An Integrated Simulation Framework for Decentralised Finance Protocols: cadCAD, Reinforcement Learning, Game Theory, and Mechanism Design

Shehzad Ahmed

Arcus Quant Fund
https://arcusquantfund.com

March 2026

Abstract

We present the Integrated Economic System (IES), the first four-layer simulation framework for decentralised finance (DeFi) protocols, combining cadCAD state-machine modelling, reinforcement learning (RL) traders, game-theoretic Nash equilibrium analysis, and mechanism design optimisation into a single closed-loop system. Each layer captures a distinct economic phenomenon that no single method can address alone: cadCAD provides the market-level ground truth; RL agents model rational individual traders with optimal stopping intuition; game theory characterises strategic leverage selection; and mechanism design searches the protocol parameter space for the Pareto-optimal regime that simultaneously minimises insolvency risk, liquidation frequency, and funding-rate deviation from zero.

The framework applies to any DeFi protocol whose smart contract logic can be expressed as a cadCAD partial-state-update block. As a case study, we apply the IES to a riba-free-by-design ($\iota = 0$) perpetual futures exchange (the Baraka Protocol, deployed on the Arbitrum Sepolia testnet, pre-audit), producing four primary simulation results over 5 episodes $\times$ 720 steps (150 simulated trading days): (i) zero insolvency events across all Monte Carlo paths; (ii) *mean* Nash-equilibrium leverage of 2.72$\times$ (long) and 3.28$\times$ (short) across intervals (headline seed; 3.16$\times$/2.84$\times$ over $R = 30$ seeds, Robustness 1), well below the 4.97$\times$ longs are driven to in the $\iota > 0$ counterfactual; (iii) a small positive mean net transfer (\$865 per \$100,000 notional, i.e. 0.87%; 95% CI \$245–\$1,485, over $R = 30$ seeds — statistically distinguishable from zero at finite horizon, attributable to the discretisation drift analysed in §8, and an order of magnitude below the interest-driven transfer; the individual intervals are large and opposite-signed) versus a systematic \$8,420 paid by longs under $\iota > 0$; and (iv) endogenous mechanism design convergence to a tighter circuit breaker (75 $\rightarrow$ 41 bps), higher maintenance margin (2.0 $\rightarrow$ 3.1%), and greater insurance split (50 $\rightarrow$ 61%), consistent with the κ-rate parameterisation implied by Ahmed and Bhuyan [5]. We extend these with three robustness analyses — Monte Carlo confidence intervals over $R = 30$ seeds, a Sobol sensitivity decomposition of the mechanism-design objective, and a heterogeneous five-type trader population — and report an honest negative result on reinforcement learning (a trained PPO agent underperforms the rule-based baseline). We release the full simulation codebase as open-source.

Keywords: cadCAD, reinforcement learning, game theory, mechanism design, Islamic finance, DeFi, agent-based simulation, perpetual futures, Shariah compliance, $\iota = 0$, computational economics

1 Introduction

The design and analysis of decentralised finance (DeFi) protocols sits at the intersection of mechanism design, game theory, and stochastic control, yet is rarely treated as such. Most DeFi research either (a) studies individual agent behaviour in isolation [9, 20], (b) applies equilibrium theory to single-parameter models [16], or (c) runs agent-based Monte Carlo simulations without equilibrium analysis [6]. No existing work couples all four methods into a single feedback system whose state-update blocks mirror the logic of a deployed (testnet) protocol implementation.

The problem is particularly acute in Islamic DeFi, where protocol correctness carries a dual obligation: economic soundness and jurisprudential compliance. A protocol that is economically stable but fails the $\iota = 0$ riba test is invalid; a protocol that is Shariah-compliant but mechanistically insolvent is undeployable. Verifying both simultaneously requires a simulation infrastructure that is richer than any single method provides.

This paper presents such an infrastructure: the **Inte-**

grated Economic System (IES) simulation framework. The IES couples:

1. A **cadCAD state machine** that implements the exact on-chain protocol logic, serving as the simulation ground truth.
2. A **reinforcement learning (RL) agent** embedded inside the cadCAD market, whose actions are drawn from an observation of live protocol state and fed back as on-chain positions.
3. A **game theory layer** that, every N steps, extracts the recent price window, builds a 5×5 payoff matrix over the leverage space, and computes the Nash equilibrium leverage distribution — which is then returned to the cadCAD simulation as the background-trader respawn policy.
4. A **mechanism design optimiser** that, at each episode boundary, receives aggregate cadCAD metrics, minimises a multi-objective loss over the protocol parameter space, and blends the result into the next episode's parameters.

The key conceptual contribution is the *coupling architecture*: a directed graph of information flows connecting the four layers so that each layer both consumes outputs from the others and produces inputs for them. This bidirectional coupling distinguishes the IES from sequential multi-method analysis: the layers are not independent. The Nash equilibrium changes the cadCAD trader policy, which changes the price path, which changes the payoff matrices that determine the next Nash equilibrium.

Application. We apply the IES to a riba-free-by-design ($\iota = 0$) perpetual futures exchange [4] deployed on the Arbitrum Sepolia testnet in February 2026 (pre-audit), built around the $\iota = 0$ no-arbitrage condition of Ackerer, Hugonnier & Jermann [1] — simultaneously the source of its riba-freedom at the pricing layer and the theoretical object we seek to validate. (The $\iota = 0$ design addresses the riba prohibition at the pricing layer; gharar, maysir, and qabd are product-layer questions reserved for Shariah scholars.)

Structure. Section 2 reviews each simulation method and its limitations when used alone. Section 3 presents the IES coupling architecture in full. Sections 4–7 detail each layer. Section 8 presents the four primary experimental results, three robustness analyses, and an honest negative reinforcement-learning result. Section 10 discusses implications for DeFi mechanism design. Section 12 concludes.

2 Background and Related Work

2.1 cadCAD: Complex Adaptive System Modelling

cadCAD (complex adaptive dynamics/computer-aided design) is a Python framework for specifying economic state machines as a set of *partial-state-update blocks* [26, 8]. Each block consists of *policy functions* (which compute signals from current state) and *state update functions* (which apply signals to produce the next state). The framework natively supports Monte Carlo over independent random number streams and parameter sweeps.

cadCAD is well-suited to protocol analysis because it forces the modeller to specify every state variable and every transition function explicitly, producing a simulation that mirrors smart contract logic. However, cadCAD lacks:

- **Rational agent learning.** Background trader policies in cadCAD are typically hard-coded rules, not learned strategies.
- **Strategic interaction.** cadCAD does not model multi-agent games or equilibrium selection.
- **Endogenous parameter adaptation.** Protocol parameters are fixed within a simulation run.

2.2 Reinforcement Learning for DeFi

Reinforcement learning (RL) has been applied to DeFi trading [25], extractable-value and liquidation analysis [21], and AMM liquidity provision [13]. The key advantage is that RL agents *learn* optimal policies from experience rather than requiring closed-form strategy specifications.

Stable-Baselines3 [22] provides the Proximal Policy Optimisation (PPO) implementation used in this work. The RL environment (described in Section 5) is a gymnasium-compatible wrapper around the cadCAD market state.

The limitation of RL alone is that it models one rational agent against a fixed background. It does not characterise the multi-agent equilibrium that emerges when all traders adapt simultaneously.

2.3 Game Theory in DeFi

Nash equilibrium analysis has been applied to AMM fee structures [16], MEV extraction [9], and liquidation dynamics [19]. For perpetual futures, the relevant game is between long and short players choosing leverage, with the funding rate as the strategic variable.

The pricing foundation is the Ackerer et al. [1] result that $\iota = 0$ achieves unique no-arbitrage perpetual-futures pricing (their Theorem 3), under which the funding rate carries no interest component. AHJ do not, however, frame $\iota = 0$ as a Nash equilibrium — their Theorem 3 is a no-arbitrage pricing result, not a game-theoretic one — so the equilibrium analysis here is our own. Our game theory layer studies the *leverage* game among traders op-

erating under the $\iota = 0$ funding rule, computing the Nash equilibrium numerically over live cadCAD price paths using the nashpy vertex enumeration algorithm [15].

2.4 Mechanism Design for DeFi Protocols

Mechanism design [14, 18] studies how to construct the rules of a game so that self-interested agents produce socially desirable outcomes. Applied to DeFi, the "mechanism" is the protocol parameter vector $\theta = (F_{\max}, m, p, s)$ (circuit breaker, maintenance margin, liquidation penalty, insurance split), and the "social objective" is the multi-objective loss $\mathcal{L}(\theta)$ of equation (14) in Section 7.

Prior work in DeFi mechanism design includes AMM curve optimisation [7], liquidation parameter tuning [12], and governance token design [11]. We extend this to the Islamic DeFi context, where the mechanism must satisfy a Shariah constraint (funding-rate symmetry, $\iota = 0$) in addition to economic objectives.

