AMM-Aware Central Counterparty Risk: Reflexive Cascades, Slow-Fast Singular Barriers, and Exogenous Settlement Severance

Raeez Lorgat

April 2025

Abstract

We develop a risk framework for a central counterparty (CCP) that clears trades executed on an automated market maker (AMM) in a hub-and-spoke topology. The paper separates the adapted clearinghouse machinery (portfolio margin, default waterfalls, haircuts, depeg controls) from four original mathematical contributions and a defense-in-depth architecture that places the correlation model as one layer among four.

The portfolio-margin, liquidation-waterfall, fund-reservation, and haircut-scheduling machinery specialises classical CCP practice (CME-SPAN style scenario margining [7], LCH/DTCC waterfall design, Eisenberg-Noe network clearing [1], Expected Shortfall in the Acerbi-Tasche formulation [44]) to AMM pool accounting, where the price impact of liquidations is a deterministic function of pool reserves. We prove existence and uniqueness of the concurrent clearing vector with capped socialisation and auto-deleveraging (Theorem 7.5, extending Eisenberg-Noe), and order-independence of the five-stage waterfall under simultaneous defaults (Theorem 7.6).

Four results are new.

(i) *AMM fire-sale contagion bound (Theorem 9.1).* The indirect-contagion framework of Cont-Wagalath [4] adapted to an AMM-price-impact kernel yields a deterministic bound on cross-pool loss amplification. For a hub-spoke topology with N spokes and hub concentration w^{hub}, the first-round contagion loss is bounded by a closed-form function of the initial liquidation volume, pool depths, and a correlation norm.

(ii) *Reflexive cascade ODE and slow-fast classification (Theorems 10.3, 11.3, Proposition 11.4, Theorem 10.4).* When the hub asset's value has an endogenous component $f(L)$ depending on locked liquidity L, the coupled price/liquidity dynamics admit a reflexivity coefficient $\varrho = \tau\lambda L f'(L)/P_M$. The locus $\{\varrho = 1\}$ is not a Sotomayor fold bifurcation: the zero eigenvalue of the Jacobian is structural on the equilibrium curve rather than parameter-triggered. The locus is, instead, the fold curve of the singular-perturbation parameter $\varepsilon = 1 - \varrho$ in a Fenichel slow-fast system [18, 19]. Trajectories entering this locus from the subcritical side encounter a canard-type blow-up rather than a structurally stable saddle-node.

(iii) *Exogenous-settlement severance (Theorem 12.2).* When margin requirements, collateral valuation, and the insurance fund are denominated in an asset whose price is exogenous to the locked-liquidity process, the first-order clearing-layer reflexivity coefficient vanishes. Residual contagion reduces to spoke-asset correlation, settlement-asset depeg (bounded per-asset under independent depegs; characterised by a copula upper bound under correlated stress), a second-order tail-copula channel when the settlement-asset return and the hub-asset return share right-tail dependence, and governance risk. We derive closed-form sufficient bounds for the gate thresholds X, Y and the decorrelation lag Δ_{lag} that preclude a reflexive spiral in the counter-cyclical buyback mechanism.

(iv) *Continuous-state-price-density Expected Shortfall and learned-correlation identifiability (Theorems 4.3, 5.3).* When the clearing venue produces a consensus state-price density μ_t over a bounded outcome space Ω, the scenario-ES margin of Definition 3.1 extends to a tail-integral $\mathrm{ES}_q[\Pi] = (1/q) \int_{\Omega_q(\Pi)} \ell(\Pi, \omega)\, \mu_t(d\omega)$ over the tail event Ω_q; quadrature on a fixed grid yields bounded discretisation error. Pairwise correlations entering the parametric ES of (3) are fit as a maximum-entropy Markov random field on the observed fill stream; we state sample-complexity and misspecification-robustness bounds closing the gap between the model's parametric ES and the realised tail loss.

The correlation-driven parametric ES of (3) sits inside a four-layer defense-in-depth: hard per-position margin floors, correlation-adjusted cross-margin (the parametric layer), discrete stress-scenario replay (Definition 8.8), and liquidation circuit-breakers halting clearing when the estimator degrades. This architecture is adopted in explicit response to the Gaussian-copula overreliance pattern diagnosed by Li's (2000) copula [43] in credit-default-swap clearing; no single correlation-model specification is load-bearing for solvency.

The main text is theorem, proof, and limitations. Controller tuning, scenario replay, and empirical parameter calibration are deferred to appendices. The framework assumes a power-weighted constant-function AMM execution layer (the g_i-weighted family of Assumption 2.1, which specialises to Uniswap-v2 at $g_i = 1$ [12, 13]); all statements hold under the assumption model of §2.

Contents

1 Introduction

1.1 Problem

A central counterparty (CCP) clearing trades executed on a hub-and-spoke AMM inherits two structural features absent from traditional CCP settings. First, the price-impact function is deterministic and known to the clearinghouse: when a position is liquidated into an AMM pool with reserves (R_M, R_B), the realised slippage is a closed-form function of volume relative to reserves, rather than a distribution over an unobservable order-book state. Second, if the hub asset's valuation depends partly on liquidity locked in the spokes, the clearing and price-formation layers are coupled: collateral value and liquidity supply co-evolve through a feedback loop with a natural dynamical-systems representation.

Traditional CCP literature treats price impact stochastically (SPAN-style [7] shock scenarios applied to exogenous markets) and treats collateral value as independent of clearing decisions (a reasonable approximation when clearing volume is small relative to the external market for the collateral asset). Neither assumption holds in our setting. The framework of Cont and Wagalath [4] for indirect contagion via fire-sale externalities, and of Eisenberg and Noe [1] for network clearing vectors, provide the appropriate starting points, but require specialisation to the AMM-price-impact kernel and extension to cover the reflexive feedback channel.

1.2 Contributions and adapted material

The paper is organised around the distinction between adapted and original content.

Adapted CCP machinery (§§2-13). Portfolio margin with correlation-aware scenario aggregation (§3), the multi-stage default waterfall as a clearing operator (§8), and the haircut-based multi-settlement-asset model (§13) specialise well-established CCP designs [7, 8, 6] to the AMM-price-impact setting. The ES measure follows Acerbi-Tasche [44]. We prove the concurrent-clearing-vector extension of Eisenberg-Noe with capped socialisation and auto-deleveraging (Theorem 7.5), and the order-independence of the five-stage execution under simultaneous defaults (Theorem 7.6). Both are new formal statements; the underlying design pattern is standard CCP practice.

Original contributions. Four results are novel.

Result 1 (§9): AMM fire-sale contagion bound. Theorem 9.1 gives a deterministic upper bound on first-round indirect-contagion losses in an N-spoke hub-and-spoke AMM, as a function of liquidation volumes, pool depths, and a correlation norm on the spoke-asset covariance. This is the AMM analogue of the Cont-Wagalath [4] fire-sale bound. The AMM structure permits a deterministic bound where the Cont-Wagalath original is stochastic.

Result 2 (§§10-11): reflexivity ODE and slow-fast singular classification. We derive the coupled price/liquidity ODE (42), introduce the reflexivity coefficient ϱ, and resolve the bifurcation type at $\{\varrho = 1\}$. Theorem 11.3 proves that $\{\varrho = 1\}$ is not a Sotomayor fold because condition (S1) fails: the zero eigenvalue of the desingularised Jacobian is present at every equilibrium on the equilibrium curve $\mathcal{E}$, rather than only at a bifurcation parameter. Proposition 11.4 identifies $\{\varrho = 1\}$ as the fold curve of the singular-perturbation parameter $\varepsilon = 1 - \varrho$ in a Fenichel slow-fast system [18], producing a canard-type blow-up [20, 21]. Theorem 10.4 proves global convergence of cascade trajectories to $(\bar{V}, 0)$ via an orbit integral and a strict Lyapunov function, closing the classical Poincaré-Bendixson hypotheses.

Result 3 (§12): exogenous-settlement severance. Theorem 12.2 proves that if the clearing layer (margin, collateral, insurance fund) is denominated in an asset exogenous to the locked-liquidity process, the first-order reflexivity coefficient of the clearing layer vanishes, independent of the hub asset's reflexivity. Corollary 12.3 translates this into a bounded-contagion theorem: the worst-case loss under any hub-asset trajectory is bounded above by the sum of the insurance fund, capped socialisation capacity, and auto-deleveraging capacity, plus a tail-copula term for correlated depegs of the settlement asset. We give closed-form sufficient bounds on the counter-cyclical buyback gate thresholds X, Y and the decorrelation lag Δ_{lag} (Propositions 12.7-12.8-12.10) as monotone functions of the haircut envelope, observation noise, and post-buyback stress horizon.

Result 4 (§4, §5, §6): continuous-density ES, learned correlations, and compliance-tier recalculation. Theorem 4.3 extends the scenario ES of Definition 3.1 to a tail-integral over a continuous state-price density produced natively by a convex-potential clearing layer; Proposition 4.5 gives a quadrature-error bound. Theorem 5.3 fits the correlation input to the parametric ES as a pairwise maximum-entropy Markov random field on the observed fill stream, states $O(p^2 \log p/\varepsilon^2)$ sample complexity for p-node consistency (Wainwright-Jordan [49]), and bounds the margin gap under misspecification. Definition 6.3 specifies the compliance-tier-triggered margin recalculation protocol when an instrument transitions between structurally distinct tiers (e.g., compliant-for-all-investors versus compliant-for-a-restricted-pool), and Proposition 6.4 bounds the instantaneous margin delta. Together, these close the gap between the correlation-model-backed parametric ES of §3 and the realised tail-loss distribution that the defense-in-depth architecture of §1.3 must contain.

1.3 Defense-in-depth against correlation-model failure

The parametric ES of (3) is a correlation-model output. A clearinghouse that relies on any single correlation specification as its sole margining input inherits the failure mode of credit-default-swap clearing prior to 2008: Gaussian-copula overreliance [43, 10, 53] produced margin levels insufficient to absorb the 2008 tail realisation because the copula's tail-dependence parameters were estimated from a calibration sample that did not contain the tail event. The clearinghouse's exposure was bounded above by a correlation model whose error exceeded its estimated value under the realised shock.

This paper architects its margin framework against that failure mode. The parametric ES enters as *one layer* of a four-layer stack; no single layer is load-bearing for solvency.

Definition 1.1 (Margin defense-in-depth stack)**.** The CCP's required collateral for position Π satisfies

$$M_{\text{req}}(\Pi) \;:=\; \max\{M^{\text{floor}}(\Pi),\, M^{\text{param}}(\Pi),\, M^{\text{stress}}(\Pi)\} \cdot \mathbb{1}_{\mathcal{H}(\Pi)} \;+\; \infty \cdot \mathbb{1}_{\neg\mathcal{H}(\Pi)}, \tag{1}$$

where:

(a) $M^{\text{floor}}(\Pi) := \sum_i |Q_i| \cdot P_i \cdot \mu_i^{\text{floor}}$ is the *hard per-position floor* with μ_i^{floor} the per-instrument minimum margin rate (structural constant independent of any correlation estimate);

(b) $M^{\text{param}}(\Pi)$ is the correlation-adjusted parametric ES of Definition 3.1, fit on live fills per Theorem 5.3;

(c) $M^{\text{stress}}(\Pi)$ is the scenario-maximum over the stress catalogue $\Omega_0 \cup \Omega_{\text{MC}}$ of Definition 8.8;

(d) $\mathcal{H}(\Pi) = \{\text{circuit-breakers all nominal}\}$ is the health predicate defined by (i) the MRF fit's condition number below a threshold, (ii) the realised short-window tail loss not exceeding ES_q by more than a tolerance factor, (iii) the state-price-density quadrature error bound of Proposition 4.5 holding, and (iv) no compliance-tier transition in flight on any leg. $\neg\mathcal{H}(\Pi)$ halts opening on Π until the predicate is restored; existing positions are re-margined at the stress layer alone.

Proposition 1.2 (Layer monotonicity)**.** *$M_{\text{req}}(\Pi)$ is monotone non-decreasing in each argument on the right-hand side of (1), and satisfies $M_{\text{req}}(\Pi) \geq M^{\text{floor}}(\Pi)$ always. Under a correlation-model breakdown in which M^{param} understates the realised tail loss by factor $\gamma \geq 1$, the clearinghouse's collateral shortfall on Π is bounded by $(\gamma - 1) \cdot M^{\text{param}}(\Pi) - \big(M^{\text{stress}}(\Pi) - M^{\text{param}}(\Pi)\big)_+$ - the stress layer absorbs the gap up to $M^{\text{stress}} - M^{\text{param}}$.*

Proof. Monotonicity is from the max in (1). When M^{param} understates by factor γ, the realised loss is $\gamma \cdot M^{\text{param}}$; the collateral posted is $\max(M^{\text{floor}}, M^{\text{param}}, M^{\text{stress}})$. If $M^{\text{stress}} \geq \gamma M^{\text{param}}$, the shortfall is zero. Otherwise, the shortfall is $\gamma M^{\text{param}} - \max(M^{\text{floor}}, M^{\text{param}}, M^{\text{stress}})$, which equals $(\gamma - 1)M^{\text{param}} - (M^{\text{stress}} - M^{\text{param}})_+$ when $M^{\text{stress}} \geq M^{\text{floor}}$ and $M^{\text{param}} \geq M^{\text{floor}}$. □

Remark 1.3 (Historical referent)**.** Li's 2000 Gaussian-copula model [43] supplied the correlation-aggregation step in CDO tranche pricing and CDS portfolio margin in the 2001-2008 period. The model's parameters were calibrated on pre-2007 default-correlation samples containing no systemic correlated-default event. When the 2008 housing-default correlation realised at the model's tail, the margin and capital provisions proved insufficient for several dealers; the failure is documented in [53, 54]. The specific failure mechanism was model-reliance, not copula choice: any single-specification correlation aggregation calibrated on a pre-shock sample would have exhibited the same tail gap. The defense-in-depth architecture of Definition 1.1 is adopted in direct response: the parametric ES is allowed to enter as a correlation-aware refinement of the floor and stress layers, never as a standalone solvency boundary.

Remark 1.4 (Relation to PFMI Principle 4)**.** The stress layer M^{stress} of Definition 1.1 corresponds to PFMI Principle 4's "extreme but plausible" scenario set [35], sized at Cover-1 or Cover-2 per Definition 7.2. The floor layer M^{floor} corresponds to margin-minimum schedules required by EMIR Article 41 [36]. The parametric layer M^{param} corresponds to the CPMI-IOSCO "model-based initial margin" discipline and is subject to the PFMI Principle 4 validation and backtesting requirements. The circuit-breaker health predicate $\mathcal{H}$ is an operational implementation of Principle 4's model-validation discipline.

