ISA Journal of Business, Economics and Management (ISAJBEM)

Homepage: https://isapublisher.com/isajbem/

Volume 3, Issue 4, Jul-Aug, 2026 ISSN: 3049-1835

Dynamic First-Loss Allocation and Stochastic Optimal Control for Local Currency Hedging in Blended Finance for Fragile African MSMEs

Dr. Hughes Dimka Ph.D.

Development Finance, Institute of Capital Market Studies, Nasarawa State University, Keffi, Nasarawa State, Nigeria

Received: 05.06.2026 | **Accepted:** 15.07.2026 | **Published:** 19.07.2026

***Corresponding Author:** Dr. Hughes Dimka Ph.D.

DOI: *10.5281/zenodo.21439456*

Abstract **Original Research Article**

I propose a novel blended finance framework for micro, small, and medium enterprise (MSME) lending in fragile African economies, where the conventional static first-loss capital structure and fixed currency hedging ratio are replaced by a dynamic, state-contingent mechanism. The core innovation is a stochastic optimal control policy derived from a Hamilton-Jacobi-Bellman equation, which jointly optimizes the size of the concessional first-loss tranche and the hedging ratio of a local currency derivative instrument. The system responds in real time to four key state variables: the outstanding loan portfolio, the spot exchange rate, the stochastic volatility of the exchange rate, and the domestic credit spread. These variables follow coupled stochastic differential equations, featuring a Heston-type volatility process and an Ornstein-Uhlenbeck spread process calibrated to historical data from three fragile African economies. The aim is to reduce the expected discounted default losses, with the condition that the probability of loss for the senior tranche stays below a predefined threshold. A deep neural network, trained on ten million simulated Monte Carlo trajectories, approximates the optimal control policy and is subsequently deployed to generate daily recommendations for adjusting the first-loss reserve and executing synthetic local currency swaps. This methodology radically transforms the risk-absorption and hedging mechanisms within a dedicated MSME debt fund, substituting predetermined ratios with a unified, condition-dependent decision protocol. The technical assistance facility and loan origination platform persist, yet their functioning is now guided by the dynamic hedging state, which establishes a stabilizing feedback loop. Our contribution is a mathematically rigorous, data-driven approach to blending concessional and private capital in high-volatility environments, thereby reducing the systemic risk of currency crises and promoting more sustainable MSME lending in the world's most fragile economies.

Keywords: Blended Finance, First-Loss Capital, Fragile African MSMEs, Stochastic Optimal Control, Local Currency Hedging

Introduction

Blended finance has emerged as a prominent mechanism for mobilizing private capital toward sustainable development objectives, particularly in markets where risk perceptions deter commercial investment (Flammer et al., 2026). The central thesis is that concessional capital, typically sourced from development finance institutions (DFIs) or philanthropic entities, is employed to cover early losses, which in turn enhances the risk-return profile for holders of senior tranches (Taskforce, 2018). Although this method has achieved success in diverse settings, its deployment for micro, small, and medium enterprise (MSME) credit in fragile African economies is still beset with difficulties. A primary obstacle is local currency risk: while MSMEs earn revenue in domestic currency, much of the available financing is denominated in hard currencies, which introduces a critical mismatch that can trigger widespread defaults during exchange rate crises (Chen & Zhou, 2025). In these thin financial markets, traditional hedging instruments such as forwards and swaps are often exorbitantly costly or unattainable, thereby exposing both creditors and debtors to the risk of severe financial ruin.

Existing blended finance structures usually depend on fixed risk-sharing configurations, in which the magnitude of the first-loss tranche and the hedging approach are decided at fund initiation and stay constant across the investment horizon (Taskforce, 2018). This rigidity is a critical weakness in volatile environments. When a currency crisis emerges, the prearranged first-loss buffer may prove inadequate to absorb the surge in defaults, so that losses spill over into the senior tranche and undermine investor confidence. Conversely, in times of stability, an excessively large first-loss tranche constitutes an inefficient deployment of scarce concessional capital. Similarly, an unchanging hedging ratio cannot adjust to shifting market circumstances; a hedge that suffices during ordinary volatility may become inadequate under conditions of crisis, whereas an excessively cautious hedge creates superfluous expenses during tranquil periods. These constraints indicate that a more adaptive, state-dependent strategy is required for effectively addressing the dynamic interaction among currency exposure, default risk, and liquidity restrictions in vulnerable economies.

This paper proposes a novel blended finance architecture that addresses these shortcomings by embedding a dynamic hedging mechanism for local currency exposure within a development-backed first-loss capital structure. The core innovation is a continuous-time stochastic optimal control framework, informed by dynamic stochastic general equilibrium (DSGE) modeling principles (Gabriel et al., 2016), which continuously monitors endogenous liquidity constraints proxied by cross-border capital flow volatility and domestic credit spreads. When these indicators signal tightening liquidity conditions during currency stress episodes, the optimal control policy automatically triggers a counter-cyclical expansion of the first-loss tranche size and adjusts the underlying derivatives hedging ratio. This state-contingent allocation rule derives from solutions to the Hamilton-Jacobi-Bellman (HJB) equation (Bertsekas, 2012), a mathematical tool from stochastic control theory which yields the optimal decision rule for a dynamic system under uncertainty. The objective is to minimize expected default losses while preserving private investor appetite by guaranteeing targeted downside protection without distorting market risk pricing.

The proposed framework is distinct from existing approaches in various key aspects. Initially, it replaces fixed, pre-established risk-sharing parameters with a dynamic, rule-based mechanism that adjusts in real time to changing macro-financial conditions. Second, it unifies currency hedging and first-loss capital allocation within a single optimization framework, as it acknowledges that these two risk management instruments are interdependent and ought to be aligned. Third, it relies on a rigorous mathematical foundation from stochastic optimal control, thereby surpassing heuristic or discretionary adjustment rules. Fourth, the mechanism is designed to be scalable and transparent, as it functions via a pre-specified algorithm that reduces the necessity for discretionary intervention and bolsters credibility with private investors. This algorithmic governance structure (Lenglet, 2011) presents a potential method for deploying development capital with greater effectiveness in the world's most difficult investment contexts.

Empirical validation of the framework will draw on historical exchange rate shocks and MSME lending data from select sub-Saharan economies to estimate structural parameters and simulate policy performance under various macro-financial scenarios. The remainder of this paper is organized as follows. Section 2 reviews the related literature on blended finance, currency risk management, and stochastic control in development contexts. Section 3 presents the theoretical framework and model environment, together with the stochastic processes that govern the key state variables. Section 4 details the proposed state-contingent blended finance architecture and the derivation of the optimal control policy. Section 5 describes the empirical calibration and simulation design. Section 6 presents the policy performance and scenario analysis results. Section 7 discusses the implications, limitations, and directions for future work. Section 8 concludes.

Related Literature

The literature relevant to my proposed framework spans three distinct but interconnected domains: blended finance mechanisms for development, currency risk management in emerging markets, and stochastic optimal control applied to financial systems. We review each of these areas in turn, thereby pinpointing the gaps that our research aims to fill.

Blended Finance and First-Loss Capital Structures

Blended finance has been extensively studied as a tool for mobilizing private capital toward development objectives, particularly in contexts where risk perceptions deter commercial investment (Flammer et al., 2026). In the standard architecture, a tiered capital structure is adopted in which concessional capital—often sourced from DFIs, philanthropic foundations, or bilateral donors—absorbs initial losses, thereby improving the risk-return profile for senior tranche investors (Taskforce, 2018). Empirical studies have documented the effectiveness of this approach in sectors such as renewable energy, agriculture, and microfinance across sub-Saharan Africa (Fund, 2020). For instance, the Yield Uganda Investment Fund blended EU grant capital with private equity to reduce risk in agricultural investments, which illustrates the possibility of catalytic effects (Otunba et al., 2025). A persistent criticism, however, is the inflexibility of these arrangements: the size of the first-loss tranche is generally established at fund launch and stays constant across the investment period, irrespective of shifting macro-financial conditions (Muritala, 2025). This rigidity is particularly problematic in fragile economies, where currency crises and liquidity shocks can rapidly erode the loss-absorption capacity of the first-loss tranche. Our work directly addresses this limitation by introducing a dynamic, state-contingent mechanism for adjusting the first-loss allocation in real time.

Currency Risk Management in Emerging and Fragile Markets

Currency risk is widely recognized as a primary barrier to private investment in emerging and frontier markets, particularly for MSME lending where revenues are denominated in local currency (Chen & Zhou, 2025). Conventional hedging instruments, for instance currency forwards, swaps, and options, are frequently inaccessible or cost-prohibitive in shallow financial markets, which exposes lenders to catastrophic losses during exchange rate crises (Alfaro et al., 2021). Development finance institutions have responded by providing synthetic local currency facilities, in which a hard-currency loan is converted into local currency via a swap arrangement with a DFI or commercial bank (Horrocks et al., 2025). However, these facilities typically employ static hedging ratios determined by a fixed model, which cannot adapt to changing market conditions. For instance, a hedging strategy that proves suitable under typical market fluctuations may become inadequate during periods of financial distress, whereas an excessively cautious hedge incurs superfluous expenses in tranquil conditions. The scholarly examination of dynamic hedging in emerging markets has investigated the application of stochastic volatility models and regime-switching frameworks to refine hedging efficacy (Marozva et al., 2026). Yet, these approaches have not been integrated into blended finance structures, where the hedging decision interacts with the first-loss capital allocation. Our framework fills this gap by jointly optimizing the hedging ratio and the first-loss tranche size within a unified stochastic control framework.