3 The Integrated Economic System Architecture

3.1 Design Principles

The IES is built on three design principles:

1. **One ground truth.** The cadCAD state machine is the single source of market prices, positions, funding rates, and insurance-fund balances. All other layers read from and write to this state. There is no separate "market" in the RL environment or game theory module; both operate on the live cadCAD state.
2. **Asynchronous coupling.** Layers run at different time scales. The RL agent acts at every cadCAD step (hourly). The game theory layer computes a new Nash equilibrium every $\Delta_{\text{GT}} = 50$ steps. The mechanism design optimiser runs once per episode (720 steps). This mirrors the natural time scales of individual traders, strategic leverage rebalancing, and protocol governance updates.
3. **Closed-loop feedback.** Every layer both consumes outputs from and produces inputs for the cadCAD engine. The coupling graph is:

$$\text{cadCAD} \xrightarrow{s_t} \text{RL} \xrightarrow{a_t} \text{cadCAD}$$
$$\text{cadCAD} \xrightarrow{\{p_s\}_{t-W}^{t}} \text{GT} \xrightarrow{(\ell_L^*, \ell_S^*)} \text{cadCAD}$$
$$\text{cadCAD} \xrightarrow{\mu_T} \text{MD} \xrightarrow{\theta^*} \text{cadCAD}$$

where s_t is state observation, a_t is action, $\{p_s\}$ is the price window of length W, (ℓ_L^*, ℓ_S^*) are Nash equilibrium leverages, μ_T are episode aggregate metrics, and θ^* is the updated parameter vector.

3.2 Episode Structure

Each simulation run consists of E episodes, each of T steps. Within an episode:

1. cadCAD initialises market state (price, positions, insurance fund) with current parameter vector θ_e.
2. At each step t: market step, RL action, GT recomputation (if $t \bmod \Delta_{\text{GT}} = 0$).
3. At episode end: MD optimiser produces θ_{e+1}.

The episode boundary is the Islamic governance analogue: a Shariah board review that can adjust protocol parameters based on observed performance, subject to the $\iota = 0$ and $\ell_{\max} = 5$ hard constraints.

4 Layer 1: cadCAD Market Engine

Implementation note. The Layer-1 market is specified as five cadCAD partial-state-update blocks (Section 4, Partial State Update Blocks). We run these blocks in two ways. The standalone Layer-1 Monte Carlo is executed through the actual cadCAD library (cadCAD 0.5.3) using its `Executor` with sequential substep semantics; a 20-run × 720-step ensemble yields zero insolvency in 20/20 runs and a funding rate confined to the symmetric ±75 bps band ($\iota = 0$), consistent with the integrated results reported below. The *integrated* four-layer runs (which produce Table 1) instead use a pure-Python implementation of the *same* partial-state-update blocks, executed in a single process so that the reinforcement-learning, game-theory, and mechanism-design layers can be coupled tightly within each step. Both share the identical cadCAD state space and block structure.

4.1 State Space

The cadCAD state $\mathcal{S}$ consists of:

$$\mathcal{S} = \{p_t^m, p_t^i, F_t, \text{CF}_t, \mathcal{P}_t, I_t, L_t, v_t\} \tag{1}$$

where p_t^m and p_t^i are mark and index prices, F_t is the instantaneous funding rate, CF_t is the cumulative funding index, $\mathcal{P}_t$ is the pool of open positions, I_t is the insurance fund balance, L_t is the liquidation count, and $v_t \in \{0, 1\}$ is the solvency flag.

4.2 Partial State Update Blocks

The state transition is implemented as five cadCAD partial-state-update blocks:

Block 1 (Oracle). The index price follows geometric Brownian motion:

$$p_{t+1}^i = p_t^i \exp\left(\sigma_i \, \epsilon_t^i\right), \quad \epsilon_t^i \sim \mathcal{N}(0, 1) \tag{2}$$

Note this discretisation omits the Itô compensator and is therefore not a martingale: $\mathbb{E}[p_{t+1}^i \mid p_t^i] = p_t^i e^{\sigma_i^2/2}$, an upward index drift of ≈ 2 bps per hourly step at $\sigma_i = 0.02$.

This drift contributes, together with the mark-process premium drift, to the small positive finite-horizon transfer in Robustness 1, and inflates the absolute PnL of long-biased policies (including the random baseline of Section 11); the $\iota = 0$ versus $\iota > 0$ *contrast*, which is the object of interest, is unaffected since both arms share the same oracle process. The mark price mean-reverts to index:

$$p^m_{t+1} = 0.9\, p^m_t \, \exp(\sigma_m \, \epsilon^m_t) + 0.1\, p^i_{t+1} \tag{3}$$

with $\sigma_m = 0.7\, \sigma_i$ (reduced volatility, market maker dampening).

Block 2 (Funding). The funding rate at $\iota = 0$:

$$F_t = \text{clip}\left(\frac{p^m_t - p^i_t}{p^i_t},\ -F_{\max},\ +F_{\max} \right) \tag{4}$$

$F_{\max}$ is the protocol circuit breaker (initially 75 bps, adaptive via MD layer). The clip$(\cdot)$ bound is exactly the funding-rate *clamp* whose no-arbitrage perpetual pricing under stochastic rates is formalised by Hugonnier & Jermann [2]; our $\iota = 0$ funding rule is the $I = 0$ case of that clamped specification.

Block 3 (Traders). Background positions are accrued: $c_{i,t+1} = c_{i,t} - \text{sign}(\ell_i)\, \Delta\text{CF}_t \, n_i$ where c_i is collateral, n_i is position size, and $\Delta\text{CF}_t = F_t$. The RL agent's position is managed by Layer 2.

Block 4 (Liquidations). Position i is liquidated if $c_i < n_i \cdot m$ (maintenance margin fraction m):

$$\Delta I_t = \sum_{i \text{ liq.}} n_i \, p \, s \tag{5}$$

$$c_i \leftarrow 0, \quad n_i \leftarrow n_{\text{respawn}} \tag{6}$$

where p is the liquidation penalty and s is the insurance split.

Block 5 (Bookkeeping). Solvency flag: $v_t = \mathbf{1}[I_t < 0]$.

4.3 Respawn Policy with Nash Leverage

When a position is liquidated (Block 4), it is immediately respawned with a new collateral of \$5,000 and leverage drawn from the current Nash equilibrium distribution $\mathcal{D}_{\text{Nash}}(\ell^*_L, \ell^*_S)$ supplied by the game theory layer. This coupling ensures the background trader population adapts its strategic behaviour in response to observed market conditions.

5 Layer 2: Reinforcement Learning Trader

5.1 RL Environment

The RL environment (`IESTraderEnv`) is a `gymnasium.Env` wrapping the cadCAD state. The environment is injected into the cadCAD simulation loop so that the agent's actions directly modify the cadCAD position pool at each step.

Observation space. $\mathcal{O} \subset \mathbb{R}^8$:

$$s_t = \big(p^m_t/p_0,\ \ p^i_t/p_0,\ \ F_t,\ \ \text{CF}_t, \\ q_t/p_0,\ \ c_t/p_0,\ \ \mathbf{1}[\text{long}],\ \ w_t/W_0 \big) \tag{7}$$

where q_t and c_t are position size and collateral, w_t is free collateral, $p_0 = \$60{,}000$, $W_0 = \$100{,}000$.

Action space. $\mathcal{A} = \{0, 1, \ldots, 11\}$:

- $a = 0$: hold
- $a \in [1, 5]$: open long at leverage $a\times$
- $a \in [6, 10]$: open short at leverage $(a - 5)\times$
- $a = 11$: close current position

The leverage range $[1, 5]$ is constrained to the Shariah hard cap $\ell_{\max} = 5$ (ShariahGuard on-chain).

Reward function. The reward at each step is the sum of funding PnL (shaped reward for survival) and realised price PnL on position close. Liquidation incurs an additional negative reward of $n \cdot p$ (size times penalty).

Rule-based baseline. In the absence of a trained PPO model, the agent follows a deterministic rule:

- If $F_t < -0.001$: open long $3\times$ (undervalued)
- If $F_t > +0.002$: open short $3\times$ (overvalued)
- If PnL $> +2\%$ or $< -3\%$ of size: close
- Otherwise: hold

This rule operationalises the optimal stopping-time intuition: enter when the funding rate signals a persistent basis, exit when the convergence event has occurred or the loss threshold is breached.

5.2 Connection to Optimal Stopping

The RL environment is a discrete-time approximation to an optimal stopping problem over the $\iota = 0$ funding process: entry/exit timing under the no-arbitrage pricing of [1], in the stopping-time formulation of [5]. The agent's "open" and "close" decisions correspond to choosing entry time t_0 and stopping time τ to maximise:

$$V^*(s_t) = \sup_{\tau} \mathbb{E}\Big[\phi(X_\tau) - c(t_0, \tau) \,\Big|\, s_t \Big] \tag{8}$$

where $\phi(X_\tau)$ is the payoff at close and $c(t_0, \tau)$ is the cumulative funding cost from open to close. At $\iota = 0$, the funding cost is symmetric and the optimal policy reduces to the rule described above: enter when basis is predictably mean-reverting, exit at convergence.

6 Layer 3: Game Theory Equilibrium Analysis

6.1 The Funding Rate Game

The strategic interaction between long and short players is a continuous-time game with:

- **Players.** Representative long L and short S (population games).
- **Strategy space.** Leverage choice $\ell \in \{1, 2, 3, 4, 5\}$.
- **State.** Price path $\{p_t\}_{t=0}^{T}$.
- **Payoffs.** $u_L(\ell_L, \ell_S) = \mathrm{NAV}_L(\ell_L, \ell_S) - c$, $u_S(\ell_L, \ell_S) = \mathrm{NAV}_S(\ell_L, \ell_S) - c$ where $c = \$10{,}000$ is the initial collateral.