1.4 Relation to the literature

CCP risk management. The applied CCP-risk literature decomposes into four strands. On margin methodology and cross-margin efficiency, Cont and Kokholm [29] analyse multilateral versus bilateral netting in a multi-CCP setting; their margin-efficiency theorem is the benchmark against which Theorem 3.2 and Proposition 3.6 below should be read. Cruz Lopez, Harris, Hurlin, and Pérignon [32] develop CoMargin for correlation-aware initial margin, a close cousin of the parametric ES aggregation of (3). Ghamami and Glasserman [33] characterise initial-margin efficiency under stressed versus non-stressed regimes. On default-fund sizing, skin-in-the-game, and waterfall design, Duffie and Zhu [30] give the foundational comparison of bilateral and centrally-cleared counterparty risk; Murphy [31] develops the canonical Cover-1 and Cover-2 default-fund sizing rules and the skin-in-the-game tranche between defaulter margin and mutualisation; Biais, Heider, and Hoerova [34] analyse incentive alignment between the CCP and its clearing members. The regulatory surface is the CPMI-IOSCO *Principles for Financial Market Infrastructures* [35] (Principle 4 on credit risk and Principle 7 on liquidity risk), EMIR Articles 43, 44, and 45 [36] (default-fund sizing, skin-in-the-game, and the default-management waterfall), and Dodd-Frank Title VIII [37] (U.S. Financial Market Utility designation and risk-management standards). Earlier theoretical treatments by Acharya and Bisin [5] and Pirrong [6] motivate the central-clearing arrangement. The SPAN family [7] is the industry benchmark for scenario margining. We take it as the reference against which our AMM-kernel margin is an in-topology specialisation.

The positioning of §§7-8 against this literature is as follows. Theorem 7.5 extends Eisenberg-Noe [1, 2] with a capped-socialisation stage and an auto-deleveraging stage that consume profit obligations rather than debt, the structure appropriate for derivative CCPs. Definition 7.3 maps onto the EMIR Art. 45 stages once the skin-in-the-game tranche of Definition 7.1 is inserted between defaulter seizure and fund draw; Definition 7.2 gives the Cover-k fund-sizing rule of PFMI Principle 4 in the AMM-CCP setting. Porting of client positions (EMIR Art. 48; PFMI Principle 13) receives the treatment of §8.2, where the AMM structure makes a client position a pool-share claim that can be transferred granularly across surviving counterparties, altering the porting-versus-liquidation trade-off.

AMM risk and LP economics. The AMM-risk literature characterises LP exposure in two complementary ways. Angeris, Kao, Chiang, Noyes, and Chitra [13] analyse Uniswap-style constant-product markets and introduce the arbitrage-equilibrium argument underlying deterministic price impact. Evans, Angeris, and Chitra [38] derive optimal fees for geometric-mean market makers, compensating LPs for the expected loss against a rebalancing benchmark. Milionis, Moallemi, Roughgarden, and Zhang [39] introduce loss-versus-rebalancing (LVR), a continuous-time baseline for LP losses that equals $\sigma^2/8$ per unit time per unit LP value in the constant-product setting and generalises to the g_i-weighted family through the weighted Black-Scholes replication of the trade-free LP payoff. The LP-as-loss-absorption-layer theorem (Theorem 9.4 in §9) integrates these strands into the CCP waterfall: the slippage kernel of Theorem 9.1 integrated over a liquidation event, net of fee revenue, is a continuous-time mutualised default-fund contribution borne by LPs proportionally to pool-share ownership. Section 3 also integrates LVR as a margin input when pool-share positions serve as collateral.

Network clearing and fire-sale contagion. The network-clearing literature begins with Eisenberg-Noe [1] and has been extended to include bankruptcy costs [2], fire-sale externalities [3], and indirect contagion [4]. Allen and Gale [40] develop the correlated-failure account of financial contagion that underlies the second-order copula remark of Corollary 12.4. Theorem 9.1 specialises Cont-Wagalath to a deterministic AMM-price-impact kernel and obtains a closed-form bound in place of the stochastic one.

Reflexivity and slow-fast classification. The reflexivity and death-spiral analysis of algorithmic stablecoins [14, 16, 15, 17] is qualitative; no prior work carries out the formal bifurcation classification. Our contribution is to place the threshold within the taxonomy of dynamical-systems bifurcations and to identify it correctly as a canard of a singular-perturbation parameter rather than as a fold. A Sotomayor fold would imply a structurally stable two-sided equilibrium picture near the threshold; the canard classification implies instead that trajectories crossing the threshold exhibit time-scale separation breakdown and the scalar reduction fails. The slow-fast dynamical-systems apparatus [18, 19, 20, 21, 22, 23] is standard; its application to reflexivity in AMM-coupled clearing systems is the contribution of §11.

1.5 Organization

Section 2 fixes notation and states the assumption model. Section 3 develops AMM-aware portfolio margin (adapted). Sections 7-8 prove the clearing-operator existence and order-independence theorems. Section 9 proves the AMM fire-sale contagion bound. Section 10 derives the reflexive cascade ODE and Theorem 10.3. Section 11 classifies the singular barrier (Theorem 11.3 and Proposition 11.4) and proves global attraction (Theorem 10.4). Section 12 proves exogenous-settlement severance. Section 13 develops the multi-settlement-asset haircut model and proves the stablecoin-agnostic bounded-contagion theorem. Section 15 lists open problems. Appendices cover controller tuning, scenario replay, and parameter tables.

Notation. Throughout, $P_M > 0$ denotes the price of a hub asset M; $L \geq 0$ denotes aggregate locked liquidity across N spoke pools; $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a C^1 increasing concave function with $f(0) = 0$ (the endogenous value function); $\bar{V} > 0$ is the exogenous mean-reversion target; $\theta > 0$ is the exogenous mean-reversion rate; $\lambda > 0$ is an LP sensitivity rate (dimension 1/time); $\tau > 0$ is a reference timescale (canonical choice $\tau = 1/\theta$). Reserves of pool i are $(R_M^{(i)}, R_B^{(i)})$; $\pi \in [0, 1]$ denotes an event-outcome probability; ES_q is Expected Shortfall at confidence q.

2 Setup and Assumption Model

2.1 The AMM-CCP market model

We fix the following structure. A central counterparty clears trades for N *spoke pools*, each trading a *spoke asset* B_i against a common *hub asset* M. Pool i has reserves $(R_M^{(i)}, R_B^{(i)})$ and a constant-function invariant $R_M^{(i)} \cdot (R_B^{(i)})^{g_i} = k_i$ with coupling exponent $g_i \geq 1$ (the $g_i = 1$ case recovers constant-product markets [12, 13]; general g_i gives the power-weighted family). The marginal price is $P_i = g_i R_M^{(i)} / R_B^{(i)}$.

Assumption 2.1 (Pricing and impact model)**.**

1. (AMM invariant.) Each pool satisfies $R_M^{(i)} (R_B^{(i)})^{g_i} = k_i$ with $g_i \geq 1$ and $k_i > 0$.
2. (Deterministic price impact.) Selling $q > 0$ units of B_i into pool i yields hub proceeds $\Delta_M(q; i) = k_i[(R_B^{(i)})^{-g_i} - (R_B^{(i)} + q)^{-g_i}]$ and post-trade pool reserves $(R_M^{(i)} - \Delta_M(q; i),\ R_B^{(i)} + q)$. This is exact and known to the clearinghouse.
3. (Linear-in-log return model.) Over a short horizon, spoke-asset log-returns admit a factor decomposition $r_i = \beta_i F + \varepsilon_i$ with common factor F and residual ε_i independent of F. Portfolio-margin formulas of §3 are exact under this model; elliptical extensions follow from scale invariance.

Assumption 2.2 (Clearinghouse policy)**.**

1. (Margin measure.) The clearinghouse uses scenario-based Expected Shortfall at confidence $q \in (0, 0.01]$ applied to a portfolio P&L function that sums spot, derivative, and event-linked instrument P&L within each stress scenario.
2. (Insurance fund.) An insurance fund $\mathrm{IF} \geq 0$ is held in a specified settlement asset (§12 treats the endogenous and exogenous cases separately).
3. (Waterfall.) On a default event with positive gross loss, the clearinghouse executes the generalized waterfall of §8: (1) defaulter margin seizure, (2) optional CCP skin-in-the-game layer (Definition 7.1), (3) insurance-fund draw (sized per Definition 7.2), (4) capped socialisation with per-position cap $\kappa \in (0, 1)$, (5) auto-deleveraging (ADL) of profitable counterparties. The skinless case $E_{\mathrm{SIG}} = 0$ collapses this to the four-stage operational waterfall used by the protocol default.

Assumption 2.3 (Endogenous-value structure (used only in §§10-12))**.** The hub asset's price decomposes as $P_M(t) = V_{\mathrm{ext}}(t) + V_{\mathrm{end}}(t)$ where:

- V_{ext} is exogenous and mean-reverts at rate $\theta > 0$ to $\bar{V} > 0$, independently of L;
- $V_{\mathrm{end}} = f(L)$ with f C^1, strictly increasing, concave, $f(0) = 0$. Canonical: $f(L) = \beta_L L^\gamma$ with $\gamma \in (0, 1]$.

When collateral and margin are denominated in an *exogenous* asset (Section 12), Assumption 2.3 does not propagate into the clearing layer; this is precisely the content of Theorem 12.2.

Remark 2.4 (Scope of the model)**.** Assumption 2.1 fixes the AMM kernel; Assumption 2.2 fixes the CCP policy surface; Assumption 2.3 enters only when we study the reflexive channel. Results that do not invoke Assumption 2.3 (all of §§3-9) hold regardless of hub-asset valuation structure.

3 Portfolio Margin under the AMM Kernel

We develop the portfolio-margin machinery that the CCP uses to compute required capital. The architecture is adapted from CME-SPAN scenario margining [7] and the parametric Expected Shortfall aggregation of [8], specialised to the AMM-price-impact kernel.

3.1 Scenario-based Expected Shortfall

Definition 3.1 (Portfolio margin)**.** Let Π be a portfolio and $\{\omega_s\}_{s=1}^{S}$ a finite set of stress scenarios, each specifying price shocks and event-outcome shifts. The *portfolio margin* is

$$M(\Pi) := \mathrm{ES}_q\Big(\{-\mathrm{PnL}(\Pi,\omega_s)\}_{s=1}^{S}\Big)\cdot(1+\zeta), \tag{2}$$

where $q \in (0, 0.01]$ is the confidence level, $\zeta \geq 0$ is a conservatism buffer calibrated to historical stress severity (typical industry range $\zeta \in [0.10, 0.25]$, reported with implementation calibrations in Appendix A), and $\mathrm{PnL}(\Pi,\omega)$ is computed using the AMM kernel of Assumption 2.1. $M(\Pi)$ is monotone non-decreasing in each of the underlying risk factors: a first-order stochastic increase in any scenario's $-\mathrm{PnL}$ raises ES_q, which propagates to $M(\Pi)$ through (2).

Under the Gaussian or elliptical return model of Assumption 2.1(3), the scenario ES formula collapses to the parametric identity [8]

$$M(\Pi)^2 = \mathbf{m}^\top \mathbf{P}\,\mathbf{m}, \qquad \mathbf{P}_{ij} = \rho_{ij}, \tag{3}$$

where m_i is the standalone ES of leg i and $\mathbf{P}$ is the return-correlation matrix.

Theorem 3.2 (Capital-efficiency ratio under a general factor model)**.** *Under the factor decomposition $r_i = \beta_i F + \varepsilon_i$ of Assumption 2.1(3) and standalone margins $m_i > 0$, the portfolio margin satisfies*

$$M^{\mathrm{port}}(\Pi)^2 \;=\; \Big(\textstyle\sum_i m_i\beta_i\Big)^2 + \textstyle\sum_i m_i^2(1-\beta_i^2). \tag{4}$$

The efficiency ratio $\eta(\Pi) := \sum_i |m_i| / M^{\mathrm{port}}(\Pi)$ is bounded by

$$\eta(\Pi) \;\leq\; \sqrt{\frac{N}{1+(N-1)\bar\rho}}, \qquad \bar\rho := \mathrm{avg}_{i\neq j}\rho_{ij}, \tag{5}$$

with equality under equal β_i and equal $|m_i|$.

Proof. Formula (4) is the quadratic form $\mathbf{m}^\top(\boldsymbol{\beta}\boldsymbol{\beta}^\top + \mathbf{D})\mathbf{m}$ for the factor-model correlation matrix $\mathbf{P} = \boldsymbol{\beta}\boldsymbol{\beta}^\top + \mathrm{diag}(1-\beta_i^2)$, which is PSD by construction. Equation (5) is the standard equicorrelation bound: under $\rho_{ij} = \bar\rho$ and $|m_i| = m$, $M^{\mathrm{port}} = m\sqrt{N(1+(N-1)\bar\rho)}$ and $\sum_i |m_i| = Nm$, so $\eta = \sqrt{N/(1+(N-1)\bar\rho)}$. Cauchy-Schwarz extends the bound to the heterogeneous case. □

Corollary 3.3 (Asymptotic ceiling)**.** *As $N \to \infty$ with $\bar\rho > 0$ fixed, $\eta \to 1/\sqrt{\bar\rho}$. Portfolio margin provides a constant-factor capital saving, not an asymptotically growing one.*

3.2 PSD projection for structural priors

Sample correlations converge slowly for new spoke markets [9]. CCPs mitigate this with structural priors (e.g., sector-based correlation matrices). Blending requires that both matrices be PSD to preserve the PSD structure of the blend.

Proposition 3.4 (PSD shrinkage)**.** *Let* $\hat{\mathbf{P}}$ *be a Pearson sample correlation matrix (PSD by construction), and* $\mathbf{P}^{\text{prior}}$ *a structural prior matrix with PSD projection* $\mathbf{P}^{\text{prior},+}$ *(eigendecomposition with negative eigenvalues floored at* $\epsilon > 0$*). For any* $w_s \in [0, 1]$*, the blended matrix*

$$\mathbf{P} := (1 - w_s)\,\hat{\mathbf{P}} + w_s\,\mathbf{P}^{\text{prior},+} \tag{6}$$

is PSD.

Proof. The PSD cone is convex. Both inputs are PSD; the convex combination inherits PSD. The projection step is required: a naive prior built from pairwise structural correlations need not be PSD (e.g., $(\rho_{12}, \rho_{13}, \rho_{23}) = (0.9, 0.9, -0.9)$ has minimum eigenvalue ≈ -0.80), so the projection step is necessary, not optional. □

Interpretation. Portfolio margin buys a constant factor $1/\sqrt{\bar{\rho}}$ over standard margin, and the advantage does not grow with N. At $\bar{\rho} = 0.2$ the ceiling is $\eta \to 1/\sqrt{0.2} \approx 2.24$: portfolio margin halves the capital requirement. The efficiency gain is a constant factor, not an asymptotic one.

3.3 Cross-margining specification: spot, perpetual, and event positions

The portfolio margin framework above treats positions abstractly. In practice the clearinghouse cross-margins across three instrument types (spot, perpetual, and event-linked positions), each with distinct risk profile and hedging relationship to the spoke asset. All offsets below are inputs to one scenario-ES generator; they are not independent coupons that may be summed without a positive-semidefinite risk model.

Definition 3.5 (Unified cross-margin computation)**.** Let a portfolio Π contain spot positions $\{(B_i, Q_i^s)\}$, perpetual positions $\{(B_i, Q_i^p)\}$, and event-linked positions $\{(E_j, S_j)\}$. The *isolated margin* treats each position independently:

$$m^{\text{iso}} = \sum_i |Q_i^s| \cdot P_i \cdot \mu_i^s \;+\; \sum_i |Q_i^p| \cdot P_i \cdot \mu_i^p \;+\; \sum_j S_j \cdot \mu_j^E \tag{7}$$

where $\mu_i^s, \mu_i^p, \mu_j^E$ are the respective margin rates (representative industry values: $\mu_i^s = 0.10$, $\mu_i^p = 0.05$, $\mu_j^E = 1.0$ for event stakes). The *cross-margin* applies three offsets.