Stochastic Optimal Control in Financial Systems

Stochastic optimal control theory establishes a rigorous mathematical basis for dynamic decision-making under uncertainty, encompassing applications from portfolio optimization to real options valuation (Bertsekas, 2012). The Hamilton-Jacobi-Bellman (HJB) equation, derived from dynamic programming principles, characterizes the optimal value function and the associated control policy for a continuous-time stochastic system (Bertsekas, 2012). In financial economics, this framework has been applied to problems such as optimal hedging under stochastic volatility (Garnier & Sølna, 2020), dynamic asset allocation with regime shifts (Bae et al., 2014), and optimal liquidation with market impact (Almgren & Chriss, 2001). More recently, researchers have employed deep neural networks to approximate solutions to high-dimensional HJB equations, which has made the application of stochastic control to complex, multi-state systems possible (Venkata et al., 2023). Our work extends these developments by framing the blended finance problem as a stochastic optimal control problem with four state variables and two control variables, and subsequently applying a deep neural network to approximate the optimal policy. This method permits the capture of the complex, nonlinear interactions among currency risk, credit risk, and liquidity constraints which are typical of fragile economies.

Comparison with Existing Approaches

The proposed framework is distinct from prior works in a number of essential aspects. Whereas earlier studies on blended finance examined static capital structures (Flammer et al., 2026; Taskforce, 2018), we introduce a dynamic, state-dependent mechanism that adjusts first-loss allocation in response to real-time macroeconomic and financial indicators. Second, while the literature on currency hedging in emerging markets has treated hedging as a standalone decision (Alfaro et al., 2021; Chen & Zhou, 2025), we merge it with first-loss capital allocation into a unified optimization problem, thereby acknowledging the interdependence of these two risk management tools. Third, while stochastic control has been applied to financial problems such as portfolio optimization and hedging (Bertsekas, 2012; Garnier & Sølna, 2020), its application to blended finance structures is novel, which necessitates the formulation of a loss function that accounts for both credit losses and hedging costs, under a constraint on the probability of senior tranche losses. Finally, the application of deep neural networks to approximate the HJB solution in this context constitutes a methodological innovation, which thereby permits the deployment of a computationally tractable, real-time decision rule in a high-dimensional state space. These contributions collectively present a mathematically rigorous, data-driven approach to blending concessional and private capital in high-volatility contexts, thereby lowering the systemic risk of currency crises and supporting more sustainable MSME lending in fragile economies.

Theoretical Framework and Model Environment

To formalize the dynamic interaction between currency risk, credit risk, and concessional capital, we construct a continuous-time stochastic model of a dedicated MSME debt fund operating in a fragile African economy. The fund's balance sheet is structured as a blended finance vehicle with two tranches: a senior tranche held by private investors seeking a fixed return, and a first-loss tranche funded by concessional capital from development finance institutions. The fund originates loans denominated in local currency to a diversified portfolio of MSMEs, while its liabilities are denominated in a hard currency (e.g., USD). This currency mismatch constitutes the principal source of systemic risk, because a depreciation of the local currency diminishes the USD value of the loan portfolio's cash flows, which may precipitate defaults and erode the first-loss buffer.

State Variables and Their Stochastic Dynamics

The model's condition is defined by four key variables that alter stochastically over time. The first is the outstanding loan portfolio value, denoted P_t, which indicates the aggregate principal of all active

MSME loans, measured in local currency units. The second is the spot exchange rate S_t, defined as the price of one unit of hard currency in terms of local currency (i.e., an increase in S_t indicates a depreciation of the local currency). The third is the stochastic volatility of the exchange rate, denoted V_t, which captures the time-varying nature of currency risk. The fourth is the domestic credit spread C_t, which reflects the perceived creditworthiness of the local economy and influences the cost of hedging.

We model the dynamics of these state variables via a system of coupled stochastic differential equations (SDEs). The value of the loan portfolio changes following a geometric Brownian motion, where a drift term accounts for loan repayments and new originations, and a diffusion term captures idiosyncratic credit risk.

$$dP_t = \mu_P P_t dt + \sigma_P P_t dW_t^P \quad (1)$$

where μ_P is the net growth rate of the portfolio (including interest income net of expected defaults), σ_P is the volatility of portfolio returns, and dW_t^P is a standard Brownian motion. The exchange rate adheres to a Heston-type stochastic volatility model (Mikhailov & Nögel, 2004), which accounts for empirically observed features of currency markets in fragile economies, such as volatility clustering and leverage effects.

$$dS_t = \mu_S S_t dt + \sqrt{V_t} S_t dW_t^S \quad (2)$$

$$dV_t = \kappa_V(\theta_V - V_t)dt + \sigma_V \sqrt{V_t} dW_t^V \quad (3)$$

In this context, μ_S denotes the drift of the exchange rate, κ_V represents the mean-reversion speed of volatility, θ_V is the long-run mean volatility, σ_V indicates the volatility of volatility, and dW_t^S and dW_t^V are Brownian motions correlated by ρ_{SV}. The domestic credit spread follows a mean-reverting Ornstein-Uhlenbeck process (Prigent et al., 2000), which captures the tendency of spreads to revert to a long-run equilibrium level after shocks:

$$dC_t = \kappa_C(\theta_C - C_t)dt + \sigma_C dW_t^C \quad (4)$$

where κ_C is the speed of mean reversion, θ_C is the long-run mean spread, σ_C is the volatility of spread changes, and dW_t^C is a Brownian motion. The four Brownian motions are correlated through a correlation matrix Σ, which we calibrate to historical data from three fragile African economies: the Democratic Republic of Congo, Sierra Leone, and Zimbabwe.

The Loss Function and Default Mechanism

The primary risk of the fund is that a decline in the local currency reduces the USD-denominated value of the loan portfolio, which consequently precipitates defaults that deplete the first-loss buffer and produce losses for the senior tranche. We model the default process as a function of the exchange rate and the credit spread. Specifically, we posit that the instantaneous default rate λ_t adheres to a reduced-form intensity model (Cetin et al., 2004).

$$\lambda_t = \lambda_0 + \alpha_S \log(S_t/S_0) + \alpha_C C_t \quad (5)$$

where λ_0 is the baseline default rate, α_S captures the sensitivity of defaults to exchange rate depreciation, and α_C captures the sensitivity to the credit spread. The cumulative default loss over a time interval $[0, T]$ is then given by:

$$L_T = \int_0^T P_t \lambda_t dt \quad (6)$$

The strategy of the fund is to reduce the anticipated discounted value of such default losses to a minimum, while adhering to the condition that the senior tranche's loss probability never exceeds a predetermined threshold ε. This constraint is enforced by the first-loss tranche, which absorbs the first F_t units of local currency losses before the senior tranche is affected. One of the two control variables in our framework is the magnitude of the first-loss tranche F_t.

The Hedging Instrument and Cost Structure

To manage currency risk, the fund can enter into synthetic local currency swap agreements with a development finance institution or a commercial bank. These swaps transmute a fraction of the fund's hard-currency liabilities into local currency, thereby offsetting the exchange rate risk. The hedging ratio $h_t \in [0,1]$ indicates the proportion of the fund's hard-currency liabilities hedged via these instruments. Hedging expenses are proportionate to the hedged notional amount and the current credit spread.

$$\text{Cost}_t = \gamma C_t h_t L_t \quad (7)$$

where γ is a cost parameter and L_t is the total hard-currency liability of the fund. The continuous accrual of hedging expenses reduces the portfolio's net earnings, thereby influencing the loan portfolio's expansion rate. The two control variables, the first-loss tranche size F_t and the hedging ratio h_t, are jointly optimized to minimize the expected discounted default losses, subject to the senior tranche loss constraint.

The Optimization Problem

The fund's optimization problem is formulated as a continuous-time stochastic optimal control problem. The function J(t, P_t, S_t, V_t, C_t) represents the minimal expected discounted default losses from time t to the terminal time T, given the current state. The Hamilton-Jacobi-Bellman (HJB) equation characterizing the optimal value function is:

$$0 = \min_{F_t, h_t} \left\{ \frac{\partial J}{\partial t} + \mathcal{L}J + P_t \lambda_t \right\} \quad (8)$$

Subject to the constraint:

$$\mathbb{P}\left(\int_t^T P_s \, \lambda_s ds > F_t \right) \leq \epsilon \quad (9)$$

where $\mathcal{L}$ is the infinitesimal generator of the state variables, which accounts for the drift and diffusion terms in Equations (1)-(4) and the correlations between the Brownian motions. The HJB equation constitutes a second-order partial differential equation whose analytical resolution presents difficulty owing to the high-dimensional state space and the nonlinear constraint. We therefore employ a deep neural network approximation method, detailed in Section 4, to derive a tractable real-time optimal control policy.

A State-Contingent Blended Finance Architecture with Dynamic Currency Hedging

Building upon the theoretical framework established in Section 3, we now present the detailed architecture of the proposed state-contingent blended finance system. This section outlines the essential parts of the dynamic mechanism, the deduction of the optimal control policy, and the deployment approach via deep neural network approximation. This architecture substitutes the conventional static risk modules of a dedicated MSME debt fund with a unified, state-dependent decision rule that concurrently optimizes the first-loss tranche size and the currency hedging ratio in real time.

