

Mathematical Modeling and Optimization of Nepal's Export–Import Dynamics for Sustainable Economic Growth

Suresh Kumar Sahani^{1,2}, Tsair-Fwu Lee¹, Bharat Kumar Sah^{2*}, Digvijay Pandey³,
Binay Kumar Pandey⁴

¹ National Kaohsiung University of Science and Technology, Taiwan

² Faculty of Science Technology and Engineering, Rajarshi Janak University, Janakpurdham,
Nepal. Email: sureshsahani54@gmail.com^{1,2}

¹ National Kaohsiung University of Science and Technology, Taiwan
Email: tflee@nkust.edu.tw¹

^{*2} Faculty of Science Technology and Engineering, Rajarshi Janak University, Janakpurdham,
Nepal, Email: bharatsah@rju.edu.np

³ Department of Technical Education Uttar Pradesh, India,
Email: digit11011989@gmail.com

⁴ Department of Information Technology, College of Technology, Govind Ballabh Pant
University of Agriculture and Technology Pantnagar Uttarakhand, India,
Email: binaydece@gmail.com

*Corresponding Author: bharatsah@gmail.com

Abstract. Nepal presents one of the most instructive external-balance puzzles in the developing world: five decades of progressive trade liberalization have coexisted with a secular deterioration of the trade balance, from a deficit of roughly five percent of GDP in the mid-1970s to nearly twenty-eight percent in 2024, while exports contracted from about fifteen percent of GDP at their turn-of-the-century peak to under seven percent. This paper constructs a fully articulated mathematical framework to explain this trajectory and to design a correction. We develop a continuous-time dynamic model of the external accounts in which the net international investment position and the real effective exchange rate co-evolve under an intertemporal solvency constraint, we characterize the equilibrium structure of the resulting planar system, and we prove that the unregulated economy lies on a saddle path along which external solvency is restored only through implausibly large real depreciations. The model is then enriched with lagged volume adjustment, which delivers a Marshall–Lerner-consistent J-curve with an impact deterioration of 1.65 percentage points of GDP and a break-even horizon of roughly eight and a half quarters. Cointegration and error-correction representations, a structural gravity block, and a Kalman-filtered stochastic extension provide the econometric and filtering scaffolding for calibration. Finally, we solve a discounted linear-quadratic optimal control problem in closed form: the Hamiltonian system decouples into stable and unstable eigenvalue pairs summing to the discount rate, the value function solves an algebraic Riccati equation exactly, and the resulting linear feedback rule cuts the discounted external-adjustment loss by about sixty-one percent over a fifty-year horizon relative to passive drift. Numerical experiments, including a sensitivity surface over structural elasticities and stochastic fan charts

under real-exchange-rate diffusion, quantify the fragility of unmanaged adjustment and the robustness of the optimized policy.

Keywords: trade balance dynamics; saddle-path stability; Marshall–Lerner condition; cointegration; structural gravity; optimal control; Riccati equation; Nepal

1. Introduction

A wedge of the Himalaya separates two of the fastest-growing large economies on earth, and yet the economy occupying that wedge has watched its own external position deteriorate almost without interruption for three decades. Nepal liberalized its current account in the early 1990s, unified and stabilized its exchange rate against the Indian rupee, acceded to the World Trade Organization in 2004, and layered regional trade agreements on top of most-favored-nation treatment — and through all of this the share of exports in GDP fell from a peak of nearly fifteen percent around the turn of the century to under seven percent in recent years, while imports climbed toward forty percent of GDP before the post-2022 correction. The trade deficit in goods and services now hovers in the neighborhood of twenty-eight percent of GDP, financed overwhelmingly by the remittances of migrant workers. No amount of institutional optimism can make that configuration permanently safe: a financing structure of this kind is exposed to the Gulf oil cycle, to Indian and East Asian labor-demand shocks, and to the goodwill of trading partners whose transit routes the country does not control.

The growing use of machine learning in the financial sector has created new possibilities for improving credit risk assessment and lending decisions. Sahani et al. (2026) explored how explainable machine learning can support credit risk management in Nepalese cooperative banks. The study focuses particularly on advanced machine learning techniques such as XGBoost and LightGBM, along with SHAP and LIME, which help explain how these models arrive at their predictions.

A major concern highlighted by Sahani et al. (2026) is that accurate predictions alone are not enough in financial decision-making. Banks also need to understand and explain why a borrower is considered risky or why a particular lending decision is made. To address this issue, the authors developed a theoretical framework that combines feature selection, SMOTE for handling imbalanced data, gradient boosting, and explainable AI techniques. Their study suggests that combining predictive performance with transparency could help Nepalese cooperative banks make more informed lending decisions, strengthen credit risk management, and build greater confidence in the use of machine learning.

The central question of this paper is what mathematics can say about such an economy. Our ambition is not merely to document, but to model: to write down a compact dynamical system whose equilibria, stability properties, and policy-controllable manifolds capture the anatomy of Nepal's external imbalance, and then to solve, exactly where possible, the problem of the social planner who must steer the system back toward sustainability. The exercise is deliberately theoretical. We take seriously the possibility that the most valuable contribution of formal modeling in a data-rich but estimate-poor environment is clarity about mechanisms

— which feedback loops dominate, which adjustments are fast and which are slow, and where policy leverage actually resides.

Three stylized mechanisms organize everything that follows. The first is external wealth accumulation: the net international investment position evolves according to the current account, which couples the trade balance to interest flows, so that today's deficit is tomorrow's debt service. The second is real exchange rate adjustment: persistent trade imbalances exert pressure on the real effective exchange rate, and the elasticity of the trade balance with respect to that rate — governed, as we show, by a Marshall–Lerner-type condition corrected for Nepal's chronic import excess — determines whether depreciation is a workable escape route. The third is structural policy: export promotion, trade facilitation, and diversification shift the autonomous component of the trade balance itself, an instrument that is costly to exercise but that operates on the level of the problem rather than on its price.

We assemble these mechanisms into a two-state continuous-time system in the net foreign asset position and the real exchange rate, study its phase portrait, and prove that it possesses a unique interior equilibrium with saddle-path structure: one eigenvalue of the linearization is positive and one is negative, so exactly one trajectory delivers convergence, and it requires the exchange rate to jump onto a specific linear manifold at the instant any disturbance strikes. The slope of that manifold, which we derive in closed form, quantifies the real depreciation an indebted Nepal must absorb per percentage point of external wealth lost — and the number is sobering. Because the equilibrium level of net foreign assets that would make the autonomous deficit consistent with solvency lies far outside any politically feasible range, the unregulated economy is, in a precise sense, unsustainable: its welfare integral diverges on the relevant horizon. This is the analytical core of the paper, and it is what motivates turning from positive analysis to optimal control.

The normative half of the paper poses a discounted linear-quadratic planning problem: choose a path of structural export-promotion effort to minimize a loss functional penalizing deviations of external wealth from its solvency norm, misalignment of the real exchange rate from its fundamental anchor, and the convex cost of policy effort. We solve this problem twice, deliberately, by the two canonical routes of dynamic optimization. Pontryagin, Boltyanskii, Gamkrelidze, and Mishchenko (1962) yield a four-dimensional Hamiltonian system whose eigenvalues occur in pairs summing to the discount rate — a symmetry we verify numerically to machine precision — and whose bounded solutions trace the optimal trajectory. Dynamic programming delivers the value function in closed form through the algebraic Riccati equation, and hence a linear feedback rule that is valid globally, not merely along a single path. The two solutions coincide, as they must, and the feedback rule is directly implementable: it prescribes an initial export-promotion push on the order of six to ten percentage points of GDP-equivalent effort in a crisis scenario — an implied trade-balance improvement of three quarters to slightly over one percentage point of GDP — decaying smoothly as the external position repairs itself. The paper contributes to several literatures at once, and it is worth being explicit about where the novelty lies. Relative to the empirical trade literature on Nepal, which has concentrated on single-equation estimates of export and import demand in the Dickey–Fuller (1979) and

Pesaran, Shin, and Smith (2001) tradition, we provide a fully articulated dynamic theory whose parameters are disciplined by that evidence rather than replacing it. Relative to the theoretical literature on external sustainability, which typically works in discrete time with ad hoc adjustment costs, our continuous-time formulation yields closed-form saddle-path and feedback solutions that make every comparative-statics result transparent. And relative to the optimal control tradition in macroeconomics, we exhibit the complete chain from primitive trade elasticities, through the Marshall–Lerner aggregation with an import-excess correction, through the Riccati equation, to a numerically implemented feedback rule — a chain whose every link can be checked against the numbers we report. The machinery is generic; only the calibration is Nepali.

The rest of the paper unfolds as follows. Section 2 assembles the stylized facts and the analytical setting. Section 3 constructs the core dynamic model and characterizes its equilibrium and saddle-path stability. Section 4 embeds lagged volume adjustment, derives the J-curve and the aggregate trade elasticity from primitives, and states the corrected Marshall–Lerner condition. Section 5 develops the cointegration, error-correction, and vector autoregressive representations that discipline the empirical content of the calibration, and Section 6 adds the structural gravity block and the measurement of export concentration. Section 7 introduces the stochastic extension and the filtering apparatus. Section 8 formulates and solves the optimal control problem by both the maximum principle and dynamic programming. Section 9 reports the numerical solution and diagnostics, Section 10 the policy experiments, and Section 11 the implications for Nepali trade policy. Section 12 concludes.

2. Stylized Facts and Analytical Setting

Four facts define the analytical problem, and each survives virtually any reasonable treatment of the underlying data. They are presented here in deliberately rounded, model-consistent form — assembled from the external-sector tables of the Ministry of Finance (2023), the balance-of-payments statistics of Nepal Rastra Bank (2023), and the World Bank’s (2023) country updates — because the argument of this paper depends on their orders of magnitude and their shapes, not on the second decimal place of any single series.

Fact one: the export share has been on a downward path for a quarter century. Merchandise and services exports together peaked near fifteen percent of GDP around 2000–2001, driven by the carpet and readymade-garment boom of the 1990s, and have since fallen to under seven percent. The timing is not accidental: the phase-out of the Multi-Fibre Arrangement in 2005 removed the quota rents on which the garment niche was built, the 2015 earthquakes and the accompanying disruption of the Indian border trade destroyed the recovery that had begun, and the appreciation of the real exchange rate within the rupee peg eroded competitiveness against both India and third markets. Nepal is thus a leading example of what the literature calls premature deindustrialization — the retreat from tradable production before the economy has exhausted its low-wage advantage.

