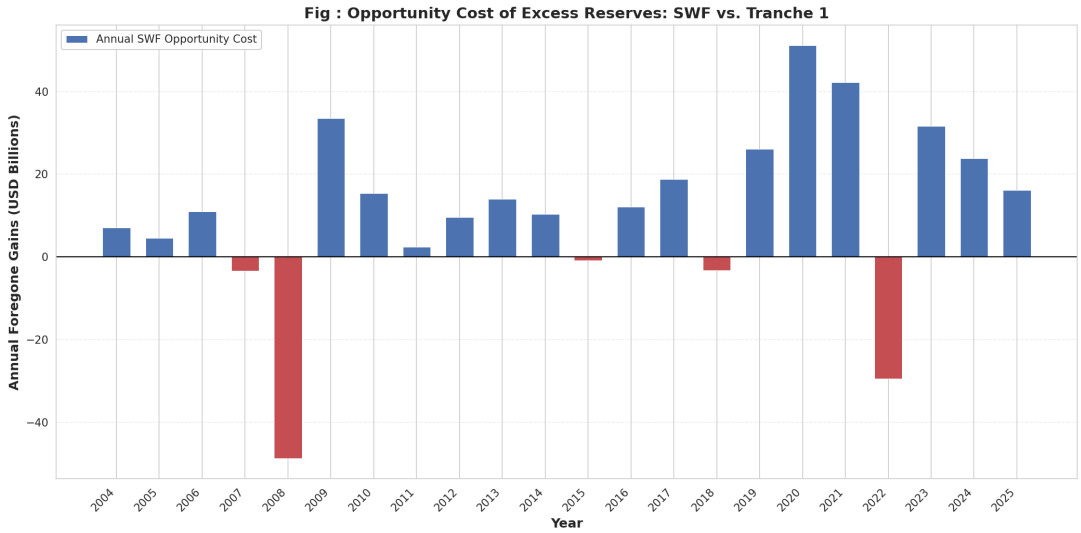


Figure 13: Opportunity Cost of Excess Reserves: SWF vs. Tranche 1.

Notes: Positive blue bars indicate foregone gains (the cost of the status quo). Negative red bars indicate periods where the SWF would have temporarily underperformed the safe asset baseline.

The annual data behind Figure 13 put a precise dollar value on what the status quo has cost India each year. The SWF strategy underperforms the traditional safe asset portfolio in only five years out of twenty-two: 2007, 2008, 2015, 2018, and 2022. Each of these years lines up with a known global shock that temporarily hit equity markets hard. In every other year, keeping excess reserves in low-yielding foreign assets meant leaving substantial returns on the table. The single largest missed opportunity came in 2020, when the sovereign foregone USD 51.17 billion as global equity markets recovered sharply from their pandemic lows while India's excess reserves sat idle in safe foreign securities. The year 2021 was nearly as costly, with USD 42.19 billion in foregone gains. Even in more recent years, as the post-pandemic interest rate cycle has normalised, the gap has remained meaningful: USD 31.68 billion in 2023, USD 23.84 billion in 2024, and USD 16.83 billion in 2025. The pattern that emerges is clear. The years where the SWF would have lost ground are few, short-lived, and fully explainable by discrete global shocks. The years where the status quo imposes a real cost are the majority, and that cost compounds over time. Since Tranche 1 already covers all of India's crisis liquidity needs, none of these foregone gains were ever necessary. They represent the recurring price of holding more insurance than the economy actually requires.

8 Policy Implications: Reserve Redeployment and Institutional Reform

The empirical findings of this paper point to a clear and quantifiable inefficiency in India's current approach to foreign exchange reserve management. India is persistently over-insured relative to the welfare-maximising benchmark derived from Jeanne and Rancière (11): actual reserves have exceeded the dynamically calibrated optimal level in every quarter from 2000 to 2025, with the gap widening materially after 2014. The excess reserve position reached $145.02 billion by 2025Q4 under the conservative baseline that assigns zero usability to the Global Financial Safety Net, and expands to $190.70 billion under a realistic 50 percent usability assumption. These excess reserves generate an annualised return of 1.96 percent while the domestic cost of the sterilisation operations used to accumulate them reflects borrowing rates that are substantially higher. The annual fiscal dividend foregone by maintaining this status quo reached $16.83 billion in 2025 alone. The GFSN further reduces the required level of self-insurance, since India's bilateral swap arrangement with the Bank of Japan ($75 billion, renewed in 2023), the unconditional BRICS Contingent Reserve Arrangement tranche ($5.4 billion), the IMF reserve tranche ($4 billion), and the SAARC swap facility ($2 billion) collectively provide supplementary liquidity that the baseline optimisation does not account for. Once these facilities are incorporated at a conservative 50 percent usability haircut, the welfare-optimal self-insurance requirement falls to $385.77 billion, making the true excess meaningfully larger than the unadjusted figure.