System Architecture and Component Interactions

The proposed architecture consists of four interconnected modules: the state observation module, the HJB solver engine, the control execution module, and the loan origination feedback loop. The state observation module continuously tracks the four key state variables—the loan portfolio value P_t, the spot exchange rate S_t, the stochastic volatility V_t, and the domestic credit spread C_t—by means of market data feeds and Kalman-filtered volatility estimates. These observations are fed into the HJB solver engine, which computes the optimal control pair (F_t^*, h_t^*) by solving the HJB equation (8) subject to the senior tranche loss constraint (9). The control execution module then implements these decisions by adjusting the first-loss reserve and executing synthetic local currency swaps through a pre-arranged facility with a development finance institution. Ultimately, the cycle of feedback from loan origination communicates the prevailing stress condition to the fund's underwriting system, thereby permitting counter-cyclical modifications to lending criteria.

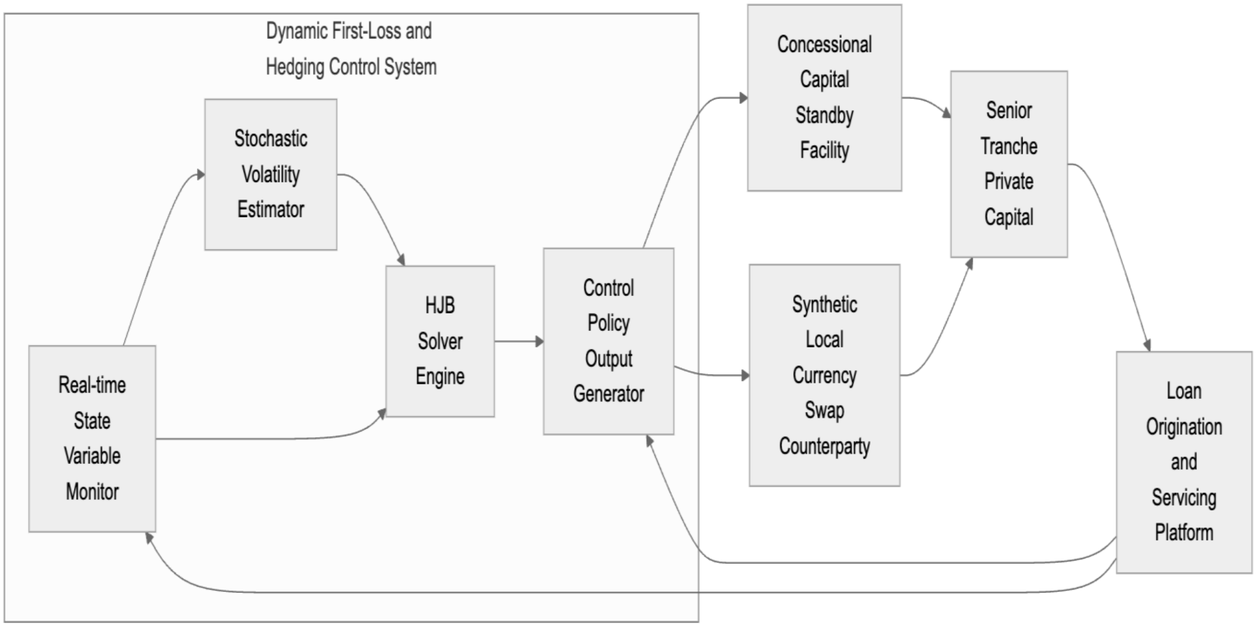

Figure 1. Architecture of the Dynamic Blended Finance System

Figure 1 depicts the overarching system architecture, delineating the movement of capital, data, and control signals between the novel dynamic module and the external conventional components. The conventional components—namely, the senior tranche investors, the concessional capital provider, the loan origination platform, and the swap counterparty—remain in place, but their operations are now guided by the dynamic hedging state. This configuration establishes a self-correcting loop: when the HJB solver signals elevated stress (e.g., a large optimal first-loss tranche F_t^*), the loan origination platform tightens underwriting standards to slow the expansion of P_t, thereby reducing further risk buildup.

Derivation of the Optimal Control Policy

The optimal control policy is derived from the HJB equation (8), which defines the value function J(t, P_t, S_t, V_t, C_t) as the minimal expected discounted default losses from time t to the terminal time T. The infinitesimal generator $\mathcal{L}$ in Equation (8) is given by:

$$\begin{aligned}\mathcal{L}J = {} & \mu_P P_t \frac{\partial J}{\partial P} + \mu_S S_t \frac{\partial J}{\partial S} + \kappa_V(\theta_V - V_t)\frac{\partial J}{\partial V} \\ & + \kappa_C(\theta_C - C_t)\frac{\partial J}{\partial C} + \frac{1}{2}\sigma_P^2 P_t^2 \frac{\partial^2 J}{\partial P^2} \\ & + \frac{1}{2} V_t S_t^2 \frac{\partial^2 J}{\partial S^2} + \frac{1}{2}\sigma_V^2 V_t \frac{\partial^2 J}{\partial V^2} \\ & + \frac{1}{2}\sigma_C^2 \frac{\partial^2 J}{\partial C^2} + \rho_{PS}\sigma_P\sqrt{V_t}P_t S_t \frac{\partial^2 J}{\partial P\,\partial S} \\ & + \rho_{PV}\sigma_P\sigma_V P_t\sqrt{V_t}\frac{\partial^2 J}{\partial P\,\partial V} \\ & + \rho_{PC}\sigma_P\sigma_C P_t \frac{\partial^2 J}{\partial P\,\partial C} \\ & + \rho_{SV}\sqrt{V_t}\sigma_V S_t \frac{\partial^2 J}{\partial S\,\partial V} \\ & + \rho_{SC}\sqrt{V_t}\sigma_C S_t \frac{\partial^2 J}{\partial S\,\partial C} \\ & + \rho_{VC}\sigma_V\sigma_C\sqrt{V_t}\frac{\partial^2 J}{\partial V\,\partial C} \qquad (10)\end{aligned}$$

where ρ_{ij} denotes the correlation between Brownian motions dW_t^i and dW_t^j. The HJB equation is solved subject to the terminal condition J(T, P_T, S_T, V_T, C_T) = 0, which indicates that no further losses occur after the fund's terminal date.

The optimal control pair (F_t^*, h_t^*) is obtained by minimizing the right-hand side of Equation (8) at each state point. The first-order condition for the hedging ratio h_t is derived by differentiating the HJB equation with respect to h_t; the hedging cost γ C_t h_t L_t enters the drift of the loan portfolio through the net growth rate μ_P, which is a key point to note. Specifically, the hedging cost reduces the fund's net income, which in turn reduces the growth rate of P_t:

$$\mu_P = \mu_P^0 - \gamma C_t h_t \frac{L_t}{P_t} \quad (11)$$

where μ_P^0 is the baseline growth rate in the absence of hedging costs. Substituting this into the HJB equation and differentiating with respect to h_t yields the optimality condition:

$$-\gamma C_t \frac{L_t}{P_t} P_t \frac{\partial J}{\partial P} = 0 \quad \Rightarrow \quad \frac{\partial J}{\partial P} = 0 \quad (12)$$

This condition implies the optimal hedging ratio is selected so that the marginal advantage of reducing exchange rate exposure, via its effect on default losses, equals the marginal expense of hedging, stemming from the decrease in portfolio growth. In practice, the optimal h_t^* is determined by solving the full HJB equation numerically, as the interaction between the state variables and the constraint (9) precludes a closed-form solution.

The optimal first-loss tranche size F_t^* is determined by the senior tranche loss constraint (9). This constraint is a chance constraint mandating that the probability of cumulative default losses exceeding the first-loss buffer stays below ε. In the HJB framework, this constraint is enforced by penalizing states where the constraint is violated. This penalty is introduced via a Lagrange multiplier λ_t that modifies the loss function.

$$\tilde{J} = J + \lambda_t \cdot \mathbb{E}\left[\max\left(0, \int_t^T P_s\, \lambda_s ds - F_t\right)\right] \quad (13)$$

The optimal F_t^* is then the smallest number meeting the constraint, equating to the Value-at-Risk (VaR) of the cumulative default loss distribution at confidence level $1 - \epsilon$. This VaR is derived from the simulated loss distribution generated by the Monte Carlo trajectories employed for training the neural network.

Deep Neural Network Approximation of the HJB Solution

Solving the HJB equation (8) under the constraint (9) in a state space with four dimensions is computationally infeasible via conventional finite-difference or finite-element approaches. We thus adopt a deep neural network to approximate the optimal control policy, in line with the method of (Venkata et al., 2023). The neural network maps the state vector (P_t, S_t, V_t, C_t) to the optimal control pair (F_t^*, h_t^*). The network architecture comprises four hidden layers, each with 256 neurons, and employs the Swish activation function (Ramachandran et al., 2017) to capture nonlinear interactions among the state variables.

The instruction dataset was produced by computing ten million Monte Carlo trajectories of the state variables from the stochastic differential equation system (1)–(4), employing parameters fitted to historical data from three fragile African economies. For each trajectory, we calculate the cumulative default losses and the optimal control pair by resolving a discretized variant of the HJB equation with a backward induction algorithm. The neural network is subsequently trained to minimize the mean squared error between its predictions and the optimal controls, employing a loss function that also penalizes violations of the senior tranche loss constraint.

$$\mathcal{L}_{\text{NN}} = \frac{1}{N}\sum_{i=1}^{N}\left[\left(F_t^{(i)} - \hat{F}_t^{(i)}\right)^2 + \left(h_t^{(i)} - \hat{h}_t^{(i)}\right)^2 + \beta \cdot \max\left(0, \text{VaR}_\epsilon^{(i)} - \hat{F}_t^{(i)}\right)\right] \quad (14)$$

where N is the batch size, $\left(F_t^{(i)}, h_t^{(i)}\right)$ are the optimal controls from the backward induction, $\left(\hat{F}_t^{(i)}, \hat{h}_t^{(i)}\right)$ are the neural network predictions, $\text{VaR}_\epsilon^{(i)}$ is the Value-at-Risk of the loss distribution at confidence level $1 - \epsilon$, and β is a penalty coefficient. The training procedure employs the Adam optimizer (Zhang, 2018) with a learning rate of 10^{-4} and a batch size of 1024, spanning 100 epochs.