Fact two: imports have grown structurally, not cyclically. The import share rose from about ten percent of GDP in the mid-1970s to twenty-three percent by 2000 and close to forty percent by

2016, dipping and surging with remittance cycles before the emergency import restrictions of fiscal year 2021/22 produced a sharp but partial compression. The composition of this growth matters as much as its level: petroleum, transport equipment, machinery, and intermediate goods dominate, which means that imports embody the capital formation and energy inputs on which any export expansion would itself depend. A model that treats import compression as a free lunch would therefore mislead; the optimization problem below prices all adjustment instruments, including their side effects.




Figure 1. Stylized external-sector trajectories of Nepal, 1975–2024. Panel (a): exports, imports, and personal remittance inflows as shares of GDP. Panel (b): the trade balance of goods and services. Series are model-consistent renderings of the aggregates reported in official publications; they preserve documented magnitudes and turning points (the garment-boom export peak near 2000, the post-2015 deterioration, the 2022 import-compression episode) rather than individual annual observations.

Fact three: remittances, not export earnings, finance the gap. Personal transfer inflows climbed from negligible levels in the 1970s to the neighborhood of a quarter of GDP — among the highest ratios recorded anywhere — and their trajectory closely shadows the import share with a lag. This financing structure has two consequences that any honest model must confront. First, it relaxes the immediate balance-of-payments constraint while leaving the trade balance itself untouched, so the current account deficit understates the structural problem. Second, it appreciates the real exchange rate through the spending channel, taxing the tradable sector that would eventually have to replace the transfers — a classic transmission of what the literature terms the Dutch disease, operating through labor migration rather than through a natural resource.

Fact four: the export basket has not merely shrunk; it has failed to diversify in the dimensions that matter. Figure 2 traces the stylized composition of merchandise exports over 2001–2024 together with the Herfindahl–Hirschman index of concentration and the effective number of products, computed formally in Section 6. The collapse and partial recovery of the two legacy clusters — carpets and garments, which together exceeded sixty-five percent of merchandise exports at the turn of the century — dominates the picture, while the arrival of cardamom, tea, and, most recently, cross-border electricity sales to India has broadened the basket only at the

margin. The index values tell a nuanced story: measured concentration first fell as the legacy clusters shrank faster than new products grew, then rose again as electricity exports scaled, so that the effective number of products ends the period only modestly above where it began.




Figure 2. Export concentration of Nepal. Panel (a): stylized composition of merchandise exports by product cluster, 2001–2024. Panel (b): the Herfindahl–Hirschman index over the same shares (left axis) and the effective number of products, the exponential of Shannon entropy (right axis). The indices are computed from the shares in panel (a) exactly as defined in equations (30) and (31).

Table 1 gathers the headline indicators. Three features deserve emphasis before any equation is written. The deterioration of the trade balance is strikingly persistent — deficits in every sub-period since 1975, with no mean reversion toward zero anywhere in the sample. The scale of the remittance offset has grown faster than the gap it finances in bad years, which is precisely why the underlying imbalance could widen for so long without a crisis of the classical kind. And the only episode of rapid deficit reduction in the record — the import compression of 2022–2023, engineered through administrative restrictions — carried visible costs in input shortages and activity, a reminder that the speed limit on external adjustment binds in practice as well as in theory.

Table 1. Stylized external-sector indicators of Nepal, selected years (percent of GDP, model-consistent rounding).

Indicator	1975	1990	2000	2010	2015	2022	2024
Exports of goods and services	5.0	11.0	14.7	10.0	8.5	6.0	5.8
Imports of goods and services	10.5	18.0	23.3	29.0	36.5	39.0	33.5
Trade balance	−5.5	−7.0	−8.6	−19.0	−28.0	−33.0	−27.8
Personal remittance inflows	0.4	2.5	10.5	21.5	27.0	25.0	26.0
Real effective exchange rate (2010 = 100)	108	104	99	100	96	94	92

Indicator	1975	1990	2000	2010	2015	2022	2024
Herfindahl–Hirschman index (merchandise)	0.28	0.27	0.27	0.25	0.24	0.30	0.30

Two institutional features complete the setting, because the model of Section 3 is built around them. The first is the exchange rate arrangement: since February 1993 the Nepalese rupee has been pegged to the Indian rupee at 1.60 per Indian rupee, and through the overwhelming weight of India in Nepal's trade this pins the nominal bilateral core of the real effective exchange rate, leaving relative prices, and hence the real rate itself, to adjust through domestic inflation differentials and administered prices. Real exchange rate management in Nepal is therefore indirect — it operates through the price level, through the pace of credit, and through the terms on which imports are permitted — which is exactly why the model treats the real rate as a slowly equilibrating state variable rather than as a policy jump variable. The second feature is geography: landlocked, transit-dependent, and bordered by the two largest economies in Asia, Nepal faces trade costs that no domestic reform fully removes. In the gravity block of Section 6 these costs enter as fixed bilateral frictions; in the optimal control problem of Section 8 they are part of what structural effort must buy down.

Throughout the paper we normalize all flow variables as ratios to GDP and all price variables in logarithms, with the real effective exchange rate q defined so that an increase is a real depreciation. The net international investment position, b , is measured as claims of residents on the rest of the world minus liabilities, again relative to GDP, so that Nepal's negative net position corresponds to $b < 0$. Sustained arithmetic on these definitions underlies everything from the current-account identity of the next section to the transversality conditions of the planning problem, and fixing the accounting now avoids a dozen small ambiguities later.

3. A Dynamic Model of the External Accounts

3.1 State Variables and the Accounting Backbone

Let Y denote real GDP; let X and M denote exports and imports of goods and services valued at current world prices converted into domestic currency and deflated by the domestic price level; and let q denote the logarithm of the real effective exchange rate, constructed so that an increase in q is a real depreciation. Two ratios organize the accounting. The trade balance ratio and the net international investment position (NIIP) are defined by

$$B \equiv \frac{X - M}{Y}, \quad b \equiv \frac{A - L}{Y}, \quad q \equiv \ln\left(\frac{E \cdot P^*}{P}\right) \quad (1)$$

where A and L are the gross foreign assets and liabilities of the resident sector, E is the nominal effective exchange rate, and P and P^* are the domestic and foreign price levels. The NIIP evolves according to the current-account identity: the change in net foreign assets equals the trade balance plus net factor income, which at the constant world real interest rate r is rb . In flow form,

$$\dot{b} = r \cdot b + B, \quad (2)$$

so the stock of external wealth today is the accumulated sum of past trade surpluses and interest flows. Two features of equation (2) carry the weight of the whole analysis. First, the interest term makes the dynamics self-reinforcing: a debtor country with $b < 0$ pays $r|b|$ each period, so the trade surplus required to hold net wealth constant grows with indebtedness itself. Second, the equation is exact rather than behavioral; all behavioral content enters through B , to which we now turn. For a creditor economy a permanent trade deficit is harmless; for a debtor it is not. The long-run solvency requirement is that the trade balance eventually cover interest obligations,

$$\lim_{t \rightarrow \infty} B(t) = -r \cdot b^*, \quad (3)$$

3.2 Behavioral Blocks

The trade balance is modeled as a linear reduced form in the real exchange rate around its long-run component, with the slope derived from primitives in Section 4:

$$B(q) = \bar{B} + \varepsilon_q \cdot q, \quad \varepsilon_q > 0, \quad (4)$$

where $\bar{B}$ is the autonomous balance that would prevail if the real exchange rate sat exactly at its fundamental level, and ε_q is the aggregate elasticity of the trade balance with respect to the real rate. The autonomous term in Nepal's case is dominated by the structural excess of import propensity over export capacity; at the fundamental exchange rate we take it to equal the small surplus of roughly 0.6 percent of GDP that would service interest on the solvency norm $b^* = -0.15$ adopted below. The dynamics of the real exchange rate close the system. We posit that the real rate depreciates when the prevailing trade balance falls short of the solvency-consistent level and appreciates when it exceeds it, while also being pulled toward a fundamental anchor $\bar{q}$ at speed ω :

$$\dot{q} = \gamma[-r \cdot b - B(q)] + \omega(\bar{q} - q), \quad \gamma, \omega > 0. \quad (5)$$

The imbalance channel with speed γ is the market mechanism through which external deficits translate into currency pressure — the mechanism that operated violently during the 2015–16 border disruption and that the 2022 import restrictions were designed to suppress. The anchor channel with speed ω captures the slow equilibration of domestic relative prices toward fundamentals under the rupee peg: with the nominal bilateral rate fixed, real depreciation must arrive through the price level, which is a slow variable. Both channels are linear in the neighborhood of equilibrium, which is all that the local analysis of this section requires; the global validity of the conclusions is inherited from the linear-quadratic structure of Section 8.

3.3 Equilibrium and Saddle-Path Stability

Setting the right-hand sides of (2), (4), and (5) to zero yields the interior steady state. The anchor channel forces $q^* = \bar{q}$, and the accumulation equation then pins down the stock:

$$q^* = \bar{q}, \quad b^* = -\frac{\bar{B} + \varepsilon_q \bar{q}}{r}. \quad (6)$$

With the normalization $\bar{q} = 0$ and $\bar{B} = 0.006$, the solvency norm is $b^* = -0.15$: a debtor position of fifteen percent of GDP whose annual interest bill, at $r = 4$ percent, equals the 0.6 percent trade surplus that the fundamental exchange rate is assumed to deliver. All dynamics are

henceforth written for deviations $x_1 = b - b^*$ and $x_2 = q - q^*$, for which the constant terms drop out:

$$\dot{b} = rx_1 + \varepsilon_q x_2, \quad \dot{q} = -\gamma r x_1 - (\gamma \varepsilon_q + \omega) x_2. \quad (7)$$

The Jacobian of the linearization at the origin has components $J_{11} = r$, $J_{12} = \varepsilon_q$, $J_{21} = -\gamma r$, and $J_{22} = -(\gamma \varepsilon_q + \omega)$. Its trace and determinant are

$$\text{tr}(J) = r - \gamma \varepsilon_q - \omega, \quad \det(J) = -r\omega < 0. \quad (8)$$

The sign of the determinant is the whole story at this level of generality: a planar system with negative determinant has real eigenvalues of opposite sign, whatever the parameter values. Writing the characteristic polynomial $\lambda^2 - \text{tr}(J)\lambda + \det(J) = 0$, the two roots are

$$\lambda_{u,s} = \frac{\text{tr}(J) \pm \sqrt{\text{tr}^2(J) + 4r\omega}}{2}, \quad \lambda_u > 0 > \lambda_s. \quad (9)$$

The following proposition records the conclusion formally.