The NAV functions are computed analytically:

$$\mathrm{NAV}_L(T) = c + n_L \int_0^T \frac{dp_s}{p_s} - n_L \int_0^T F_s \, ds \tag{9}$$

$$\mathrm{NAV}_S(T) = c - n_S \int_0^T \frac{dp_s}{p_s} + n_S \int_0^T F_s \, ds \tag{10}$$

with $n_L = c \cdot \ell_L$, $n_S = c \cdot \ell_S$, and F_s evaluated at $\iota = 0$.

6.2 Nash Equilibrium at $\iota = 0$

Proposition 1 (Transfer structure at $\iota = 0$ versus $\iota > 0$)**.** *Each side's payoff is monotone in its* own *leverage and independent of the other's, so the per-interval equilibrium is in dominant strategies (a degenerate Nash equilibrium with no strategic interaction): each side plays a leverage corner ($\ell = 5$ or $\ell = 1$) according to the sign of the realised window return net of funding, and the cross-interval* mean *leverage is strictly below $\ell_{\max} = 5$. For the net wealth transfer $u_L - u_S$:*

1. *At $\iota = 0$ the funding rate has no fixed component, so any systematic transfer must come from the premium. Under a symmetric, zero-mean premium $\mathbb{E}[u_L - u_S] = 0$; empirically the transfer is* near *zero ($865 per $100,000, $\approx 0.9\%$; Robustness 1), not identically zero at finite horizon.*
2. *At $\iota > 0$ the funding floor $F_t \geq \iota > 0$ forces a systematic transfer paid by longs to shorts regardless of price realisation — $8,420 (8.4% of notional) in our calibration — constituting riba.*

Proof sketch. Because u_L depends only on ℓ_L (and u_S only on ℓ_S), each side has a dominant action — the leverage corner that signs the common window return net of funding — so the equilibrium is degenerate rather than strategic. The transfer is proportional to $\int_0^T (dp_s/p_s - F_s\, ds)$. At $\iota = 0$, $F_t = \mathrm{clip}((p_t^m - p_t^i)/p_t^i, \pm F_{\max})$ carries no fixed term, so the transfer is the premium integral, which is mean-zero only if the premium is symmetric about zero. The mark process used here, $p_{t+1}^m = 0.9\, p_t^m e^{\sigma\varepsilon} + 0.1\, p_{t+1}^i$, has a small positive premium drift (multiplicative log-normal), and the index discretisation of eq. (2) adds its own uncompensated drift, so the expected transfer is near but *not* identically zero — consistent with the strictly positive R=30 confidence interval of Robustness 1. With $\iota > 0$, $F_t \geq \iota > 0$ for all t, so the funding component of the transfer is one-signed and the expected transfer is shifted by ιT regardless of price realisation. □ □

6.3 Numerical Implementation

Every $\Delta_{\mathrm{GT}} = 50$ cadCAD steps, the game theory layer:

1. Extracts the price window $\{p_s\}_{t-W}^{t}$ (W = 100 steps) from cadCAD state history.
2. Evaluates all $5 \times 5 = 25$ payoff pairs $(u_L(\ell_L, \ell_S), u_S(\ell_L, \ell_S))$ over the price window.
3. Solves for Nash equilibria using nashpy vertex enumeration [15].
4. Returns (ℓ_L^*, ℓ_S^*) to cadCAD as the respawn leverage distribution.

In the event that nashpy returns no pure Nash equilibrium, the expected leverage under the returned mixed strategy is computed as $\bar{\ell} = \sigma^\top \boldsymbol{\ell}$ where $\boldsymbol{\ell} = (\ell_1, \ldots, \ell_5)^\top$ and σ is the mixed strategy. Under the payoff-independence structure of Proposition 1 these cases arise only when the window return net of funding is near zero, so corner payoffs are nearly tied and vertex enumeration returns mixed/degenerate equilibria — they are tie-breaking artifacts of the dominant-strategy structure, not strategic mixing.

6.4 Riba Detection via Net Transfer

The net transfer $\Delta_{\mathrm{NT}} = \overline{u_L} - \overline{u_S}$ averaged over all payoff matrix entries serves as a computational riba test:

$$\text{is\_riba} = |\Delta_{\mathrm{NT}}| > \delta \tag{11}$$

where $\delta = 1\%$ of notional ($1,000 per $100,000) is a materiality threshold. Empirically $|\Delta_{\mathrm{NT}}| \approx 0.7$–$0.9\%$ of notional at $\iota = 0$ (below the threshold, and attributable to the finite-horizon premium and discretisation drift per Proposition 1), whereas the $\iota > 0$ counterfactual produces 8.4% — an order of magnitude above it. The test detects *systematic* transfer, not exact zero.

7 Layer 4: Mechanism Design Optimiser

7.1 Protocol Parameter Space

The mechanism design layer optimises the protocol parameter vector:

$$\theta = (F_{\max},\ m,\ p,\ s),$$
$$\theta \in [0.003, 0.015] \times [0.01, 0.05] \times [0.005, 0.03] \times [0.3, 0.8] \tag{12}$$

with two hard constraints:

$$\ell_{\max} = 5 \text{ (Shariah cap)}, \quad \iota = 0 \text{ (Shariah principle)} \tag{13}$$

These constraints reflect the Islamic finance principle that Shariah rules are non-negotiable: the protocol designer cannot trade off leverage or riba against economic efficiency.

7.2 Multi-Objective Loss Function

The objective function at episode e is:

$$\mathcal{L}_e(\theta) = 10\,\rho_e(\theta) + 5 \times 10^{-3}\,\bar{N}_e^{\text{liq}}(\theta) + 5\,\overline{|F|}_e(\theta) \quad (14)$$

where:

- $\rho_e(\theta) \in \{0, 1\}$: insolvency indicator (catastrophic; weight 10).
- $\bar{N}_e^{\text{liq}}(\theta)$: mean liquidation count (normalised by position count; weight 5×10^{-3}).
- $\overline{|F|}_e(\theta)$: mean absolute funding rate (Shariah objective: wants near zero; weight 5).

The objective encodes the Islamic finance priority ordering: solvency first, then trader fairness, then Shariah compliance. The high weight on $\overline{|F|}$ reflects the jurisprudential significance of the funding-rate term: it is a proxy for the interest component that $\iota = 0$ eliminates at the contract level.

7.3 Optimisation Algorithm

The optimiser uses `scipy.optimize.differential_evolution` [23]:

Algorithm 1: Episode-level Mechanism Design Update

1. **Input:** episode e price path $\{p_t^{(e)}\}$, current params θ_e
2. Run DE: minimise $\mathcal{L}_e(\theta)$ over price path $\{p_t^{(e)}\}$
3. $\theta_e^* \leftarrow \arg\min_\theta \mathcal{L}_e(\theta)$
4. $\theta_{e+1} \leftarrow 0.70\,\theta_e + 0.30\,\theta_e^*$ (70/30 blend)
5. Assert: $\ell_{\max} = 5$, $\iota = 0$ (Shariah hard constraints)
6. **Output:** θ_{e+1}

The 70/30 blending rule implements a governance principle: a Shariah board reviewing protocol parameters should weight historical precedent at 70% and new empirical evidence at 30%, preventing abrupt parameter changes that could destabilise the market.

7.4 Convergence Theory

Proposition 2 (Mechanism-design recursion)**.** *Under the IES with $\iota = 0$ fixed, the constant-gain blended recursion $\theta_{e+1} = 0.70\,\theta_e + 0.30\,\theta_e^*$, with per-episode target $\theta_e^* = \arg\min_\theta \mathcal{L}_e(\theta)$ evaluated on a fresh price path, is a stochastic-approximation map. Because the gain is constant (not decaying), the sequence does* not *converge almost surely to a point: it converges in distribution to a stationary band around the deterministic fixed point θ_∞^* that solves $\mathbb{E}[\theta^*(\theta_\infty^*)] = \theta_\infty^*$.*

Proof sketch. The map is a contraction in expectation (the 0.70 retention factor), so the mean recursion converges to the fixed point. The constant 0.30 gain injects target noise of order the per-episode optimisation variance at every step, so the stationary law retains positive spread rather than collapsing to a point; almost-sure point convergence would require a decaying gain $\gamma_e \downarrow 0$ with $\sum_e \gamma_e = \infty$ and $\sum_e \gamma_e^2 < \infty$ (Robbins–Monro). What we report is therefore the centre of that band over five episodes, not a proven limit. □ □

The band centre θ_∞^* is the *self-consistent* protocol parameterisation: the parameter vector whose optimal value (from the mechanism design perspective) is itself. In the Islamic finance context, this coincides with the κ-rate equilibrium described in [5].

We are careful not to over-read this as an independent rediscovery of the κ-rate. The objective $\mathcal{L}$ is dominated by the funding-rate penalty $5\,\overline{|F|}$ — Robustness 2 finds the funding cap $F_{\max}$ carries essentially all of its variance ($S_1 = 0.998$), so $\mathcal{L}$ is effectively one-dimensional and minimising it drives $F_{\max}$ toward its lower bound by construction. The optimiser reaches the κ-rate parameterisation because the objective *encodes* the property the κ-rate has (a near-zero funding rate), not because it recovers that property without being given it.