8.1 A Two-Tranche Reserve Framework

The central policy recommendation of this paper is the establishment of a two-tranche reserve management framework. Under this model, the Reserve Bank of India retains full control over a dynamically calibrated liquidity buffer, which we call Tranche 1, sized in each quarter using the Jeanne and Rancière (11) optimal reserve formula with time-varying crisis probabilities and empirically grounded shock parameters. As of 2025Q4, this liquidity floor stands at $431.45 billion, sufficient to absorb a 1.5 standard deviation sudden stop that would generate a combined macroeconomic resource shortfall of $444.36 billion comprising a $143.64 billion external financing gap and a $300.72 billion collapse in output. Tranche 1 is invested exclusively in highly liquid, short-duration foreign sovereign securities, exactly as the RBI currently manages its reserve portfolio, with no change to the central bank's balance sheet risk or its ability to intervene in currency markets. The structural excess above this floor, currently $145.02 billion under the conservative unadjusted baseline, is separated and transferred to a newly established sovereign wealth vehicle, which we call Tranche 2, with a global investment mandate targeting the binding financial hurdle rate of 8.25 percent nominal derived from India's sovereign cost of capital. The key institutional insight is that the separation is not merely a return-enhancement

device. Because the central bank's crisis-fighting liquidity is already guaranteed by Tranche 1, the sovereign wealth fund is structurally prevented from being forced to liquidate assets at distressed prices to fund crisis intervention. This structural patience is what allows the equity risk premium in Tranche 2 to compound over time rather than being eroded by forced selling during precisely the episodes when asset prices are lowest. Under an expanding-window, out-of-sample optimisation that targets this 8.25 percent mandate, the model allocates 60.1 percent to global equities, 34.9 percent to global bonds, and 5.0 percent to real assets by the end of the sample, closely tracking the roughly 60/35/5 strategic allocation used in the policy recommendations below.

8.2 International Lessons

Three sovereign wealth fund precedents are directly relevant for India. Norway's Government Pension Fund Global, with assets exceeding $1.7 trillion as of 2024, demonstrates that a rules-based spending cap, full public transparency, and strict separation from day-to-day fiscal policy are achievable at scale and are the primary source of the fund's political durability. Its long-term real return of approximately 4 to 5 percent above global inflation over the past two decades provides an external benchmark for what a well-governed sovereign fund can deliver, and its strategic allocation of approximately 70 percent global equities and 25 percent fixed income is broadly consistent with the allocation our mean-variance optimisation identifies as the minimum-variance portfolio capable of clearing the 8.25 percent return mandate. Singapore's GIC, funded by government reserves, is perhaps the most directly relevant precedent for India. GIC operates with a long-term investment mandate, invests across global asset classes, maintains a strict separation between its investment decisions and government fiscal policy, and has generated a 20-year real return of approximately 4.6 percent above global inflation under realistic institutional constraints. China's State Administration of Foreign Exchange and its subsidiary China Investment Corporation demonstrate that reserve-funded sovereign investment is operationally feasible at very large scale. However, the opacity of CIC's governance and its entanglement with broader strategic objectives represent features that India would want to explicitly avoid. The lesson from all three cases is the same: institutional independence, a rules-based mandate written into statute, and transparent reporting are not optional governance refinements but supportive conditions for political sustainability.

8.3 Institutional and Political Constraints

Acknowledging the political economy barriers to implementation is necessary for the proposal to be credible. The Reserve Bank of India has strong institutional incentives to maintain a large reserve stock. A large reserve balance enhances the central bank's perceived external credibility and operational autonomy in foreign exchange markets. Transferring a significant share of that stock to a new entity, even one with a clear statutory mandate, would reduce the RBI's

balance sheet and could be perceived, incorrectly but with political consequences, as a weakening of India's external defences. Coordination between the RBI and the Ministry of Finance would be required at every stage of the transition, and the historical record of central bank and fiscal authority coordination in India on balance sheet matters is mixed. The governance risks associated with a foreign-currency sovereign wealth fund are real. A poorly designed fund is exposed to political interference in investment decisions, procyclical withdrawal during domestic fiscal stress, and reputational sensitivity to individual investment losses. Exchange rate management presents a further constraint. The proposed framework does not mechanically generate appreciation pressure since the transfer is executed in foreign currency and the fund's investment mandate is strictly external, but a reduction in the pace of future reserve accumulation would reduce the scale of RBI foreign exchange purchases over time and could allow some gradual rupee appreciation relative to the counterfactual of continued aggressive accumulation. To the extent that the RBI has used reserve accumulation partly to manage the exchange rate, as our DOLS estimation suggests through the sensitivity of reserve demand to the real effective exchange rate, some loss of that tool is a genuine trade-off worth acknowledging rather than dismissing. Capital flow volatility adds a final layer of complexity. The optimal reserve floor must be recalibrated in real time as crisis probabilities evolve, and the transition framework would benefit from a fail-safe mechanism that automatically pauses transfers to Tranche 2 whenever Tranche 1 falls below a defined threshold of its dynamically computed optimal level.

8.4 Policy Recommendations

Based on the empirical findings and institutional analysis above, this paper advances five specific policy recommendations. First, the Reserve Bank of India could move toward dynamic reserve targeting, complementing the current practice of accumulation with a formal reserve floor calibrated each quarter using the Jeanne and Rancière (11) welfare maximisation formula and updated crisis probability estimates. Second, the Global Financial Safety Net could be incorporated more explicitly into India's reserve strategy. India's bilateral swap arrangements, contingent reserve access, and IMF facilities represent real, if imperfect, sources of supplementary liquidity that reduce the required self-insurance buffer and could justify a somewhat lower reserve floor than the unadjusted baseline implies. Third, the scale of sterilised intervention could be tied more closely to what is required to maintain the dynamically computed liquidity floor, with accumulation beyond that threshold treated as a fiscal cost warranting periodic review rather than a default policy posture. Fourth, a gradual and rules-based reserve deployment framework could be established to transfer part of the structural excess to a dedicated investment vehicle over six to eight quarters, using a systematic deployment schedule that mitigates market impact and entry timing risk. Fifth, India could consider establishing a dedicated sovereign wealth fund, legally and operationally independent from both the RBI and the Ministry of Finance, with an investment mandate informed by the 8.25 percent nominal return

target derived in this paper, a strategic allocation of approximately 60 percent global equities, 35 percent global fixed income, and 5 percent real assets, quarterly public reporting broadly consistent with the Santiago Principles, and a fiscal rule that limits repatriation of returns to the government budget in any quarter in which Tranche 1 falls below 95 percent of its dynamically calibrated floor.