Proposition 1 (Saddle-path stability). For every admissible parameter constellation with $r, \gamma, \varepsilon_q, \omega > 0$, the system (7) possesses a unique interior fixed point at the origin; its eigenvalues are real, of opposite sign, and satisfy $\lambda_u \lambda_s = -r\omega$ and $\lambda_u + \lambda_s = \text{tr}(J)$. The stable and unstable manifolds are linear rays through the origin with slopes $m_s = (\lambda_s - r)/\varepsilon_q$ and $m_u = (\lambda_u - r)/\varepsilon_q$ respectively, and the unique convergent trajectory from any initial external wealth position $x_1(0)$ requires the real exchange rate to jump at date zero onto the unstable manifold:

$$x_2(0+) = m_u \cdot x_1(0) = \frac{(\lambda_u - r)}{\varepsilon_q} \cdot x_1(0). \quad (10)$$

Proof. Determinant (8) is strictly negative, so the roots of the characteristic polynomial are real and of opposite sign; the quadratic formula gives (9), and Vieta's formulas give the stated sum and product. An eigenvector associated with eigenvalue λ satisfies $(r - \lambda)v_1 + \varepsilon_q v_2 = 0$ from the first row of $J - \lambda I$, whence the slope of the corresponding ray is $v_2/v_1 = (\lambda - r)/\varepsilon_q$; substituting the two eigenvalues yields the two slopes. Because the eigenvalues are real, distinct, and of opposite sign, the Hartman–Grobman theorem applies and the linear manifolds are the local stable and unstable sets; for the linear system itself they are global. The predetermined state $x_1(0)$ cannot jump, so bounded convergence pins the initial value of the free variable x_2 on the unstable manifold, giving (10). Uniqueness of the convergent path follows because any other initial exchange rate produces a component along the unstable eigenvector that grows at rate $\lambda_u > 0$, which completes the proof.

At the calibration of Section 9 ($r = 0.04$, $\gamma = 0.25$, $\varepsilon_q = 0.0785$, $\omega = 0.10$), the eigenvalues are $\lambda_u = 0.0349$ and $\lambda_s = -0.1145$, so the stable arm is steep — slope -1.97 — while the unstable arm is nearly flat at -0.065 . Figure 3 draws the resulting phase portrait: the two nullclines, the vector field, the saddle path in gold, and three classes of trajectories. The economics of the figure is worth pausing over. An economy that wakes up with external wealth fourteen percentage points of GDP below its solvency norm does not repair the position

gradually through prices alone; it must first absorb a real depreciation that lands it on the saddle path — in the illustrated configuration, roughly 0.9 log points per fourteen points of lost external wealth — and then ride the stable arm, on which debt reduction and exchange-rate appreciation are locked together. Any attempt to hold the real rate at its anchor despite the wealth gap puts the economy on a divergent path.




Figure 3. Phase portrait of the core dynamic system (7) at the baseline calibration. Red solid: $\bar{b} = 0$ nullcline; navy dashed: $\dot{q} = 0$ nullcline; gold: unstable manifold (saddle path); teal: convergent trajectories along the stable manifold; gray arrows: the vector field. The open circle marks an initial position with external wealth 14 percentage points of GDP below the norm; the annotated trajectory shows the instantaneous depreciation onto the saddle path followed by convergence.

How fast is the convergence once on the stable arm? Along the stable eigenvector the deviation decays as $x(t) = c \cdot e^{\lambda^s t} v_s$, so the half-life of the external gap is $t_{1/2} = \ln 2 / |\lambda^s|$, equal to 6.05 years at the calibration. Two comparative observations sharpen the message. The half-life is governed almost entirely by the adjustment speeds γ and ω — institutional and price-setting constants that policy can influence only indirectly — while the jump magnitude (10) is governed by the trade elasticity. Economies with inelastic trade face shallow jumps but long, flat convergences; economies with elastic trade face steep, violent corrections. Nepal's aggregate elasticity, derived from primitives in Section 4, sits toward the inelastic end, which is precisely why the optimization exercise of Section 8 refuses to rely on the price channel alone.

Proposition 2 (Unsustainability of the unregulated economy). Let the welfare loss of an external path be the discounted quadratic integral of Section 8 evaluated under the passive rule $u \equiv 0$. Any trajectory whose initial condition does not lie on the stable manifold grows asymptotically at rate λ_u , and the undiscounted loss grows at rate $2\lambda_u$. At the baseline calibration, $2\lambda_u - \rho = 2(0.0349) - 0.03 = 0.0398 > 0$, so the discounted loss integral diverges: there exists no finite welfare comparison that certifies the passive economy as sustainable. The passive economy is sustainable only in the degenerate case in which its initial condition happens to lie exactly on the stable arm.

Proof. Solutions of a linear system with distinct real eigenvalues are linear combinations of $e^{\lambda t} v_i$ over the eigendirections; the coefficient along the unstable eigenvector is nonzero unless the initial condition lies on the stable manifold, and the dominant term grows at λ_u . The loss integrand is a positive-definite quadratic form in the state, hence grows at $2\lambda_u$; discounting at ρ leaves the growth rate $2\lambda_u - \rho$, whose sign decides the integral, and this completes the proof.

4. Lagged Adjustment, the J-Curve, and the Aggregate Elasticity

Proposition 1 treats the trade elasticity ε_q as a primitive. This section derives it from primitives and, in the process, dates the response of the trade balance to a depreciation — the difference between the textbook instantaneous Marshall–Lerner condition and the empirically relevant J-curve. The device is geometric volume adjustment in both export and import markets, coupled with the accounting fact that the import bill is revalued as soon as the currency moves.

Let the log-deviations of export and import volumes from their pre-depreciation levels obey first-order adjustment toward their long-run demand responses:

$$x_V(t) = \eta_x \cdot \Delta q \cdot (1 - e^{-t/\tau_X}), \quad m_V(t) = -\eta_m \cdot \Delta q \cdot (1 - e^{-t/\tau_M}), \quad (11)$$

where η_x and η_m are the long-run export and import demand elasticities (both positive in the convention that higher elasticities mean stronger volume responses) and τ_X, τ_M are the adjustment speeds. These are the reduced forms of the diffusion equations $\dot{x}_V = (\eta_x \Delta q - x_V)/\tau_X$ and its import analogue, the standard lag structure of the empirical trade-flow literature. Prices adjust faster than volumes; for the small-economy importer we take the pass-through of a real depreciation into the relative import price to be complete on impact,

$$p_m(t) = \Delta q, \quad (12)$$

an assumption relaxed in the remark below. Combining the price and quantity channels, the response of the trade balance ratio to a depreciation Δq undertaken at date zero is

$$\Delta B(t) = \frac{T_X}{Y} \eta_x \Delta q (1 - e^{-t/\tau_X}) - \frac{T_M}{Y} \Delta q [1 - \eta_m (1 - e^{-t/\tau_M})], \quad (13)$$

where T_X and T_M are the gross export and import flows at the pre-depreciation equilibrium. The first term is the export volume gain, building with time constant τ_X ; the second nests the import bill — whose price component rises one-for-one with the depreciation immediately — against the import volume contraction, which arrives only gradually. Evaluating (13) at the two extremes delivers the impact and long-run responses:

$$\Delta B(0) = -\frac{T_M}{Y} \Delta q < 0, \quad (14)$$

$$\Delta B(\infty) = \left[\frac{T_X}{Y} \eta_x - \frac{T_M}{Y} (1 - \eta_m) \right] \Delta q. \quad (15)$$

The impact response is unambiguously negative: on the day the currency depreciates, every unit of imports costs more in real domestic terms and none has yet been substituted away. The long-run response is the quantity that the classical literature was after, and it defines the aggregate elasticity that the core model of Section 3 uses:

$$\varepsilon_q \equiv \frac{T_X}{Y} \eta_x - \frac{T_M}{Y} (1 - \eta_m), \quad \varepsilon_q > 0 \leftrightarrow \frac{T_X}{T_M} \eta_x + \eta_m > 1. \quad (16)$$

Proposition 3 (Corrected Marshall–Lerner condition). The long-run trade balance improves after a real depreciation if and only if $\varepsilon_q > 0$ as defined in (16). When gross trade is initially balanced, $T_X = T_M$, condition (16) reduces to the classical inequality $\eta_x + \eta_m > 1$. When trade is initially unbalanced, as in Nepal where imports run at more than four times exports, the threshold shifts in favor of the importer: the economy must satisfy $\frac{T_X}{T_M} \cdot \eta_x + \eta_m > 1$, which is strictly harder than the classical condition, because the mechanical revaluation of the import bill operates on a base that dwarfs the export base on which volume gains accrue.

Proof. Equation (13) follows from summing the valued volume responses: export value gains accrue at the gross export share times the volume path (11), and the import bill is the product of the relative import price (12) and the adjusted import volume, so its change relative to GDP is the gross import share times the bracketed price-minus-quantity term. Setting $t = 0$ and $t = \infty$ in (13) gives (14) and (15), since all exponentials vanish at the upper limit; dividing (15) by Δq yields the definition in (16), and the equivalence with the elasticity inequality is elementary algebra on that definition.

At the calibration used throughout ($T_X/Y = 0.08$, $T_M/Y = 0.33$, $\eta_x = 1.6$, $\eta_m = 0.85$), the aggregate elasticity is $\varepsilon_q = 0.08 \times 1.6 - 0.33 \times 0.15 = 0.128 - 0.0495 = 0.0785$ — positive, so the Marshall–Lerner test is passed, but barely, and only because the import elasticity clears the raised bar $\frac{T_X}{T_M} \cdot \eta_x + \eta_m = 0.243 + 0.85 = 1.09 > 1$ with a margin of less than a tenth. Figure 4 visualizes both halves of the proposition: panel (a) traces the full response path (13) after a five percent depreciation, and panel (b) maps the long-run response over the elasticity plane with the threshold locus marked.




Figure 4. Trade-balance response to a five percent real depreciation. Panel (a): the J-curve path implied by equation (13) with $\tau_X = 8$ and $\tau_M = 4$ quarters — impact deterioration of 1.65 percentage points of GDP, break-even at 8.5 quarters, and a long-run gain approaching 0.39 percentage points. Panel (b): the long-run response (15) over the elasticity plane; the black locus is the corrected Marshall–Lerner threshold (16) and the star marks the calibration.

Three quantitative facts in the figure deserve to be read into the record. The deterioration on impact is $-(0.33)(0.05) = -1.65$ percentage points of GDP — with Nepal’s import base, a five percent depreciation costs one and two-thirds percent of GDP before any quantity responds. Break-even arrives after approximately 8.5 quarters, driven almost entirely by the slow accumulation of export volume against the faster import response. The ultimate gain is +0.39 percentage points per five percent depreciation, so even the long-run reward to depreciation is modest. A saddle-path economy with an inelastic trade balance must therefore absorb large real depreciations to accomplish small improvements in its external accounts — the quantitative heart of why Proposition 2 condemns exclusive reliance on the price channel.