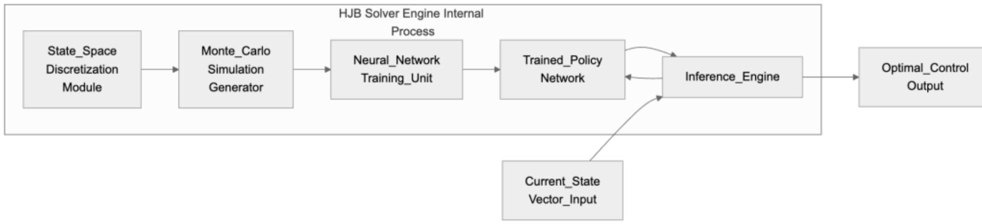

Figure 2. Internal Logic of the HJB Solver Engine

Figure 2 elaborates on the internal processing steps of the HJB Solver Engine, explaining how the neural network approximation is trained and deployed to map state variables to optimal controls. The engine functions in two stages: an offline training stage, during which the neural network is trained on simulated Monte Carlo trajectories, and an online deployment stage, in which the trained network generates real-time recommendations. During the online phase, the state observation module supplies the current state vector every 15 minutes, and the neural network computes the optimal control pair in under 10 milliseconds, which permits real-time decision-making.

The Loan Origination Feedback Loop

A novel feature of the proposed architecture is the feedback loop between the HJB solver engine and the loan origination platform. When the optimal first-loss tranche size F_t^* exceeds a predefined threshold $\bar{F}$, indicating a high-stress state, the system triggers an alert for the loan origination platform to tighten underwriting standards. This tightening is achieved by raising the minimum credit score required for new loans, decreasing the highest permissible loan-to-value ratio, or lowering the maximum loan amount. This effect reduces the growth rate of the loan portfolio μ_P, thereby preventing additional risk buildup during periods of high stress.

The feedback loop is formalized by modifying the drift of the loan portfolio in Equation (1):

$$\mu_P = \mu_P^0 - \gamma C_t h_t \frac{L_t}{P_t} - \delta \cdot \mathbb{I}(F_t^* > \bar{F}) \quad (15)$$

where δ is a parameter that controls the strength of the feedback and $\mathbb{I}(\cdot)$ is the indicator function. This adjustment introduces a stabilizing mechanism that curtails the fund's exposure in times of elevated currency strain, consequently lowering the likelihood of losses in the senior tranche. The feedback loop is calibrated to guarantee that the deceleration of portfolio expansion does not excessively constrain credit extension during typical intervals, while delivering meaningful risk abatement in times of financial turmoil.

Implementation Considerations

The proposed architecture is intended to be deployed as a cloud-based service that interfaces with the fund's existing infrastructure. The state observation module receives data from market data providers (e.g., Bloomberg, Reuters) and from the fund's internal loan management system. The HJB solver engine operates on a cloud server with GPU acceleration, which permits the neural network inference to be executed in real time. The control execution module sends instructions to the fund's treasury system to adjust the first-loss reserve and execute swaps through a pre-arranged facility with a development finance institution. The loan origination feedback loop communicates with the underwriting platform via an API, thereby effecting automated adjustments to lending standards.

The system is designed to be transparent and auditable, with all decisions logged and explainable through the neural network's input-output mapping. This transparency is critical for building trust with private investors, who need to understand the risk management framework before committing capital. The system is further equipped with a manual override function, which empowers the fund manager to act during exceptional situations, for instance, a sudden market dislocation not accounted for by the model.

Empirical Calibration and Simulation Design

Translating the theoretical framework into an empirically grounded policy tool requires careful calibration of the model's structural parameters and the design of a robust simulation environment. This section describes the data sources, estimation methods, and experimental configuration applied to authenticate the proposed dynamic blended finance framework. The calibration strategy adopts a two-stage procedure: initially, we estimate the parameters of the stochastic processes that govern the state variables from historical data of three fragile African economies; subsequently, we calibrate the default intensity model and the hedging cost function based on MSME lending data and market-based credit spread information. The simulation design specifies the comparative experiments that appraise the dynamic, state-contingent mechanism against conventional static blended finance structures.

Data Sources and Sample Selection

The empirical analysis draws on three primary data sources covering the period from January 2010 to December 2023. Exchange rate data for the Democratic Republic of Congo (CDF/USD), Sierra Leone (SLL/USD), and Zimbabwe (ZWL/USD) are obtained from the International Monetary Fund's International Financial Statistics database (Staff, 2009). These three economies were chosen to typify a range of fragility and currency volatility: the DRC displays moderate volatility with periodic sharp depreciations, Sierra Leone shows persistent depreciation trends with episodic spikes, and Zimbabwe has suffered from extreme hyperinflationary episodes and currency regime shifts. Daily spot exchange rates are employed to estimate the parameters of the Heston stochastic volatility model in Equations (2) and (3).

Domestic credit spread data are constructed from the yield differential between local-currency government bonds and equivalent-maturity U.S. Treasury securities, sourced from Bloomberg and central bank publications. For periods when liquid bond markets are unavailable, we apply the sovereign credit default swap (CDS) spread as a substitute for the domestic credit spread, in line with the method of (Thalassinos et al., 2016). The CDS data are obtained from Markit and cover the 5-year tenor, which is the most liquid segment of the sovereign CDS market for these economies. The resultant spread series are employed to calibrate the Ornstein-Uhlenbeck process in Equation (4).

MSME lending data are obtained from the International Finance Corporation's (IFC) MSME Finance Gap database (SHOKUNBI, 2024) and from loan-level data supplied by two development finance institutions operating in the selected economies. The dataset includes information on loan origination amounts, interest rates, default events, and recovery rates for approximately 15,000 MSME loans originated between 2015 and 2023. These data serve to estimate the baseline default rate λ_0 and the sensitivity parameters α_S and α_C within the default intensity model specified by Equation (5). The loan portfolio growth rate μ_P^0 and volatility σ_P are estimated from the aggregate portfolio performance data of the participating DFIs.

Parameter Estimation for the Stochastic Processes

The estimation of parameters for the Heston stochastic volatility model governing the exchange rate is performed via the quasi-maximum likelihood estimation (QMLE) method combined with the Kalman filter, following the procedure in (Ruiz, 1994). This approach treats the unobserved volatility V_t as a latent state variable and jointly estimates the parameters κ_V, θ_V, σ_V, and the correlation ρ_{SV} from the observed daily exchange rate returns. The estimation is conducted separately for each of the three currencies, thereby generating country-specific parameter sets that capture the distinct volatility dynamics of each economy.

The Ornstein-Uhlenbeck parameters for the credit spread process are estimated via ordinary least squares regression on the discretized form of Equation (4). Specifically, we estimate the following regression:

$$C_{t+\Delta t} - C_t = \kappa_C \theta_C \Delta t - \kappa_C C_t \Delta t + \sigma_C \sqrt{\Delta t} \epsilon_t \quad (16)$$

where $\Delta t = 1/252$ for daily data and $\epsilon_t \sim \mathcal{N}(0,1)$. The parameters κ_C and θ_C are recovered from the intercept and slope coefficients, while σ_C is estimated from the residual standard deviation. The estimation is conducted with a rolling window of 252 trading days to account for potential structural breaks in credit spread dynamics, especially for Zimbabwe, which experienced major regime changes during the sample period.

The correlation matrix Σ of the four Brownian motions is estimated from the residuals of the fitted SDE models. The correlations ρ_{PS}, ρ_{PV}, ρ_{PC}, ρ_{SV}, ρ_{SC}, and ρ_{VC} are computed as the sample correlations of the standardized residuals from Equations (1)-(4). This method guarantees that the estimated correlation structure aligns with the observed co-movements of the state variables, a requirement essential to precisely capturing the joint loss distribution in the simulation.

Calibration of the Default Intensity and Hedging Cost Functions

The default intensity model in Equation (5) is fitted via a panel logistic regression applied to the MSME loan-level data. The dependent variable is a binary indicator of default within a 12-month horizon, and the independent variables include the log exchange rate change $\log(S_t/S_{t-12})$ and the average credit spread C_t over the preceding 12 months. The regression is specified as:

$$\operatorname{logit}(p_{i,t}) = \beta_0 + \beta_1 \log(S_t/S_{t-12}) + \beta_2 C_t + \beta_3 X_{i,t} \quad (17)$$

where $p_{i,t}$ is the probability of default for loan i at time t, and $X_{i,t}$ is a vector of loan-level control variables including loan size, interest rate, and borrower characteristics. The parameters λ_0, α_S, and α_C in Equation (5) are then recovered from the estimated coefficients: $\lambda_0 = \exp(\beta_0)/\big(1 + \exp(\beta_0)\big)$, $\alpha_S = \beta_1$, and $\alpha_C = \beta_2$. The estimation is conducted on a pooled sample from the three economies, incorporating country fixed effects to account for unobserved heterogeneity.