(i) Spot/perpetual hedge discount. When a trader holds offsetting spot and perpetual positions in the same spoke asset, the overlapping notional receives a basis-risk-adjusted margin discount:

$$D_i^{sp} = d_{\text{sp}}(b_i) \times \min(|Q_i^s| \cdot P_i,\; |Q_i^p| \cdot P_i) \times \mu_i^p \times \mathbf{1}[\text{sgn}(Q_i^s) \neq \text{sgn}(Q_i^p)] \tag{8}$$

where the indicator ensures the discount applies only to opposing directions (long spot / short perp, or vice versa), and the discount rate is a function of the observed basis volatility b_i:

$$d_{\text{sp}}(b_i) = \max\left(0.50,\;\; 1 - \frac{\text{ES}_{0.995}[\,|P_i^{\text{spot}} - P_i^{\text{perp}}|\,]}{\mu_i^p \cdot P_i}\right). \tag{9}$$

A default value applies when insufficient basis data exists; the floor ensures the hedge always receives some margin benefit.

Equation (9) selects the discount so that the retained margin covers the estimated basis ES. A static discount d combined with a realised tail basis b_q at tail quantile q leaves residual exposure

$$L_{\text{basis}}^{\max} = \left(b_q - (1-d)\cdot \mu_i^p\right)_+ \times N, \tag{10}$$

which vanishes under (9) by construction.

(ii) Event/spot correlation offset. Event-linked positions whose outcome correlates with spoke asset i's spot price (through a coupling $G(E) = E^\kappa$) earn a margin offset of 40% of the lesser notional:

$$D_i^{pE} = 0.40 \times \min(|Q_i^s| \cdot P_i,\ S_i) \times \mu_i^s, \tag{11}$$

where S_i is the event-linked stake on spoke asset i's outcomes. The coefficient 0.40 is an admissible cap on one covariance term, not an additional discount if the same risk vector was already credited by the spot/perpetual hedge.

(iii) Margin floor. To prevent zero-margin portfolios even under perfect hedges:

$$m^{\text{floor}} = 0.05 \times m^{\text{iso}}. \tag{12}$$

Let $D^{\text{PSD}}(\Pi)$ be the total hedge credit produced by a single scenario-ES or PSD covariance generator that contains the spot/perpetual and event/spot risk factors once each. The unified cross-margin is:

$$m^{\text{cross}} = \max\left(m^{\text{iso}} - D^{\text{PSD}}(\Pi),\ \ m^{\text{floor}}\right), \tag{13}$$

with $0 \le D^{\text{PSD}}(\Pi) \le \sum_i D_i^{sp} + \sum_i D_i^{pE}$ and with each covariance term counted once.

Proposition 3.6 (Capital efficiency of cross-margining)**.** *Under the factor model of Assumption 2.1(3) and the cross-margin formula of Definition 3.5, for a balanced portfolio (equal per-spoke notional across spot, perpetual, and event positions) over N_{br} spoke-asset clusters, the capital-reduction ratio admits the closed form*

$$\frac{m^{\text{cross}}}{m^{\text{iso}}} = \frac{m^{(1)}\sqrt{N_{\text{br}}(1+(N_{\text{br}}-1)\bar{\rho}_{\text{eff}})}}{N_{\text{br}}\, m^{\text{iso}}_{\text{per spoke}}}, \qquad m^{(1)} := m^{\text{iso}}_{\text{per spoke}} - D^{sp} - D^{pE}, \tag{14}$$

whenever the per-spoke floor (12) is non-binding, where $\bar{\rho}_{\text{eff}} := \bar{\rho} + \delta_\rho$ is the shock-adjusted equicorrelation. The ratio is monotone decreasing in each of the hedge-offset rates (D^{sp}, D^{pE}) and monotone increasing in $\bar{\rho}_{\text{eff}}$ and N_{br}. No calibration assumption enters (14) beyond equal per-spoke notionals and a non-binding floor.

Proof. With per-instrument notional N and rates (μ^s, μ^p, μ^E), the per-spoke isolated margin is $m^{\text{iso}}_{\text{per spoke}} = N(\mu^s + \mu^p + \mu^E)$ and $m^{(1)} = m^{\text{iso}}_{\text{per spoke}} - D^{sp} - D^{pE}$ by (8)-(11). The isolated total is $N_{\text{br}} \cdot m^{\text{iso}}_{\text{per spoke}}$. Under equicorrelation $\bar{\rho}_{\text{eff}}$ with per-spoke margin $m^{(1)}$, Theorem 3.2 yields $m^{\text{port}}_{\text{total}} = m^{(1)}\sqrt{N_{\text{br}}(1+(N_{\text{br}}-1)\bar{\rho}_{\text{eff}})}$, and (14) follows. Monotonicity in $\bar{\rho}_{\text{eff}}$ and N_{br} is immediate from the square root; monotonicity in (D^{sp}, D^{pE}) follows from the linear dependence of $m^{(1)}$ on these variables. □

Remark 3.7 (Sensitivity surface, not a point estimate)**.** Equation (14) is a closed form in $(N_{\text{br}}, \bar{\rho}, \delta_\rho, D^{sp}, D^{pE})$. The paper's formal content is the functional form and its monotonicities; specific numerical evaluations at representative parameter choices appear in Appendix A.

Remark 3.8 (SPAN as benchmark)**.** CME-SPAN [7] is the industry scenario-margin engine for listed derivatives, serving in the Cont-Kokholm [29] framework as the benchmark against which the parametric ES formula of (3) is evaluated. The cross-margin formula of Definition 3.5 is an in-topology specialisation of SPAN scenario aggregation to the spoke-AMM kernel; the correlation credits in (14) are of the same functional form as SPAN inter-spread credits. A benchmark against published SPAN tables for a reference book is an empirical calibration, deferred to Appendix A.

Remark 3.9 (Design rationale for offset rates)**.** The spot-perp offset d_{sp} is the complement of the basis-risk ES reserved by the margin engine; the monotonic dependence on observed basis volatility in (9) is a classical risk-of-ruin argument. The prediction-spot offset rate in (11) is set below the regression beta of spot against event probability to accommodate coupling-model misspecification. The floor in (12) reserves capital against residual operational risks orthogonal to the hedge.

Remark 3.10 (Hedge-leg removal and re-margining)**.** The offset D_i^{sp} in (8) depends on the indicator $\mathbf{1}[\text{sgn}(Q_i^s) \neq \text{sgn}(Q_i^p)]$. If margin is recomputed only at block boundaries, an adversary can close one leg within a block and retain the discount until the next block. The AMM-deterministic price-impact kernel permits deterministic intra-block re-margining (Remark 14.13), making the indicator update synchronous with the leg removal and eliminating the exploit.

Remark 3.11 (Synthetic-directional decorrelation under events)**.** Portfolios engineered to appear hedged under the pre-event correlation estimate may decorrelate after the event, leaving the clearinghouse under-margined by (61). Setting the margin-engine correlation input to $\max(\rho^{\text{pre}}, \rho^{\text{stress}})$ for a stress floor $\rho^{\text{stress}} \geq \rho^{\text{post}}$ bounds the exposure to pre-event correlation overstatement; forward-looking event-aware correlation updates close the residual gap.

3.4 Impermanent loss and LVR as margin inputs

When LP pool-share positions serve as collateral at the CCP, the collateral value drains continuously under ambient volatility at the loss-versus-rebalancing (LVR) rate of Milionis-Moallemi-Roughgarden-Zhang [39] and jumps downward under liquidation events. The margin computation of §3 must account for both.

Definition 3.12 (LP-collateral process)**.** Let $V_k^{\text{LP}}(t) := \sum_i \alpha_{k,i}\big(R_M^{(i)}(t) + R_B^{(i)}(t) \cdot P_i(t)\big)$ be LP k's pool-denominated collateral value at time t, where $\alpha_{k,i}$ is k's share of pool i. The collateral available to the CCP is $C_k^{\text{LP}}(t) = (1 - h_{\text{LP}}) \cdot V_k^{\text{LP}}(t)$ where $h_{\text{LP}} \in [0, 1)$ is an LP-share haircut.

Proposition 3.13 (LVR-induced drain on LP collateral)**.** *Under the constant-product invariant ($g_i = 1$) with hub-asset volatility σ_M, the expected drain rate on V_k^{LP} is*

$$\mathbb{E}\big[\dot{V}_k^{\text{LP}}(t)\big] \;=\; -\frac{\sigma_M^2}{8} \cdot V_k^{\text{LP}}(t). \tag{15}$$

For general $g_i \geq 1$, the drain rate is $-(\sigma_M^2/2) \cdot G(g_i) \cdot V_k^{\text{LP}}(t)$ with $G(g_i) = g_i/(4(1 + g_i))$, recovering $G(1) = 1/8$ at constant product.

Proof. By [39], the LVR of a constant-product pool against an exogenous-price rebalancer is $\sigma_M^2/8$ per unit time per unit LP value. For the g_i-weighted invariant $R_M^{g_i^{-1}} R_B = k_i^{1/g_i}$, the same argument applied to the weighted Black-Scholes replication of the trade-free LP payoff yields $G(g_i) = g_i/(4(1 + g_i))$. □

Corollary 3.14 (LVR-adjusted margin)**.** *A CCP accepting LP collateral with margin horizon T_m must augment the portfolio margin of Definition 3.1 by a time-decay factor*

$$M^{\text{LVR}}(\Pi) \;=\; M(\Pi) \cdot \Big(1 + \tfrac{\sigma_M^2}{8} \cdot T_m\Big) \tag{16}$$

(constant product; the g_i-weighted analogue uses $G(g_i)$ in place of $1/8$). When LP shares are not accepted as collateral (hub-asset-denominated collateral only), the LVR drain operates on the pool itself (Theorem 9.4) rather than on the margin ledger, and Corollary 3.14 does not apply.

Proof. The margin buffer must cover the expected collateral drain $\sigma_M^2 T_m/8$ per unit LP value in addition to the scenario-ES of Definition 3.1; the two contributions are additive at leading order under the independent-volatility assumption of Assumption 2.1(3). □

3.5 Scope of the elliptical return model

Remark 3.15 (Scope of Assumption 2.1(3))**.** Theorem 3.2 and (14) are exact under the Gaussian or elliptical factor model. For heavy-tailed return distributions (a Student-t family with finite tail index, or a copula-based aggregation in the sense of Embrechts, McNeil, and Straumann [10]), (3) fails. The ES aggregation then requires either a Student-t ES identity with explicit tail index or a copula-tail bound. The $R_e^2/\sqrt{1-\bar{\rho}}$ structure is preserved up to tail-dependent corrections; extensions to fat-tailed settings are outlined in Open Problem 10.

4 Continuous-Density Expected Shortfall

The scenario ES of Definition 3.1 takes a finite stress-scenario set $\{\omega_s\}_{s=1}^S$ as input. When the clearing layer produces a consensus state-price density μ_t over a bounded outcome space $\Omega \subseteq \mathbb{R}^d$ - as arises from a convex-potential clearing mechanism that solves for market-clearing measures directly, following the potential-functional clearing formulation of Peters and Ye [47] and the measure-valued extension to continuous outcomes - the scenario set is replaced by a continuous distribution, and the ES reduces to a tail integral. This section states the continuous-density analogue of Definition 3.1 and bounds its quadrature-discretisation error.

4.1 Definition and tail-integral form

Definition 4.1 (State-price density)**.** At clearing time t, the *state-price density* is a non-negative measure $\mu_t \in \mathcal{M}_{\geq 0}(\Omega)$ on the bounded outcome space Ω with $\mu_t(\Omega) = 1$. The price of a contingent payoff $\ell : \Omega \to \mathbb{R}$ at clearing is $\langle \ell, \mu_t \rangle := \int_\Omega \ell(\omega)\,\mu_t(d\omega)$. Under the Peters-Ye [47] convex-potential formulation, μ_t is the unique minimiser of a strictly convex potential functional subject to market-clearing constraints; its existence and uniqueness are standard [52, 47].

Definition 4.2 (Continuous-density portfolio margin)**.** Let Π be a portfolio with realised loss function $\ell_\Pi : \Omega \to \mathbb{R}$, with $\ell_\Pi(\omega) = -\mathrm{PnL}(\Pi, \omega)$. The *continuous-density Expected Shortfall* of Π at confidence $q \in (0,1)$ is

$$\mathrm{ES}_q^{\mu_t}[\Pi] := \frac{1}{q}\int_{\Omega_q(\Pi)} \ell_\Pi(\omega)\,\mu_t(d\omega), \qquad \Omega_q(\Pi) := \{\omega \in \Omega : \ell_\Pi(\omega) \geq \mathrm{VaR}_q^{\mu_t}[\Pi]\}, \quad (17)$$

where $\mathrm{VaR}_q^{\mu_t}[\Pi] := \inf\{v : \mu_t(\{\ell_\Pi \leq v\}) \geq 1-q\}$. The continuous-density portfolio margin is $M^{\mathrm{cont}}(\Pi) := \mathrm{ES}_q^{\mu_t}[\Pi]\cdot(1+\zeta)$ with ζ as in Definition 3.1.

Theorem 4.3 (Continuous-density ES exists and equals the tail integral)**.** *Let $\mu_t \in \mathcal{M}_{\geq 0}(\Omega)$ be a state-price density with $\mu_t(\Omega) = 1$ on a bounded $\Omega \subseteq \mathbb{R}^d$, and $\ell_\Pi : \Omega \to \mathbb{R}$ a continuous, bounded loss function. Then $\mathrm{ES}_q^{\mu_t}[\Pi]$ is well-defined, finite, and satisfies $\mathrm{ES}_q^{\mu_t}[\Pi] \geq \mathrm{VaR}_q^{\mu_t}[\Pi]$ with equality iff μ_t assigns no mass to $\Omega_q(\Pi)$'s interior. The Acerbi-Tasche [44] coherence properties (subadditivity, monotonicity, translation invariance, positive homogeneity) hold.*

Proof. Boundedness of Ω and continuity of ℓ_Π make ℓ_Π bounded; the integral in (17) exists and is finite. $\mathrm{VaR}_q^{\mu_t}$ is well-defined by the completeness of the distribution function; subadditivity and the other coherence properties transfer directly from the Acerbi-Tasche [44] treatment, which is stated for a general probability measure and applies here with μ_t playing the role of the distribution. □

Remark 4.4 (Breeden-Litzenberger baseline)**.** In a discrete-scenario setting, the state-price density must be reconstructed post-hoc from option prices [52]. The consensus-output setting of Definition 4.1 makes μ_t a direct output of the clearing layer; no Breeden-Litzenberger reconstruction is needed. Theorem 4.3 holds regardless of origin, provided μ_t satisfies the definitional constraints.

4.2 Quadrature error bound

In implementation, $\mathrm{ES}_q^{\mu_t}$ is evaluated on a finite quadrature grid $\{\omega_j\}_{j=1}^J$ with weights $\{w_j\}$ satisfying $\sum_j w_j = 1$. The discretisation error is bounded by standard quadrature theory.