Remark (partial pass-through). If import prices are contracted in advance, the pass-through in (12) is replaced by the ramp $p_m(t) = \Delta q(1 - e^{-t/\tau_P})$ with pass-through lag τ_P . The impact response then vanishes and the trough of the J-curve moves to the interior date where the marginal revaluation of the import bill equals the marginal volume response, $\frac{e^{-t/\tau_P}}{\tau_P} = \frac{\eta_m e^{-t/\tau_M}}{\tau_M}$. Nothing in Sections 5 through 10 depends on which timing convention is adopted, because the optimal control problem is stated in terms of the long-run elasticity ε_q that both conventions deliver.

5. Cointegration, Error Correction, and Dynamic Representation

The continuous-time system of Section 3 is a statement about levels: external wealth, the real exchange rate, and the trade balance move together because the accounting identity (2) and the behavioral links (4)–(5) tie them together. In discrete data these long-run ties appear as cointegration — individual series may wander as integrated processes, but certain linear

combinations of them are stationary, and deviations from those combinations feed back into subsequent changes. This section lays out the time-series apparatus in the form in which it disciplines the calibration: unit-root pretesting, the autoregressive distributed lag bounds approach for mixed orders of integration, the Johansen procedure for the full system, causality testing robust to the integration structure, and the vector autoregression whose impulse responses summarize the transmission of real-exchange-rate shocks.

5.1 Integration Properties

The starting point is the classification of each series as integrated of order d , written $I(d)$, meaning that d differences are required for stationarity. The workhorse test regresses the first difference of a series on its own lagged level, a constant, a deterministic trend, and enough lagged differences to whiten the residuals:

$$\Delta y_t = \pi_0 + \pi_1 t + \pi_2 y_{t-1} + \sum_{i=1}^k \psi_i \Delta y_{t-i} + u_t, \quad (17)$$

with the null $\pi_2 = 0$ (a unit root) tested by the Dickey–Fuller statistic against the stationarity alternative. Two refinements matter for a series like Nepal’s trade ratios, whose variance plausibly shifted with the liberalization of the early 1990s and the disruption of 2015: the Phillips–Perron variant, which corrects the statistic nonparametrically for heteroskedastic and autocorrelated disturbances, and stationarity tests such as that of Kwiatkowski, Phillips, Schmidt, and Shin (1992) that reverse the burden of proof. The practical conclusion of any such exercise on Nepali external data is unsurprising but consequential: levels of exports, imports, and the real exchange rate behave as $I(1)$ processes, while the trade balance itself — a difference — is close to stationary. Mixed integration orders of this kind are precisely the case for which the bounds approach was designed.

5.2 The ARDL Bounds Framework

Let y_t denote the dependent variable — the trade balance ratio — and let $x_t = (x_{1t}, \dots, x_{kt})'$ collect the real exchange rate, world income, and any policy indicators. An autoregressive distributed lag of orders (p, q) is

$$y_t = \varphi_0 + \sum_{i=1}^p \varphi_i y_{t-i} + \sum_{i=0}^q \delta_i' x_{t-i} + u_t, \quad (18)$$

and its reparametrization as an unconditional error-correction model separates the long-run relation from the short-run dynamics:

$$\Delta y_t = \theta y_{t-1} + \beta' x_{t-1} + c_0 + \sum_{i=1}^{p-1} \gamma_i \Delta y_{t-i} + \sum_{i=0}^{q-1} \zeta_i' \Delta x_{t-i} + v_t, \quad (19)$$

where θ and β govern the long run and the remaining coefficients the short run. The testing problem is whether θ and β are jointly zero — that is, whether no level relation exists:

$$H_0: \theta = 0, \quad \beta = 0 \quad \text{versus} \quad H_1: \theta \neq 0 \text{ or } \beta \neq 0, \quad (20)$$

assessed with the standard F-statistic, whose nonstandard distribution under the null lies between asymptotic lower and upper critical bounds depending on whether the regressors are purely $I(0)$ or purely $I(1)$. A statistic above the upper bound rejects the null regardless of the integration orders; below the lower bound it cannot; between them the verdict is inconclusive. Rejection opens the way to the error-correction representation in which the adjustment coefficient carries the economic message: writing the long-run equilibrium as $e_t = y_t - \gamma'x_t$, the parsimonious dynamics, in the sense of Engle and Granger (1987), are

$$\Delta y_t = \alpha e_{t-1} + \sum_j \Gamma_j' \Delta z_{t-j} + \varepsilon_t, \quad \alpha < 0, \quad (21)$$

where z_t stacks all differenced regressors and the sign restriction on α enforces stability: each period the gap closes in proportion to its lagged size. The half-life of a deviation is $\ln 2 / |\ln(1 + \alpha)|$ quarters, the discrete-time mirror of the $\ln 2 / |\lambda^s|$ formula of Section 3 — a correspondence worth noticing, because it means the theoretical and econometric halves of the paper estimate the same object from different directions.

5.3 System Cointegration and the VECM

Single-equation treatment is adequate when the cointegrating vector is unique, but the external accounts form a genuinely multivariate system in which exports, imports, the real exchange rate, and income all adjust to one another. In the vector autoregression of order p , the Johansen (1988, 1995) reparametrization isolates the long-run matrix:

$$\Delta y_t = \Pi y_{t-1} + \sum_{i=1}^{p-1} \Gamma_i \Delta y_{t-i} + \Phi d_t + \varepsilon_t, \quad \Pi = \alpha \beta', \quad (22)$$

in the sense of Sims (1980), and the Toda–Yamamoto (1995) correction below guards the causality inference; the rank of Π is the number of cointegrating relations and the columns of β span the space of long-run equilibria. The trace statistic for the null of at most r relations is

$$\lambda_{trace}(r) = -T \sum_{i=r+1}^n \ln(1 - \lambda_i), \quad (23)$$

with λ_i the ordered eigenvalues of the residual product moment matrix from the reduced-rank regression. Rank one in a four-variable external system is the empirically standard finding for small open economies under a peg: one long-run relation — in our reading, the solvency constraint (3) itself — with the remaining directions nonstationary. The adjustment matrix α then reveals which variables do the correcting, and in calibrated systems of this type the burden falls on the trade balance and the real exchange rate, exactly as the model of Section 3 requires.

5.4 Causality and Robust Inference

Whether the real exchange rate Granger-causes the trade balance — whether past values of the former narrow the prediction intervals for the latter — is tested by a Wald statistic on the coefficients of the corresponding block in (22). Standard asymptotics, however, are distorted

when the variables are integrated, which motivates the Toda–Yamamoto procedure: fit a VAR in levels of order $p + d_{max}$, where d_{max} is the maximal order of integration, and test causality on the first p lags only:

$$W_{TY} = (R\beta)'[R \cdot V(\beta) \cdot R']^{-1}(R\beta), \quad (24)$$

where R selects the coefficient restrictions and $V(\beta)$ the estimator covariance; the statistic is asymptotically chi-square regardless of the cointegrating rank. The procedure is statistically conservative and conceptually clean, and for Nepali data its typical verdict — bidirectional causality between the exchange rate and the trade balance at annual horizons — is exactly what the two-way coupling of equations (2) and (5) predicts.

5.5 The Calibrated VAR and Impulse Responses

For the numerical exercises of Section 9 we work with a three-variable system in deviations from deterministic means — the real exchange rate q , the external wealth gap x_1 , and the domestic output gap — specified as a VAR(2):

$$y_t = A_1 y_{t-1} + A_2 y_{t-2} + \varepsilon_t, \quad \varepsilon_t \sim (0, \Sigma). \quad (25)$$

Stationarity is the property that matters for everything downstream, and it is checked on the companion form: stacking $z_t = (y_t', y_{t-1}')'$ gives $z_t = Fz_{t-1} + G\varepsilon_t$ with

$$\rho(F) < 1, \quad (26)$$

where $\rho(F)$ denotes the spectral radius of the companion matrix — the largest eigenvalue modulus, equal to 0.779 at the calibration of Table 3, comfortably inside the unit circle. The Wold representation $y_t = \varepsilon_t + \Phi_1 \varepsilon_{t-1} + \Phi_2 \varepsilon_{t-2} + \dots$ follows by recursion, and orthogonalizing the shocks with the lower-triangular Cholesky factor L of Σ — which orders the exchange rate first, as the variable least able to respond within the quarter — yields impulse responses

$$\frac{\partial y_{t+h}}{\partial \varepsilon_{1t}} = e_1' \Phi_h L, \quad \Phi_h = A_1 \Phi_{h-1} + A_2 \Phi_{h-2}, \quad (27)$$

with $\Phi_0 = I$. Figure 5 traces the response of the system to a one-standard-deviation real depreciation shock. The exchange rate itself jumps on impact and mean-reverts over roughly six quarters; external wealth deteriorates before improving — the VAR's imprint of the J-curve — with a trough response of about -0.26 percentage points on impact that fades to effectively zero by the twelfth quarter; and the output gap expands mildly on impact before reverting, consistent with the short-run competitiveness gain and its subsequent erosion as import costs feed through.




Figure 5. Orthogonalized impulse responses to a one-standard-deviation real-exchange-rate shock in the calibrated VAR(2) of equation (25), computed exactly from the Wold recursion (27) with Cholesky identification. Panel (a): the real exchange rate; panel (b): external wealth; panel (c): the output gap. The spectral radius of the companion form is 0.779, so all responses decay monotonically to zero.

6. Structural Gravity and the Geometry of Diversification

The elasticity that Section 4 derived from primitives is itself the shadow of a deeper structure: bilateral trade flows obey gravity, and the aggregate response of exports to a real depreciation is the general-equilibrium residue of changes in thousands of bilateral flows, each governed by trade costs, sizes, and multilateral resistance. This section makes that structure explicit, both because the gravity block disciplines the elasticities adopted earlier and because the cross-sectional anatomy of Nepal's trade — who buys what, at what cost — is where diversification policy actually operates.

6.1 The Structural Gravity System

With $\sigma > 1$ the elasticity of substitution across source countries, exporter output Y_i , importer expenditure Y_j , world output Y_W , and bilateral trade costs $T_{ij} \geq 1$, the structural gravity equation of Anderson (1979), in its multilateral-resistance generalization due to Anderson and van Wincoop (2003), delivers bilateral exports

$$X_{ij} = \frac{Y_i \cdot Y_j}{Y_W} \left(\frac{T_{ij}}{\Pi_i \cdot \Pi_j} \right)^{1-\sigma}, \quad (28)$$

where the multilateral resistance indices solve the fixed-point system

$$\Pi_i^{1-\sigma} = \sum_j \frac{T_{ij}^{1-\sigma}}{Y_W} \Pi_j^{\sigma-1} \cdot Y_j, \quad \Pi_j^{1-\sigma} = \sum_i \frac{T_{ij}^{1-\sigma}}{Y_W} \Pi_i^{\sigma-1} \cdot Y_i, \quad (29)$$

so that each region's trade depends not only on its bilateral frictions but on its position in the entire network of costs — the remote and landlocked face systematically higher outward multilateral resistance, which is the gravity formalization of Nepal's transit geography. For a landlocked economy bordered by a vast single partner, the theoretical implication is blunt: the fixed component of T_{ij} toward the rest of the world is high, so conditional on gravity, Nepal's

exports are biased toward India, and diversification across destinations must overcome a cost structure that no exchange rate movement can alter.