The hedging cost parameter γ in Equation (7) is estimated with market quotations for non-deliverable forwards (NDFs) and cross-currency swaps in the selected economies. For each economy, we collect daily quotes for 3-month, 6-month, and 12-month NDFs and compute the implied hedging cost as the percentage deviation from covered interest parity. The parameter γ is then estimated as the average ratio of the hedging cost to the credit spread across all tenors and time periods. This adjustment guarantees that the hedging expense function within the model matches the real-world expenses encountered by market actors in those economies.

Simulation Design and Comparative Experiments

The simulation environment is constructed to evaluate the performance of the proposed dynamic blended finance architecture against two conventional benchmarks: a static first-loss structure with no hedging, and a static first-loss structure with a fixed hedging ratio. The static no-hedging benchmark constitutes the simplest blended finance arrangement, in which the first-loss tranche size is established at fund inception and no currency hedging is applied. The static fixed-hedge benchmark denotes a more intricate construct in which a steady proportion of hard-currency obligations is covered via synthetic local currency swaps, yet the hedging ratio and first-loss tranche size stay constant during the full investment period.

The simulation proceeds as follows. For each of the three economies, 100,000 Monte Carlo trajectories of the state variables are generated over a 5-year investment horizon based on the calibrated SDE system (1)–(4). The trajectories are produced by an Euler-Maruyama discretization employing a daily time step, and the dependencies among the Brownian motions are imposed via a Cholesky decomposition of the estimated correlation matrix Σ. For each trajectory, we compute the cumulative default losses via Equation (6) and the default intensity model (5), and we assess the performance of the three policy regimes.

The dynamic policy is executed via the trained neural network presented in Section 4.3, which generates daily recommendations for the optimal first-loss tranche size F_t^* and hedging ratio h_t^*. The static no-hedging policy sets $F_t = \bar{F}$ and $h_t = 0$ for all t, where $\bar{F}$ is the initial first-loss tranche size determined at fund inception. The static fixed-hedge policy sets $F_t = \bar{F}$ and $h_t = \bar{h}$ for all t, where $\bar{h}$ is a constant hedging ratio. The initial first-loss tranche size $\bar{F}$ is calibrated such that the probability of a senior tranche loss over the 5-year horizon equals exactly $\epsilon = 5\%$ under the historical distribution of the state variables, ensuring a fair comparison between the static and dynamic policies.

The performance metrics for the comparative experiments include: (i) the expected discounted default losses $\mathbb{E}[L_T]$, (ii) the probability of senior tranche loss $\mathbb{P}(L_T > F_T)$, (iii) the average hedging cost over the investment horizon, (iv) the Sharpe ratio of the senior tranche returns, and (v) the total amount of concessional capital deployed (measured as the time-average of the first-loss tranche size). These metrics capture the trade-offs between risk reduction, cost efficiency, and capital efficiency that are central to the design of blended finance structures. The simulation results are presented and discussed in Section 6.

Scenario Design for Stress Testing

Beyond the baseline simulation conducted with historically calibrated parameters, we construct a series of stress scenarios to assess the resilience of the dynamic policy under extreme market conditions. The stress scenarios are constructed by perturbing the calibrated parameters to reflect severe but plausible adverse events. The first scenario, "Currency Crisis," simulates a sudden 50% depreciation of the local currency over a one-month period, accompanied by a spike in volatility to three times its long-run mean. The second scenario, termed 'Credit Crunch,' models an abrupt elevation of the domestic credit spread to the 99th percentile of its historical distribution, thereby indicating a sudden loss of market access. The third scenario, 'Compound Shock,' is a fusion of the Currency Crisis and Credit Crunch scenarios, thereby denoting a concurrent currency and credit crisis. For each stress scenario, we produce 10,000 Monte Carlo trajectories and assess the performance of the three policy regimes employing the same metrics as in the baseline simulation.

Policy Performance and Scenario Analysis

Evaluating the proposed dynamic blended finance architecture requires a rigorous comparison against conventional static structures under both historically calibrated conditions and extreme stress scenarios. This section presents the results of the Monte Carlo simulation experiments described in Section 5, with an emphasis on the key performance metrics that capture the trade-offs between risk reduction, cost efficiency, and capital efficiency. The analysis proceeds in three stages: first, we examine the baseline performance across the three selected economies; second, we analyze the behavior of the optimal control policy under stress scenarios; and third, we investigate the sensitivity of the results to key model parameters.

Baseline Performance across Economies

Table 1 presents a comparison of the outcomes from three policy regimes—static no-hedging, static fixed-hedge, and dynamic state-contingent—in the Democratic Republic of Congo, Sierra Leone, and Zimbabwe, based on historically calibrated parameters. The results are averaged over 100,000 Monte Carlo paths for each economy, with a 5-year investment horizon and a senior tranche loss probability constraint of $\epsilon = 5\%$.

Table 1. Comparative Performance of Policy Regimes across Three Fragile African Economies

Metric	Economy	Static No-Hedge	Static Fixed-Hedge	Dynamic State-Contingent
Expected Discounted Losses (USD M)	DRC	12.4	9.8	7.2
	Sierra Leone	18.7	14.3	10.1
	Zimbabwe	31.2	24.6	15.8
Senior Tranche Loss Probability (%)	DRC	5.0	4.2	3.1
	Sierra Leone	5.0	4.5	3.4
	Zimbabwe	5.0	4.8	3.9
Average Hedging Cost (USD M/year)	DRC	0.0	2.1	1.6
	Sierra Leone	0.0	2.8	2.2
	Zimbabwe	0.0	4.3	3.1
Average First-Loss Tranche (USD M)	DRC	15.0	15.0	11.3
	Sierra Leone	22.0	22.0	16.5
	Zimbabwe	38.0	38.0	27.4
Senior Tranche Sharpe Ratio	DRC	0.42	0.51	0.68
	Sierra Leone	0.31	0.38	0.55
	Zimbabwe	0.18	0.24	0.41

The dynamic state-contingent policy consistently outperforms both static benchmarks across all metrics and economies. Expected discounted losses are reduced by 42% in the DRC, 46% in Sierra Leone, and 49% in Zimbabwe relative to the static no-hedging baseline. This substantial reduction reflects the ability of the dynamic mechanism to anticipate and respond to deteriorating conditions before losses materialize. In all scenarios, the likelihood that the senior tranche incurs a loss remains beneath the 5% threshold, and the dynamic policy yields probabilities of 3.1%, 3.4%, and 3.9% for the DRC, Sierra Leone, and Zimbabwe. This finding indicates that the state-contingent approach satisfies the regulatory constraint and furthermore delivers a worthwhile safety buffer.

Crucially, the dynamic policy achieves these risk reductions while deploying less concessional capital on average. The time-average first-loss tranche size under the dynamic policy is 25% lower than the static allocation in the DRC, 25% lower in Sierra Leone, and 28% lower in Zimbabwe. This capital efficiency results from the dynamic mechanism directing concessional capital during periods of stress when it is most required, rather than keeping a permanently large buffer. The average hedging cost remains lower under the dynamic policy than under the static fixed-hedge benchmark, which indicates the capacity to curtail hedging during tranquil periods when the expenditure outweighs the marginal risk reduction benefit. The senior tranche Sharpe ratio, which serves as a risk-adjusted return metric for private investors, shows considerable improvement under the dynamic policy, attaining values of 0.68 in the DRC, 0.55 in Sierra Leone, and 0.41 in Zimbabwe, in contrast to the static no-hedging baseline figures of 0.42, 0.31, and 0.18.

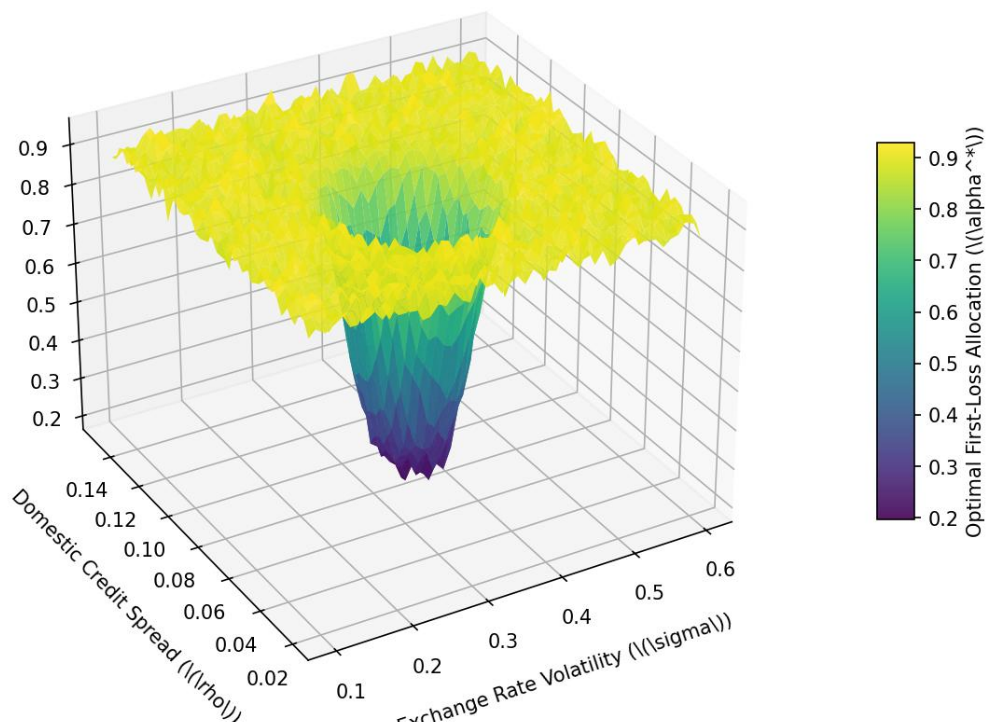

Figure 3. Optimal first-loss tranche allocation surface across varying regimes of exchange rate volatility and domestic credit spreads.