Proposition 4.5 (Quadrature error bound)**.** *Let μ_t be Lipschitz with constant $\mathrm{Lip}(\mu_t) \le L_\mu$ and ℓ_Π be Lipschitz with constant $\mathrm{Lip}(\ell_\Pi) \le L_\ell$. Let $\{\omega_j, w_j\}_{j=1}^J$ be a Gauss-Legendre quadrature on a tensor-product grid over $\Omega \subseteq [a,b]^d$. Write $\hat{\mathrm{ES}}_q^J[\Pi] := (1/q)\sum_{j:\omega_j\in\hat{\Omega}_q(\Pi)} w_j\,\ell_\Pi(\omega_j)$ for the grid estimator of (17), with $\hat{\Omega}_q$ the grid-discretised tail event. Then*

$$\left|\mathrm{ES}_q^{\mu_t}[\Pi] - \hat{\mathrm{ES}}_q^J[\Pi]\right| \;\le\; C_d \cdot \frac{L_\ell \cdot \|\ell_\Pi\|_\infty}{q} \cdot \left(L_\mu + \|\mu_t\|_\infty\right) \cdot J^{-2/d}, \tag{18}$$

for a dimension-dependent constant $C_d > 0$ depending on the smoothness of $\ell_\Pi \mu_t$ restricted to Ω_q.

Proof sketch. Separate the error into two components: (a) the quadrature error of $\int_{\Omega_q} \ell_\Pi\,\mu_t\,d\omega$ on a known domain, and (b) the error from discretising the tail event $\Omega_q \to \hat{\Omega}_q$. (a) is the standard Gauss-Legendre error [51] applied to the Lipschitz product $\ell_\Pi\mu_t$, giving $O(L_\ell L_\mu J^{-2/d})$ on a tensor grid of total size J in dimension d. (b) scales as the quadrature spacing times the Lipschitz constant of ℓ_Π and the density μ_t on the tail boundary, giving $O(L_\ell(L_\mu + \|\mu_t\|_\infty)J^{-1/d})$; the worst of the two components dominates (18) up to the dimension-dependent constant. □

Remark 4.6 (Tail resolution)**.** The $1/q$ factor in (18) reflects the sharp tail-event localisation at small q: resolving the q-tail of a smooth density with fixed error tolerance requires quadrature nodes on the order of $J = O(q^{-d/2})$ concentrated near the tail boundary. Adaptive quadrature or importance-sampled grids relax this cost; the stress catalogue Ω_0 of Definition 8.8 supplies a non-adaptive tail-biased grid.

4.3 Compatibility with the parametric ES

Corollary 4.7 (Continuous-density to parametric reduction)**.** *When μ_t is an elliptical (Gaussian or Student-t) measure specified by mean $\boldsymbol{\mu}$, covariance $\boldsymbol{\Sigma}$, and (for Student-t) tail index ν, and the loss ℓ_Π is linear in the outcome, the continuous-density ES of Theorem 4.3 reduces to the parametric formula (3) exactly for Gaussian and up to a tail-index correction factor for Student-t [10, 8].*

Proof. Direct substitution of the elliptical density into (17) yields the Gaussian/elliptical-ES closed form; the Student-t correction follows from the standard tail-index identity [8]. □

Remark 4.8 (When to use which)**.** The continuous-density ES is the appropriate measure when the clearing layer outputs a full μ_t directly; the parametric ES is the appropriate reduction when μ_t is fit to a factor model with elliptical residuals (Assumption 2.1(3)). Defense-in-depth (Definition 1.1) uses both: the continuous-density evaluation on μ_t drives M^{param} when the quadrature-error bound of Proposition 4.5 holds, and the scenario catalogue drives M^{stress} as an independent cross-check.

4.4 Robustness to state-price-density manipulation

The state-price density μ_t is an output of the clearing layer, and the clearing layer observes trade flow. An adversary controlling a fraction $\alpha \in [0, 1)$ of clearing volume can therefore influence μ_t up to that fraction. We bound the influence on the continuous-density margin.

Proposition 4.9 (Margin sensitivity to adversarial density perturbation)**.** *Let μ_t^* be the true state-price density absent adversarial manipulation, and $\mu_t = (1-\alpha)\mu_t^* + \alpha\,\nu$ the density produced when an adversary controlling fraction α of clearing volume biases the output toward an adversarial measure ν. The continuous-density ES satisfies*

$$\left|\mathrm{ES}_q^{\mu_t}[\Pi] - \mathrm{ES}_q^{\mu_t^*}[\Pi]\right| \;\le\; \frac{\alpha \cdot \|\ell_\Pi\|_\infty}{q}. \tag{19}$$

For $q = 0.01$ (99-ES) and a bounded loss $\|\ell_\Pi\|_\infty \le L$, the per-unit-loss margin sensitivity to adversary-controlled fraction α is $100\alpha L$ - linear in α.

Proof. The ES is a Lipschitz functional of the density with Lipschitz constant $\|\ell_\Pi\|_\infty/q$ in total-variation (TV) distance [44]. $\mathrm{TV}(\mu_t, \mu_t^*) = \alpha \cdot \mathrm{TV}(\mu_t^*, \nu) \le \alpha$. Combining gives (19). □

Remark 4.10 (Defence via volume cap and adversarial-detection circuit breaker)**.** Equation (19) implies: capping the per-account volume share at $\alpha_{\max}$ caps the density-manipulation margin error at $\alpha_{\max}\|\ell_\Pi\|_\infty/q$. When the manipulation-caused margin gap exceeds the stress-layer buffer $M^{\text{stress}} - M^{\text{param}}$, the backtest-rejection circuit breaker (Remark 5.7) drops the parametric layer and the defense-in-depth stack binds at M^{stress}; the density output itself is not a standalone safety surface. The health predicate $\mathcal{H}$ includes the quadrature-error bound of Proposition 4.5 as a validity check on μ_t and drops the parametric layer when the Lipschitz-constant estimate grows anomalously (a proxy for density distortion).

5 Learned Correlations: Pairwise Max-Entropy MRF

The parametric ES of (3) requires a correlation matrix $\mathbf{P}$. The PSD-shrinkage prescription of Proposition 3.4 blends a sample estimate with a structural prior. This section replaces the ad hoc sample-plus-prior construction with a principled estimator: $\mathbf{P}$ is fit on the observed fill stream as a pairwise maximum-entropy Markov random field (MRF), with an identifiability statement, a sample-complexity bound, and a margin-gap bound under misspecification.

Positioning relative to Definition 1.1: the MRF fit is the correlation source *inside* the M^{param} layer. It does not itself provide solvency protection; the stress and floor layers do that. The MRF's role is to make M^{param} a more accurate reflection of realised tail dependence in normal regimes, so that $M^{\text{req}}(\Pi) = \max(M^{\text{floor}}, M^{\text{param}}, M^{\text{stress}})$ binds at M^{param} in the regime where the correlation estimate tracks reality, and binds at M^{stress} or the floor when the estimator degrades (caught by the circuit-breaker health predicate $\mathcal{H}$).

5.1 Model and identifiability

Definition 5.1 (Pairwise max-entropy MRF on fill returns)**.** Let $\{r_i\}_{i=1}^{p}$ denote the return processes of p spoke assets sampled at the per-fill cadence. The *pairwise maximum-entropy MRF* on $(r_1, \ldots, r_p)$ with observed first moments $\boldsymbol{\mu}$ and second moments $\mathbf{S}$ is the distribution

$$p_{\boldsymbol{\theta}}(r) \;=\; Z(\boldsymbol{\theta})^{-1} \exp\big(\boldsymbol{\theta}_1^{\top} r + r^{\top} \boldsymbol{\Theta}_2\, r/2\big), \tag{20}$$

where $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\Theta}_2)$ are the natural parameters (linear and pairwise-quadratic coupling). $\boldsymbol{\theta}$ is identified by the moment-matching condition $\mathbb{E}_{p_{\boldsymbol{\theta}}}[r] = \boldsymbol{\mu}$, $\mathbb{E}_{p_{\boldsymbol{\theta}}}[rr^{\top}] = \mathbf{S}$.

Remark 5.2 (Max-entropy interpretation)**.** Among all distributions on $\mathbb{R}^p$ matching the observed first and second moments, $p_{\boldsymbol{\theta}}$ is the unique maximum-entropy distribution (Jaynes [48]; Wainwright-Jordan [49]). It is honest about what the data pin down: pairwise correlations are the data-constrained content; higher moments are the least-informative completion consistent with that pinning. The MRF formulation is a principled replacement for the "blend sample with structural prior" prescription of Proposition 3.4: the pairwise max-entropy has no hidden structural-prior influence.

Theorem 5.3 (Sample-complexity bound for consistent MRF identification)**.** *Let $\{r^{(t)}\}_{t=1}^{T}$ be an i.i.d. sample of size T from a pairwise MRF on p nodes with true parameter $\boldsymbol{\theta}^*$ and maximum degree d in the induced dependence graph. Let $\hat{\boldsymbol{\theta}}_T$ denote the moment-matching (equivalently, maximum-likelihood) estimator. Then for any $\varepsilon > 0$ and $\delta \in (0,1)$, there exists a constant $C(d)$ depending only on the graph degree such that*

$$T \;\geq\; C(d) \cdot \frac{p^2 \log(p/\delta)}{\varepsilon^2} \quad\Longrightarrow\quad \mathbb{P}\big(\|\hat{\boldsymbol{\theta}}_T - \boldsymbol{\theta}^*\|_2 \leq \varepsilon\big) \;\geq\; 1 - \delta. \tag{21}$$

Under ε-consistent estimation, the covariance-matrix estimation error satisfies $\|\hat{\mathbf{P}}_T - \mathbf{P}^*\|_F \leq C'\,\varepsilon\sqrt{p}$ for a constant C' depending on the condition number of $\boldsymbol{\Sigma}^*$.

Proof sketch. The moment-matching estimator is the Kullback-Leibler projection of the empirical distribution onto the exponential family [49]. For pairwise MRFs with bounded-degree dependence graphs, the moment-matching sample complexity is governed by the standard sub-Gaussian concentration of $p \times p$ empirical covariance matrices, which gives $T = \Omega(p^2 \log p/\varepsilon^2)$ for ε-Frobenius-norm consistency; the dependence-graph-degree factor $C(d)$ absorbs the local-Markov-property cost. The mapping from $\boldsymbol{\theta}$ to $\mathbf{P}$ is smooth on the interior of the natural-parameter domain, and the $\sqrt{p}$ factor follows from applying the Frobenius-norm bound element-wise. □

Remark 5.4 (Online / streaming estimation)**.** A production CCP computes $\hat{\mathbf{P}}_T$ in an online setting: fills arrive sequentially; the estimator must update without storing the full fill history. Welford's online algorithm [46] maintains running first and second moments with $O(1)$ memory per element and $O(p^2)$ memory for the full moment matrix; the moment-matching MRF update reduces to a quasi-Newton step on the natural parameters. Welford's algorithm is chosen because it is numerically stable under catastrophic cancellation in floating-point arithmetic, which matters when two fills produce near-cancelling contributions to a variance estimate.

Remark 5.5 (Single-market LMSR calibration baseline)**.** The bid-ask reservation-price model of Avellaneda and Stoikov [45] sets a single-market market-maker target from inventory dynamics. In the portfolio-margin setting of this paper, the Avellaneda-Stoikov target is a single-node baseline for the per-spoke margin contribution m_i in (4); the MRF fit supplies the cross-asset coupling in $\boldsymbol{\Theta}_2$. The two are orthogonal inputs to the margin calculation: Avellaneda-Stoikov pins per-spoke m_i, the MRF pins pairwise ρ_{ij}.

5.2 Misspecification-robustness bound

The MRF is fit on the assumption of pairwise-only dependence (no higher-order coupling). Real fill streams may exhibit higher-order dependence; the pairwise fit misspecifies the joint distribution when this occurs. We bound the margin gap this introduces.

Proposition 5.6 (Margin gap under higher-order misspecification)**.** *Let p^* denote the true joint distribution with pairwise-plus-higher-order dependence, and $p_{\boldsymbol{\theta}^*}$ the best-approximating pairwise MRF (the KL projection onto the pairwise family). Let $\varepsilon_{\mathrm{KL}} := \mathrm{KL}(p^* \| p_{\boldsymbol{\theta}^*})$ measure the misspecification. Then the parametric ES under $p_{\boldsymbol{\theta}^*}$ and the true p^* satisfy*

$$\left|\mathrm{ES}_q^{p^*}[\Pi] - \mathrm{ES}_q^{p_{\boldsymbol{\theta}^*}}[\Pi]\right| \;\le\; \frac{\|\ell_\Pi\|_\infty}{q} \cdot \sqrt{2\varepsilon_{\mathrm{KL}}}, \tag{22}$$

by Pinsker's inequality [50]. When the true distribution is close in KL to the pairwise family, the margin gap is small; when $\varepsilon_{\mathrm{KL}}$ grows (e.g., at regime transitions that introduce triplet-level dependence), the parametric-layer margin becomes a biased under- or over-estimate of the realised tail loss, and the defense-in-depth stack (Definition 1.1) falls back to the stress layer M^{stress}.

Proof. Expected Shortfall is a Lipschitz functional of the loss distribution with Lipschitz constant $\|\ell_\Pi\|_\infty/q$ in total-variation distance [44]. Pinsker bounds the total-variation distance by $\sqrt{2\varepsilon_{\mathrm{KL}}}$, which gives (22). □

Remark 5.7 (Estimator breakdown is caught by the health predicate)**.** The health predicate $\mathcal{H}$ in Definition 1.1 tests (i) the MRF's condition number below a threshold (singular $\mathbf{P}$ is rejected), and (ii) the realised short-window tail loss against $\mathrm{ES}_q^{p_{\boldsymbol{\theta}^*}}[\Pi]$. When the realised tail loss exceeds ES_q by a tolerance factor (e.g., Kupiec backtest rejection at 95% confidence), the circuit-breaker halts new-position opening on the affected leg, and existing positions re-margin at M^{stress} alone. The Li-copula failure mode of Remark 1.3 is caught by (ii): a correlation-model understating the 2008 shock would be rejected after the first backtest-window realised loss, and the stress layer would bind thereafter.

Remark 5.8 (Relation to Ledoit-Wolf shrinkage)**.** The PSD-shrinkage construction of Proposition 3.4 is a particular case of the MRF fit when the structural prior $\mathbf{P}^{\mathrm{prior}}$ is treated as an entropy-regularisation term on $\boldsymbol{\Theta}_2$: blending $\mathbf{P}$ toward $\mathbf{P}^{\mathrm{prior}}$ with weight w_s is the KL-projection onto the pairwise family with a centred prior. The MRF formulation makes the regularisation explicit and makes the sample-complexity trade-off transparent.

5.3 Regime-shift detection and estimator reset

Theorem 5.3 assumes i.i.d. samples. Real fill streams exhibit regime shifts - discrete changes in the underlying correlation structure driven by macro, micro, or protocol events - that violate the i.i.d. assumption and degrade the MRF estimator consistency.