8.5 Limitations

The simulation results reported in this paper embed assumptions that bound the scope of the conclusions. The out-of-sample portfolio performance is computed using highly liquid index fund proxies that assume frictionless quarterly rebalancing at closing prices. A real sovereign wealth fund managing over $100 billion would face transaction costs, meaningful market impact on large trades, position size constraints in smaller asset classes, and redemption periods in private asset allocations that the simulation does not capture. The 8.25 percent target return mandate is derived from historical asset class returns over 2001 to 2025, a period that includes two extended equity bull markets. Future returns may be lower if the global equity risk premium compresses as institutional capital continues to flow into passive index strategies. The optimal reserve estimates are sensitive to the calibrated shock parameters, and smaller assumed shocks would lower the optimal floor and increase the measured excess. The crisis probability model achieves an AUC-ROC of 0.734, which reflects strong but not perfect discriminatory power, and the smoothed probability series may understate or overstate true sudden stop risk in any given quarter. These limitations suggest that the headline figure of approximately $204 billion in additional sovereign wealth generated over two decades should be treated as an indicative upper-bound estimate rather than a precise forecast. The directional finding, that deploying excess reserves into a diversified global portfolio generates higher risk-adjusted returns than the status quo, appears reasonably robust to downward adjustments in assumed returns. Even at a meaningfully lower equity risk premium than observed historically, the two-tranche framework would still be expected to generate more cumulative wealth than the status quo over a twenty-year horizon, though the margin would narrow.

9 Conclusion

This paper has examined whether India's large foreign exchange reserve stock is economically efficient and what the fiscal consequences of the current reserve management strategy are. Using quarterly data from 2000Q1 to 2025Q4, we have shown that India holds reserves persistently in excess of what a welfare-based optimality criterion requires. The Jeanne and Rancière (2011) framework, calibrated with time-varying crisis probabilities from an asymmetric complementary log-log model and empirically grounded shock parameters, estimates an optimal reserve floor of $431.45 billion as of 2025Q4 against actual holdings of $576.47

billion, implying an excess of $145.02 billion under the conservative baseline that assigns no usability to the Global Financial Safety Net. Once the GFSN is incorporated at a realistic 50 percent usability rate, the estimated excess rises to $190.70 billion. The Dynamic OLS reserve demand estimation suggests that this accumulation pattern is not random: it reflects systematic responses to financial deepening, exchange rate movements, sterilisation costs, and global risk conditions, but the level of accumulation it predicts is best understood as a structural outcome of these drivers rather than a welfare optimum in itself.

The fiscal cost of maintaining this excess appears large and has grown over time. The annual opportunity cost of holding reserves reached $14.624 billion in 2025, while the additional sterilisation cost was $0.064 billion, implying a total fiscal cost of $14.688 billion. In the counterfactual where only the optimal reserve level is held and the excess is deployed into a diversified global portfolio, the foregone returns reached approximately $51.17 billion in 2020 alone. Over the full 2004 to 2025 simulation horizon, the two-tranche framework generates roughly $204 billion more in cumulative sovereign wealth than the status quo, at a blended annualised return of 4.68 percent against 1.96 percent for the status quo. While these figures are simulation-based and subject to the limitations discussed above, they point to a sustained and compounding fiscal drag that has received relatively little attention in the existing literature on India's external sector policy.

The policy implications that follow are directional rather than prescriptive in every detail. India could consider moving toward more dynamic reserve targeting anchored to a welfare-informed floor along the lines derived here, incorporating the Global Financial Safety Net more explicitly into its reserve adequacy assessments, and exploring whether a legally independent sovereign wealth vehicle could manage part of the structural excess under a global investment mandate. The international precedents from Norway, Singapore, and China suggest that reserve-funded sovereign investment is institutionally feasible at scale, provided governance is rules-based, transparent, and reasonably insulated from short-term fiscal pressures. The barriers to implementation are real but appear to be manageable medium-term institutional challenges rather than fundamental obstacles: a legislative and governance design phase, a gradual and carefully sequenced capital transfer, and a steady-state operational review framework would all be needed.

The broader contribution of this paper is to suggest that the question of reserve efficiency deserves a more central place in India's macroeconomic policy agenda than it currently occupies. A reserve stock that appears to cost over one percent of GDP annually to maintain in excess of estimated optimal levels, at a time when the economy faces real constraints on public investment, infrastructure, and development financing, is a question worth taking seriously. The analytical tools to examine it exist, and this paper has applied them to India's case. The findings suggest that a carefully designed reform of reserve management could, in principle, help preserve macroeconomic stability while reducing the fiscal burden of sterilisation and generating

meaningful intergenerational wealth, without materially compromising the central bank's core crisis-fighting capacity.

References

[1] Ben-Bassat, A., & Gottlieb, D. (1992). *On the Effect of Opportunity Cost on International Reserve Holdings*. The Review of Economics and Statistics, 74(2), 327-332.

[2] Frenkel, J. A., & Jovanovic, B. (1981). *Optimal International Reserves: A Stochastic Framework*. The Economic Journal, 91(362), 507-514.