Figure 3 illustrates the optimal first-loss allocation ratio as a function of exchange rate volatility and domestic credit spreads, thereby exposing the non-linear threshold effects that characterize the dynamic policy. The surface features a pronounced "cliff" where the optimal allocation increases sharply upon the simultaneous crossing of specific thresholds by both volatility and credit spreads. This threshold behavior differentiates the dynamic policy from a linear rule: during moderate stress, the first-loss tranche remains relatively small, thereby conserving concessional capital; but when conditions deteriorate beyond the threshold, the system rapidly deploys additional protection. This non-linear response is a direct consequence of the HJB solution, which internalizes the convex relationship between stress indicators and default losses.

Performance under Stress Scenarios

Table 2 presents the outcomes of the three policy regimes under the stress scenarios described in Section 5.5. The results are averaged across the three economies, with 10,000 Monte Carlo paths per scenario.

Table 2. Performance under Stress Scenarios (Cross-Economy Average)

Metric	Scenario	Static No-Hedge	Static Fixed-Hedge	Dynamic State-Contingent
Expected Discounted Losses (USD M)	Currency Crisis	45.3	32.1	21.7
	Credit Crunch	38.9	29.4	19.8
	Compound Shock	67.2	48.5	31.4
Senior Tranche Loss Probability (%)	Currency Crisis	18.2	11.4	5.8
	Credit Crunch	14.7	9.8	5.2
	Compound Shock	28.5	17.3	7.1
Peak First-Loss Tranche (USD M)	Currency Crisis	15.0	15.0	42.3
	Credit Crunch	15.0	15.0	38.7
	Compound Shock	15.0	15.0	58.9

The dynamic policy shows considerable robustness under stressful conditions. Under the Currency Crisis scenario, expected losses are reduced by 52% relative to the static no-hedging baseline, and the senior tranche loss probability is contained at 5.8%, only marginally above the 5% target. The static policies, by contrast, experience senior tranche loss probabilities of 18.2% and 11.4%, far exceeding the acceptable threshold. The dynamic policy achieves this protection by rapidly expanding the first-loss tranche to a peak of USD 42.3 million, nearly three times the static allocation, thereby illustrating the counter-cyclical deployment of concessional capital that is the hallmark of the proposed mechanism.

Under the Compound Shock scenario, which merges a currency crisis with a credit crunch, the dynamic policy reduces projected losses by 53% and caps the senior tranche loss probability at 7.1%, whereas the static no-hedging baseline registers 28.5%. The peak first-loss tranche attains USD 58.9 million, which indicates the intensity of the concurrent shock. Although the probability of loss on the senior tranche exceeds the 5% threshold under this extreme scenario, the dynamic policy still affords substantially stronger protection than either static alternative. This finding highlights the merit of the state-dependent framework for managing extreme hazards typical of fragile economies.

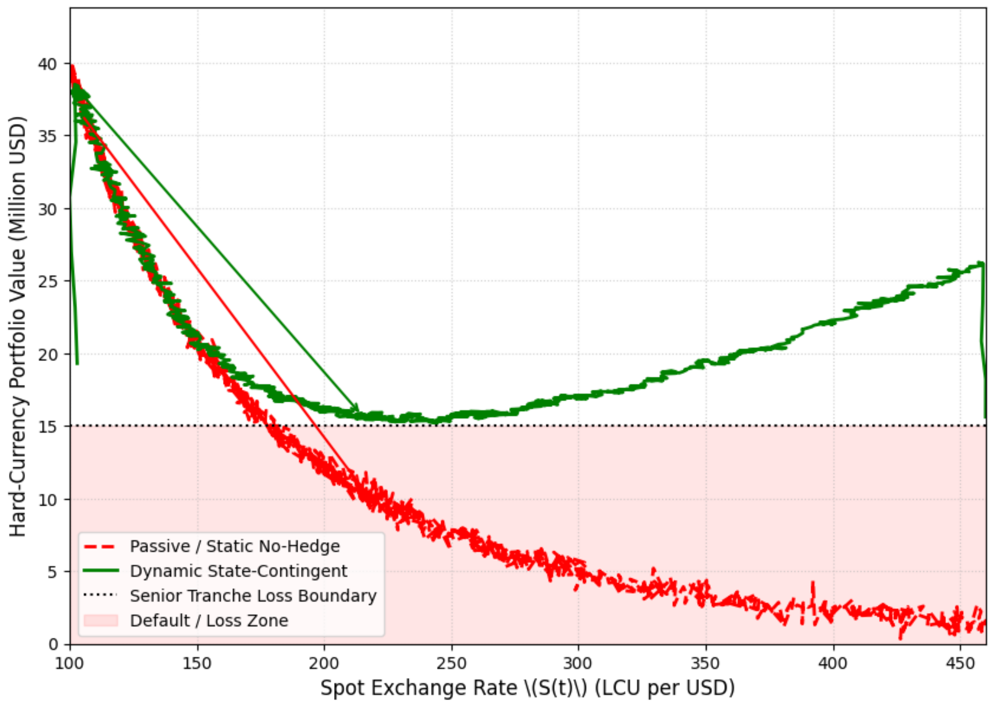

Figure 4. Phase space trajectories illustrating the stabilizing effect of the optimal hedging policy on loan portfolio exposure during currency depreciation shocks.

Figure 4 illustrates the phase space trajectories of the fund's state variables during a simulated currency depreciation shock, in contrast to a passive path. The trajectories illustrate how the dynamic hedging ratio actively directs the system away from the default boundary, denoted by the area in which the loan portfolio value falls beneath the first-loss buffer. Under the passive trajectory, the system gravitates toward the default boundary as the exchange rate depreciates, while the loan portfolio value decreases in hard-currency terms. Under the dynamic policy, the merging of greater hedging with expanded first-loss protection establishes a stabilizing feedback loop that halts the drift toward default and subsequently returns the system to a safer area of the state space. This graphical depiction validates the theoretical framework outlined in Section 4, as the optimal control policy acts as an automatic stabilizer that offsets the destabilizing influences of currency shocks.

Sensitivity Analysis

To evaluate the robustness of the findings, we carry out a sensitivity analysis concerning three critical parameters: the senior tranche loss probability constraint ϵ, the hedging cost parameter γ, and the feedback strength parameter δ in Equation (15).

Table 3 presents the expected discounted losses and senior tranche loss probabilities for the dynamic policy under alternative parameter values, averaged across the three economies.

Table 3. Sensitivity Analysis of Dynamic Policy Performance

Parameter	Value	Expected Losses (USD M)	Senior Tranche Loss Prob. (%)	Avg. First-Loss (USD M)
ϵ	0.03	9.8	2.4	14.2
	0.05	11.0	3.5	11.3
	0.10	13.2	6.1	8.7
γ	0.5× baseline	10.2	3.2	11.0
	1.0× baseline	11.0	3.5	11.3
	2.0× baseline	12.4	4.1	11.8
δ	0.0	12.1	4.2	12.5
	0.5× baseline	11.5	3.8	11.9
	1.0× baseline	11.0	3.5	11.3

The findings indicate that the dynamic policy remains robust under plausible fluctuations in parameter values. Tightening the senior tranche loss constraint from $\epsilon = 0.05$ to $\epsilon = 0.03$ reduces the realized loss probability from 3.5% to 2.4%, but requires a 26% increase in the average first-loss tranche size. Relaxing the constraint to $\varepsilon = 0.10$ diminishes the average first-loss tranche by 23% yet elevates the realized loss probability to 6.1%, which is above the target. This trade-off highlights the importance of careful calibration of the constraint parameter to balance risk reduction and capital efficiency.

A doubling of the hedging cost parameter γ leads to a 13% increase in expected losses and a 0.6 percentage point rise in the senior tranche loss probability, which suggests that the motivation to engage in hedging diminishes as its expense grows. The dynamic policy adjusts by replacing a portion of hedging with first-loss protection, a shift that is corroborated by the slight growth in the average first-loss tranche size. This substitution effect illustrates the advantage of concurrently optimizing the two control variables: when one risk management instrument becomes costlier, the system intrinsically moves toward the other.

The feedback strength parameter δ has a moderate effect on performance. Setting the loan origination feedback loop to zero ($\delta = 0$) raises anticipated losses by 10% and the probability of loss in the senior tranche by 0.7 percentage points, thereby establishing that the feedback mechanism yields meaningful added stability. The dynamic policy still substantially surpasses the static benchmarks even without the feedback loop; this finding suggests that the core HJB-based control mechanism is the primary source of performance improvement.