Definition 5.9 (Regime-shift detector)**.** A *regime-shift detector* on a rolling window of size W tests the null hypothesis that the samples $\{r^{(t-W+1)}, \ldots, r^{(t)}\}$ and $\{r^{(t-2W+1)}, \ldots, r^{(t-W)}\}$ are drawn from a common MRF. Under the MRF parameter $\boldsymbol{\theta}$, the generalised-likelihood-ratio test statistic

$$\Lambda_t \;:=\; 2\big[\ell(\hat{\boldsymbol{\theta}}^{(1)}; \mathrm{window}_1) + \ell(\hat{\boldsymbol{\theta}}^{(2)}; \mathrm{window}_2) - \ell(\hat{\boldsymbol{\theta}}^{\mathrm{pool}}; \mathrm{pooled})\big], \tag{23}$$

is distributed as $\chi^2_{d_\theta}$ under the null, where d_θ is the effective dimension of $\boldsymbol{\theta}$. A regime-shift detection at level α rejects the null when $\Lambda_t > \chi^2_{d_\theta, 1-\alpha}$.

Proposition 5.10 (Estimator reset on regime shift)**.** *Under the regime-shift detector of Definition 5.9 with detection probability at least $1-\beta$ for a true-regime-shift magnitude $\|\Delta\boldsymbol{\theta}\|_2 \ge \kappa_{\mathrm{reg}}$:*

(i) *On detection at time t, the MRF estimator is reset: the pre-t fill-window data are excluded from the new fit.*

(ii) *The health predicate $\mathcal{H}$ of Definition 1.1 holds the parametric layer M^{param} off until the reset-window length exceeds the sample-complexity requirement of Theorem 5.3: $W^{\text{reset}} \geq C(d) \cdot p^2 \log(p/\delta)/\varepsilon^2$.*

(iii) *During the hold-off period, the defense-in-depth stack binds at $\max(M^{\text{floor}}, M^{\text{stress}})$; the correlation-aware layer is deferred.*

Proof. (i) is a standard change-point reset. (ii) and (iii) follow from the MRF sample-complexity bound (Theorem 5.3): the reset window must accumulate enough post-shift fills for consistent estimation before the parametric layer re-engages. □

Remark 5.11 (Regime-shift interacts with tier transition)**.** A compliance-tier transition (Definition 6.3) is a structural regime-shift driver: the eligible-counterparty set changes, which changes the realised correlation structure. The regime-shift detector of Definition 5.9 flags the tier transition, and the MRF reset of Proposition 5.10 initiates automatically. The atomic margin recalculation of Definition 6.3 composes with the estimator-reset hold-off: the tier change is fully propagated before the parametric layer re-engages.

Remark 5.12 (Sensitivity to detector mis-tuning)**.** A detector tuned too sensitively (small α) triggers spurious resets, causing the parametric layer to be held off excessively and pushing the margin binding to the stress layer permanently; this is conservative (solvency-preserving) but capital-inefficient. A detector tuned too loosely (large α) misses real regime shifts, leaving the parametric layer running on stale data; this is caught by the backtest-rejection condition of Remark 5.7. The two detectors compose: the regime-shift detector is a leading indicator; the backtest-rejection is a lagging indicator; both feed $\mathcal{H}$ and either can drop the parametric layer.

Proposition 5.13 (Health-predicate adversarial robustness)**.** *Let $\mathcal{A} \in \mathcal{A}_{\text{AM}}(T)$ be an adaptive adversary attempting to keep $\mathcal{H}$ nominal (parametric layer engaged) while biasing the MRF fit toward a structure that understates the realised tail loss by factor $\gamma \geq 1$. Define the adversary's influence on the parametric layer as $\Delta^{\text{adv}}_{\text{param}} := M^{\text{param,adv}}(\Pi) - M^{\text{param,true}}(\Pi)$. Then the adversary's maximum solvency-relevant impact, subject to $\mathcal{H}$ remaining nominal throughout $[0, T]$, is bounded by*

$$|\Delta^{\text{adv}}_{\text{param}}| \;\leq\; (M^{\text{stress}} - M^{\text{param,true}})_+, \tag{24}$$

because the defense-in-depth stack of Definition 1.1 binds at $\max(M^{\text{floor}}, M^{\text{param}}, M^{\text{stress}})$; any adversarial understatement of M^{param} is absorbed up to the stress-layer buffer.

Proof. The required collateral $M^{\text{req}} = \max\{M^{\text{floor}}, M^{\text{param}}, M^{\text{stress}}\}$ is monotone non-decreasing in each argument. An adversary lowering M^{param} by $\Delta^{\text{adv}}_{\text{param}}$ lowers M^{req} by at most $\Delta^{\text{adv}}_{\text{param}}$ - and only when the max was already binding at M^{param}. Whenever $M^{\text{stress}} \geq M^{\text{param}}$, the max binds at M^{stress} regardless of M^{param}'s value; the adversary has no impact on M^{req}. Whenever $M^{\text{stress}} < M^{\text{param,true}}$, the adversary's understatement can move the binding down by at most $M^{\text{param,true}} - M^{\text{stress}}$, yielding (24). □

Remark 5.14 (Composition of defense-in-depth with MRF adversarial robustness)**.** Proposition 5.13 formalises the structural purpose of the stress layer: even absent the circuit-breaker activation, even with an adversary-engineered MRF fit that passes the health predicate, solvency is preserved to the extent of M^{stress}. This is the concrete sense in which the architecture rejects the Li-copula failure mode (Remark 1.3): the 2008 CDS clearing breakdown was load-bearing on its Gaussian-copula margin layer; the defense-in-depth stack here is load-bearing on the max, not on any single layer.

6 Compliance-Tier-Dependent Margin

A CCP clearing instruments whose compliance state varies across investor classes faces margin-state transitions that cannot be captured by the return-correlation structure alone. An instrument that is compliant for all investors has a different liquidity and hedging profile than an instrument restricted to a subset of investors under a jurisdictional compliance gate. Tier transitions between these states change the margin requirement instantaneously, independently of any price or correlation change. This section states the tier-transition protocol and bounds the instantaneous margin delta.

6.1 Instrument tiers and gated pools

Definition 6.1 (Instrument compliance tier)**.** Each instrument I at clearing time t carries a *compliance tier* $\tau_I(t) \in \{\text{Cold}, \text{Warm}, \text{Hot}\}$ drawn from the instrument's composed compliance state across all applicable jurisdictional domains. Informally:

- Cold: open to issuance and still uncleared for public secondary trade; held in a restricted book.
- Warm: cleared for a gated pool of eligible investors (qualified by compliance-domain constraints); positions are isolated to the gated pool with separate margin and default capacity.
- Hot: fully compliant across all applicable domains; cleared into the main book with the margin framework of §3 and the waterfall of §7.

Definition 6.2 (Gated-pool margin floor)**.** A Warm-tier instrument held in the gated pool carries a margin floor

$$\mu_I^{\text{Warm}-\text{floor}} := \max(\mu_I^{\text{floor}}, \mu^{\text{gated}}), \tag{25}$$

where $\mu^{\text{gated}} > \mu_I^{\text{floor}}$ is a protocol-set conservative floor reflecting the reduced hedge-book depth available in the gated pool relative to the main book. Warm-tier positions do not inherit correlation credits from cross-pool positions in the main book.

6.2 Tier-transition recalculation protocol

Definition 6.3 (Tier-transition margin recalculation)**.** Let I be an instrument with current tier $\tau_I(t^-)$ and new tier $\tau_I(t^+)$ following a compliance-state change at time t. The CCP re-margins positions in I atomically at time t per the tier-transition matrix of Table 1.

Transition	Margin formula	Pool handling	Halt condition
Cold → Warm	$M^{\text{cross}}_{\text{gated}}$ with $\mu^{\text{Warm}-\text{floor}}$	Move to gated pool	None
Warm → Cold	M^{iso} with $\mu^{\text{Cold}-\text{floor}} \geq \mu^{\text{Warm}-\text{floor}}$	Freeze positions	Stop new opens
Warm → Hot	$M^{\text{req}}(\Pi)$ per Definition 1.1	Move to main book	None
Hot → Warm	Main → gated pool re-isolation	$M^{\text{cross}}_{\text{gated}}$	Drain correlation credits
Hot → Cold	M^{iso} with $\mu^{\text{Cold}-\text{floor}}$	Freeze + liquidate	Revoke market-making

Table 1: Tier-transition margin recalculation matrix. Each transition triggers an atomic margin recomputation at time t; the new margin binds before any subsequent position change or mark-to-market. Downgrades (Hot → Warm, Warm → Cold, Hot → Cold) freeze or liquidate positions to respect the tighter pool boundary; upgrades (Cold → Warm, Warm → Hot) release reserved collateral smoothly.

Proposition 6.4 (Instantaneous margin delta at tier transition)**.** *Let Π_I be a position in instrument I transitioning from $\tau_I(t^-)$ to $\tau_I(t^+)$. The instantaneous margin change $\Delta M := M^{\text{new}}(\Pi_I) - M^{\text{old}}(\Pi_I)$ satisfies:*

(i) Downgrade direction. *For Hot → Warm or Hot → Cold, $\Delta M \geq 0$; the position's required collateral increases by $(\mu^{\text{new-floor}} - \mu^{\text{old-floor}}) \cdot |Q_I| \cdot P_I$ plus the drain of cross-book correlation credits.*
(ii) Upgrade direction. *For Cold → Warm or Warm → Hot, $\Delta M \leq 0$; collateral is released up to the difference in floors.*
(iii) Lateral Warm ↔ Hot. *The delta equals the difference between gated-pool correlation structure and main-book correlation structure; both are positive since gated pools carry fewer hedging instruments.*

Proof. (i) Downgrades both raise the per-instrument floor and remove positions from the main book's cross-pool correlation credits; ΔM is a sum of two non-negative terms. (ii) Upgrades reverse both. (iii) Follows from Definition 6.2: the gated pool's reduced eligible-counterparty set always tightens the margin beyond the main-book equivalent. □

Remark 6.5 (Health-predicate interaction)**.** The health predicate $\mathcal{H}$ of Definition 1.1 includes "no compliance-tier transition in flight on any leg" to prevent margin-recalculation race conditions: during the atomic re-margin at tier transition, the defense-in-depth stack falls back to the stress layer alone, and new-position opening on the affected leg is halted. The atomic re-margin completes within a single clearing step.

Remark 6.6 (Transition-announcement front-running)**.** An adversary who learns that a downgrade (Hot → Warm, Warm → Cold, Hot → Cold) is imminent can unwind positions at the main-book correlation structure before the transition, avoiding the higher post-transition floor of Proposition 6.4. Two protocol choices mitigate:
(1) *Embargo window.* The compliance-state change that drives a downgrade is announced only at the clearing step of the transition itself. Between pre-detection and clearing-step commit, the transition is held in a private queue inaccessible to counterparties; no position may be closed on the affected instrument during the window.
(2) *Retroactive re-margin.* Position closures within an embargo window of length W_{emb} blocks before a detected downgrade are re-margined at the post-transition floor; the additional collateral is clawed back from the closing counterparty through the main-book waterfall.

(1) is preferred when the downgrade signal arrives synchronously with the clearing step (the embargo reduces to zero length); (2) is preferred when compliance-state changes are detected asynchronously and the protocol must reconcile the gap. Both choices preserve the margin invariant of Proposition 6.4.

Remark 6.7 (Separate default capacity for gated pools)**.** Warm-tier positions in gated pools do not share the insurance fund IF of the main book. A separate gated-pool fund IF^{gated} is sized under the same Cover-k rule (Definition 7.2) against the gated-pool's stress-scenario set, and the waterfall of §8 applies within-pool. Cross-pool contagion from a gated-pool default into the main book is blocked by the tier-isolation discipline; the LP loss-absorption layer (Theorem 9.4) is similarly isolated. This isolation is the tier-transition analogue of the exogenous-settlement severance of Theorem 12.2: compliance-tier isolation breaks a contagion channel at the architectural level.

6.3 Margin-breach response and cross-pool unwind

A margin-breach event occurs when a position's realised loss exceeds its posted collateral, whether because the parametric ES estimator understated the tail (correlation-model breakdown caught by the health predicate $\mathcal{H}$), because a tier transition drained correlation credits to below the new-tier floor, or because an execution slippage exceeded the scenario catalogue. The response protocol must propagate the breach through the waterfall of §8 and, under compliance-tier interaction, respect tier-boundary isolation.

Definition 6.8 (Margin-breach response protocol)**.** On detection of a margin breach on position Π at time t:

(1) *Freeze.* New-position opening on Π's leg is halted and the health predicate $\mathcal{H}$ is set to false for the affected leg.
(2) *Re-margin at stress.* Existing positions are re-margined at M^{stress} alone per Definition 1.1; the correlation-aware M^{param} is bypassed until circuit-breaker conditions are re-satisfied.
(3) *Default handling.* If the re-margined requirement exceeds posted collateral, the position is a default and the five-stage waterfall of §8 applies with the defaulter's collateral seized first.
(4) *Tier-isolated unwind.* If Π spans multiple tiers (e.g., a Warm-tier leg and a Hot-tier leg of a synthetic position), the breach is first allocated to the tier of the primary leg, then propagates only through tier-preserving channels. Hot-tier gained correlation credits that the position relied on are drained before any Warm-pool socialisation.
(5) *Insurance-fund selection.* Default allocations to the insurance fund draw from the IF corresponding to the position's tier at breach-detection time; Warm-tier breaches draw from IF^{gated}, Hot-tier breaches from the main IF.

Proposition 6.9 (Single-position-breach containment)**.** *Let Π be a single breaching position at time t, with shortfall $\Delta_\Pi := M^{\text{req}}(\Pi) - C_\Pi$ where C_Π is posted collateral and M^{req} is the defense-in-depth requirement of (1). Under the protocol of Definition 6.8, the portfolio-level shortfall propagation to counterparties $k \neq \text{owner}(\Pi)$ satisfies*

$$\Delta_k^{\text{prop}} \le \kappa \cdot C_k \cdot \frac{|\Pi_k|}{\sum_{k' \neq \text{owner}(\Pi)} |\Pi_{k'}|} \cdot \mathbb{1}_{\Delta_\Pi > C_\Pi + E_{\text{SIG}} + \text{IF}^{\text{tier}}}, \tag{26}$$

where IF^{tier} is the tier-appropriate insurance fund. Socialisation engages only when the stage-1, stage-2, and stage-3 capacities are all exhausted; the cap κ bounds per-counterparty exposure to the breach even in the worst case.

Proof. The five-stage waterfall of Theorem 7.5 consumes stages in order; propagation to a non-defaulting counterparty occurs only at stage 4 (socialisation), and the stage-4 allocation is capped by $\kappa \cdot C_k$ per Definition 7.3. The tier-selection rule of Definition 6.8(5) ensures that the fund drawn is the one matching the breaching position's tier; cross-tier fund draws are excluded by the tier-isolation of Remark 6.7. □

Proposition 6.10 (Portfolio-level breach cascade bound under tier isolation)**.** *Let $D_t \subseteq \mathcal{A}$ be the set of counterparties whose positions breach in a single clearing step at time t. Denote by $D_t^{\text{Hot}}, D_t^{\text{Warm}}, D_t^{\text{Cold}}$ their partition by tier. Under the protocol of Definition 6.8, the aggregate shortfall absorbed by each tier is bounded independently:*

$$\Delta_t^{\text{Hot}} \le E_{\text{SIG}}^{\text{Hot}} + \text{IF}^{\text{Hot}} + \kappa \textstyle\sum_{k \notin D_t^{\text{Hot}}, k \in \text{Hot}} C_k + \sum_{k \in \mathcal{P}^{\text{Hot}}} \pi_k, \tag{27}$$

$$\Delta_t^{\text{Warm}} \le E_{\text{SIG}}^{\text{Warm}} + \text{IF}^{\text{gated}} + \kappa \textstyle\sum_{k \notin D_t^{\text{Warm}}, k \in \text{Warm}} C_k + \sum_{k \in \mathcal{P}^{\text{Warm}}} \pi_k, \tag{28}$$

$$\Delta_t^{\text{Cold}} \le C_{\text{total}}^{\text{Cold}-\text{frozen}}. \tag{29}$$

No cross-tier socialisation occurs; a Warm-tier cascade that exhausts IF^{gated} does not access the main IF.