[3] Heller, H. R. (1966). *Optimal International Reserves*. The Economic Journal, 76(302), 296-311.

[4] International Monetary Fund. (2011). *Assessing Reserve Adequacy in Emerging Market Economies*. IMF Policy Paper.

[5] Ramachandran, K. (2006). *On the Upsurge of Foreign Exchange Reserves in India*. Reserve Bank of India Occasional Papers.

[6] Rodrik, D. (2006). *The Social Cost of Foreign Exchange Reserves*. International Economic Association.

[7] Das, & Nath, S. (2015). Assessing Reserve Adequacy in India. *Reserve Bank of India Occasional Papers*, 35 & 36(1 & 2), ISSN 0972-7493.

[8] Stock, J. H., & Watson, M. W. (1993). A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems. *Econometrica*, 61(4), 783-820.

[9] International Monetary Fund (2011). *Assessing Reserve Adequacy*. IMF Policy Paper. Washington, DC.

[10] International Monetary Fund (2016). *Adequacy of the Global Financial Safety Net*. IMF Policy Paper. Washington, DC.

[11] Jeanne, O., & Rancière, R. (2011). The Optimal Level of International Reserves for Emerging Market Countries: A New Formula and Some Applications. *The Economic Journal*, 121(555), 905-930.

[12] Aguilera, R. V., Capapé, J., & Santiso, J. (2023). The Heterogeneous and Evolving Roles of Sovereign Wealth Funds. *Journal of International Business Policy*, 6, 264-278.

[13] Gupta, A., & Van Nieuwerburgh, L. (2021). Valuing Private Equity and Real Estate Investments Strip by Strip. *The Journal of Finance*, 76(6), 2727-2777.

[14] Das, U. S., Mazarei, A., & van der Hoorn, H. (Eds.). (2010). *Economics of Sovereign Wealth Funds: Issues for Policymakers*. International Monetary Fund.

[15] Norges Bank Investment Management (NBIM). (2024). *Government Pension Fund Global Annual Report 2023*. Oslo, Norway.

A Data Appendix

The empirical analysis in this study relies on a comprehensive dataset spanning the period from 2000 to 2025. The data is broadly categorized into two segments: macroeconomic fundamentals used for the econometric estimations and reserve optimization, and financial market data used to calibrate the sovereign wealth fund proxies and benchmark returns. To ensure uniform periodicity, all higher frequency data was rigorously aggregated or resampled to a quarterly frequency (end of period).

Table 2: Macroeconomic and External Sector Variables (Part 1 of 2)

Variable	**Definition**	**Frequency**	**Source**
Foreign Exchange Reserves	Total official reserve assets excluding gold (USD Billions). Components include FCA, SDR, and IMF Reserve Tranche.	Quarterly	Haver Analytics / RBI
Real GDP	Gross Domestic Product at constant prices, converted to USD Billions.	Quarterly	Haver Analytics
Nominal GDP	Gross Domestic Product at current prices (USD Billions). Annualized and used as the denominator to scale macroeconomic flow ratios.	Quarterly	Haver Analytics
NEER	Nominal Effective Exchange Rate index.	Quarterly	Haver Analytics
REER	Real Effective Exchange Rate index. A decrease indicates depreciation.	Quarterly	Haver Analytics
Short Term External Debt	STED (USD Billions) based on residual maturity.	Quarterly	Haver Analytics
Portfolio Liabilities	Total outstanding foreign portfolio liabilities (USD Billions).	Quarterly	Haver Analytics

Table 3: Macroeconomic and External Sector Variables (Part 2 of 2)

Variable	Definition	Frequency	Source
Net FPI to GDP	Net Foreign Portfolio Investment flows scaled by annualized Nominal GDP (%).	Quarterly	Haver Analytics
Exports to GDP	Total exports of goods and services scaled by annualized Nominal GDP (%).	Quarterly	Haver Analytics
Imports to GDP	Total imports of goods and services scaled by annualized Nominal GDP (%).	Quarterly	Haver Analytics
Current Account Receipts to GDP	Total current account receipts scaled by annualized Nominal GDP (%).	Quarterly	Haver Analytics
Broad Money (M3)	Broad money supply, converted to USD Billions.	Quarterly	Haver Analytics
Opportunity Cost	The opportunity cost rate of holding reserves (percentage points) and the absolute cost (USD Billions).	Quarterly	Haver Analytics
US Federal Funds Rate	Effective Federal Funds Rate (%). Used as a macroeconomic proxy for global push factors driving external capital flows.	Quarterly	FRED (FEDFUNDS)