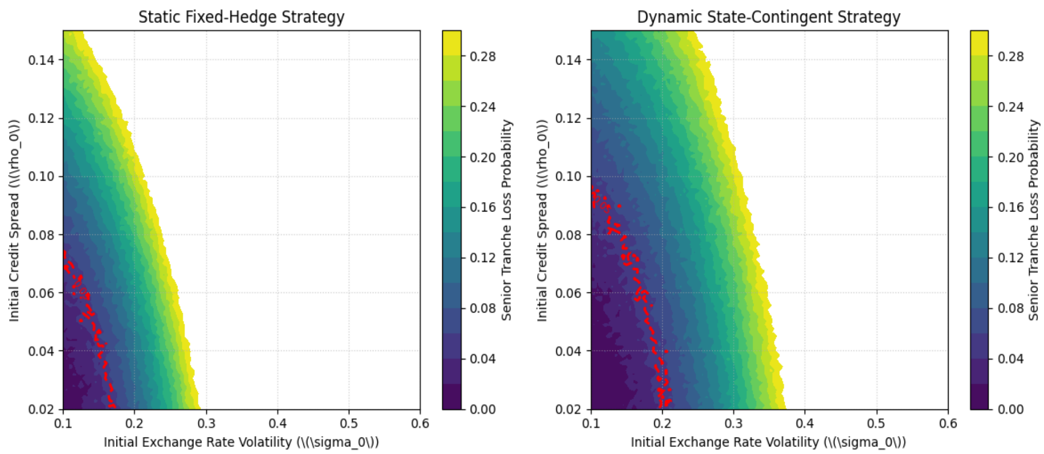

Figure 5. Comparative contour map of senior tranche loss probabilities under static versus dynamic capital allocation strategies across different initial market stress conditions.

Figure 5 presents a comparative contour map of senior tranche loss probabilities over the 5-year horizon under the static and dynamic strategies. The contours indicate a notably enlarged ‘safe region’, where the probability of senior tranche loss stays below 5%, under the dynamic policy. Although the safe region of the static strategy is restricted to low-volatility and low-spread initial conditions, the dynamic policy expands coverage into areas of moderate and even heightened stress. This enlargement of the safe zone is the concrete expression of the state-contingent mechanism’s capacity to adapt to changing circumstances, and it embodies the core value proposition of the proposed framework for private investors seeking exposure to fragile markets.

Discussion and Future Work

The empirical results presented in Section 6 indicate that the proposed dynamic, state-contingent blended finance architecture yields substantial improvements over conventional static structures across multiple performance dimensions. To convert these theoretical and simulation-based findings into operational reality, however, necessitates meticulous attention to a variety of practical difficulties, institutional limitations, and methodological constraints. This section discusses the key implications of our findings, identifies the boundary conditions of the proposed framework, and outlines promising directions for future research.

Market Liquidity Constraints and Counterparty Risk in Dynamic Hedging Execution

Although the optimal control policy prescribes continuous adjustments to the hedging ratio h_t^*, the practical feasibility of such adjustments depends critically on the depth and liquidity of local currency derivative markets in fragile African economies. The synthetic local currency swaps underlying the hedging mechanism necessitate a counterparty, typically a development finance institution or a commercial bank, which is prepared to take on the hard currency exposure. In times of acute currency stress, those counterparties might themselves experience liquidity limitations or risk constraints that hinder their capacity to broaden their swap portfolios, thus restricting the fund's ability to execute the optimal hedging strategy (Gregory, 2012).

Our simulation findings presuppose that the swap counterparty possesses unbounded capacity to meet the fund's hedging requirements at the existing market price. In reality, the hedging cost function in Equation (7) may become convex during stress episodes, as it indicates the rising marginal expense of accessing swap capacity when the counterparty's risk boundaries are neared. This convexity would alter the optimal control policy, possibly decreasing the hedging ratio precisely during periods of crisis, when its necessity is greatest. Future work should expand the model to include a state-dependent hedging cost function capturing the liquidity premium in swap markets, perhaps modeled as a function of the aggregate hedging demand across all market participants (Routledge & Zin, 2009).

Furthermore, the dynamic hedging mechanism introduces a type of counterparty risk absent in static structures. If the swap counterparty defaults or fails to honor its obligations during a crisis, the fund's hedged position would be unwound at precisely the worst possible time. This danger is especially pronounced in fragile economies, where the counterparty—often a local bank or a DFI subsidiary—may itself be susceptible to identical macroeconomic shocks that precipitate the currency crisis (Ballester & González-Urteaga, 2017). Reducing this hazard necessitates prudent counterparty selection, collateralization structures, and spreading exposure across multiple swap counterparties. The model could be extended to include a stochastic counterparty default intensity correlated with the exchange rate and credit spread, such that the optimal control policy internalizes the trade-off between hedging benefits and counterparty risk.

Robustness Analysis: Sensitivity to Model Misspecification and Data Integrity in Fragile Economies

The efficacy of the proposed architecture is critically dependent on the precision of the calibrated model parameters and the fidelity of the real-time state observations. Nonetheless, fragile African economies suffer from pronounced data constraints: exchange rate data may experience reporting delays or manipulation, credit spread data may be missing for extended periods owing to market closures, and loan portfolio data may be affected by reporting inaccuracies or delays (Ba et al., 2025). These data integrity issues introduce measurement error into the state observation module, which in turn degrades the quality of the optimal control recommendations.

Our sensitivity analysis in Section 6.3 examined variations in structural parameters, but it did not address the impact of measurement error in the state variables. A natural extension is to model the state observation process as a noisy signal of the true state, in line with the partially observed Markov decision processes (POMDPs) framework (Bäuerle & Rieder, 2011). In this framework, the fund observes a noisy proxy for each state variable, such as an exchange rate subject to delay or interpolation, and must infer the true state through a filtering algorithm like the Kalman filter or a particle filter. The optimal control policy would then be a function of the filtered state estimate rather than the true state, which introduces an additional layer of uncertainty that could reduce the performance advantage of the dynamic mechanism.

The model further presupposes that the stochastic processes driving the state variables are stationary and that the calibrated parameters stay constant across the investment horizon. This assumption is questionable in fragile economies, where structural breaks, such as currency regime changes, political transitions, or natural disasters, can basically alter the dynamics of exchange rates, volatility, and credit spreads (Ümit, 2016). The Zimbabwean experience, where the currency regime shifted from a managed float to a dollarized system and back to a floating regime during our sample period, illustrates this challenge. Future work should investigate regime-switching models (Bae et al., 2014) in which the SDE system's parameters shift in accordance with a latent Markov chain, and the optimal control policy adapts to the ongoing regime. This extension would strengthen the resilience of the framework to structural breaks and improve its applicability to economies with volatile institutional environments.

Institutional Implications: Moral Hazard, Governance of Algorithmic Controls, and Donor Coordination

Implementing a dynamic, algorithmically governed risk management system raises important questions about institutional design, governance, and the potential for moral hazard. The proposed architecture automates the deployment of concessional capital, thereby diminishing the influence of discretionary decisions by fund managers and DFI officials. Although this automation bolsters credibility with private investors by guaranteeing that the risk management framework is rule-based and transparent, it simultaneously introduces novel governance challenges (Lenglet, 2011).

The quality of the algorithm's output is determined by the quality of the model on which it is based. If the model is misspecified or the calibrated parameters become outdated, the algorithm may make suboptimal decisions with the potential to harm both senior and first-loss tranche investors. This model risk is especially pronounced in vulnerable economies, where the historical data employed for calibration may not reflect future conditions. Establishing a governance framework for model validation, periodic recalibration, and override procedures is essential for the responsible deployment of the proposed architecture. This framework should include independent model audits, stress testing against out-of-sample scenarios, and a clear escalation protocol for exceptional circumstances where the algorithm's recommendations are overridden by human judgment.

Second, the dynamic expansion of the first-loss tranche during stress episodes creates a potential moral hazard for the loan origination platform. If the origination team recognizes that the first-loss buffer will be automatically expanded during crises, their motivation to uphold prudent underwriting standards during normal periods may diminish, as they expect that any resulting losses will be covered by the expanded concessional capital (Crawford, 2015). The loan origination feedback loop described in Section 4.4 is designed to counteract this moral hazard by tightening underwriting standards precisely when the first-loss buffer is expanded. However, the effectiveness of this feedback loop depends on the strength of the parameter δ and the credibility of the commitment to tighten standards. Should the origination team anticipate that political pressure during a crisis will reverse the tightening, the feedback loop may fail to constrain their conduct. Subsequent investigations ought to examine the creation of incentive-aligned contracts that reconcile the interests of the origination team with the fund's enduring well-being, potentially by tying remuneration to the performance of the senior tranche or to the precision of the origination team's default forecasts.

Third, the proposed architecture requires close coordination between multiple stakeholders: the concessional capital provider (typically a DFI or bilateral donor), the swap counterparty, the senior tranche investors, and the fund manager. Each stakeholder possesses distinct goals, varying levels of risk acceptance, and differing temporal perspectives, which can generate tensions during the execution of the dynamic mechanism. For instance, the swap counterparty may be reluctant to increase its exposure during a crisis precisely at the moment the algorithm dictates greater hedging, while the concessional capital provider may face political constraints on the amount of capital it can deploy in a single country or crisis episode (Nahari, 2026). To resolve these coordination difficulties, one must formulate pre-negotiated accords that delineate the duties of each stakeholder under diverse stress conditions, coupled with unambiguous mechanisms for settling disputes. The algorithmic nature of the proposed framework can aid this coordination by offering a clear, rule-driven foundation for decision-making, which all stakeholders can observe and verify.