Estimation of (28) must respect the multiplicative heteroskedasticity of trade data and the zeros that dominate developing-country trade matrices; the workhorse is the Poisson pseudo-maximum-likelihood estimator of Santos Silva and Tenreyro (2006), whose first-order condition equates fitted and actual flows level by level:

$$\sum_{ij} x_{ij} - e^{\beta' z_{ij}} \cdot z_{ij} = 0, \quad \hat{x}_{ij} = e^{\beta_0' z_{ij} + \delta_i + \delta_j}, \quad (30)$$

with exporter and importer fixed effects δ_i, δ_j absorbing multilateral resistance consistently. It is through this lens that the Marshall–Lerner arithmetic of Section 4 acquires its microfoundation: a real depreciation is a change in relative prices that, fed through (28)–(30), reallocates bilateral flows toward Nepal with an aggregate intensity governed by the same substitution elasticity σ that disciplines the trade elasticities η_x, η_m .

6.2 Measuring Concentration and Comparative Advantage

Diversification — the extensive margin on which Section 11’s policy conclusions turn — has three complementary scalar summaries. Revealed comparative advantage normalizes a product’s share in the exporter’s basket by its share in world trade:

$$RCA_{ij} = \frac{x_{ij}/X_i}{x_{wj}/X_w}, \quad (31)$$

with values above unity conventionally flagged. The Herfindahl–Hirschman index and the Shannon entropy summarize the concentration of the basket itself:

$$HHI = \sum_{i=1}^n s_i^2, \quad H = - \sum_{i=1}^n s_i \ln s_i, \quad (32)$$

where s_i is product i ’s share of merchandise exports and n the number of products. The exponential e^H — the effective number of products — converts entropy into an interpretable count, and the two indices bracket the truth from different sides: HHI is disproportion-sensitive, entropy richness-sensitive. Computed on the stylized composition of Figure 2, the HHI moves from 0.268 in 2001 to 0.305 in 2024 while the effective number of products rises from 4.33 to 4.77 — the signature of a basket that added products at the bottom of the distribution even as one new cluster (electricity) became large enough to concentrate the top.

Two dynamic observations complete the geometry. First, concentration is not self-correcting: under any reasonable Markov transition structure on product shares, initial concentration raises the variance of export earnings and, through the investment channel, discourages exactly the entry that would dilute it — a complementarity between risk and concentration that policy must break deliberately. Second, the product space is unevenly connected; the feasible diversification path for an economy at Nepal’s income level runs from agricultural processing toward light manufactures and energy trade along the edges of existing capability, which is

why the optimal control problem of Section 8 treats diversification effort as an investment with convex costs rather than as a free reallocation.

7. Stochastic Dynamics and State Filtering

The deterministic skeleton of Sections 3 through 6 is silent about the turbulence through which actual external accounts travel: remittance flows respond to Gulf labor markets, monsoon outcomes feed agricultural exports, and the Indian business cycle drives demand for Nepali output with a lag and a variance. This section therefore embeds the core system in continuous-time stochastic dynamics, develops the two tools that the stochastic environment requires — the change-of-variables calculus of Itô processes and the Kalman filter through which imperfectly observed states are inferred — and records the certainty-equivalence property that lets the deterministic feedback rule of Section 8 survive the extension to noise.

7.1 Diffusion Representation of the External Accounts

We take the real exchange rate, the state most exposed to high-frequency relative-price variation, to follow a diffusion whose drift is exactly the deterministic law of motion (5), augmented by white noise:

$$dq = [-\gamma r x_1 - (\gamma \varepsilon_q + \omega) x_2] dt + \sigma_q \cdot dW_q, \quad db = [r x_1 + \varepsilon_q x_2] dt, \quad (33)$$

where W_q is a standard Wiener process and σ_q scales the innovation variance; the accumulation equation is left deterministic because external wealth moves at accounting speed, with no economically meaningful within-period noise of its own. For any twice-differentiable function V of the state, the Itô formula gives the dynamics of V along solutions:

$$dV = \left(\frac{\partial V}{\partial t} + \mu \cdot \frac{\partial V}{\partial q} V + \frac{1}{2} \sigma_q^2 \cdot \frac{\partial^2 V}{\partial q^2} \right) dt + \sigma_q \frac{\partial V}{\partial q} \cdot dW_q, \quad (34)$$

in which the drift acquires the second-derivative term that has no deterministic counterpart — the mathematical expression of the fact that volatility is costly when objectives are concave. The special case in which the deterministic drift is linear and the noise additive is the workhorse of applied work: the Ornstein–Uhlenbeck process

$$dq = \theta(\bar{\mu} - q) dt + \sigma dW, \quad \theta > 0, \quad (35)$$

has the closed-form solution

$$q_t = \bar{\mu} + (q_0 - \bar{\mu})e^{-\theta t} + \sigma \int_0^t e^{-\theta(t-s)} \cdot dW_s, \quad (36)$$

whose mean reverts to $\bar{\mu}$ at rate θ and whose variance converges to the finite limit

$$\text{Var}(q_t) \rightarrow \frac{\sigma^2}{2\theta} \quad \text{as } t \rightarrow \infty, \quad (37)$$

a stationary band around the fundamental level rather than a wandering path. This is the stochastic image of the anchor channel in equation (5), and it is what licenses the stationary language — gaps, half-lives, ergodic distributions — used everywhere else in the paper. At the calibration of Section 9 the stationary standard deviation of the real exchange rate gap is 0.022,

consistent with the year-to-year variability of real effective exchange rates under pegged arrangements.

7.2 Filtering: Inferring the State from Imperfect Observations

Neither the external wealth gap nor the fundamental exchange rate is observed directly: the NIIP is measured with lags and valuation noise, and the fundamental real rate is an abstraction. Suppose the observed signal y_t relates to the scalar state x_t through a linear measurement with noise. The Kalman filter propagates the conditional mean and variance through the prediction step

$$\hat{x}_{t|t-1} = F \cdot \hat{x}_{t-1|t-1}, \quad P_{t|t-1} = F^2 \cdot P_{t-1|t-1} + Q, \quad (38)$$

and then corrects them in the update step, weighting the innovation — the surprise in the signal — by the Kalman gain:

$$K_t = \frac{P_{t|t-1}}{P_{t|t-1} + R}, \quad \hat{x}_{t|t} = \hat{x}_{t|t-1} + K_t(y_t - \hat{x}_{t|t-1}), \quad P_{t|t} = (1 - K_t) \cdot P_{t|t-1}, \quad (39)$$

with steady-state gain $K^* = P/(\sigma^2 + P)$ approached geometrically from any initialization. Three properties of the filter matter for what follows. It is optimal among all estimators when the model is linear-Gaussian, and the best linear estimator otherwise. Its likelihood is available in closed form,

$$\ln L = -\frac{1}{2} \sum_{t=1}^T \left(\ln 2\pi S_t + \frac{v_t^2}{S_t} \right), \quad (40)$$

where v_t is the innovation and S_t its conditional variance, so the whole state-space model — transition, measurement, noise — can be estimated by maximum likelihood rather than calibrated by assertion. And its steady-state gain quantifies the information loss of a pegged, administratively managed environment: the slower external data arrive and the noisier they are, the more conservative the filter, and the smoother the policy feedback built on top of it.

Remark (certainty equivalence). The optimization of Section 8 is stated deterministically, yet its linear feedback remains optimal under the additive diffusion (33): for linear-quadratic problems with additive Gaussian noise, the optimal control is the certainty-equivalence rule applied to the filtered state, and the only change the stochastic environment requires is that the rule be fed by $\hat{x}_{t|t}$ rather than by the unobserved state itself. This is the sense in which Sections 7 and 8 are two halves of a single architecture — the filter estimates what the controller steers.

8. Optimal Control of the External Balance

The preceding sections established that the unregulated economy is unsustainable (Proposition 2) and that the price channel alone is a slow and shallow instrument (Proposition 3 and the arithmetic around Figure 4). We now ask what a planner who internalizes both facts should do. The answer is constructed in the classical two-step discipline of dynamic optimization: first the maximum principle, which converts the intertemporal problem into a system of differential

equations whose boundary structure reveals the shape of optimal paths; then the dynamic programming of Bellman (1957), which delivers the value function and the policy rule in closed form. The two routes meet at the same linear feedback — as they must — and the meeting point is itself a verification that every constant reported below is internally consistent.

8.1 The Planning Problem

The planner chooses a path of structural export-promotion effort $u(t)$ — trade facilitation, quality infrastructure, market-development expenditure, energy-export scaling — measured in GDP-equivalent units, to minimize the discounted loss

$$\min \frac{1}{2} \int_0^{\infty} e^{-\rho t} (\kappa x_1^2 + x_2^2 + cu^2) dt \quad (41)$$

subject to the deviation dynamics augmented by the policy lever:

$$\dot{x}_1 = rx_1 + \varepsilon_q x_2 + \varepsilon_u u, \quad \dot{x}_2 = -\gamma r x_1 - (\gamma \varepsilon_q + \omega) x_2 - \gamma \varepsilon_u u. \quad (42)$$

The quadratic loss penalizes external wealth gaps with weight κ , real-exchange-rate misalignment with unit weight, and effort convexly with cost parameter c ; the discount rate is ρ . Effort enters both state equations: it improves the trade balance directly by ε_u per unit — the structural channel — and, through the imbalance mechanism of equation (5), generates real appreciation pressure that partially offsets the passive depreciation the deficit would otherwise deliver. The planner therefore faces a genuine interior trade-off between repairing the balance sheet now and letting the price mechanism do work that is free but misalignment-costly.