Table 4: Financial Market Data and Asset Proxies

Variable	Definition	Frequency	Source
US 10 Year Treasury Yield	10 Year Treasury Constant Maturity Rate (%). Used to calculate the structural term premium and risk free rate.	Quarterly	FRED (GS10)
Federal Funds Rate	Effective Federal Funds Rate (%). Base risk free rate for term premium calculation.	Quarterly	FRED (FEDFUNDS)
US 3 Month T-Bill Yield	3 Month T Bill Secondary Market Rate (%). Used as the baseline risk free rate and proxy for central bank deposit returns.	Quarterly	FRED (TB3MS)
US Consumer Price Index	CPI for All Urban Consumers. Used to calculate the baseline US inflation rate for wealth fund targets.	Quarterly	FRED (CPIAUCSL)
India Sovereign Spread Proxy	ICE BofA BBB US Emerging Markets Option-Adjusted Spread (%). Used as a proxy for the sovereign credit spread.	Quarterly	FRED (BAMLC0A4CBBB)
India 10 Year Bond Yield	Long term domestic borrowing rate (%). Used to calculate the opportunity cost of reserves.	Quarterly	FRED (INDIRLTLT01STM)
India 91 Day T-Bill Yield	Short term domestic policy rate (%). Used to calculate the sterilization cost of active reserve accumulation.	Quarterly	RBI Data
S&P 500 Index	Broad US equity market index. Used to calculate the historical global equity risk premium.	Quarterly	Yahoo Finance (^GSPC)
Global Equities Proxy	Vanguard Total Stock Market Index (VTSMX). Broad proxy for the global equity allocation.	Quarterly	Yahoo Finance
Global Bonds Proxy	Vanguard Total Bond Market Index (VBMFX). Broad proxy for the global fixed income allocation.	Quarterly	Yahoo Finance
Real Assets Proxy	Vanguard Real Estate Index (VGSLX). Proxy for the unlisted and real asset allocation.	Quarterly	Yahoo Finance
US Short Term Securities Proxy	iShares 1 to 3 Year Treasury Bond ETF (SHY). Proxy for the central bank's foreign currency asset securities deployment.	Quarterly	Yahoo Finance

B Summary Statistics and Correlation Analysis

Table 5 presents the descriptive statistics for the primary macroeconomic variables utilized in the cointegration and optimal reserve models. Figure 14 provides the cross-correlation matrix to check for multicollinearity among the core predictors of exchange market pressure.

Table 5: Descriptive Statistics of Key Macroeconomic Variables (2000:Q1 – 2025:Q4)

Variable	**Mean**	**Std. Dev.**	**Min**	**Max**
Reserves / GDP	0.1677	0.0363	0.0711	0.2552
Broad Money (M3) / GDP	0.8079	0.0809	0.5199	0.9769
Imports / GDP	0.0642	0.0154	0.0352	0.0931
Exports / GDP	0.0437	0.0075	0.0257	0.0584
Short-Term External Debt (USD Bn)	68.69	44.48	2.74	140.04
Portfolio Liabilities / GDP	0.0082	0.0085	-0.0124	0.0299
Opportunity Cost of Reserves (%)	4.5002	2.2729	1.0777	10.2727
Nominal Effective Exchange Rate (NEER)	89.21	18.79	62.43	121.72
Real Effective Exchange Rate (REER)	95.90	8.56	81.54	110.25
IMF Assessing Reserve Adequacy (ARA) Ratio	1.4825	0.2801	0.9428	2.1647
Current Account Volatility	0.0843	0.0355	0.0239	0.1513
VIX Index	19.8387	7.3928	10.3079	58.5959