Proof. Direct from the tier-selection rule of Definition 6.8(5) and the tier-isolated waterfall of Remark 6.7. The bounds are the Theorem 7.5 capacity bound applied within-tier. □

Remark 6.11 (Cross-pool unwind on Hot → Warm downgrade)**.** A Hot → Warm downgrade (Table 1) drains the instrument's correlation credits in the main book. The unwind proceeds in two steps: (i) close all hedge legs in the main book that referenced the downgrading instrument, crystallising their P&L into the book-wide mark-to-market; (ii) re-open the position in the gated pool at the $M^{\text{cross}}_{\text{gated}}$ level. The inter-step window during which the position is in neither book is a clearing-step atomic operation (zero latency by construction); the crystallised P&L and the gated-pool re-entry are committed together. The atomicity blocks any adversary from opening a synthetic position across the unwind window.

Remark 6.12 (Tier-isolation is orthogonal to exogenous settlement)**.** Theorem 12.2 breaks the reflexive coupling between hub-asset dynamics and clearing-layer capacity; tier isolation (Proposition 6.10) breaks the reflexive coupling between compliance-state shifts and cross-tier solvency. The two are architecturally independent: a non-exogenous-settlement Warm pool still benefits from tier isolation, and an exogenous-settlement Hot book still benefits from cross-tier fund separation. Defense-in-depth layers the two without overlap.

7 Clearing Operator and the Waterfall

We formalise the CCP's default-handling procedure as a clearing operator on a network of position accounts, extending the Eisenberg-Noe framework [1, 2] to cover the five-stage waterfall of Assumption 2.2(3), including the CCP skin-in-the-game tranche of Definition 7.1 and the Cover-k fund-sizing rule of Definition 7.2.

7.1 Setup

Let $\mathcal{A} = \{1, \dots, K\}$ index accounts. Account k has *gross obligation* $o_k \geq 0$, *gross claim* $c_k \geq 0$, and *outside collateral* $C_k \geq 0$ posted with the CCP. Write

$$L_k := (o_k - c_k)_+$$

for the gross loss before collateral. An account *defaults* at clearing time if $L_k > C_k$. The waterfall first seizes collateral $s_k = \min(L_k, C_k)$ and then runs the remaining loss $r_k := (L_k - C_k)_+$ through the mutualised layers. Let $E_{\text{SIG}} \geq 0$ denote the CCP's skin-in-the-game tranche, IF ≥ 0 the mutualised insurance fund, $\kappa \in (0, 1)$ the per-position socialisation cap, and $\mathcal{P} \subseteq \mathcal{A}$ the set of profitable accounts (those with $c_k > o_k$, with unrealised profit $\pi_k := c_k - o_k$).

Definition 7.1 (CCP skin-in-the-game tranche)**.** Let $E_{\text{SIG}} \geq 0$ denote the CCP's own equity posted as a dedicated default-management layer, ring-fenced from the mutualised insurance fund IF. The tranche is consumed in full before any draw on IF, aligning the CCP's loss-absorption incentives with members' (EMIR Art. 45(4), Dodd-Frank §5464 [36, 37, 31, 34]). The design choice $E_{\text{SIG}} = 0$ ("skinless" CCP) is admissible and is discussed in Remark 7.4.

Definition 7.2 (Cover-k insurance-fund sizing)**.** Let $q \in (0, 1)$ be an extreme-but-plausible stress quantile and $k \geq 1$ an integer. The insurance fund is *Cover-k sized* at quantile q if

$$\text{IF} \;\geq\; \sup_{\substack{D \subseteq \mathcal{A} \\ |D| \leq k}} \Big(\sum_{j \in D} (\Delta_j^{(q)} - C_j)_+ \Big) - E_{\text{SIG}}, \tag{30}$$

where $\Delta_j^{(q)}$ is the q-quantile gross pre-collateral loss L_j under the clearinghouse's stress-scenario distribution. Cover-1 and Cover-2 are the PFMI Principle 4 / EMIR Art. 43 standards for systemically important CCPs [35, 36, 31].

Definition 7.3 (Waterfall clearing vector)**.** A *waterfall clearing vector* is a tuple $(\mathbf{s}, e, d, \boldsymbol{\ell}, \boldsymbol{\delta})$ with:

- $\mathbf{s} \in \mathbb{R}^K_{\geq 0}$ the *seizure vector* ($s_k \in [0, C_k]$ for defaulting k);
- $e \in [0, E_{\text{SIG}}]$ the *skin-in-the-game draw*;
- $d \in [0, \text{IF}]$ the *fund draw*;
- $\boldsymbol{\ell} \in \mathbb{R}^K_{\geq 0}$ the *socialisation vector* with $\ell_k \leq \kappa\, C_k$;
- $\boldsymbol{\delta} \in \mathbb{R}^K_{\geq 0}$ the *ADL allocation* with $\delta_k \leq \pi_k$ and $\delta_k = 0$ for $k \notin \mathcal{P}$.

The vector *clears* the gross loss $\sum_k L_k$ if

$$\sum_k s_k + e + d + \sum_k \ell_k + \sum_k \delta_k \;=\; \sum_k L_k \;-\; (\text{terminal residual})_+, \tag{31}$$

and satisfies the *stage-priority constraints*: $e > 0 \Rightarrow \sum_k s_k = \sum_k \min(L_k, C_k)$ (defaulter seizure exhausted before SIG); $d > 0 \Rightarrow e = E_{\text{SIG}}$ (SIG exhausted before fund draw); $\ell_k > 0 \Rightarrow d = \text{IF}$ (fund exhausted before socialisation); $\delta_k > 0 \Rightarrow \ell_{k'} = \kappa\, C_{k'}$ for some k' (socialisation cap binds before ADL). This realises the five-stage EMIR Art. 45 structure [36].

Remark 7.4 (Skinless design)**.** Setting $E_{\text{SIG}} = 0$ reduces Definition 7.3 to a four-stage waterfall without CCP first-loss. The trade-off is studied in [6, 34]: positive SIG aligns the CCP's incentives with those of members at the cost of CCP-equity commitment, and a skinless CCP externalises first-loss to the mutualised fund. Jurisdictional requirements (EMIR Art. 45(4); Dodd-Frank §5464) mandate $E_{\text{SIG}} > 0$ for systemically important CCPs.

Theorem 7.5 (Existence and uniqueness of the clearing vector)**.** *For any configuration* $(\mathbf{o}, \mathbf{c}, \mathbf{C}, E_{\text{SIG}}, \text{IF}, \kappa)$ *with* $\kappa \in (0, 1)$ *there exists a unique waterfall clearing vector* $(\mathbf{s}^*, e^*, d^*, \boldsymbol{\ell}^*, \boldsymbol{\delta}^*)$ *satisfying Definition 7.3. The terminal residual*

$$\Delta^{\text{res}} := \sum_k L_k - \sum_k s^*_k - e^* - d^* - \sum_k \ell^*_k - \sum_k \delta^*_k$$

is non-negative, and $\Delta^{\text{res}} = 0$ *iff the total available capacity* $\sum_k C_k + E_{\text{SIG}} + \text{IF} + \kappa \sum_k C_k + \sum_{k \in \mathcal{P}} \pi_k$ *exceeds* $\sum_k L_k$.

Proof. The stages are decoupled by the priority constraints. Stage 1: $s^*_k = \min(L_k, C_k)$ is uniquely determined; let $\Delta_1 := \sum_k (L_k - s^*_k)_+ = \sum_k r_k$. Stage 2 (SIG): $e^* = \min(\Delta_1, E_{\text{SIG}})$, residual $\Delta_2 := (\Delta_1 - e^*)_+$. Stage 3 (fund): $d^* = \min(\Delta_2, \text{IF})$, residual $\Delta_3 := (\Delta_2 - d^*)_+$. Stage 4 is the capped water-filling problem:

$$\boldsymbol{\ell}^* = \arg\min\Big\{ \textstyle\sum_k \max(0, \ell_k - \ell^{\text{prorata}}_k) \;:\; \sum_k \ell_k = \min(\Delta_3, \kappa \sum_k C_k),\; 0 \leq \ell_k \leq \kappa\, C_k \Big\}, \tag{32}$$

where ℓ^{prorata}_k is the notional-proportional share. This is a convex program with linear constraints: existence follows from compactness, uniqueness from strict convexity in the interior. Residual $\Delta_4 := \Delta_3 - \sum_k \ell^*_k \geq 0$; $\Delta_4 = 0$ iff the cap κC_k is slack on at least one account with positive notional. Stage 5: $\delta^*_k = \pi_k \cdot \Delta_4 / \sum_{k' \in \mathcal{P}} \pi_{k'}$ scaled to $\min(\Delta_4, \sum_{k' \in \mathcal{P}} \pi_{k'})$; the terminal residual is $\Delta^{\text{res}} := (\Delta_4 - \sum_{k' \in \mathcal{P}} \pi_{k'})_+$. Uniqueness is immediate at each stage; non-negativity holds by construction. □

Theorem 7.6 (Order-independence under simultaneous defaults)**.** *Let $D \subseteq \mathcal{A}$ be the set of accounts defaulting at a single clearing time (concurrent defaults). Any two permutations of D yield the same clearing vector* $(\mathbf{s}^*, e^*, d^*, \boldsymbol{\ell}^*, \boldsymbol{\delta}^*)$ *provided the stages are applied* level by level*: all seizures before any SIG draw, all SIG draws before any fund draw, all fund draws before any socialisation, all socialisation before any ADL.*

Proof. Stage 1: $s_k^* = \min(L_k, C_k)$ depends only on defaulter k's own state; $\sum_{k \in D} s_k^*$ is permutation-invariant. Stages 2 and 3: the aggregate draws $e^* = \min(\Delta_1, E_{\text{SIG}})$ and $d^* = \min(\Delta_2, \text{IF})$ depend only on the aggregates Δ_1, Δ_2. Stage 4: the water-filling problem (32) is a convex program depending only on Δ_3 and $\{C_{k'}\}_{k' \notin D}$. Stage 5: identical argument. Applying stages per-account-sequentially would violate order-independence because subsequent stages' inputs would then depend on prior accounts' activity across levels; the level-by-level discipline makes each stage's output a function of the aggregate input only. □

Remark 7.7 (AMM-price-impact gives deterministic C_k adjustment)**.** In the classical Eisenberg-Noe setting, C_k is an exogenous collateral pool. In the AMM-CCP setting, C_k includes the proceeds from liquidating k's underwater positions into the AMM pool. Under Assumption 2.1(2), these proceeds are a closed-form function of pool reserves; the Stage-1 seizure operator is deterministic, not stochastic. This is the key structural advantage over order-book-liquidation settings where C_k depends on counterparty auction behaviour.

8 The Five-Stage Waterfall

We make the five stages of Definition 7.3 explicit and prove the conservation invariant. The stage structure maps onto EMIR Art. 45 [36]: defaulter margin, CCP skin-in-the-game, mutualised fund, member socialisation, auto-deleveraging.

8.1 Stage definitions

Stage 1: Defaulter margin seizure. The clearinghouse seizes $s_k = \min(L_k, C_k)$ from each defaulting account's collateral, using the AMM-price-impact kernel to convert non-settlement-asset collateral to settlement-asset value.

Stage 2: CCP skin-in-the-game. The residual $\Delta_1 := \sum_k (\Delta_k - s_k)_+$ is drawn against the CCP's dedicated tranche: $e = \min(\Delta_1, E_{\text{SIG}})$; see Definition 7.1.

Stage 3: Insurance fund draw. Residual $\Delta_2 := (\Delta_1 - e)_+$ is drawn from IF: $d = \min(\Delta_2, \text{IF})$. The fund balance is reconciled at block end to prevent intra-block oracle manipulation of fund state.

Stage 4: Capped socialisation. Residual $\Delta_3 := (\Delta_2 - d)_+$ is allocated across non-defaulting counterparties by water-filling with per-position cap $\kappa \in (0, 1)$:

$$\ell_k = \min\left(\frac{|Q_k| P_k}{\sum_{k' \in \mathcal{A}^{(n)}} |Q_{k'}| P_{k'}} \, \Delta_3^{(n)}, \ \kappa \, C_k \right),$$

iterating with active set $\mathcal{A}^{(n)}$ until all residual is allocated or all caps bind.

Stage 5: Auto-deleveraging (ADL). Residual Δ_4 is absorbed by reducing profitable positions proportional to unrealised profit π_k:

$$\delta_k = \pi_k \cdot \frac{\min(\Delta_4, \sum_{k' \in \mathcal{P}} \pi_{k'})}{\sum_{k' \in \mathcal{P}} \pi_{k'}}.$$

Proposition 8.1 (Conservation identity)**.** *The clearing vector of Theorem 7.5 satisfies, for every* $t \geq 0$,

$$\sum_k C_k(t) + E_{\text{SIG}}(t) + \text{IF}(t) + \text{ER}(t) \ = \ \textit{const} \ - \ \Delta^{\text{res}}(t), \tag{33}$$

where $\text{ER}(t)$ *is the clearinghouse equity reserve and* $\Delta^{\text{res}}(t) \geq 0$ *is the cumulative terminal residual through time* t*. When* $\Delta^{\text{res}} = 0$ *the system is closed; when* $\Delta^{\text{res}} > 0$ *the shortfall equals the cumulative gap between obligations and absorbed capital.*

Proof. Stages 1, 2, 3, 5 each transfer capital between balance-sheet entries, an accounting identity. Stage 4's cap introduces the only mechanism by which $\Delta^{\text{res}} > 0$: if the cap binds on every counterparty before the residual is allocated, the remainder passes to Stage 5; if Stage 5's ADL capacity $\sum_{k' \in \mathcal{P}} \pi_{k'}$ is exhausted, the residual is $\Delta^{\text{res}} := (\Delta_4 - \sum_{k' \in \mathcal{P}} \pi_{k'})_+$. Summing over time gives (33). □

8.2 Porting of client positions

Standard CCP machinery ports client positions to a surviving member on default (EMIR Art. 48; PFMI Principle 13 [36, 35]). In the AMM-CCP setting, a client position decomposes into pool-share claims and derivative bookings against surviving counterparties, so porting is structurally more granular than the traditional single-member carry-over.

Definition 8.2 (Portable position)**.** Let k denote a clearing member with client accounts $\{j : \text{parent}(j) = k\}$. The *portable component* of client j's position is the pool-share claim $\alpha_j \in \mathbb{R}_{\geq 0}$ on each pool i (the share of pool reserves attributable to j); the *non-portable component* is the (Q_j^p, S_j) vector of perpetual and event positions, which may be ported to any surviving member k' willing to book them.