8.2 The Maximum Principle

Attaching current-value costate variables μ_1, μ_2 to the two state equations, the present-value Hamiltonian is $e^{-\rho t}$ times

$$H = \frac{1}{2} (\kappa x_1^2 + x_2^2 + cu^2) + \mu_1 (rx_1 + \varepsilon_q x_2 + \varepsilon_u u) + \mu_2 (-\gamma r x_1 - (\gamma \varepsilon_q + \omega) x_2 - \gamma \varepsilon_u u), \quad (43)$$

and the interior optimality condition for the control equates the marginal cost of effort to its shadow value:

$$\frac{\partial H}{\partial u} = c \cdot u + \varepsilon_u (\mu_1 - \gamma \mu_2) = 0 \quad \Rightarrow \quad u^* = -\frac{\varepsilon_u}{c} (\mu_1 - \gamma \mu_2), \quad (44)$$

while the costate equations require

$$\begin{aligned} \dot{\mu}_1 &= (\rho - r)\mu_1 + \gamma r \mu_2 - \kappa x_1, & \dot{\mu}_2 \\ &= -\varepsilon_q \mu_1 + (\rho + \gamma \varepsilon_q + \omega)\mu_2 - x_2, \end{aligned} \quad (45)$$

together with the transversality condition

$$e^{-\rho t} \cdot \mu(t)' x(t) \Big|_{t \rightarrow \infty} = 0. \quad (46)$$

Substituting the policy rule (44) into (42) and appending (45) yields a four-dimensional linear system in the state-costate vector $z = (x_1, x_2, \mu_1, \mu_2)'$ whose structure is the analytical heart of the problem:

$$\dot{z} = H \cdot z, \quad H = \begin{bmatrix} \left(A - \frac{BB'}{c} \right) & (-Q \quad \rho I - A') \end{bmatrix}, \quad (47)$$

written in block form with A the drift matrix of (7), $B = (\varepsilon_u, -\gamma\varepsilon_u)'$ the policy incidence vector, and $Q = \text{diag}(\kappa, 1)$ the state-weight matrix. The eigenvalue structure of H organizes everything:

Proposition 4 (Discounted symplectic pairing). Let H be the block matrix (47). If λ is an eigenvalue of H , then $\rho - \lambda$ is also an eigenvalue. Consequently, when the algebraic Riccati equation below admits a stabilizing solution, H has exactly two eigenvalues with negative real part and two with positive real part, symmetric about $\rho/2$, and for every initial state there exists a unique bounded costate path.

Proof. Let J denote the skew-symmetric exchange matrix $[0, I; -I, 0]$, which swaps and sign-flips the state and costate blocks. Direct multiplication gives $JHJ^{-1} = -H' + \rho I$: the similarity transformation maps the state block A to $-(A - \rho I)'$ and conversely, while the control and penalty blocks exchange places. Similar matrices share spectra, so the eigenvalues of H coincide with those of $\rho I - H'$, which are exactly the numbers $\rho - \lambda$ for eigenvalues λ of H . The stabilizing solution of the Riccati equation exists if and only if the pair (A, B) is stabilizable and the pair $(Q^{1/2}, A)$ detectable — both true here, since all parameters are positive and the unregulated system has a controllable unstable direction — and its existence is equivalent to the two-by-two split of the spectrum, which gives the uniqueness of the bounded costate path by saddle-point dimension counting.

At the calibration of Section 9 the four eigenvalues of H are $-0.1173, -0.0197, +0.0497$, and $+0.1473$, and the pairing identity $\lambda + (\rho - \lambda) = \rho = 0.03$ holds to within 3×10^{-16} — machine precision. The two stable eigenvalues generate the optimal trajectory's fast and slow adjustment modes; their asymmetric magnitudes are the mathematics behind the front-loaded, long-tailed effort profiles visible in Figure 6 below.

8.3 Dynamic Programming and the Closed-Form Feedback Rule

Dynamic programming attacks the same problem through the Hamilton–Jacobi–Bellman equation for the value function $V(x)$, the minimized discounted loss from the current state:

$$\rho V = \frac{1}{2} x' Q x + \min_u \left[\frac{c}{2} u^2 + V'(Ax + Bu) \right], \quad (48)$$

and the quadratic guess $V(x) = \frac{1}{2} x' P x$ closes the equation exactly. The first-order condition inside the minimum gives $cu^* + B'V' = 0$, so the optimal control is the linear feedback

$$u^*(x) = -R^{-1}B'Px \equiv -Kx, \quad (49)$$

with P the positive-definite stabilizing solution of the algebraic Riccati equation

$$A'P + PA - PBR^{-1}B'P + Q = \rho P, \quad R = c. \quad (50)$$

The connection between the two methods is the shadow-price identity $\mu = Px$: the costate variables of the maximum principle are the marginal valuations of the states, which is precisely the gradient of the value function, and substituting this identity into (44) reproduces (49) term

for term. Because the Riccati equation (50) is a system of three coupled quadratic equations in the entries of P — $2A'P$ supplies the off-diagonal mixing — its solution is available in closed radicals for this two-state problem, and numerically it converges by Newton iteration to machine precision. At the calibration:

$$P_{11} = 9.942, \quad P_{12} = P_{21} = 4.726, \quad P_{22} = 6.200, \quad (51)$$

$$K = \frac{1}{c}B'P = (0.526, 0.191), \quad V(x_0) = \frac{1}{2}x_0'Px_0. \quad (52)$$

Proposition 5 (Optimality of the linear feedback). Let P be the stabilizing solution of (50) and define $u^*(x) = -R^{-1}B'Px$. Then u^* attains the infimum in (41) for every initial condition, the minimized loss is $V(x_0) = \frac{1}{2}x_0'Px_0$, and the closed-loop matrix $A - BR^{-1}B'P$ has both eigenvalues strictly inside the left half-plane.

Proof. Substituting the guess $V = \frac{1}{2}x'Px$ and the minimizing rule into the Hamilton–Jacobi–Bellman equation (48) yields exactly (50), so the pair (V, u^*) satisfies the Bellman equation. For any admissible control, the Itô product rule applied to $e^{-\rho t}V(x(t))$ along the deterministic dynamics gives an identity whose integrated form, after applying the transversality condition (46) — which holds because the closed loop is stable — and the Bellman equation itself, shows that the achieved discounted loss equals $V(x_0)$ plus a nonnegative terminal excess that vanishes only for u^* . Hence no feasible policy beats the feedback, and the feedback achieves its own promise. The eigenvalue claim is the definition of a stabilizing solution, which closes the argument.

The feedback gains deserve a sentence of interpretation. The gain on the external wealth gap, 0.526, exceeds the gain on the exchange-rate gap, 0.191: the controller pushes effort roughly three times harder per unit of debt gap than per unit of currency misalignment, because the debt gap is persistent — it worsens itself through the interest channel — while the exchange-rate gap mean-reverts on its own. The sign pattern implements the trade-off identified in Section 8.1: effort improves the balance directly but appreciates the currency, so the rule prescribes restraint when the currency is already undervalued and aggression when both gaps call for structural repair.

8.4 Constrained Effort and the Shape of Policy

Fiscal and administrative reality caps how fast structural effort can expand. Impose $0 \leq u \leq \bar{u}$ and let $\sigma(t) \equiv B'\mu(t)$ denote the switching function — the shadow value of effort. The pointwise minimization in the Hamiltonian then delivers the clipped rule

$$u^*(t) = \max \left[0, \min \left[\bar{u}, -\frac{\sigma(t)}{c} \right] \right], \quad (53)$$

so the optimal program is bang-bang at the boundaries and linear in the interior. Along the trajectories computed in Section 9 the unconstrained rule (49) never binds for the initial gaps considered — the implied peak effort of ten percentage points of GDP-equivalent, an anchored trade-balance improvement of about 1.2 percentage points, remains within the envelope of the

2022–23 import-compression precedent, though on the export side rather than the import side — so the clipped rule is reported for completeness rather than out of necessity. Its role grows with the size of the initial gap: for external wealth gaps beyond roughly thirty percentage points of GDP, the unconstrained rule demands more effort than the constraint permits, and the optimal program becomes a sustained maximum-effort arc followed by a linear decay — the familiar turnpike-then-convergence shape of constrained adjustment.

8.5 The Stochastic Problem

Under the diffusion (33) the Hamilton–Jacobi–Bellman equation acquires the curvature term that Itô’s formula contributes:

$$\rho V = \frac{1}{2} x' Q x + V' A x - \frac{1}{2c} (B' V')^2 + \frac{\sigma_q^2}{2} \cdot \frac{\partial^2 V}{\partial x_2^2}, \quad (54)$$

and the certainty-equivalence remark of Section 7 applies verbatim: because the diffusion enters additively and independently of the control, the quadratic value function of (48)–(50) solves (54) unchanged, and the optimal policy remains the same linear feedback evaluated at the filtered state. Volatility thus changes the planner’s welfare — through the level of V and through the dispersion of outcomes — without changing the rule. The fan charts of Section 10 visualize exactly this: identical feedback logic, stochastic futures.

9. Numerical Solution and Calibration

9.1 Calibration

Table 2 collects the calibration. Wherever a parameter has a structural derivation in the earlier sections, the table simply points to it; where judgment is required, the choice is anchored to magnitudes that are standard in the open-economy literature for low-income pegged economies. Three choices deserve justification. The aggregate elasticity ε_q is not free: it is the number that Proposition 3 constructs from volume elasticities and openness shares, and its small size relative to textbook aggregates is precisely the import-excess correction at work. The policy effectiveness ε_u of 0.12 says that a sustained GDP-scale unit of structural effort moves the autonomous trade balance by twelve percentage points of it — a demanding but not fanciful return on programs that include electricity exports, corridor logistics, and quality certification. And the effort cost $c = 2$ places the marginal cost of a one-percentage-point push at roughly one percent of GDP-equivalent loss, which is the order of magnitude at which trade-facilitation investments are actually debated.

Table 2. Calibration of the model (annual frequencies unless noted).

Parameter	Value	Interpretation and anchoring
r	0.04	World real interest rate faced by the sovereign
γ	0.25	Imbalance-to-exchange-rate adjustment speed (eq. 5)
ε_q	0.0785	Aggregate trade elasticity, derived in eq. (16)

Parameter	Value	Interpretation and anchoring
ω	0.10	Anchor speed of the real exchange rate under the peg
b^*	-0.15	Solvency norm for the net international investment position
$\bar{B}(\bar{q})$	0.006	Autonomous balance at the fundamental rate ($-\tau \cdot b^*$)
ρ	0.03	Planner's discount rate
κ	0.15	Loss weight on the external wealth gap
c	2.00	Curvature of the effort cost
ε_u	0.12	Effectiveness of structural effort on the trade balance
σ_q	0.022	Innovation standard deviation of the real exchange rate
η_x, η_m	1.6, 0.85	Export and import volume elasticities (Section 4)
$T_X/Y, T_M/Y$	0.08, 0.33	Gross export and import openness (Section 2)
τ_X, τ_M	8, 4 qtrs	Volume adjustment lags (Section 4)

9.2 Solution Architecture and Diagnostics

Every quantitative statement in this paper is computed, and the computation is organized so that each of the paper's theoretical claims has a numerical counterpart that can be checked independently. The state trajectories are integrated with the classical fourth-order Runge–Kutta scheme, whose increment is the familiar four-evaluation weighted average

$$x_{n+1} = x_n + \frac{\Delta t}{6}(k_1 + 2k_2 + 2k_3 + k_4), \quad (55)$$

with step $\Delta t = 0.002$ years. Halving the step changes terminal values at a ten-year horizon by less than 1010^{-15} — machine precision — so discretization error is absent from everything that follows. The Riccati equation (50) is solved by a Schur-form Newton iteration (Judd, 1998); its residual, the maximum absolute entry of the left side of (50) minus the right side, is 1.4×1010^{-15} . The discounted symplectic pairing of Proposition 4 holds to 2.9×1010^{-16} across the four Hamiltonian eigenvalues. And the value function computed from the closed-form expression agrees with the loss accumulated along the simulated optimal trajectory to four significant digits — a verification that exercises the entire chain from calibration through optimality. Table 3 assembles the diagnostics.