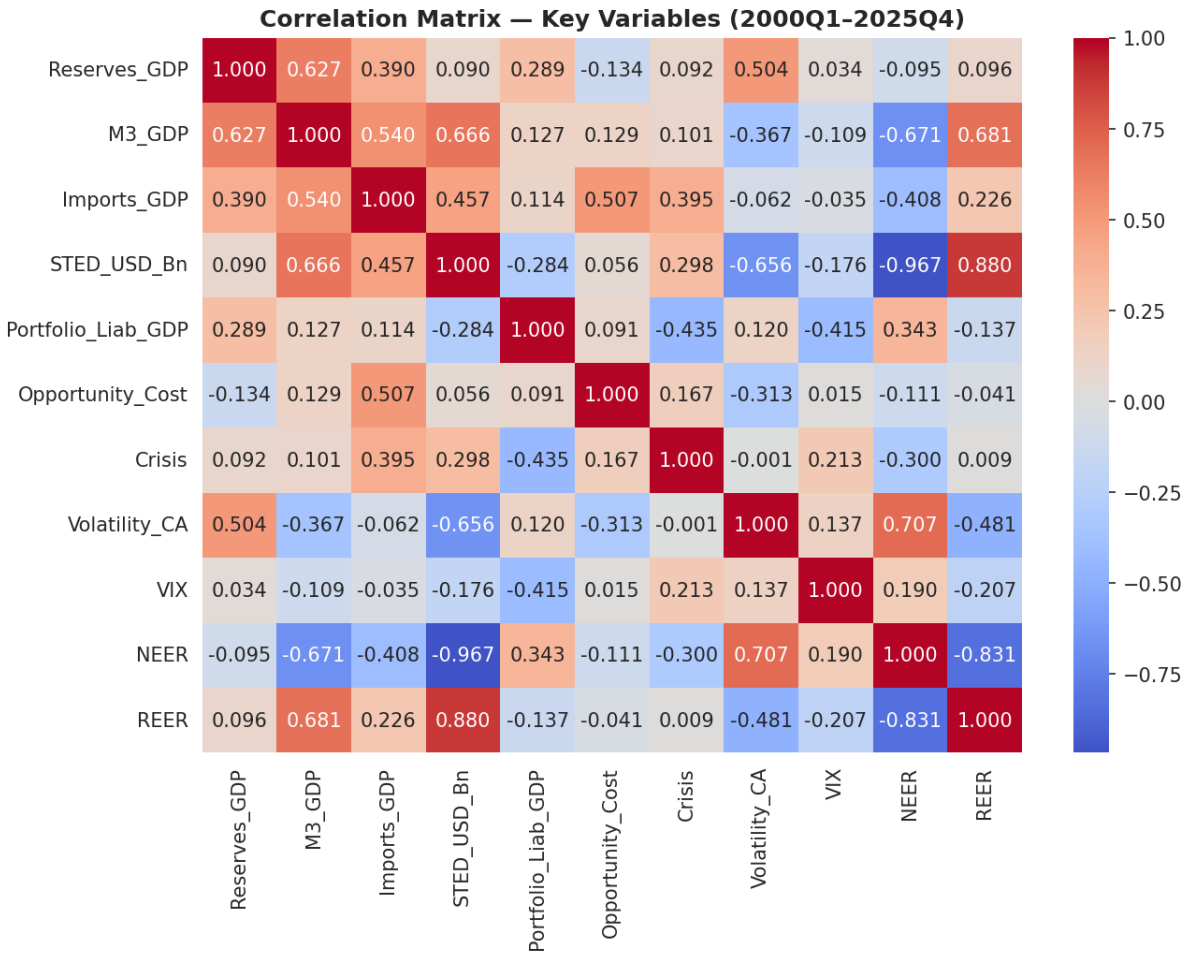


Figure 14: Figure A1: Correlation Matrix of Key Variables. *Notes:* Heatmap of Pearson correlation coefficients.

C Financial Optimization and Comparative Performance Metrics

The following table summarizes the primary empirical inputs, asset class parameters, and full sample performance results discussed in this paper. The underlying data spans from 2000 to 2025, with the out-of-sample simulation covering the 2004 to 2025 period.

Table 6: Financial Optimization and Comparative Performance Metrics

Metric	**Value / Return**	**Volatility**
Panel A: Target Return Mandate Derivation (2001 to 2025)		
US 10-Year Treasury Yield (Risk-Free Rate)	3.20%	
India Sovereign Spread (BBB EM Proxy)	1.19%	
Global Equity Risk Premium	3.76%	
Institutional Management Frictions	0.10%	
Financial Hurdle Rate (Sovereign Cost of Capital)	8.25%	
Macroeconomic Hurdle (10-Year ARA CAGR)	5.77%	
Binding Target Return Mandate	**8.25%**	
Panel B: Asset Class Historical Parameters (2001 to 2025)		
Global Equities	10.91%	16.65%
Global Bonds	3.43%	4.31%
Real Assets	10.00%	21.24%
Panel C: Full Sample Performance Simulation (2004 to 2025)		
Status Quo (100% Liquidity)	1.96%	1.63%
Proposed 2-Tranche Blended	4.68%	5.12%
Total Cumulative Wealth Generated (Status Quo)	\$140.80B	
Total Cumulative Wealth Generated (Proposed)	\$344.58B	
Excess Sovereign Wealth Generated	**+\$203.78B**	
Panel D: 2025 Single Year Sanity Check		
2025 Yield (Status Quo)	4.83%	
2025 Yield (Proposed 2-Tranche Mechanism)	7.68%	
2025 Total Wealth Generated (Status Quo)	\$28.10B	
2025 Total Wealth Generated (Proposed)	\$44.93B	
2025 Excess Wealth Generated	**+\$16.83B**	

D Econometric Estimations and Diagnostics

This section details the progression from the baseline Ordinary Least Squares (OLS) estimation to the robust Dynamic Ordinary Least Squares (DOLS) specification. The baseline OLS model exhibits severe serial correlation (Breusch-Godfrey $p < 0.001$) and utilizes non-stationary variables, rendering standard inference invalid and necessitating the cointegration framework used in the main text.

Table 7: Baseline OLS Reserve Demand Estimation

Dependent Variable: Reserves / GDP	**Coefficient**	**Robust Std. Err.**	**P-Value**
Constant	-0.1514***	0.0490	0.002
Imports / GDP	-0.0328	0.3650	0.928
NEER	0.0011***	0.0000	0.003
Volatility (Current Account)	0.0851	0.1390	0.540
Volatility (Portfolio Liabilities)	-0.0021	0.0110	0.853
M3 / GDP	1.1614***	0.1550	0.000
REER Undervaluation	0.0087**	0.0040	0.016
Opportunity Cost	-0.0033	0.0030	0.191
Observations	81	R-squared: 0.686	
Breusch-Godfrey LM Stat (4 lags)	52.27	P-Value: 0.000	

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are HC3 robust.

Severe serial correlation and unit roots invalidate standard OLS inference here.

Table 8: Dynamic Ordinary Least Squares (DOLS) Structural Elasticities

Dependent Variable: ln(Reserves / GDP)	**Coefficient**	**P-Value**	**Significance**
Constant	3.1010	0.0004	***
ln(Imports / GDP)	-0.1085	0.2867	
ln(Volatility CA)	-0.0193	0.5917	
ln(Volatility FPI)	-0.1523	0.0006	***
ln(M3 / GDP)	1.4530	0.0000	***
Opportunity Cost	-0.0362	0.0000	***
ln(REER)	-1.0683	0.0000	***
VIX	0.0073	0.0000	***
Observations	80	Model R-squared: 0.8579	

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Differenced lead and lag terms are included in the estimation to correct for endogeneity but are suppressed from this table for brevity.

Table 9: DOLS Robustness Diagnostics and Cointegration Tests

Panel A: Variance Inflation Factor (Multicollinearity Check)	
Variable	**VIF Score**
ln(Imports / GDP)	3.0270
ln(Volatility CA)	2.2808
ln(Volatility FPI)	3.4980
ln(M3 / GDP)	1.3441
Opportunity Cost	3.5921
ln(REER)	3.0297
VIX	1.1273
Panel B: System Cointegration Verification	
Test Type	**Statistic / Result**
Engle-Granger (Residual ADF)	-4.6658 ($p = 0.0001$, Cointegrated)
Johansen Trace Statistic ($r = 0$)	263.0493 $>$ 159.5290 (5% Critical Value)

Notes: Panel A confirms the absence of fatal multicollinearity (all VIFs < 10).