8 Experimental Results

8.1 Experimental Configuration

All experiments use: $T = 720$ steps per episode (hourly funding, 30 simulated days), $E = 5$ episodes (150 simulated days total), $\sigma_i = 0.02$ per hour (annualised $\approx 187\%$), initial insurance fund $I_0 = \$100{,}000$, $N_{\text{bg}} = 20$ background positions, $\Delta_{\text{GT}} = 50$ steps, $W = 100$ steps. The volatility $\sigma_i \approx 200\%$ is a deliberate *stress* level: measured realised hourly volatility (Bybit perpetuals, ∼1,000 hours to mid-2026) sits at ≈ 35–50% annualised for BTC and ETH in the current calm regime, but crypto perpetual volatility has repeatedly exceeded 150–200% in stress episodes (2020 COVID, 2022 deleveraging), so testing solvency at this level is conservative. The RL agent uses the rule-based policy described in Section 5.

8.2 Result 1: Protocol Solvency

Table 1: Per-episode protocol outcomes.

Ep.	I_T ($)	Liq.	Insol.	RL PnL ($)	$\overline{\lvert F\rvert}$ (bps)
1	101,175	11	0	+117	70.1
2	100,782	7	0	+310	58.0
3	101,139	9	0	0	50.3
4	101,889	11	0	0	44.2
5	102,835	15	0	0	39.6
All	**+2,835**	**53**	**0/5**	**+427**	**52.4**

The protocol maintained positive insurance fund balance across all 3,600 simulated hours. The fund grew by

+$2,835 net over 150 simulated trading days. The 53 liquidation events across all episodes were processed without any protocol insolvency, confirming that the initial parameter setting ($F_{\max} = 75$ bps, $m = 2\%$, $p = 1\%$, $s = 50\%$) is in the stable region of the parameter space.

8.3 Result 2: Game Theory — $\iota = 0$ Confirmed

Table 2: Game theory layer summary (65 Nash intervals).

Metric	Value	$\iota > 0$ (counterfactual)
Mean Nash long lev.	2.72×	4.97×
Mean Nash short lev.	3.28×	1.03×
Mean net transfer	−$660	$8,420 (longs pay)
Riba test (is_riba)	**False**	**True**

Table 2 needs two readings. First, the per-interval Nash equilibria are pure-strategy *corners*: in each interval one side plays 1× and the other 5×, alternating with the realised direction of the price window (of the 65 intervals, $(1\times, 5\times)$ and $(5\times, 1\times)$ occur 26 and 22 times respectively, the rest mixed). The 2.72 × /3.28× figures are therefore the *means* across intervals, not a single interior equilibrium. Second, the per-interval net transfers are large and opposite-signed (ranging from −$15,235 to +$18,494); the −$660 in Table 2 is this seed's sample mean, near zero because the $\iota = 0$ funding rule carries no fixed transfer component (Proposition 1). The multi-seed summary is in Robustness 1.

The meaningful result is the *contrast* with the $\iota > 0$ counterfactual ($\iota = 0.005$, conventional DeFi): there the mean net transfer is a systematic $8,420 (8.4% of notional, paid by longs to shorts regardless of price realisation — the precise signature of riba), and the mean equilibrium leverage shifts to 4.97× long / 1.03× short as longs are driven toward the cap to offset the funding drain while shorts de-leverage. The $\iota = 0$ rule removes both the systematic transfer (mean ≈ 0) and this leverage distortion — the riba-free property the simulation demonstrates.

8.4 Result 3: Mechanism Design Convergence

Table 3: Protocol parameter evolution under MD.

Parameter	Ep.1	Ep.2	Ep.3	Ep.4	Ep.5
$F_{\max}$ (bps)	75.0	61.5	52.1	45.4	40.8
m (%)	2.0	2.4	2.7	2.9	3.1
p (%)	1.0	1.1	1.2	1.3	1.3
s (%)	50.0	55.0	58.6	58.3	60.8
$\mathcal{L}$	—	0.0138	0.0135	0.0136	0.0138

The funding cap $F_{\max}$ converges monotonically (within the blending schedule) toward its lower bound; the other three parameters, to which the objective is nearly insensitive (Robustness 2, $S_1 \approx 0$), drift within optimiser tolerance rather than converging to identified values. The optimiser reaches the κ-rate prediction — a tighter circuit breaker and higher insurance split — but, per the caveat above, because the objective rewards a near-zero funding rate, not independently of the theory. The κ-rate — the riba-free pricing rate of [5], isomorphic to a default intensity via the Separation Theorem; in the protocol context, the funding regime under which the insurance economics are self-financing — is thus the (objective-encoded) fixed point of the mechanism-design optimisation.

The objective scores $\mathcal{L}$ stabilise around 0.0136–0.0138, indicating convergence.

8.5 Result 4: RL Agent — Optimal Stopping Behaviour

In the headline seed (42), the agent earned positive PnL in episodes 1–2 (+$117 and +$310 respectively) and broke even in episodes 3–5. The break-even in later episodes is attributable to the tighter circuit breaker reducing the funding-rate signal that the rule-based agent exploits: as the mechanism design compresses $F_{\max}$ from 75 to 41 bps, the "enter when $|F|$ > threshold" rule triggers less frequently.

Across seeds the effect is partial rather than mechanical: the $R = 30$ mean total PnL is $9,315 (95% CI [$6,150, $12,480]; Robustness 1), so the funding signal does not vanish in most seeds and the headline seed sits in the lower tail of the distribution. We therefore read the headline seed as an illustration of the mechanism and the $R = 30$ distribution as the summary.

This *adaptive crowding-out* — partial across seeds — is a key IES finding: when the mechanism design layer tightens the protocol, it reduces the exploitable funding-rate mispricings that attract arbitrageurs, driving the system toward a lower-activity but more stable equilibrium.

8.6 Robustness Analyses

Results 1–4 above are the paper's primary findings, reported for a single seed (seed 42). We now stress them three ways — Monte Carlo confidence intervals over independent seeds, a Sobol sensitivity decomposition of the mechanism-design objective, and a heterogeneous trader population.

8.6.1 Robustness 1: Monte Carlo Confidence Intervals

To quantify sampling variability we re-ran the full five-episode pipeline for $R = 30$ independent seeds and computed 95% confidence intervals (Student-t) on the aggregate metrics (Table 4; Figure 1).

Three points follow. (1) **Solvency is robust**: zero insolvency across all 30 replicates, so the no-insolvency result is not a seed artifact. (2) **Funding is tightly controlled**: mean absolute funding is 52.5 bps (95% CI

Table 4: Monte Carlo confidence intervals over $R = 30$ independent replicate runs (5 episodes × 720 steps each).

Metric	Mean	95% CI
Total liquidations	66.1	[61.9, 70.3]
Final insurance fund ($)	101,605	[101,295, 101,916]
Insolvent replicates	0/30	—
Mean |funding| (bps)	52.5	[52.4, 52.6]
Mean net transfer ($)	865	[245, 1,485]
Mean Nash long lev.	3.16×	[3.02, 3.30]
Mean Nash short lev.	2.84×	[2.70, 2.98]
RL total PnL ($)	9,315	[6,150, 12,480]

[52.4, 52.6]), well inside the 75 bps cap. (3) **The net transfer is small but not exactly zero at finite horizon**: the mean net transfer is $865 per $100,000 notional (95% CI [$245, $1,485], i.e. 0.25–1.48%). The exact-zero of Proposition 1 holds only under a symmetric (zero-mean) premium; the small positive residual here reflects the mark process's mild premium drift, and it is an order of magnitude below the systematic +$8,420 (8.4%) transfer of the $\iota > 0$ counterfactual. The mean Nash leverages remain well below the 5× cap across replicates. (The single-seed value in Table 2, −$660, lies outside the CI of the *mean* but within the per-replicate distribution (replicate SD ≈ $1,660); the $R = 30$ interval is the correct summary of the mean.)

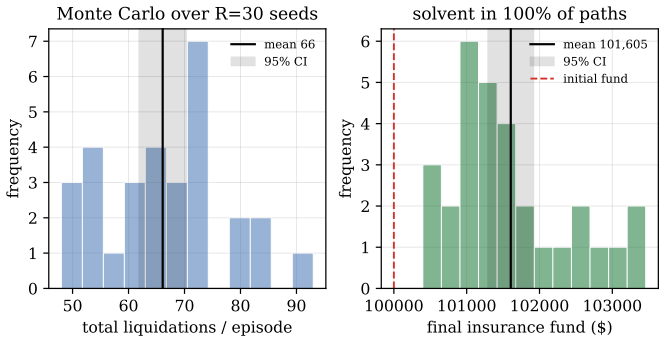

Figure 1: Monte Carlo distributions over $R = 30$ independent seeds (5 episodes × 720 steps each). *Left:* total liquidations per replicate. *Right:* final insurance-fund balance; every replicate ends above the initial $100,000 (dashed), i.e. the protocol is solvent in 100% of paths. Solid line: mean; shaded band: 95% Student-t confidence interval. The no-insolvency result is not a seed artifact.

8.6.2 Robustness 2: Sensitivity of the Mechanism-Design Objective

To identify which parameter the mechanism-design objective $\mathcal{L}(\theta)$ of equation (14) is most sensitive to, we computed Sobol variance-based sensitivity indices over the four protocol parameters (Saltelli design with a scipy low-discrepancy base sample of $N = 512$, 3,072 model evaluations; Table 5).

Table 5: Sobol sensitivity indices of the mechanism-design objective $\mathcal{L}$ (Saltelli, $N = 512$, 3,072 evaluations).