Proposition 8.3 (Porting versus liquidation)**.** *Let D be a default set and $|D|^c = |\mathcal{A} \setminus D|$ the count of surviving members with capacity. Write* $\text{Port}(D)$ *for the aggregate residual shortfall under full porting (surviving members absorb positions without mutualisation) and* $\text{Liq}(D)$ *for the residual under the liquidation path (defaulter positions sold into the AMM at the slippage cost of Theorem 9.1). Then*

$$\text{Port}(D) \;\leq\; \text{Liq}(D) \qquad \Longleftrightarrow \qquad \sum_{k' \notin D} (C_{k'}^{\text{slack}}) \;\geq\; \mathcal{L}_{\text{slip}}(D), \tag{34}$$

where $C_{k'}^{\text{slack}}$ is surviving member k''s post-port collateral capacity above maintenance, and $\mathcal{L}_{\text{slip}}(D)$ is the AMM slippage loss from liquidating D's positions. Porting is preferred when surviving-member slack exceeds the slippage loss; liquidation is preferred otherwise.

Proof. Porting transfers the defaulter's position to surviving members without crystallising the slippage loss $\mathcal{L}_{\text{slip}}$ into the pool; the cost is the consumption of surviving-member slack. Liquidation crystallises $\mathcal{L}_{\text{slip}}$ but preserves surviving members' capacity. The "$\leq$" direction of (34) follows immediately from this accounting; the "$\geq$" direction is the capacity constraint for the porting path to clear without Stage 4 socialisation. □

Remark 8.4 (Porting under hub-asset concentration)**.** When all surviving members are exposed to the same hub asset M, porting concentrates rather than distributes risk: the surviving-member slack $C_{k'}^{\text{slack}}$ is reduced in proportion to the hub-asset exposure. In the limit of full hub exposure across the membership, $\text{Port}(D) \to \text{Liq}(D)$ and the porting advantage vanishes; exogenous settlement (Theorem 12.2) restores the gap by decoupling surviving-member capacity from hub-asset dynamics.

Corollary 8.5 (Cascade termination under correlated defaults)**.** *Let $\{D_n\}_{n \geq 1}$ be the sequence of defaulter sets arising in successive clearing rounds (round $n+1$'s defaulters being those counterparties whose Stage-4 socialisation in round n pushed them below maintenance). Write $\Delta_n^{\text{res}} \geq 0$ for the terminal residual of round n and $\Delta_\infty^{\text{res}} := \sum_{n \geq 1} \Delta_n^{\text{res}}$ for the cumulative residual summed over all rounds. Then:*

(i) Cumulative loss bound. *The cumulative clearing-layer equity destruction across the entire cascade is bounded by*

$$\sum_{n \geq 1} (\Delta_n^{(\text{abs})}) \;\leq\; E_{\text{SIG},0} + \text{IF}_0 + \kappa \sum_{k \in \mathcal{A}_0} C_k^{(0)} + \sum_{k \in \mathcal{P}_0} \pi_k^{(0)} + \Delta_\infty^{\text{res}}, \tag{35}$$

where IF_0, $\{C_k^{(0)}\}$, $\{\pi_k^{(0)}\}$ *are the pre-cascade values, and* $\Delta_n^{(\mathrm{abs})}$ *is the absorbed obligation in round* n*.*

(ii) Finite termination. *The cascade terminates in finitely many rounds: since each round consumes a non-negative quantum of the absorptive capacity on the RHS of* (35)*, and the RHS is finite, at most* $O(|\mathcal{A}_0|)$ *rounds can contribute absorbed obligation before the absorptive capacity is exhausted and the residual is carried to* $\Delta_\infty^{\mathrm{res}}$*.*

Proof. (i) Sum Proposition 8.1 over rounds. The LHS is the cumulative absorbed obligation; the RHS is a telescoping sum of the initial capacity minus the residual carried forward, which gives (35).

(ii) Each round consumes at least one counterparty's worth of collateral or ADL capacity (or else the residual increments and the round is absorbed into $\Delta_\infty^{\mathrm{res}}$). Because $|\mathcal{A}_0|$ is finite, the set of counterparties is exhausted in $O(|\mathcal{A}_0|)$ rounds. □

Remark 8.6 (Cascade depth is buffer-distribution dependent)**.** Corollary 8.5 bounds cumulative equity destruction. The number of rounds cannot be derived from the socialisation cap κ alone. The cascade depth depends on the joint distribution of post-Stage-4 margin buffers across counterparties: when all counterparties have post-socialisation margin above maintenance, the cascade terminates in one round. A sharp bound on cascade depth requires a distributional assumption on the buffer density; we leave this as Open Problem 7.

8.3 Stress scenario framework

The waterfall capacity bound of Corollary 8.5 is evaluated against a scenario distribution. PFMI Principle 4 [35] and EMIR Art. 49 [36] require CCPs to stress the waterfall against extreme but plausible events; Cover-1 and Cover-2 sizing (Definition 7.2) is stated with respect to the quantile q of this distribution. We fix the minimal scenario template used throughout the paper.

Definition 8.7 (Stress scenario)**.** A *stress scenario* $\omega = (\mathbf{\Delta}P, \delta, \mathbf{\Delta R_B}, \Delta\bar{\rho})$ is a tuple specifying: (i) a joint price shock $\mathbf{\Delta}P \in \mathbb{R}^N$ on spoke assets; (ii) settlement-asset depeg magnitudes $\delta \in [0,1]^{|\mathcal{S}|}$; (iii) pool-depth shocks $\mathbf{\Delta R_B} \in \mathbb{R}^N_{\le 0}$; (iv) a correlation shift $\Delta\bar{\rho} \in [0, 1-\bar{\rho}]$. The joint distribution $\pi(\omega)$ is characterised by a marginal-plus-copula specification.

Definition 8.8 (Scenario catalogue)**.** The minimal catalogue Ω_0 replays historical tail events rescaled to current spoke and hub exposures: (H1) October 1987 equity crash (single-day -20% broad-index shock); (H2) September 2008 cross-asset credit freeze (three-day joint -30% spot, widening basis, funding-liquidity withdrawal); (H3) March 2020 pandemic shock (one-week -25% spot, elevated stablecoin depeg probability); (H4) May 2022 Terra/Luna collapse (three-day full-depeg of an endogenous settlement asset, elevated spoke correlation); (H5) June 2021 Iron Finance collapse (one-day reflexive cascade in a minority-backed algorithmic asset). Ω_0 is supplemented by a Monte Carlo envelope Ω_{MC} that resamples joint shocks at the $q = 10^{-3}$ quantile under the copula of Definition 8.7.

Remark 8.9 (Reverse stress testing)**.** Reverse stress testing inverts the direction: given a residual shortfall threshold $\bar{\Delta} > 0$, identify the minimum scenario (in a norm on $(\mathbf{\Delta}P, \delta, \mathbf{\Delta R_B}, \Delta\bar{\rho})$) under which $\Delta^{\mathrm{res}}(\omega) \ge \bar{\Delta}$. The scenario template of Definition 8.7 admits this inversion as a finite-dimensional optimisation; we defer its empirical implementation to Appendix A.

9 AMM Fire-Sale Contagion Bound

Traditional fire-sale analysis [3, 4] treats the price-impact kernel as an empirically estimated function. The AMM-CCP setting permits a closed-form deterministic bound because the kernel is exact.

9.1 Setup

Consider a first-round liquidation event in which defaulting accounts collectively sell volumes $q_i \geq 0$ of spoke asset B_i into pool i for $i = 1, \ldots, N$. Let $P_i(q_i)$ denote the realised average price and $\Delta P_i(q_i) := P_i - P_i(q_i)$ the slippage. Under Assumption 2.1, for constant-product pools ($g_i = 1$),

$$\Delta P_i(q_i) \;=\; P_i \cdot \frac{q_i}{R_B^{(i)} + q_i}. \tag{36}$$

For a general $g_i \geq 1$, $\Delta P_i(q_i) = P_i\big(1 - (1 + q_i/R_B^{(i)})^{-g_i}\big)$, which reduces to (36) at $g_i = 1$ and is increasing in g_i for fixed $q_i/R_B^{(i)}$.

Fire-sale contagion captures the second-order effect: a slippage ΔP_i on pool i reduces the mark-to-market value of all portfolios holding B_i, which triggers additional margin calls. The magnitude depends on the spoke-asset correlation structure and on portfolio concentrations.

Theorem 9.1 (AMM fire-sale contagion bound)**.** *Let $\mathbf{V} \in \mathbb{R}^N_{\geq 0}$ be the vector of mark-to-market portfolio values indexed by spoke asset, and let $\boldsymbol{\Sigma}$ be the spoke-return covariance matrix with spectral norm $\|\boldsymbol{\Sigma}\|_2$. Let $q_i \geq 0$ be the first-round liquidation volume on pool i, and define the slippage vector $\boldsymbol{\psi} = (\Delta P_i/P_i)_i$. The second-round loss amplification satisfies*

$$\mathcal{L}_{\text{2nd}} \;\leq\; \|\boldsymbol{\Sigma}\|_2 \cdot \|\mathbf{V}\|_2 \cdot \|\boldsymbol{\psi}\|_2, \tag{37}$$

where $\mathcal{L}_{\text{2nd}}$ is the aggregate portfolio-value loss from the second-round mark-to-market propagation. The bound is attained when $\mathbf{V}$ and $\boldsymbol{\psi}$ are aligned with the leading eigenvector of $\boldsymbol{\Sigma}$.

Proof. The second-round loss is $\mathcal{L}_{\text{2nd}} = \mathbf{V}^\top \boldsymbol{\Sigma}\, \boldsymbol{\psi}$ (portfolio exposures times price shocks weighted by the covariance). By Cauchy-Schwarz, $|\mathbf{V}^\top \boldsymbol{\Sigma}\, \boldsymbol{\psi}| \leq \|\mathbf{V}\|_2 \, \|\boldsymbol{\Sigma}\, \boldsymbol{\psi}\|_2 \leq \|\boldsymbol{\Sigma}\|_2 \, \|\mathbf{V}\|_2 \, \|\boldsymbol{\psi}\|_2$. Attainment is achieved when $\mathbf{V} \parallel \boldsymbol{\Sigma}\boldsymbol{\psi}$ and $\boldsymbol{\psi}$ is the leading eigenvector of $\boldsymbol{\Sigma}$. The slippage vector $\boldsymbol{\psi}$ is deterministic given $\{q_i, R_B^{(i)}\}$ via (36); hence the bound is a closed-form function of the liquidation event and pool depths. □

Corollary 9.2 (Bound for the hub-spoke topology)**.** *Under equicorrelation $\rho_{ij} = \bar{\rho}$ with unit individual variances (standardised returns) and equal pool depths $R_B^{(i)} = R_B$, the bound reduces to*

$$\mathcal{L}_{\text{2nd}} \;\leq\; \big(1 + (N-1)\bar{\rho}\big) \cdot \|\mathbf{V}\|_2 \cdot \frac{\|q\|_2}{R_B}, \tag{38}$$

where $\|q\|_2$ is the aggregate liquidation norm. The hub-spoke topology's spectral norm is $1 + (N - 1)\bar{\rho}$, linear in $\bar{\rho}$.

Proof. $\|\boldsymbol{\Sigma}\|_2 = 1 + (N-1)\bar{\rho}$ for equicorrelation with unit variances. The slippage $\psi_i = q_i/(R_B + q_i)$ has the conservative bound $\|\boldsymbol{\psi}\|_2 \leq \|q\|_2/R_B$ because $q_i/(R_B + q_i) \leq q_i/R_B$ for every i. □

Remark 9.3 (Deterministic versus stochastic contagion)**.** Cont and Wagalath [4] give a stochastic contagion bound in which the price-impact kernel is specified by a distribution. The bound here is deterministic: once $\{q_i, R_B^{(i)}\}$ are observed, the slippage vector is a closed-form quantity and the contagion loss is bounded by a matrix-norm inequality. This is a strict improvement, enabled by the AMM-invariant structure.

9.2 AMM LPs as a distributed loss-absorption layer

The slippage of Theorem 9.1 is borne by the pool's liquidity providers: when the clearinghouse liquidates a defaulter's position into pool i, LPs experience a loss equal to the slippage integral, partially offset by the fee revenue generated on that trade. The pool is therefore a continuous-time mutualised default-fund layer distributed across LPs by pool-share ownership. This is the AMM-CCP analogue of the EMIR Art. 45 mutualisation tranche [36, 31], made continuous by the AMM invariant and quantifiable in closed form.

Theorem 9.4 (LP loss-absorption layer). *Let $\{q_i(s)\}_{s\in[0,T]}$ be the cumulative liquidation schedule through a default event of duration T, let f_i denote pool i's trading fee, and let $\alpha_{j,i} \geq 0$ be LP j's share of pool i. The cumulative LP loss absorbed into the event is*

$$\mathcal{L}_j^{\mathrm{LP}} \;=\; \sum_i \alpha_{j,i} \int_0^T \left[\Delta P_i(q_i(s)) \cdot \dot{q}_i(s) \;-\; f_i \cdot |\dot{q}_i(s)| \cdot P_i(q_i(s))\right] ds, \tag{39}$$

with ΔP_i as in (36). The aggregate $\sum_j \mathcal{L}_j^{\mathrm{LP}}$ is a continuous-time default-fund contribution, mutualised proportionally to pool-share ownership, and is $O(\|q\|_2^2/R_B)$ as pool depth scales.

Proof. The AMM invariant assigns each infinitesimal trade dq a slippage $\Delta P_i(q)$ per (36) and a fee $f_i \cdot |dq| \cdot P_i$. LP j holds fraction $\alpha_{j,i}$ of pool i; by the pool-share accounting of the invariant, j's share of the slippage net of the fee share is $\alpha_{j,i}$ times the integrand, and integrating over the liquidation schedule yields (39). The $\|q\|_2^2/R_B$ scaling follows from the Taylor expansion of $\Delta P_i(q) = P_i \cdot q/(R_B + q)$ at small q/R_B: the slippage is $P_i q^2/R_B + O(q^3/R_B^2)$, and integrating against dq over $[0, q_i]$ gives $P_i q_i^3/(2R_B) + O(q_i^4/R_B^2)$ which at fixed liquidation fraction q_i/R_B is $O(\|q\|_2^2/R_B)$. □

Remark 9.5 (LVR baseline and event-driven absorption). Under Milionis, Moallemi, Roughgarden, and Zhang [39], the continuous-time LP loss decomposes into two components: a loss-versus-rebalancing baseline $\mathrm{LVR}(t) = (\sigma_M^2/8)V^{\mathrm{LP}}(t)$ per unit time (constant-product; g_i-weighted analogue for the power-weighted family) proportional to ambient volatility, plus an event-driven component given by Theorem 9.4. Optimal-fee calibration [38] compensates LPs for the baseline; the event-driven component is absorbed by the pool and feeds back into the margin and waterfall accounting of §§3-8.

Remark 9.6 (Position of LPs in the waterfall). When LP pool shares are not accepted as collateral at the CCP (hub-asset-denominated collateral only), the LP loss of Theorem 9.4 is a Stage 0 pre-waterfall absorption: it enters through the slippage cost before the defaulter-margin seizure is computed, reducing the notional recovered by Stage 1. When LP shares are accepted as collateral, the LP loss also enters the collateral value in Stages 1, 4, and 5, and LVR becomes a margin input (§3.4).