Panel B verifies that a stable, long-run structural relationship exists.

E Crisis Definition Audit and Threshold Robustness

To calculate the probability of a sudden stop (π) for the optimal reserve calibration, we construct a composite Exchange Market Pressure index. This index relies on six distinct criteria spanning currency depreciation, reserve depletion, and capital flow reversals. A specific quarter is classified as a crisis if at least two of the six criteria trigger simultaneously. Following standard literature, consecutive crisis quarters are clustered into a single episode, and a new episode can only begin after at least two tranquil, non-crisis quarters.

Under this baseline definition, the model identifies 27 crisis quarters out of the 104 quarter sample (26.0 percent), grouped into 11 distinct stress episodes. Table 10 details the specific thresholds for each individual criterion and how frequently they triggered over the sample period.

Table 10: Individual Crisis Criteria and Trigger Frequencies (2000 to 2025)

Criterion (Threshold)	**Quarters Flagged**	**Frequency (%)**
C1: REER deviation (< -1.0 Standard Deviation)	16	15.4%
C2: Reserve drawdown ($<$ Mean $- 1.5$ S.D.)	6	5.8%
C3: Net FPI to GDP ($<$ Mean $- 0.75$ S.D.)	24	23.1%
C4: REER quarter over quarter change ($< -2.0\%$)	16	15.4%
C5: STED to Reserves Ratio ($>$ Mean $+ 1.0$ S.D.)	13	12.5%
C6: Portfolio outflow absolute change ($<$ Mean $- 1.0$ S.D.)	15	14.4%

To validate the accuracy of this mechanical filter, we conduct a sanity check against known macroeconomic stress events. Table 11 demonstrates that the baseline composite index successfully captures all major global shocks affecting India. The model does not flag the 2000 Dot-com bust, which accurately reflects India's relatively insulated capital account during that specific period.

Table 11: Sanity Check: Model Identification of Known Historical Episodes

Historical Event	**Onset Quarter**	**Criteria Score**	**Assessment**
Dot-com Bust	2000Q4	0 out of 6	Correctly Missed (Insulated)
Global Financial Crisis	2008Q4	3 out of 6	Correctly Flagged
Taper Tantrum	2013Q3	3 out of 6	Correctly Flagged
COVID-19 Capital Flight	2020Q1	2 out of 6	Correctly Flagged
Global Rate Shock	2022Q1	3 out of 6	Correctly Flagged

We also identify potential false positives at the end of the sample (late 2024 to 2025). During this period, the index flags several quarters due to reserve drawdowns and REER movements. However, a manual narrative audit suggests these movements were driven by active central bank

foreign exchange interventions to manage localized volatility, rather than a true exogenous sudden stop. Consequently, we create a "Conservative" variant that manually overrides and excludes these specific intervention quarters.

Finally, Table 12 presents a robustness check on the index thresholds. We test a "Strict" variant requiring three simultaneous criteria to trigger, which restricts the sample to only the most severe systemic shocks (7 quarters). We also test tightening the REER deviation parameter. The baseline approach (≥ 2 criteria, 1.0 S.D. REER deviation) remains the preferred specification, as it optimally balances capturing severe sudden stops alongside the moderate capital flow reversals that dictate emerging market reserve adequacy.

Table 12: Threshold Robustness and Sample Variations

Model Specification	**Total Crisis Quarters**	**Sample Frequency (%)**
Panel A: Composite Score Variations		
Baseline Composite (≥ 2 out of 6 criteria)	27	26.0%
Conservative Variant (Excluding 2024 to 2025)	23	22.1%
Strict Composite (≥ 3 out of 6 criteria)	7	6.7%
Panel B: REER Deviation (Criterion 1) Sensitivity		
REER Deviation at 1.0 S.D. (Baseline)	16	15.4%
REER Deviation at 1.5 S.D.	8	7.7%
REER Deviation at 2.0 S.D.	3	2.9%

F Cloglog Model Estimates for Crisis Probability

For completeness and to document the methodological progression described in Section 5.2, Appendix Table 13 reports the baseline Probit model results estimated prior to the transition to the asymmetric cloglog specification. The Probit model is estimated on the same sample and predictor set as the cloglog, and its poor fit, evidenced by a McFadden Pseudo R^2 of just 0.089 and two of three slope coefficients failing to reach conventional significance levels, directly motivates the methodological contribution of this paper.

Table 13: Baseline Probit Model Results (Prior to Cloglog Upgrade)

Dependent Variable: Crisis Indicator $(0/1)$*. Estimator: Maximum Likelihood Probit. Observations: 96. Standard errors are non-robust (MLE default).*

Variable	Coef.	Std. Err.	z	$P>\|z\|$	95% Conf.	Interval
Constant	-2.1720^{***}	0.642	-3.382	0.001	-3.431	-0.913
REER Overvaluation (Lag Avg)	0.1045	0.068	1.537	0.124	-0.029	0.238
Financial Openness (Lag Avg)	102.8779	84.100	1.223	0.221	-61.954	267.710
Ext. Liab. / M3 (Lag Avg)	28.7861^{**}	12.725	2.262	0.024	3.845	53.727
Log-Likelihood: -51.953	Null Log-Likelihood: -57.036					
McFadden Pseudo R^2: 0.0891	LLR p-value: 0.0172					

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Predictors are averaged across the fourth and eighth quarterly lags to match the annual identification horizon of Das and Nath (2015) and Jeanne and Rancìere (2011). The weak fit and near-zero significance of two of three slope coefficients reflects the structural limitation of the symmetric Probit link function when applied to rare, asymmetric crisis events, motivating the transition to the cloglog specification reported in Table 14.