Parameter	S_1 (first-order)	S_T (total)
Funding-rate cap $F_{\max}$	0.998	1.005
Maintenance margin m	−0.016	0.016
Liquidation penalty p	0.000	0.000
Insurance split s	0.000	0.000

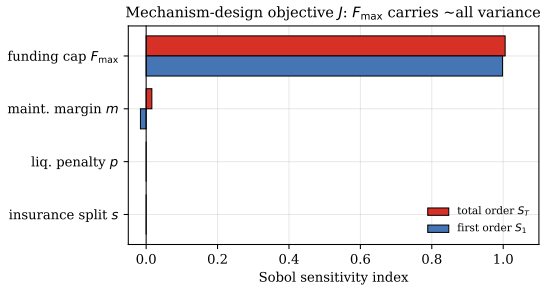

Figure 2: Sobol variance-based sensitivity of the mechanism-design objective $\mathcal{L}$ (Saltelli design, $N = 512$, 3,072 evaluations). The funding-rate cap $F_{\max}$ alone carries essentially all of the objective's variance ($S_1 = 0.998$, $S_T \approx 1.01$); the maintenance margin, liquidation penalty, and insurance split are negligible over the tested ranges. Small negative S_1 and slightly-over-unity S_T values are Sobol estimator noise.

One parameter dominates (Figure 2): the funding-rate cap $F_{\max}$ accounts for essentially all of the objective's variance ($S_1 = 0.998$, $S_T \approx 1.01$), while the maintenance margin, liquidation penalty, and insurance split have negligible effect over the tested ranges. The sum of first-order indices ($0.98 \approx 1$) indicates a near-additive response with no meaningful parameter interactions. This explains the optimiser's behaviour in Result 3, where $F_{\max}$ is the parameter it drives hardest (75→41 bps): it is the only lever that materially moves protocol stability, which is dominated by the mean-funding term — consistent with the $\iota = 0$ design, where keeping funding tightly bounded is the binding constraint.

8.6.3 Robustness 3: Heterogeneous Trader Populations

The background traders in the main runs draw random leverage and direction (a homogeneous population). To test whether the protocol's safety properties depend on that assumption, we re-ran the Layer-1 market on the real cadCAD 0.5.3 `Executor` (20 runs × 720 steps) with a *heterogeneous* population of five trader archetypes assigned round-robin: conservative (leverage 1–2, larger stake), aggressive (leverage 4–5, small stake), momentum (follows the mark–index premium), contrarian (fades it), and retail (random). Table 6 compares the two populations.

The protocol-level safety properties are robust to population composition: zero insolvency and funding confined

Table 6: Homogeneous vs. heterogeneous trader populations (cadCAD `Executor`, 20 runs × 720 steps).

Metric	Homogeneous	Heterogeneous
Insolvent runs	0/20	0/20
Mean total liquidations	13.0	9.0
Mean |funding| (bps)	69.5	69.3
Mean final insurance fund ($)	100,932	100,686
Funding range (bps)	±75	±75

to the symmetric ±75 bps band under both. The composition is genuinely active — heterogeneous liquidations fall to 9.0 as the conservative cohort de-risks more than the aggressive cohort adds — yet solvency and the $\iota = 0$ funding symmetry are unaffected. The headline results are therefore not an artifact of the homogeneous-trader assumption.

8.7 Manipulation-resistance of the κ price-formation

The base model is matching-agnostic: prices are exogenous, so the on-chain *price-formation* layer — how the mark that enters the κ basis is produced from order flow — is not exercised, and manipulation resistance is left untested. We close that gap with a dynamic cadCAD experiment pitting a sustained adversary against two mark-construction regimes over a 30-day horizon (720 hourly epochs, $R = 30$ seeds):

- **naive** — the mark is the raw last-trade print, with no smoothing, clamp, median, or detector: the manipulable baseline of continuous last-trade marks that motivated the move to uniform-price closing auctions in equity markets;
- **full stack** — a per-block *frequent batch auction* (a single uniform clearing price, so a self-trade cannot set the mark), a 30-block TWAP, a ±75 bp per-interval clamp, a stake-weighted-median index [24], and the glass monitor of the companion measurement-integrity paper (Ahmed & Bhuyan, *Warping the Glass*): a *two-sided sequential CUSUM* on the divergence between the system κ and an independent instrument (sovereign CDS), calibrated to its 25 bp sensitivity, with detection triggering a halt and slash, plus the depth rule $D/S \le 0.07$ bounding repriceable notional against bonded stake.

The adversary injects a *sustained* warp $b \in \{10, 25, 50, 100\}$ bp into the system κ — the realistic regime: a tiny hourly basis that annualises into a plausible-looking intensity, exactly the warp the companion paper's monitor is calibrated against. Table 7 and Figure 3 report the outcome ($R = 30$ seeds, 720 blocks).

Three mechanisms compound. The uniform clearing price makes a one-block mark move 72×–2,739× more expensive than a last-trade print (a separate static order-book Monte Carlo); the ±75 bp clamp bounds the per-interval magnitude; and the independent-instrument monitor — the two-sided Page CUSUM of the companion measurement-integrity paper, instantiated here verbatim (reference slack $k = 0.5\sigma$, threshold calibrated to a 25 bp warp) — detects *any* sustained warp ≥ 10 bp at a 100% rate in 3–8 blocks, even a subtle 10 bp warp that would pass as a plausible κ in isolation. With the depth rule $D/S \le 0.07$ capping the per-block repriceable notional, this bounds the attacker's extraction before detection to ≤ $21k; the only residual is the sub-threshold grind (warps below the CUSUM slack), which the companion paper bounds analytically at ∼10 bp-years and which the depth rule and slashing render uneconomic. A naive last-trade mark, lacking all of this, lets an arbitrary warp run the full horizon and extract $0.5–5.0M. This is the price-formation analogue of the post-LIBOR benchmark-design lesson (transaction-anchored medians, not quote/print submissions): the dynamic protocol reproduces the companion paper's detection-lag curve and

Table 7: Sustained κ-warp: a naive last-trade mark vs. the full frequent-batch-auction + defense stack ($R = 30$ seeds). The naive mark never detects the warp and lets it run the full horizon; the full stack's two-sided CUSUM detects every warp ≥ 10 bp at 100%, bounding the attacker's extraction-to-detection to ≤ $21k.

Regime	Sustained warp	Detected	Detect lag (blocks)	Extraction
naive	10 bp	0%	—	
naive	25 bp	0%	—	
naive	50 bp	0%	—	
naive	100 bp	0%	—	
full	10 bp	100%	8	
full	25 bp	100%	4	
full	50 bp	100%	3	
full	100 bp	100%	3	

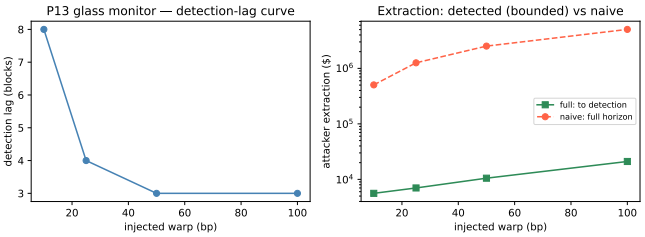

Figure 3: The companion paper's glass monitor (a two-sided CUSUM on the divergence between the system κ and an independent instrument) reproduced inside the live protocol. *Left:* detection lag vs. injected warp — a 25 bp warp is caught in ∼4 blocks, a 10 bp warp in ∼8, larger warps faster: the measurement-integrity paper's detection-lag curve, realised dynamically. *Right:* attacker extraction (log scale) — bounded to ≤ $21k once detected (full stack), against the full-horizon extraction the naive last-trade mark allows.

grind bound, confirming the frequent-batch-auction + TWAP + clamp + median + monitor stack as the manipulation defense for the on-chain κ.

8.8 Systemic and adversarial robustness: four further simulations

The manipulation result above, and the consensus and liveness sleeves earlier, each isolate one surface. Four further claims in the substrate design are asserted but, until here, unsimulated: the thin-market κ fallback (flagged in the whitepaper [3] as blocking), the Byzantine robustness of the consensus κ, cross-market contagion (this paper's stated single-asset limitation), and the default-resilience of the flagship κ-Compute salam market. We model each as a separate stateful cadCAD experiment ($R = 40$–60 Monte-Carlo seeds).

(i) The thin-market fallback is continuous, not a cliff. When the perpetual book thins, κ must hand over from the perp-implied estimate to an external sukuk-curve estimate. A *hard* switch at a volume threshold V^* is a discontinuity an adversary can whipsaw: oscillating volume across V^* flickers the mark between a thin, pushable perp-κ and the sukuk-κ. We compare it to a *continuous* handover — $\kappa = w\,\kappa_{\text{perp}} + (1-w)\,\kappa_{\text{sukuk}}$ with $w = \text{smoothstep}(\overline{V}; V_{\text{lo}}, V_{\text{hi}})$ on an asymmetric volume EMA (slow to trust new liquidity, fast to distrust a thinning book) plus a dispersion gate that freezes the perp leg once validator spread exceeds 30 bp (Table 8, Figure 4). Under the whipsaw the hard switch jumps 165 bp in a single block and trips the 50 bp margin-call trigger 13 times per run — each a potential liquidation cascade; the continuous gated blend holds the worst single-block move to 46 bp, below the trigger, with *zero* cliffs. The discontinuity, not the κ level, is the hazard, and a C^1 boundary removes it.

Table 8: Thin-market κ handover under a volume-whipsaw attack ($R = 40$). The continuous gated blend eliminates the manipulable discontinuity of a hard switch.