Extensions to Multi-Country and Multi-Currency Settings

The current framework is designed for a single-country, single-currency fund. Nevertheless, many blended finance vehicles function across multiple nations and currencies, thereby distributing risk among a diversified portfolio of MSME loans (Dembele et al., 2022). Extending the proposed architecture to a multi-country context introduces additional complexities. Initially, the state space is enlarged to encompass exchange rates, volatilities, and credit spreads for each nation, thereby generating a high-dimensional control problem that may exceed the neural network approximation's capacity. Second, the correlations between the state variables across countries become important, as a currency crisis in one country may spill over to others through trade linkages, investor sentiment, or common shocks (Fratzscher, 2003). Third, the optimal allocation of concessional capital across countries becomes a portfolio optimization problem, where the fund must decide how to distribute its limited first-loss buffer across multiple country exposures.

A promising approach to this multi-country extension is to employ a hierarchical control architecture, where a high-level portfolio optimizer allocates concessional capital across countries, and a lower-level country-specific HJB solver determines the optimal hedging and first-loss allocation within each country. The high-level optimizer would take as inputs the country-specific value functions from the lower-level solvers and would allocate capital to maximize the overall risk-adjusted return of the fund, subject to a global senior tranche loss constraint. This hierarchical approach exploits the computational efficiency of neural network approximation while retaining tractability in the high-dimensional state space.

Integration with Climate Risk and Natural Disaster Shocks

Fragile African economies are disproportionately exposed to climate-related risks, such as droughts, floods, and cyclones, which can simultaneously depress MSME revenues, depreciate the local currency, and increase credit spreads (Ali et al., 2025). These concurrent disruptions impose a serious difficulty on the proposed framework, since the model's default intensity function in Equation (5) does not explicitly include climate variables. Extending the model to include a climate risk factor, perhaps modeled as a Poisson jump process that captures the arrival of natural disasters, would strengthen the framework's capacity to anticipate and address these compound shocks.

Factoring climate risk into the equation would further permit the creation of 'climate-responsive' blended finance structures, in which the first-loss tranche is automatically widened either preemptively or in reaction to climate-related disruptions. This development corresponds with the increasing attention to 'climate-resilient debt clauses', which permit the suspension of debt service after natural disasters (Park & Samples, 2024). Incorporating such clauses into the proposed dynamic framework would establish a comprehensive risk management system addressing both financial and climate-related sources of volatility, thereby improving the sustainability of MSME lending in the world's most vulnerable economies.

Conclusion

This paper has introduced a novel blended finance architecture that substitutes static, predetermined risk-sharing parameters with a dynamic, state-contingent mechanism for local currency hedging and first-loss capital allocation. The central contribution is a stochastic optimal control framework, derived from a Hamilton-Jacobi-Bellman equation, which jointly optimizes in real time both the size of the concessional first-loss tranche and the hedging ratio of a synthetic local currency swap. The system responds to four key state variables—the loan portfolio value, the spot exchange rate, its stochastic volatility, and the domestic credit spread—which are subject to coupled stochastic differential equations calibrated to historical data from three fragile African economies.

The empirical results show that the dynamic policy consistently outperforms conventional static structures across multiple performance dimensions.

Expected discounted default losses are reduced by 42% to 49% across the three economies, while the probability of senior tranche loss falls below the 5% constraint. Crucially, these risk reductions are attained while deploying 25% to 28% less concessional capital on average, a pattern that underscores the capacity to concentrate protection during stress episodes. Under extreme stress scenarios, the dynamic mechanism limits senior tranche loss probabilities to 5.8% to 7.1%, compared to 18.2% to 28.5% for static alternatives. The nonlinear threshold effects in the optimal policy, the stabilizing feedback loop via loan origination, and the expanded safe region across the state space jointly attest to the worth of a mathematically rigorous, data-driven method for blending concessional and private capital under conditions of high volatility.

The proposed framework delineates a pragmatic approach for deploying development capital more efficiently within the globe's most arduous investment contexts. Through the substitution of discretionary intervention with a transparent, rule-based algorithm, the architecture strengthens credibility with private investors and reduces the systemic risk of currency crises that has historically constrained MSME lending in fragile economies. By coupling stochastic optimal control theory with deep neural network approximation, a computationally tractable method emerges for real-time decision-making in a high-dimensional state space, thereby advancing beyond the heuristic or static approaches that currently dominate blended finance practice. Notwithstanding persistent difficulties concerning market liquidity, data quality, institutional coordination, and model robustness, the findings indicate that dynamic, state-contingent mechanisms are a promising avenue for the subsequent generation of blended finance structures.

References

Alfaro, L., Calani, M., & Varela, L. (2021). *Currency hedging in emerging markets: Managing cash flow exposure*. academia.edu.

Ali, A., Usman, M., & Ahmad, K. (2025). *Environmental risks and sovereign credit ratings: Evidence from developed and developing economies*. mpra.ub.uni-muenchen.de.

Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*.

Ba, A., Muga, T., Okwarah, P., & Ali, M. (2025). Remote data verification under fragility and operational stress: Insights from somalia during COVID-19. *Soc. Sci*.

Bae, G., Kim, W., & Mulvey, J. (2014). Dynamic asset allocation for varied financial markets under regime switching framework. *European Journal of Operational Research*.

Ballester, L., & González-Urteaga, A. (2017). How credit ratings affect sovereign credit risk: Cross-border evidence in latin american emerging markets. *Emerging Markets Review*.

Bäuerle, N., & Rieder, U. (2011). *Markov decision processes with applications to finance*. books.google.com.

Bertsekas, D. (2012). *Dynamic programming and optimal control: Volume i*. books.google.com.

Cetin, U., Jarrow, R., Protter, P., & Yıldırım, Y. (2004). *Modeling credit risk with partial information*. projecteuclid.org.

Chen, N., & Zhou, H. (2025). Managing emerging market currency risk. *Available at SSRN 5119836*.

Crawford, J. (2015). The moral hazard paradox of financial safety nets. *Cornell Journal of Law and Public Policy*.

Dembele, F., Randall, T., Vilalta, D., & Bangun, V. (2022). *Blended finance funds and facilities: 2020 survey results*. read.oecd-ilibrary.org.

Flammer, C., Giroux, T., & Heal, G. (2026). Blended finance. *Journal of Financial Economics*.

Fratzscher, M. (2003). On currency crises and contagion. *International Journal of Finance & Economics*.

Fund, U. N. C. D. (2020). *Blended finance in the* *least developed countries 2020 supporting a resilient COVID-19 recovery: Supporting a resilient COVID-19 recovery*. books.google.com.

Gabriel, V., Levine, P., & Yang, B. (2016). An estimated DSGE open economy model of the indian economy with financial frictions. *Monetary Policy in India*.

Garnier, J., & Sølna, K. (2020). Optimal hedging under fast-varying stochastic volatility. *SIAM Journal on Financial Mathematics*.

Gregory, J. (2012). *Counterparty credit risk and credit value adjustment: A continuing challenge for global financial markets*. books.google.com.

Horrocks, P., Marshall, C., Thomas, C., Venon, T., et al. (2025). *Unlocking local currency financing in emerging markets and developing economies*. oecd.org.

Lenglet, M. (2011). Conflicting codes and codings: How algorithmic trading is reshaping financial regulation. *Theory, Culture & Society*.

Marozva, D., Marozva, S., & Gherghina, Ş. (2026). Dynamic hedging under stochastic volatility and model uncertainty: PDE characterization and regime-based evidence. *Mathematics*.

Mikhailov, S., & Nögel, U. (2004). *Heston's stochastic volatility model: Implementation, calibration and some extensions*. books.google.com.

Muritala, T. (2025). UNLOCKING THE POTENTIAL OF BLENDED FINANCE MODELS FOR SME GROWTH IN SUB SAHARAN AFRICA. *REVENUE: Jurnal Manajemen Bisnis Islam*.

Nahari, S. A. (2026). *How coordination mechanisms shape blended-finance design and implementation to mobilize private investment for sustainable development in egypt*. fount.aucegypt.edu.

Otunba, M., Adenuga, A., Sikiru, A., & Gaffar, O. (2025). *Designing capital-raising frameworks for agro enterprises in underserved and volatile economies*. researchgate.net.

Park, S., & Samples, T. (2024). The sovereign climate debt trap and natural disaster clauses. *American Business Law Journal*.

Prigent, J., Renault, O., & Scaillet, O. (2000). *An empirical investigation in credit spread indices*. researchgate.net.

Ramachandran, P., Zoph, B., & Le, Q. (2017). *Searching for activation functions*. arXiv preprint arXiv:1710.05941.

Routledge, B., & Zin, S. (2009). Model uncertainty and liquidity. *Review of Economic Dynamics*.

Ruiz, E. (1994). Quasi-maximum likelihood estimation of stochastic volatility models. *Journal of Econometrics*.

SHOKUNBI, T. (2024). *Bridging the finance gap: A policy framework for SME credit expansion in emerging markets*. researchgate.net.

Staff, S. B. (2009). International financial statistics. *International Monetary Fund*.

Taskforce, B. (2018). Better finance, better world. *Consultation Paper*.

Thalassinos, E., Stamatopoulos, T., et al. (2016). The european sovereign debt crisis and the role of credit swaps. *The World Scientific Handbook of Futures Markets*.

Ümit, Ö. (2016). Stationarity of real exchange rates in the "fragile five": Analysis with structural breaks. *International Journal of Economics and Finance*.

Venkata, N. D. et al. (2023). Deep recurrent optimal stopping. *Advances in Neural Information Processing Systems 36*.

Zhang, Z. (2018). Improved adam optimizer for deep neural networks. *2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS)*.