Appendix Table 14 presents the preferred cloglog model results, which are discussed in full in Section 5.2 of the main text. The model is estimated using a Binomial family with a complementary log-log link function and HC1 robust standard errors. The inclusion of the lagged US Federal Funds rate as a global push factor, combined with the asymmetric link function, raises the McFadden Pseudo R^2 to 0.139 and delivers an AUC-ROC of 0.737, confirming strong discriminatory power between crisis and tranquil quarters.

Table 14: Cloglog Model Regression Results

Dependent Variable: Crisis	**Coefficient**	**Robust S.E.**	**z-statistic**	$P > \|z\|$
Constant	-5.2164^{***}	1.388	-3.758	0.000
REER Overvaluation (Lag Avg)	0.1566^{**}	0.078	2.019	0.043
Financial Openness (Lag Avg)	135.9732	99.969	1.360	0.174
External Liab. / M3 (Lag Avg)	69.0168^{***}	24.501	2.817	0.005
US Federal Funds Rate (Lag Avg)	0.3286^{**}	0.130	2.526	0.012
Model Diagnostics				
Observations	96			
Log-Likelihood	-49.131			
Null Log-Likelihood	-57.036			
McFadden's Pseudo R^2	0.1386			
AUC-ROC Score	0.7370			

Note: The model is estimated using a Binomial family with a cloglog link function. Standard errors are HC1 robust. Predictor variables are smoothed over a two-year lag average ($t-4$ and $t-8$) to capture structural buildup prior to the crisis onset. Significance levels are denoted as: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

G Jeanne and Rancière Calibration: Parameter Transparency and Uncertainty

To ensure complete methodological transparency and address standard reviewer requirements for theoretical models, this section details the exact parameter calibration for the Jeanne and Rancière optimization framework. Table 15 provides the baseline derivations. Table 16 conducts a one at a time sensitivity analysis to isolate the marginal impact of each parameter on the optimal reserve threshold (ρ^*). Finally, Table 17 models the aggregate uncertainty bands driven by time-varying crisis probabilities and varying stress magnitude assumptions.

Table 15: Baseline Parameter Calibration for the Jeanne and Rancière Model

Parameter	**Symbol**	**Value**	**Derivation / Source**
Risk aversion	σ	2.0000	Standard Jeanne and Rancière (2011) assumption
Crisis probability	π	0.2563	Mean of Cloglog estimated crisis probabilities
Financing gap	λ	0.0370	1.5 standard deviation drop in Exports/GDP + FPI/GDP
Output loss	γ	0.0775	1.5 standard deviation drop in HP filtered real GDP growth
Depreciation	ΔQ	0.1286	1.5 standard deviation depreciation in REER index
Opportunity Cost wedge	$r - g$	-0.0472	Mean US 3M yield proxy minus mean potential GDP growth
Term premium	δ	0.0131	Calculated from 10Y vs 3M US Treasury spread
Baseline Optimal Reserves	ρ^*	**0.1088**	**10.88% of GDP**
Actual Reserves to GDP		**0.1677**	**16.77% of GDP (Sample Average)**

Table 16: Sensitivity of Optimal Reserves (ρ^*) to Parameter Variation

	Parameter Inputs			**Resulting Optimal Threshold**		
Parameter	**Low**	**Base**	**High**	ρ^* **(Low)**	ρ^* **(Base)**	ρ^* **(High)**
Risk aversion (σ)	1.0000	2.0000	3.0000	0.1087	0.1088	0.1088
Crisis probability (π)	0.1282	0.2563	0.3845	0.1071	0.1088	0.1104
Financing gap (λ)	0.0185	0.0370	0.0555	0.0898	0.1088	0.1277
Output loss (γ)	0.0387	0.0775	0.1162	0.0733	0.1088	0.1442
Depreciation (ΔQ)	0.0643	0.1286	0.1929	0.1090	0.1088	0.1063
Opportunity Cost ($r - g$)	-0.0672	-0.0472	-0.0272	0.1087	0.1088	0.1088

Table 17: Uncertainty Bands for Optimal Reserves (ρ^*)

Panel A: Uncertainty Induced by Estimated Cloglog Probabilities (π)

Percentile	**Probability (π)**	**Optimal (ρ^*)**	**Excess vs Actual**
10th Percentile	0.0574	0.0635	+0.1042
25th Percentile	0.1750	0.1077	+0.0599
50th Percentile	0.2753	0.1090	+0.0587
75th Percentile	0.3462	0.1099	+0.0577
90th Percentile	0.4123	0.1108	+0.0569

Panel B: Uncertainty Induced by Stress Scenario Magnitudes

Shock Magnitude	**Gap (λ)**	**Loss (γ)**	**Optimal (ρ^*)**	**Excess vs Actual**
1.00 S.D. Shock	0.0319	0.0645	0.0923	+0.0754
1.25 S.D. Shock	0.0345	0.0710	0.1005	+0.0671
1.50 S.D. Shock	0.0370	0.0775	0.1088	+0.0589
1.75 S.D. Shock	0.0395	0.0840	0.1169	+0.0508
2.00 S.D. Shock	0.0421	0.0904	0.1250	+0.0427