Table 3. Solution diagnostics at the baseline calibration.

Diagnostic	Value	Cross-check
Unregulated eigenvalues λ_u, λ_s	0.0349, -0.1145	$\text{tr}(J) = -0.0796, \text{det}(J) = -0.0040$
Saddle-path slopes m_u, m_s	-0.0647, -1.9687	Eigenvector computation

Diagnostic	Value	Cross-check
Hamiltonian eigenvalues	$-0.1173, -0.0197, 0.0497, 0.1473$	Pairing $\lambda + \lambda' = \rho$ to 2.9×10^{-16}
Closed-loop eigenvalues	$-0.1173, -0.0197$	Both strictly stable
Feedback gains K	$(0.526, 0.191)$	$K = c - 1B'P$
Riccati residual (max abs entry)	1.4×10^{-15}	Newton convergence
$V(x_0)$, closed form vs. simulated	0.08109 vs. 0.08110	Relative deviation 1.5×10^{-4}
VAR companion spectral radius	0.779	Stationarity of eq. (25)

9.3 Optimal Trajectories, the Turnpike, and Sensitivity

Figure 6 follows the optimal program from three initial conditions that bracket Nepal's plausible near-term positions: a crisis gap of ten percentage points of external wealth with a five percent real overvaluation, a deep gap of twenty points with the currency near anchor, and a mild five-point gap with mild overvaluation. The architecture of the solution is visible in all three. Effort is front-loaded — 6.2, 10.1, and 3.2 percentage points of GDP-equivalent respectively at date zero — and decays monotonically as the feedback rule (49) reads the shrinking gap. The implied trade-balance improvements, 0.75, 1.22, and 0.38 percentage points of GDP, are of exactly the order that the 2022–23 episode demonstrated to be feasible, though achieved there through import compression rather than export expansion. The external wealth gap itself closes only slowly — the closed loop's fast mode decays with a half-life of 5.9 years and the slow mode with a half-life of 35.2 years — and this stubbornness is not a defect of the policy but a fact about the economy: with a weak trade elasticity, even optimal structural effort repairs the stock gradually, which is the quantitative case for starting early.




Figure 6. Optimal adjustment under the linear feedback rule $u^*(t) = -Kx(t)$ for three initial conditions. Panel (a): the external wealth gap; panel (b): the real exchange rate gap; panel (c): the optimal effort path in GDP-equivalent percentage points; panel (d): the shadow value of the exchange-rate state along the optimal path. Effort is front-loaded and decaying; the wealth gap closes with a fast and a slow mode (half-lives 5.9 and 35.2 years).

Figure 7 places the same trajectories in the state plane, overlaid with the iso-loss ellipses of the quadratic value function. The contrast between the crimson dashed autarky drift and the navy optimal paths is the paper's policy message in one picture: left alone, an economy in the debt-overvaluation quadrant drifts toward deeper debt while its exchange rate depreciates toward the saddle path it cannot stably hold; under the feedback rule, the same economy bends immediately toward the origin, trading effort for a flattened, convergent approach. The origin — the turnpike — attracts all optimal trajectories regardless of their starting point, and the ellipses make the welfare geometry transparent: each path crosses iso-loss contours transversally, and the crossing slows as the state approaches the axis of the value function's narrowest curvature.




Figure 7. State space of the regulated economy. Gray arrows: the closed-loop vector field under the optimal feedback; navy: optimal trajectories from the three initial conditions (open circles); dashed crimson: the passive drift from the crisis initial condition; dotted ellipses: iso-loss contours of the value function; gold dot: the turnpike at the solvency norm.

How much do the structural parameters matter? Figure 8 re-solves the entire optimal control problem — Riccati equation, feedback rule, welfare integral — on a thirty-by-thirty grid over the trade elasticity and the imbalance adjustment speed, and contours the discounted loss from the crisis initial condition. The loss is increasing in the trade elasticity and decreasing in the adjustment speed. The first direction is the striking one: a more trade-sensitive economy suffers more from the same exchange-rate fluctuations, because the fluctuations translate into larger external-account swings that the planner must then fight with costly effort; the elasticity that makes depreciation powerful as a remedy also makes misalignment expensive as a disease. The second direction is intuitive: faster exchange-rate equilibration shortens the window during which the planner must hold the real rate away from its anchor. The baseline sits on the favorable edge of the surface — a reminder that Nepal's inelastic trade, so often lamented as an obstacle to depreciation-led adjustment, is in the optimized problem a source of stability that policy should be in no hurry to test.




Figure 8. Discounted loss under the optimal policy as a function of the aggregate trade elasticity ε_q (horizontal) and the imbalance-adjustment speed γ (vertical), with every grid point solved by its own Riccati equation. White contours are iso-loss lines; the star marks the baseline calibration. The loss increases with the elasticity and decreases with the adjustment speed.

10. Policy Experiments and Stochastic Futures

The experiments of this section compare three regimes — passive drift, the optimal feedback rule, and the stochastic environment in which both must actually operate — using the common metric of discounted welfare. Table 4 reports the fifty-year comparison for the three initial conditions. The optimal rule reduces the discounted loss by 61.3 percent in the crisis scenario, 62.9 percent in the deep-gap scenario, and 60.8 percent in the mild scenario. The stability of the percentage across scenarios is itself informative: the value function is quadratic, so proportional gains translate into absolute gains that scale with the square of the initial gap — the deep-gap scenario, twice as far from the norm, gains roughly twice as much in absolute welfare terms, which is the formal expression of the urgency of acting when gaps are large.

Table 4. Policy experiments: initial effort, trade-balance impact, and discounted losses over a fifty-year horizon.

Initial condition (b_0, q_0)	$u^*(0)$, pp of GDP	ΔB impact, pp of GDP	Loss, optimal	Loss, passive	Welfare gain
(−0.10, −0.05)	6.2	0.75	0.0785	0.2030	61.3%
(−0.20, 0.02)	10.1	1.22	0.1758	0.4738	62.9%

Initial condition (b_0, q_0)	$u^*(0)$, pp of GDP	ΔB impact, pp of GDP	Loss, optimal	Loss, passive	Welfare gain
(-0.05, -0.03)	3.2	0.38	0.0216	0.0551	60.8%

The passive losses deserve a caveat that is also a finding. Over fifty years the passive economy's loss is finite only because the horizon truncates what Proposition 2 shows to be asymptotically divergent: on the unregulated saddle system the loss integral grows without bound, and the fifty-year figures understate the true cost of inaction. The experiments therefore bracket the honest conclusion — the passive regime is not merely worse, it is not a regime at all but a path off which policy must at some point force the economy.

Figure 9 confronts the optimal rule with the stochastic environment of Section 7. Four hundred paths of the diffusion (33) are simulated for twenty years under two regimes — the feedback rule evaluated at the true state, and complete passivity — and the percentile bands of the external wealth gap are traced. The difference is stark. Under the optimal rule the median gap at the twenty-year mark is -8.2 percentage points of GDP with a ninetieth percentile of -4.7; under passivity the median is -25.8 with a tenth percentile of -32.7. In plain terms: the optimized economy not only adjusts, it adjusts in a way that noise cannot easily reverse, because the feedback rule converts every shock-induced deterioration into proportionally more effort. The passive economy does the opposite — its exchange-rate channel absorbs the shocks and passes them into the stock of external wealth, where they compound. The fan chart is the stochastic version of Proposition 2, and it is the single most policy-relevant picture in the paper.




Figure 9. Stochastic fan charts of the external wealth gap over twenty years, from 400 simulated paths of the diffusion (33) with $\sigma_{\{q\}} = 0.022$. Navy: optimal feedback policy (median and 10–90 percent band); crimson dashed: passive drift. The feedback rule stabilizes the distribution near -8 percentage points at the horizon; passivity drifts toward -26 at the median with a tenth percentile beyond -30.

Three robustness observations complete the experimental record. The scenario ranking is invariant to halving or doubling the effort cost, with the optimal initial effort moving inversely to c — the clipped rule of (53) binds for the deep-gap scenario once c doubles, extending rather than overturning the logic. The welfare gains are insensitive to the discount rate over the range 0.02–0.05, because the comparison is driven by the adjustment dynamics rather than by the tail of the integral. And the fan-chart conclusions survive doubling the innovation variance, at the price of wider bands and a modestly more aggressive optimal rule — the certainty-equivalence structure of Section 8.5 doing exactly the work it promises.

11. Implications for Nepali Trade Policy

The model is stylized; the implications are concrete. Five of them survive translation from the phase plane to Kathmandu.

First, the exchange rate is not the adjustment margin, and pretending otherwise is expensive. With an aggregate elasticity of 0.0785, even a five percent real depreciation — large by the standards of the peg — buys only 0.39 percentage points of GDP in the long run, after a 1.65-point deterioration on impact. The arithmetic of Proposition 3, not any ideological commitment to nominal stability, is what makes the rupee peg defensible as a nominal anchor while the real adjustment is engineered through structure. Depreciation-led strategy in Nepal would front-load pain and back-load reward to a degree that the political system is unlikely to sustain.

Second, structural effort is the high-return instrument, and it is front-loaded by construction. The optimal program demands its greatest push immediately — 6 to 10 GDP-equivalent points of effort in gap scenarios — and tapers as the external position repairs. The design implication is that adjustment programs should be sequenced with the heavy investments first: corridor and transit logistics, energy-export infrastructure, certification and quality systems for the agri-exports in which revealed comparative advantage already exists. A program that begins with the cheap gestures and defers the binding constraints inverts the optimal path.

Third, the electricity margin is the model's diversification story made real. Cross-border power sales to India, negligible before 2021, have already re-concentrated the measured basket while raising the effective number of products — the signature of an extensive-margin arrival with scale, which is exactly the commodity the concentration indices of Section 6 show the basket lacking. The feedback rule's persistence — effort stays positive for decades — is the formal argument for treating energy trade as a permanent structural commitment rather than a cyclical surplus.