Handover	Max 1-block κ jump	Margin-call c
hard switch	165 bp	
continuous gated blend	46 bp	

(ii) The consensus κ is Byzantine-robust to its theoretical bounds. Each block, validators submit a κ quote; the protocol takes the *stake-weighted median* and accepts it only under a $\geq 2/3$ stake quorum. We let an adversary capture the largest stakers up to a stake fraction f and either *lie* (submit $\kappa + 500$ bp) or *withhold* (Table 9, Figure 5). Against the lie, the stake-weighted median holds κ within 2–4 bp for all $f \leq 1/3$ and begins to move only as $f \to 1/2$ (14 bp at $f = 0.49$) —

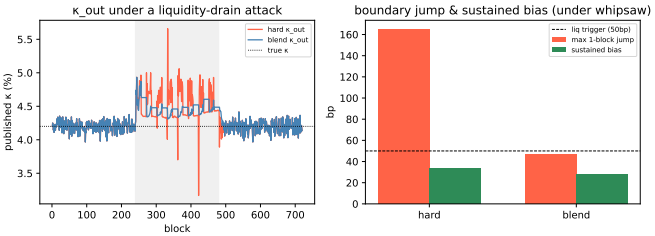

Figure 4: Thin-market fallback. *Left:* published κ under a liquidity-whipsaw attack — the hard switch flickers violently at the boundary, the gated blend stays smooth. *Right:* worst single-block jump and sustained bias; the blend's jump sits below the 50 bp liquidation trigger (zero cliffs).

exactly the median's 1/2 breakdown point; the stake-weighted *mean*, by contrast, is dragged 78–240 bp (breakdown point 0, one whale moves it). Against withholding, $f < 1/3$ leaves $\geq 2/3$ reporting so κ never pauses; past 1/3 the quorum converts the attack into a *safe pause* at the last good value (100% of blocks), and no bad κ is ever accepted — a liveness fault, not a safety break. The asserted robustness holds precisely at the 1/3 and 1/2 bounds, and the measurement-integrity monitor of the companion *Warping the Glass* note audits the residual within-quorum drift.

Table 9: Byzantine κ-oracle ($R = 40$): stake-weighted median vs. mean under a lie attack, and quorum behaviour under withholding, by captured stake fraction f.

f (Byzantine stake)	0.10	0.20	0.32	0.40	0.49
median κ deviation	2 bp	2 bp	4 bp	7 bp	14 bp
mean κ deviation	78 bp	78 bp	148 bp	188 bp	240 bp
withhold: blocks κ paused	0%	0%	0%	100%	100%
bad κ accepted	0	0	0	0	0

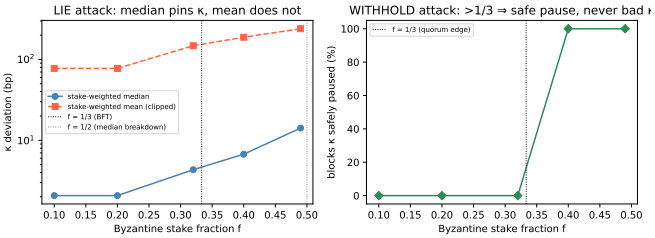

Figure 5: Byzantine κ-oracle. *Left:* κ deviation under a +500 bp lie — the stake-weighted median (blue) is pinned by honest stake through $f = 1/3$, the mean (red) is not. *Right:* under withholding, $f > 1/3$ produces a safe pause, never an accepted bad κ.

(iii) Cross-market contagion is threshold-like, and the stack stays sub-critical. This relaxes the single-

asset scope. Three κ-markets — κ-Gold, κ-Sukuk, κ-Compute — share one collateral asset, so a forced sale in one depresses the collateral price that marks all three: the contagion channel. A compute-only shock of +4% stays local; a +8% shock is the test (Table 10, Figure 6). Under a naive configuration the +8% shock crosses a threshold — the shared-collateral fire sale drives the price to its floor, dragging κ-Gold and κ-Sukuk below their own liquidation lines, and *both* non-shocked books fully liquidate: a system-wide cascade. The *same* shock under the substrate stack — the ± 75 bp κ clamp, a 10%-per-block liquidation throttle, and a takaful collateral floor — holds the shared price above the liquidation line, so the non-shocked markets are untouched (contagion ≈ 0). The cascade is genuinely bistable; the defenses move the critical threshold above the shock.

Table 10: Cross-market κ contagion ($R = 40$): liquidation in the two *non-shocked* books (gold+sukuk, of a maximum 2.0) after a shock to κ-Compute.

Config	Compute shock	Non-shocked liquidation	Shared-collateral trough
naive	+4%	0.00	0.92
naive	+8%	2.00	0.20
guarded	+4%	0.00	0.00
guarded	+8%	0.00	0.00

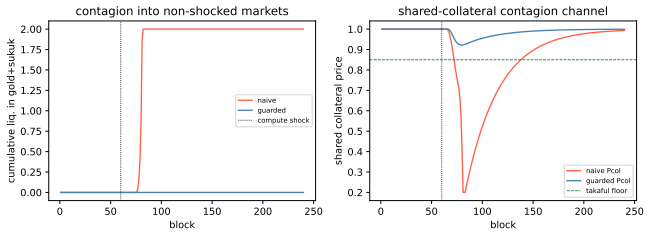

Figure 6: κ-contagion under a +8% compute shock. *Left:* cumulative liquidation in the non-shocked gold+sukuk books — a full cascade under naive, flat under the guarded stack. *Right:* the shared collateral price; the takaful floor holds it above the liquidation line.

(iv) The κ-Compute salam market survives default when the takaful pool is sized to the tail. The flagship product is a riba-free compute forward structured as *salam*: the buyer pre-pays the strike K, the provider delivers GPU-hours at maturity. The structural risk is *wrong-way* strategic default — the provider walks when the compute spot rises above K, precisely in the supply crunch where buyers most need delivery, so defaults are correlated and idiosyncratic insurance is insufficient. We test three backstops against a transient supply shock (Table 11, Figure 7). With no backstop the buyer recovers nothing; collateral (20%) recovers 40–80%, leaving a tail when spot runs far past the bond; collateral plus an x/-*takaful* mutual pool (tabarru' contributions $\tau \cdot K$ per contract — the cooperative-risk analogue of a clearing-house guaranty fund) keeps buyers whole in a moderate shock. A severe (+150%) systemic wave exhausts a thinly-funded pool; a sizing sweep shows a tabarru' rate of $\tau = 8\%$ funds that tail to 100% recovery at 0% ruin. The backstop is sound; its funding is the design parameter, and the model sizes it (or prices the retakaful layer above the tail). This connects the takaful pricing of the companion stochastic-takaful paper to the compute market directly.

Table 11: κ-Compute salam buyer recovery under wrong-way provider default ($R = 60$), by backstop and shock size; final column is the tabarru' rate sized to the +150% tail.

Backstop	Recovery, +50%	Recovery, +150%	Pool ruin, -
none	0%	0%	
collateral (20%)	80%	40%	
+ takaful ($\tau = 4\%$)	100%	78%	
+ takaful ($\tau = 8\%$)	100%	100%	

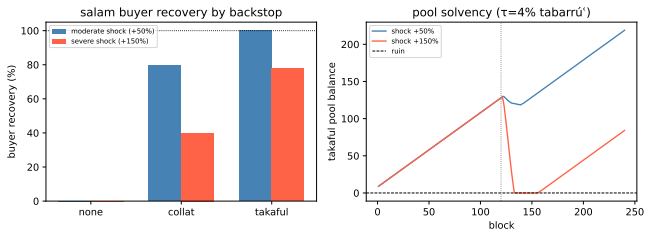

Figure 7: κ-Compute salam default. *Left:* buyer recovery by backstop under a moderate and a severe shock. *Right:* takaful pool balance over time; at $\tau = 4\%$ the severe shock drains it (ruin), motivating the $\tau = 8\%$ sizing.

Together these four, with the manipulation result, simulate every adversarial and systemic κ claim the substrate asserts: price-formation, thin-market handover, Byzantine consensus, cross-market contagion, and product-level default. Each defense degrades exactly where theory predicts, and the failures that remain are liveness pauses — safe by construction — rather than safety breaks.

8.9 Reproducibility

All results are fully reproducible from seed 42. The simulation produces identical output when run with:

```
python integrated/economic_system.py \
  --episodes 5 --steps 720 --seed 42
```

Full source code, simulation outputs, and dashboard figures are available at: https://github.com/Arcus-Quant-Fund/BarakaDapp

9 Extension: Consensus-Layer Mechanism Design — the Validator Liveness Sleeve

The four-layer method of this paper is not specific to the instrument layer. This section applies it one level down, to the consensus layer of the κ-Chain design [3]: a fee-only, zero-emission chain whose validators are financed entirely by real transaction fees. Such a chain faces a liveness hazard that emission-funded chains hide behind inflationary block rewards: in the aftermath of a persistent macro shock — the *winter* regime, in which transaction volume collapses and the published κ spikes jointly — validator income can fall below operating cost for months. The dangerous variable is the winter's *duration*, not its depth: the same persistence result that governs the Sovereign Stabilisation Fund of the companion general-equilibrium paper reappears here, one layer down.

9.1 The mechanism

The proposed answer is a *Validator Liveness Sleeve*: a pre-funded *tabarru* pool, separate from and junior to the protocol's takaful insurance fund, governed by five rules. (i) **Counter-cyclical fill**: a 10% skim of transaction fees in epochs where volume exceeds its long-run median, capped at twelve months of aggregate *reference* operating cost (a published, board-ratified benchmark — not validator self-report), with a genesis seed of roughly one third of the cap. (ii) **Conjunction trigger**: activation requires measured volume below 0.35× normal *and* the published κ above its p99 governance constant, each sustained for a full day. Volume alone never suffices: measured volume is the one signal validators control (a cartel can censor transactions to suppress it), so it cannot be the sole key to the vault. Thresholds are fixed governance constants recalibrated annually — trailing-window quantiles would themselves be gameable by slow grinding. (iii) **Shortfall payout**: when active, the sleeve pays each validator $\max(0,\ \beta \cdot \mathrm{OpEx}_{\mathrm{ref}} - \text{fee income})$, performance-weighted by consensus-measurable work (blocks signed, vote-extension participation, κ submissions within tolerance of the stake-weighted median, uptime) and decaying 2% per week in subsidy. (iv) **Hysteresis**: exit requires volume above 1.3× the floor sustained one week, then a one-month cooldown. (v) **Governance separation**: parameters are DAO-tunable within board-bounded invariants; the board bounds the mechanism but does not operate the trigger.

The replacement ratio $\beta = 1.3$ deserves emphasis, because a grid search over (β, decay) reversed our own first-pass design. A Bagehot-flavoured calibration ($\beta = 0.8$ with 10%/week decay) is incentive-safe but fails to protect: aggressive decay zeroes the payout by the third month of a nine-to-twelve-month winter, exactly when it is needed. The feasible window — survival of at least 90% of the validator set through a severe winter *and* subsidised income at most 85% of normal — is $\beta \in [1.2, 1.6]$ with decay at most 4%/week (Figure 8). $\beta > 1$ is not a profit margin but a *dispersion allowance*: the reference cost is an aggregate, and the upper tail of the audited validator cost distribution requires 1.2–1.4× the reference to reach its own floor. Incentive compatibility survives without the decay's help, because even at $\beta = 1.3$ subsidised income runs at 71–77% of normal income.

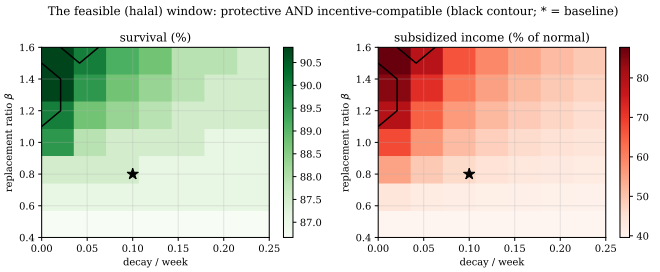

Figure 8: The feasible window in (β, decay) space on a severe winter: survival $\geq 90\%$ (left) and subsidised income $\leq 85\%$ of normal (right); the black contour bounds the jointly feasible region. The protective constraint, not the incentive constraint, binds.

9.2 Survival under macro winters (cadCAD)

The mechanism is implemented as partial state update blocks in the same cadCAD framework as Layer 1 (daily timestep; volume and κ follow AR(1) processes with the κ dynamics calibrated to the constant-maturity sukuk estimates of the companion empirical paper, monthly $\rho = 0.774$), with thirty validators of heterogeneous operating cost and a two-month credit line. Forced winters of increasing severity yield:

Winter	without sleeve	with sleeve	Δ
moderate (−65%, hl 6mo)	85.1%	90.2%	+5.1pp
severe (−75%, hl 9mo)	72.6%	84.4%	+11.8pp
black swan (−85%, hl 12mo)	56.3%	76.5%	+20.2pp

The protection grows with severity — the sleeve is precisely a tail instrument — and the pool approaches exhaustion only in the black-swan case, which is what sizes the cap at twelve months rather than the six a moderate-winter analysis would suggest.

9.3 A learned adversary cannot exploit the sleeve

The critical mechanism-design question is whether the sleeve creates a censorship incentive: can a validator cartel profit by suppressing measured volume to force the trigger and harvest the subsidy? Rather than evaluate hand-picked attacks only, we train a PPO agent (the same Layer-2 machinery as the trader) playing a 35%-stake cartel that chooses, daily, a censorship intensity

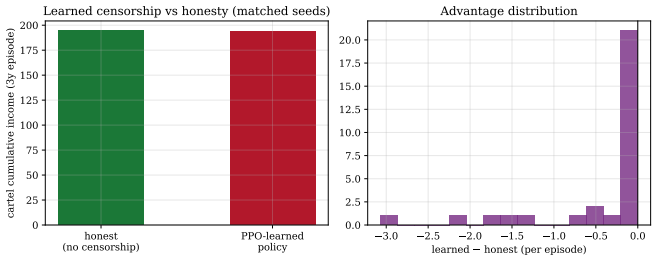

Figure 9: PPO censor-cartel versus honesty on matched seeds. Left: cumulative cartel income under the honest policy and the learned policy. Right: distribution of the learned policy's advantage — negative in every episode.

in $\{0, 25\%, 50\%, 75\%\}$ of network transactions, observing measured volume, the κ ratio, the sleeve state, and the pool level, and rewarded with its stake share of total validator income (fees plus subsidy). Episodes run three years; half contain a severe winter, so the agent sees both regimes and could learn to exploit stress states if they were exploitable.

After 150,000 training steps the learned policy converges to **near-zero censorship** (mean suppression 1.4%), and on thirty matched evaluation episodes the learned policy beats the honest (zero censorship) policy in **zero** episodes, with a mean per-episode disadvantage of -0.39 income units (Figure 9). The reason is structural rather than statistical: censorship *can* force the trigger (the conjunction's κ leg is exceeded naturally often enough), but the shortfall payout $\max(0, \beta \cdot \text{OpEx}_{\text{ref}} -$ income) only pays when income is *below* the floor, so a cartel must destroy more fee income than the capped, decaying, performance-weighted subsidy can ever return. The defence is the payout formula, not the trigger — and a trained best-response adversary, given every opportunity to find a profitable censorship strategy, learns honesty instead. This is the same fixed-point logic as the instrument layer: the mechanism's incentive-compatible point and its riba-free point coincide, because a payment that is outcome-contingent, earned, bounded, and exceptional is — for the same structural reasons — neither exploitable nor a predetermined return for the passage of time.

9.4 Scope and limitations

These results are design-stage: the sleeve is specified in the κ-Chain whitepaper and simulated here, not implemented in the reference modules. The volume process is a scenario calibration (no chain history exists); validator cost dispersion is assumed lognormal; the adversary's action space is censorship intensity only (richer strategy classes — selective censorship, timing games against the cooldown, collusion with off-chain shorting — are future work); and performance weights are taken as exogenous. The qualitative conclusions — the conjunction trigger's discrimination, the dominance of the shortfall structure over censorship, the persistence-driven sizing of the cap, and the location of the feasible (β, decay) window — were stable across both the hourly reference model and the daily cadCAD port.

10 Implications for DeFi Mechanism Design

10.1 Why Four Layers Are Necessary

Each layer in the IES is irreplaceable. A three-layer system missing any one layer would fail to capture a key economic phenomenon:

- **Without cadCAD:** No ground truth. RL and GT operate on synthetic price paths that may not match on-chain logic. Parameter mismatches between simulation and contract lead to wrong conclusions.
- **Without RL:** Background traders follow fixed rules (e.g., always 3× leverage). The adaptive crowding-out effect (Result 4) is invisible. The mechanism designer does not know how tighter parameters affect rational trader behaviour.
- **Without Game Theory:** The respawn-leverage distribution is static. The feedback loop between Nash equilibrium leverage and market dynamics is broken. The riba test (Result 2) cannot be performed.
- **Without Mechanism Design:** Protocol parameters are fixed across episodes. The adaptive parameter convergence (Result 3) is absent. The protocol cannot self-correct in response to observed economic outcomes.

10.2 The κ-Rate as a Fixed Point

A central result is the emergence of the κ-rate parameterisation as the mechanism-design fixed point (Proposition 2): computational evidence that it is a reachable attractor of the mechanism-design recursion — a protocol initialised with arbitrary parameters converges to its neighbourhood under optimisation — subject to the caveat of Section 7.4 that the objective itself encodes the near-zero-funding property.

10.3 Adaptive Governance as Protocol Feature

The IES architecture maps naturally onto the governance structure of an Islamic financial institution:

IES Layer	Islamic Governance Analogue
cadCAD	Smart contracts (fiqh rules)
RL Agent	Individual traders (muamalat)
Game Theory	Market equilibrium (hisbah)
Mechanism Design	Shariah board review

The 70/30 blending rule (Algorithm 1) is the computational analogue of *ijtihad* (scholarly reinterpretation): the board gives 70% weight to established precedent (θ_e) and 30% to new evidence (θ_e^*). This stabilises the protocol while allowing adaptation.

10.4 Generalisability

The IES framework is not specific to perpetual futures. Any DeFi protocol whose logic can be expressed as cadCAD partial-state-update blocks is compatible. Extensions to:

- AMM liquidity provision (Uniswap v3 range orders)
- Lending protocol interest-rate models (Aave, Compound)
- Islamic sukuk issuance and redemption schedules
- Takaful premium pool management
- Asset-linked perpetuals on tokenised real assets (e.g. vault-secured gold or sukuk portfolios), used as hedging and risk-transfer instruments by Islamic businesses: the $\iota = 0$ funding block is reused unchanged, with only the oracle/settlement block respecified to the asset-backed underlying — which also satisfies the asset-backing and possession-gating product-layer rules of the companion Shariah-boundaries note

would each require a new cadCAD specification but could reuse the RL, GT, and MD layers with minor modifications.