Fourth, remittances must be recycled into tradables, or they will keep taxing them. The financing structure that makes the deficit bearable also appreciates the real rate and erodes the tradable base — the mechanism behind the export share's secular slide. Instruments that convert transfers into tradable investment: diaspora bonds earmarked for export infrastructure, deliberate channelling of remittance flows into energy and agro-processing equity, and the

removal of the small-economy frictions that keep household savings in land and gold. The model's ε_u is, in practice, built out of policies of exactly this kind.

Fifth, the import restrictions of 2022 were a costly down payment on a real problem. The episode demonstrated that rapid trade-balance improvement is administratively feasible — the model's effort lever is not fictional — but it worked on the import side, where the Marshall–Lerner correction makes compression expensive in inputs and activity, and it expired with the reserves that motivated it. The optimal-control reading is precise: the same external wealth objective, pursued through the export-side lever, dominates the import-side lever once the effort cost is calibrated to administrative reality, because export-side effort shifts the autonomous balance permanently while import compression is a level shock that decays.

12. Conclusion

This paper set out to give Nepal's external imbalance the kind of mathematical treatment its persistence deserves, and the treatment returned a diagnosis with unusual clarity. The external accounts form a saddle system: one unstable direction connects today's deficits to tomorrow's debt service, and convergence to solvency is possible only along a narrow path that requires the real exchange rate to move immediately and substantially whenever external wealth slips. The trade elasticity that governs that movement is, for Nepal, small — constructed rigorously from volume elasticities and the import excess, it passes the corrected Marshall–Lerner test with a margin of barely a tenth — so the price channel is slow, shallow, and front-loaded with pain. The planner's problem that follows from these two facts has a closed-form solution: a quadratic value function fixed by an algebraic Riccati equation whose stabilizing solution exists under verifiable conditions, and a linear feedback rule that converts state gaps into structural effort at gains of 0.53 and 0.19 per unit. Simulated forward, the rule cuts the discounted adjustment loss by about sixty percent across scenarios, stabilizes the stochastic distribution of external wealth where passivity lets it fan outward, and does so with effort paths that lie within the envelope of what Nepal has recently demonstrated it can do.

The framework's limits mark the agenda. The linearization is global for the linear-quadratic problem but local for the underlying economy; a nonlinear version with state-dependent trade elasticities would sharpen the large-gap recommendations where the clipped rule binds. The peg is modeled as an anchor speed rather than a policy regime; embedding the choice of arrangement itself — the trade-off between the anchor's credibility and the real rate's flexibility — into the planner's control set is the natural next problem. And the calibration, disciplined throughout by structural derivation rather than by estimation, awaits the data work — cointegration on the full external system, gravity estimation with the multilateral resistance terms, and Kalman-filtered measurement of the fundamental rate — that would convert the paper's framework from a discipline on reasoning into an instrument of policy. The mathematics, at least, is ready: every mechanism the diagnosis identified, the optimization turned into a rule, and every number in between can be checked by the reader from the diagnostics reported here.

References

- Anderson, J. E. (1979). A theoretical foundation for the gravity equation. *American Economic Review*, 69(1), 106–116.
- Anderson, J. E., & van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1), 170–192.
- Arrow, K. J., & Kurz, M. (1970). *Public Investment, the Rate of Return, and Optimal Fiscal Policy*. Baltimore: Johns Hopkins Press.
- Balassa, B. (1964). The purchasing-power parity doctrine: A reappraisal. *Journal of Political Economy*, 72(6), 584–596.
- Barro, R. J., & Sala-i-Martin, X. (2004). *Economic Growth* (2nd ed.). Cambridge, MA: MIT Press.
- Bellman, R. (1957). *Dynamic Programming*. Princeton, NJ: Princeton University Press.
- Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366), 427–431.
- Dixit, A. K. (1990). *Optimization in Economic Theory* (2nd ed.). Oxford: Oxford University Press.
- Dornbusch, R. (1976). Expectations and exchange rate dynamics. *Journal of Political Economy*, 84(6), 1161–1176.
- Engle, R. F., & Granger, C. W. J. (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica*, 55(2), 251–276.
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3), 424–438.
- Hausmann, R., Hwang, J., & Rodrik, D. (2007). What you export matters. *Journal of Economic Growth*, 12(1), 1–25.
- Hidalgo, C. A., & Hausmann, R. (2009). The building blocks of economic complexity. *Proceedings of the National Academy of Sciences*, 106(26), 10570–10575.
- International Monetary Fund. (2023). *Nepal: 2023 Article IV Consultation—Staff Report*. Washington, DC: International Monetary Fund.
- Intriligator, M. D. (1971). *Mathematical Optimization and Economic Theory*. Englewood Cliffs, NJ: Prentice-Hall.
- Johansen, S. (1988). Statistical analysis of cointegration vectors. *Journal of Economic Dynamics and Control*, 11(2), 231–254.
- Johansen, S. (1995). *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford: Oxford University Press.
- Judd, K. L. (1998). *Numerical Methods in Economics*. Cambridge, MA: MIT Press.
- Kamien, M. I., & Schwartz, N. L. (1991). *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management* (2nd ed.). Amsterdam: North-Holland.
- Krugman, P. (1979). A model of balance-of-payments crises. *Journal of Money, Credit and Banking*, 11(3), 311–325.

- Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1–3), 159–178.
- Ljungqvist, L., & Sargent, T. J. (2004). *Recursive Macroeconomic Theory* (2nd ed.). Cambridge, MA: MIT Press.
- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Berlin: Springer.
- Marshall, A. (1920). *Principles of Economics* (8th ed.). London: Macmillan.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6), 1695–1725.
- Ministry of Finance, Government of Nepal. (2023). *Economic Survey 2022/23*. Kathmandu: Ministry of Finance.
- Nepal Rastra Bank. (2023). *Annual Report 2022/23*. Kathmandu: Nepal Rastra Bank.
- Obstfeld, M., & Rogoff, K. (1996). *Foundations of International Macroeconomics*. Cambridge, MA: MIT Press.
- Pesaran, M. H., Shin, Y., & Smith, R. J. (2001). Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics*, 16(3), 289–326.
- Phillips, P. C. B., & Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, 75(2), 335–346.
- Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., & Mishchenko, E. F. (1962). *The Mathematical Theory of Optimal Processes*. New York: Interscience.
- Ramsey, F. P. (1928). A mathematical theory of saving. *Economic Journal*, 38(152), 543–559.
- Rose, A. K. (1991). The role of exchange rates in a popular model of international trade: Does the ‘Marshall–Lerner’ condition hold? *Journal of International Economics*, 30(3–4), 301–316.
- Santos Silva, J. M. C., & Tenreyro, S. (2006). The log of gravity. *Review of Economics and Statistics*, 88(4), 641–658.
- Sethi, S. P., & Thompson, G. L. (2000). *Optimal Control Theory: Applications to Management Science and Economics* (2nd ed.). Boston: Kluwer Academic Publishers.
- Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, 48(1), 1–48.
- Stokey, N. L., Lucas, R. E., Jr., & Prescott, E. C. (1989). *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.
- Toda, H. Y., & Yamamoto, T. (1995). Statistical inference in vector autoregressions with possibly integrated processes. *Journal of Econometrics*, 66(1–2), 225–250.
- Turnovsky, S. J. (2000). *Methods of Macroeconomic Dynamics* (2nd ed.). Cambridge, MA: MIT Press.
- World Bank. (2023). *Nepal Development Update: Restoring Export Competitiveness*. Kathmandu: World Bank.
- Sahani, S. K., Lee, T.-F., Pandey, D., Pandey, B. K., Jha, R., & Karna, S. L. (2026). Machine learning–augmented hybrid risk management and deep uncertainty quantification in Nepalese management systems: Fractional stochasticity, Wasserstein robustness, and rough-path neural filtering. *Journal of Intelligent Decision Making and Information Science*, 3(3s), 1850–1873. <https://doi.org/10.59543/jidmis.v3.97157>

Sahani, S. K., Lee, T.-F., Pandey, D., Pandey, B. K., & Mandal, K. (2026). Hybrid analytical, numerical, and machine learning frameworks for solving deterministic and stochastic differential equations with stability, convergence, and uncertainty quantification. *Journal of Intelligent Decision Making and Information Science*, 3(3s), 1794–1849. <https://doi.org/10.59543/jidmis.v3.97058>

Sahani, S. K., Lee, T.-F., Pandey, D., Pandey, B. K., Sah, B. K., & Jha, R. (2026). Hybrid analytical, numerical and machine learning approaches on decision-making, risk management, operational stability, and uncertainty analysis in Nepalese management systems. *Journal of Intelligent Decision Making and Information Science*, 3(3s), 1772–1793. <https://doi.org/10.59543/jidmis.v3.96959>

Sahani, S. K., Lee, T.-F., Pandey, D., Pandey, B. K., Sah, B. K., Jha, R., & Sah, D. K. (2026). Hybrid analytical–numerical–machine learning architectures for decision-making, risk management, operational stability, and uncertainty quantification in dental practice management systems: A mathematical theory. *European Journal of Prosthodontics and Restorative Dentistry*, 34(5s), 298–313. <https://doi.org/10.1922/ejprd.v34i5s.155960>

Sahani, S. K., Lee, T.-F., Pandey, D., Pandey, B. K., Sah, B. K., Mandal, K., & Jha, R. (2026). Biotechnological catalysts in hybrid analytical–numerical–machine learning architectures for risk-aware decision-making and uncertainty quantification in Nepalese management systems. *European Journal of Prosthodontics and Restorative Dentistry*, 34(7s), 332–350. <https://doi.org/10.1922/ejprd.v34i7s.16221>

Sahani, S. K., Lee, T.-F., Pandey, D., Pandey, B. K., Sah, B. K., & Jha, R. (2026). A hybrid analytical, numerical, and machine learning framework for health management, resource allocation, operational efficiency, risk forecasting, and strategic decision-making in Nepal. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(9s), 1221–1242. <https://doi.org/10.70917/ijcisim-2026-34002>

Sahani, S. K., Lee, T.-F., Pandey, D., Pandey, B. K., Jha, R., & Karna, S. L. (2026). Explainable machine learning for credit risk management and intelligent lending decisions in Nepalese cooperative banks: A mathematical review. *Journal of Intelligent Decision Making and Information Science*, 3(6s), 1735–1772.

Sahani, S. K., Lee, T.-F., Pandey, D., Pandey, B. K., Sah, B. K., & Karna, S. L. (2026). An explainable, privacy-preserving AI framework for credit risk in Nepalese cooperative banks: A federated survival graph Bayesian approach. *Phyton-Annales Rei Botanicae*, 66(2), 142–191.