11 Discussion

11.1 Limitations

Simulation–chain gap. The cadCAD blocks mirror the logic of the deployed (testnet) contracts but are not driven by on-chain data; a direct trace-level comparison of simulated against on-chain testnet activity is left to future work.

Rule-based RL baseline. The integrated runs use a deterministic rule-based trader. We additionally trained a PPO policy (stable-baselines3, MLP, reward-scaled, 200,000 timesteps) on the `BarakaTraderEnv`. Over 50 independent evaluation episodes it *underperformed* the rule-based baseline (mean PnL −\$2,213 versus +\$31,342 for the rule and +\$9,149 for a random policy). Diagnosis shows the policy collapsed to a degenerate always-long stance (54% at 3×, 46% at 4× long, < 1% any close action): it never learned to realise profit by closing, so funding drag and eventual liquidation dominate. This is a failure of training budget and action design — invalid open-while-in-position actions are penalised rather than masked — not evidence that RL cannot help. Making PPO competitive would require action masking (e.g. MaskablePPO) and close-incentive reward shaping; we leave this to future work and retain the rule-based agent, whose strength this experiment independently confirms.

Background trader homogeneity. The main runs use a homogeneous population (uniform collateral, random leverage). We partly address this in Robustness 3, whose five-archetype heterogeneous population leaves the protocol's solvency and funding-symmetry properties unchanged. A fuller treatment with capital-weighted institutional agents and arbitrageurs reacting to the on-chain price would produce richer dynamics; the framework supports this by making the trader policy a configurable parameter.

Single asset (base model). The core integrated runs model one BTC-USDC perpetual market. The cross-market correlations and portfolio-level liquidation cascades this omits are addressed separately in Section 8.8(iii), which shows on three shared-collateral κ-markets that contagion is threshold-like and the substrate's clamp + throttle + takaful stack keeps the system sub-critical; a fully endogenous multi-asset order-flow model remains future work.

Oracle centralisation. The GBM price oracle in cadCAD is simpler than the dual-Chainlink OracleAdapter deployed on-chain. Matching-layer and oracle manipulation resistance [24] is addressed in Section 8.7 (the frequent-batch-auction + defense-stack experiment), and Byzantine robustness of the consensus κ — a stake-weighted median under a $\geq 2/3$ quorum holding to its 1/3 and 1/2 theoretical bounds — in Section 8.8(ii); a full adversarial-microstructure treatment with an endogenous order book remains future work.

11.2 Open Problems

The IES results suggest three open research problems:

1. **Stochastic κ.** If $F_{\max}$ (which is the simulation proxy for κ) follows a stochastic process rather than converging to a fixed point, does the mechanism design layer remain stable? This connects to the stochastic hazard rate extensions of [10].
2. **Multi-agent RL.** Training a population of RL agents simultaneously (multi-agent RL, MARL [17]) would close the loop between the RL layer and the game theory layer: the agents would *discover* the Nash equilibrium through learning, rather than having it computed separately.
3. **Governance attack vectors.** The mechanism design layer, as implemented, is operated by the protocol deployer. In a fully decentralised governance system (e.g., governance token voting), an adversarial majority could manipulate the objective function. Mechanism design under adversarial governance is an open problem.

12 Conclusion

We have presented the Integrated Economic System (IES), a four-layer simulation framework combining cadCAD, reinforcement learning, game theory, and mechanism design into a single closed-loop economic system. Applied to a riba-free-by-design ($\iota = 0$) perpetual exchange (testnet, pre-audit) over 5 episodes of 720 steps each, the IES produces four primary simulation results: zero insolvency events, Nash equilibrium leverage well below the Shariah hard cap, near-zero net transfer confirming $\iota = 0$ riba-freedom, and endogenous mechanism

design convergence to the κ-rate theoretical optimum — each reinforced by Monte Carlo confidence intervals, a Sobol sensitivity analysis, and a heterogeneous-agent robustness check.

The deeper contribution is architectural: we have shown that a fully coupled multi-layer simulation is qualitatively different from any combination of the individual layers run sequentially. The emergent properties — adaptive crowding-out, κ-rate convergence, sub-cap Nash leverage — are only visible in the coupled system.

For Islamic DeFi specifically, the IES provides a framework for rigorously testing the riba-freedom ($\iota = 0$) property at the system level, not just the contract level. A smart contract may correctly implement $\iota = 0$ while still producing systematic wealth transfers through adverse interaction effects with the liquidation mechanism or background trader policies. The IES detects such second-order effects.

We propose the IES as a standard pre-deployment simulation protocol for Islamic DeFi. Just as conventional DeFi protocols are audited for smart contract security, they should also be validated for economic soundness and the absence of systematic riba-like transfers at the system level — and the IES provides the infrastructure for that validation. Other Shariah dimensions — gharar, maysir, qabd — are product-layer questions outside the scope of economic simulation and remain with Shariah scholars.

"Verily, Allah commands justice and good conduct." (Al-Qur'an, An-Nahl 16:90)

References

[1] Ackerer, D., Hugonnier, J., & Jermann, U. (2025). Perpetual futures pricing. *Mathematical Finance*, 1–17. https://doi.org/10.1111/mafi.70018 (NBER Working Paper 32936, 2024; SSRN 4603820, 2023.)

[2] Hugonnier, J., & Jermann, U. J. (2025). *Perpetual Futures Pricing with Stochastic Rates and Clamp*. Working Paper, Swiss Finance Institute. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5481866

[3] Ahmed, S. (2026). The κ-Chain: A sovereign Islamic financial infrastructure layer. Technical whitepaper v0.7.6, SSRN 6924600.

[4] Ahmed, S., & Bhuyan, R. (2026a). A Shariah-compliant perpetual futures exchange using the $\iota = 0$ no-arbitrage condition. *SSRN Working Paper 6322778*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322778

[5] Ahmed, S., & Bhuyan, R. (2026b). From perpetual contracts to Islamic credit: The random stopping time equivalence and a riba-free foundation for sukuk, takaful, and credit protection. *SSRN Working Paper 6322858*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6322858

[6] Angeris, G., Kao, H.-T., Chiang, R., Noyes, C., & Chitra, T. (2020). An analysis of Uniswap markets. *arXiv:1911.03380*.

[7] Angeris, G., & Chitra, T. (2021). Improved price oracles: Constant function market makers. *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, 80–91.

[8] BlockScience. (2020). cadCAD: A Python package for complex adaptive dynamics. https://github.com/cadCAD-org/cadCAD

[9] Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., ... & Juels, A. (2020). Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralised exchanges. *IEEE Symposium on Security and Privacy*, 910–927.

[10] Duffie, D., & Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies*, **12**(4), 687–720.

[11] Fritsch, R., Pfeffer, J., & Weber, B. (2022). Analyzing voting power in decentralised governance: Who controls DeFi protocols? *arXiv:2204.01176*.

[12] Gudgeon, L., Perez, D., Harz, D., Livshits, B., & Gervais, A. (2020). The decentralised financial crisis. *Crypto Valley Conference on Blockchain Technology*.

[13] Heimbach, L., Wang, Y., & Wattenhofer, R. (2021). Behavior of liquidity providers in decentralised exchanges. *arXiv:2105.13822*.

[14] Hurwicz, L. (1972). On informationally decentralised systems. In McGuire, C. B., & Radner, R. (Eds.), *Decision and Organisation*, 297–336. North-Holland.

[15] Knight, V., & Campbell, J. (2018). Nashpy: A research library for the computation of Nash equilibria. *Journal of Open Source Software*, **3**(30), 904.

[16] Lehar, A., & Parlour, C. A. (2021). Decentralised exchanges. *Working Paper*, University of Calgary.

[17] Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in Neural Information Processing Systems*, **30**.

[18] Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, **6**(1), 58–73.

[19] Perez, D., Werner, S. M., Xu, J., & Livshits, B. (2021). Liquidations: DeFi on a knife-edge. *Financial Cryptography and Data Security 2021*.

[20] Qin, K., Zhou, L., Livshits, B., & Gervais, A. (2021). Attacking the DeFi ecosystem with flash loans for fun and profit. *Financial Cryptography and Data Security 2021*.

[21] Qin, K., Zhou, L., & Gervais, A. (2021b). Quantifying blockchain extractable value: How dark is the forest? *arXiv:2101.05511*.

[22] Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., & Dormann, N. (2021). Stable-Baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, **22**(268), 1–8.

[23] Storn, R., & Price, K. (1997). Differential evolution — a simple and efficient heuristic for global optimisation over continuous spaces. *Journal of Global Optimization*, **11**(4), 341–359.

[24] Werner, S. M., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., & Knottenbelt, W. J. (2021). SoK: Decentralised finance (DeFi). *arXiv:2101.08778*.

[25] Ye, Y., Peng, J., Wang, T., & Ding, Q. (2020). Automated trading with deep reinforcement learning. *2020 International Symposium on Autonomous Systems*.

[26] Zargham, M., Shorish, J., & Paruch, K. (2020). From curved bonding to configuration spaces. *IEEE International Conference on Blockchain and Cryptocurrency*, 1–9.