10 Reflexive Cascade Dynamics

From this section forward we invoke Assumption 2.3. The hub asset's price P_M has an endogenous component $f(L)$ depending on locked liquidity. This creates a feedback channel not present in the adapted CCP material of §§3-9.

10.1 Coupled price-liquidity ODE

Definition 10.1 (LP withdrawal dynamics). The aggregate LP withdrawal rate responds to price declines with a strictly positive sensitivity λ:

$$\dot{L}(t) \;=\; -\lambda\, L(t) \cdot \left[-\dot{P}_M(t)/P_M(t)\right]^+, \tag{40}$$

where $[x]^+ := \max(0, x)$.

Definition 10.2 (Reflexivity coefficient)**.** Fix a reference timescale $\tau > 0$ (canonical: $\tau = 1/\theta$). The *reflexivity coefficient* is

$$\varrho(P_M, L) \;:=\; \tau \cdot \frac{\lambda\, L\, f'(L)}{P_M}. \tag{41}$$

With λ a rate ($[\lambda] = 1/\text{time}$), ϱ is dimensionless.

In the declining-price regime, substituting Assumption 2.3 and Definitions 10.1-10.2 into the price equation yields the coupled ODE

$$\dot{P}_M \;=\; \frac{\theta(\bar{V} - P_M + f(L))}{1-\varrho}, \qquad \dot{L} \;=\; \frac{\lambda L}{P_M}\,\dot{P}_M. \tag{42}$$

10.2 Critical slowdown and subcritical stability

Theorem 10.3 (Critical slowdown of the reduced flow)**.** *Consider (42) with $f(L) = \beta_L L^\gamma$, $\gamma \in (0, 1]$.*

(a) Subcritical *($\varrho < 1$): perturbations of P_M about the equilibrium curve $\mathcal{E} = \{P_M = \bar{V} + f(L)\}$ decay at the amplified exponential rate $\theta/(1-\varrho) > \theta$.*

(b) Critical *($\varrho \to 1^-$): the linearised relaxation rate diverges; this is the signature of a singular limit. The full classification of the two-dimensional dynamics at $\{\varrho = 1\}$ is given in Section 11.*

(c) Supercritical (linear instability) *($\varrho > 1$): the effective coefficient $\theta/(1-\varrho) < 0$; the reduced scalar ODE has a positive linearised eigenvalue and perturbations grow exponentially on the fast timescale of Proposition 11.4.*

(d) Equilibrium subcriticality condition. *At equilibrium $(P_M^*, L^*) \in \mathcal{E}$,*

$$\tau\lambda\gamma < \frac{P_M^*}{V_{\text{end}}^*} = 1 + \frac{\bar{V}}{V_{\text{end}}^*}, \tag{43}$$

with $V_{\text{end}}^* = f(L^*)$.

Proof. The equilibrium condition is $\bar{V} - P_M^* + f(L^*) = 0$. Linearising (42) at (P_M^*, L^*): $\dot{P}_M \approx -[\theta/(1-\varrho^*)](P_M - P_M^*)$. The sign of $1 - \varrho^*$ determines the stability of the linearisation. For (d): $\varrho^* = \tau\lambda L^* f'(L^*)/P_M^* = \tau\lambda\gamma V_{\text{end}}^*/P_M^*$ using $Lf'(L) = \gamma f(L)$ for $f = \beta_L L^\gamma$. Setting $\varrho^* < 1$ gives (43). □

10.3 Global attraction on the supercritical branch

Theorem 10.4 (Global attractor under supercritical cascade)**.** *Consider (42) with $f(L) = \beta_L L^\gamma$, $\gamma \in (0, 1]$, restricted to supercritical cascade trajectories ($\varrho(P_{M,0}, L_0) > 1$, $\dot{P}_M < 0$) on the open quadrant $\Omega := \{P_M > 0,\ L \geq 0\}$. Then:*

(i) Orbit integral. *$L(t) = L_0(P_M(t)/P_{M,0})^\lambda$; the orbit is confined to a one-dimensional curve.*

(ii) Compact invariance. *$\mathcal{K} := \{(P_M, L) \in \Omega : P_M \leq P_{M,0},\ L \leq L_0\}$ is forward-invariant and compact.*

(iii) Strict Lyapunov function. *$V(P_M, L) := (P_M - \bar{V})^2 + c\, L^{2/\lambda}$ is a strict Lyapunov function on $\mathcal{K} \setminus \{(\bar{V}, 0)\}$ for $c > 0$ sufficiently large.*

(iv) Convergence. *$(P_M(t), L(t)) \to (\bar{V}, 0)$ as $t \to \infty$.*

Proof. (i) From (42), $dL/dP_M = \lambda L/P_M$ (the singular prefactor $1/(1-\varrho)$ cancels). Separation of variables gives $L(P_M) = L_0(P_M/P_{M,0})^\lambda$.

(ii) By hypothesis $\dot{P}_M < 0$ and $\dot{L} = (\lambda L/P_M)\dot{P}_M < 0$. The coordinate bounds are preserved; compactness follows.

(iii) Differentiating, $\dot{V} = 2(P_M - \bar{V})\dot{P}_M + c \cdot (2/\lambda)L^{2/\lambda-1}\dot{L} = 2\dot{P}_M\big[(P_M - \bar{V}) + c\,L^{2/\lambda}/P_M\big]$. On the cascade branch $\dot{P}_M < 0$; we need the bracket positive. For $P_M \geq \bar{V}$ the first term is non-negative. For $P_M < \bar{V}$, pick c large enough on the compact $\mathcal{K}$: $c \geq \bar{V}\,P_{M,0}/L_*^{2/\lambda}$ for a lower bound $L_* > 0$ on the boundary of $\mathcal{K} \setminus \{(\bar{V}, 0)\}$. Strictness $\dot{V} < 0$ holds off $(\bar{V}, 0)$ because $\dot{P}_M = 0$ only on $\mathcal{E}$, and the supercritical cascade trajectory intersects $\mathcal{E}$ only at $(\bar{V}, 0)$.

(iv) $\mathcal{K}$ is compact; the system is autonomous and C^1 away from $\{\varrho = 1\}$, which is not encountered on the supercritical branch. Poincaré-Bendixson [23, Thm. 7.4.1] implies the ω-limit set is an equilibrium, a periodic orbit, or a union of heteroclinic connections. By (iii), V strictly decreases off $(\bar{V}, 0)$; periodic orbits and heteroclinic connections are ruled out. The unique equilibrium on the supercritical branch's closure is $(\bar{V}, 0)$, and convergence follows. □

Remark 10.5 (Poincaré-Bendixson hypotheses are verified)**.** Previous analyses of reflexive collapse in algorithmic stablecoins [14, 15] are qualitative and do not formalise global attraction. The integrability $L = L_0(P_M/P_{M,0})^\lambda$ reduces the two-dimensional flow to a one-dimensional orbit equation, closing the required hypotheses for Poincaré-Bendixson.

Remark 10.6 (Scope of Theorem 10.4)**.** Theorem 10.4 does not claim global attraction from arbitrary initial conditions in Ω: the statement is restricted to the supercritical cascade branch $\{\varrho(P_{M,0}, L_0) > 1, \dot{P}_M(0) < 0\}$, on which the system is C^1 away from Σ (which the cascade branch does not encounter). The Lyapunov construction requires $\lambda > 0$; for small λ (slow LP withdrawal) the term $L^{2/\lambda}$ becomes degenerate, and the Lyapunov-constant c grows accordingly on compact $\mathcal{K}$.

10.4 LP concentration and single-node failure

The withdrawal dynamics of Definition 10.1 treat $L(t)$ as aggregate locked liquidity. When the LP distribution is concentrated (a small number of LPs hold most of L), a single-LP withdrawal can collapse L discontinuously, inducing the fast transition across Σ that Proposition 11.4 identifies as a canard blow-up.

Definition 10.7 (LP concentration)**.** Let $\alpha_j \geq 0$ denote LP j's share of aggregate L with $\sum_j \alpha_j = 1$. The *LP Herfindahl index* is $H := \sum_j \alpha_j^2 \in [0, 1]$; the *maximum share* is $s_{\max} := \max_j \alpha_j \in [0, 1]$. A pool is *LP-diffuse* if $s_{\max}$ is bounded below some threshold $\bar{s}$; otherwise it is *LP-concentrated*.

Proposition 10.8 (Time-to-Σ under single-LP withdrawal)**.** *Suppose a single LP with share $s_{\max}$ withdraws at some time t_0, causing $L(t_0^+) = (1 - s_{\max})L(t_0^-)$ and a simultaneous hub-asset price drop $P_M(t_0^+) = (1 - \delta_P)P_M(t_0^-)$ with $\delta_P \in [0, 1)$. If $\varrho(P_M(t_0^-), L(t_0^-)) < 1$ (subcritical) and $s_{\max}$ satisfies*

$$s_{\max} > 1 - \left((1 - \delta_P)\,\frac{P_M(t_0^-)}{\tau\lambda L(t_0^-) f'(L(t_0^-))}\right)^{1/\gamma} \tag{44}$$

(with $f(L) = \beta_L L^\gamma$, $\gamma < 1$), then the post-withdrawal state satisfies $\varrho(P_M(t_0^+), L(t_0^+)) > 1$; the system is pushed onto the supercritical cascade branch discontinuously. A pure liquidity withdrawal with $\delta_P = 0$ decreases ϱ by the factor $(1 - s_{\max})^\gamma$ and cannot by itself push a subcritical state across Σ.

Proof. The reflexivity coefficient is $\varrho = \tau\lambda L f'(L)/P_M$. Substituting $L(t_0^+) = (1 - s_{\max})L(t_0^-)$ and $P_M(t_0^+) = (1 - \delta_P)P_M(t_0^-)$ gives

$$\varrho(t_0^+) = \frac{(1 - s_{\max})^\gamma}{1 - \delta_P}\varrho(t_0^-).$$

Rearranging $\varrho(t_0^+) > 1$ gives (44). The price-jump term is load-bearing: without it the multiplicative factor is $(1 - s_{\max})^\gamma < 1$, so the withdrawal alone moves away from the critical locus. □

Remark 10.9 (Liquidity concentration and funding liquidity)**.** Proposition 10.8 is the AMM-CCP analogue of the Brunnermeier-Pedersen [25] funding-liquidity spiral: concentrated LP capital provides market liquidity that is vulnerable to funding-liquidity shocks on the dominant LP. An LP-concentration cap $s_{\max} \le \bar{s}$ and a per-LP withdrawal-rate limit together bound the time-to-Σ from below and preserve the subcritical stability of Theorem 10.3. Optimising the trade-off between LP-scale efficiency (which encourages concentration via fee revenue) and systemic-safety concentration cap is Open Problem 9.

11 Bifurcation Classification at $\varrho = 1$

We resolve the dynamical-systems classification of the critical locus $\Sigma := \{\varrho = 1\}$. The main result is that Σ is *not* a Sotomayor fold bifurcation: condition (S1) fails. We then identify the correct classification.

11.1 Desingularization

Equation (42) is singular on Σ. Rescale time: $ds := dt/(1-\varrho)$ on the subcritical sheet $\{\varrho < 1\}$. Multiplying through gives the smooth two-dimensional field

$$\frac{dP_M}{ds} = \theta(\bar{V} - P_M + f(L)), \qquad \frac{dL}{ds} = \frac{\lambda L}{P_M}\theta(\bar{V} - P_M + f(L)). \tag{45}$$

Proposition 11.1 (Equilibria of the desingularized flow)**.** *The equilibria of* (45) *in* $\{P_M > 0, L \ge 0\}$ *are exactly* $\mathcal{E} = \{(P_M, L) : P_M = \bar{V} + f(L),\ L \ge 0\}$.

Proof. Setting both components to zero requires $\theta(\bar{V} - P_M + f(L)) = 0$; the second component then vanishes automatically. □

11.2 Sotomayor conditions on the desingularized flow

Sotomayor's theorem [23, Thm. 3.4.1][22, Thm. 3.4] classifies a parametrised equilibrium as a fold if the Jacobian $A = D_x F$ has a simple zero eigenvalue with right null v and left null $w^\top$, and the following conditions hold:

(S1) $\dim\ker A = 1$; non-zero eigenvalues have non-zero real part; the zero eigenvalue appears at the bifurcation parameter value $\mu = \mu_0$ (and not for nearby μ).

(S2) $w^\top D^2_{xx} F(x^*, \mu_0)[v, v] \neq 0$.

(S3) $w^\top D_\mu F(x^*, \mu_0) \neq 0$ (transversality).

We instantiate on (45). The Jacobian at an equilibrium $(P_M^*, L^*) \in \mathcal{E}$ is

$$A = \begin{pmatrix} -\theta & \theta f'(L^*) \\ (\lambda L^*/P_M^*)(-\theta) & (\lambda L^*/P_M^*)\,\theta f'(L^*) \end{pmatrix}. \tag{46}$$

Lemma 11.2 (Rank deficiency on $\mathcal{E}$)**.** *At every equilibrium* $(P_M^*, L^*) \in \mathcal{E}$ *with* $L^* > 0$, *A has rank one: row 2* $= (\lambda L^*/P_M^*)\cdot$ *row 1. Consequently,* $\det A = 0$ *and* $0 \in \sigma(A)$ *for all such equilibria, independent of any bifurcation parameter.*

Proof. Direct from (46): each entry of row 2 equals $(\lambda L^*/P_M^*)$ times the corresponding entry of row 1. □

Theorem 11.3 ($\Sigma = \{\varrho = 1\}$ is not a Sotomayor fold)**.** *Let $\mu := \tau\lambda\gamma$ parametrise the one-parameter family (45). At any equilibrium $(P_M^*, L^*) \in \mathcal{E}$ with $\varrho^* = 1$, Sotomayor's condition (S1) fails: the zero eigenvalue of A is present at every μ in a neighbourhood of the putative bifurcation point, not isolated to $\mu = \mu_0$. The fold bifurcation does not occur.*

Proof. By Lemma 11.2, $0 \in \sigma(A)$ at every $(P_M^*, L^*) \in \mathcal{E}$ for every value of μ. Sotomayor requires the zero eigenvalue to be an isolated event at $\mu = \mu_0$; here it is persistent. Independently, the transversality condition (S3) also fails: $D_\mu F = 0$ identically along $\mathcal{E}$ after desingularization, since μ enters the original ODE only through the singular prefactor $(1 - \varrho)$, which has been removed. Both (S1) and (S3) fail; the classification is categorically not a fold. □

11.3 Correct classification: Fenichel canard of a slow-fast system

Proposition 11.4 (Σ is the fold curve of a singular-perturbation parameter)**.** *Write (42) as a slow-fast system with singular parameter $\varepsilon := 1 - \varrho$:*

$$\varepsilon \dot{P}_M = \theta(\bar{V} - P_M + f(L)), \qquad \dot{L} = \frac{\lambda L}{P_M} \dot{P}_M.$$

The slow manifold is $\mathcal{E}$. On $\Sigma = \{\varepsilon = 0\}$, normal hyperbolicity of $\mathcal{E}$ fails in the Fenichel sense [18]; slow-fast reduction is invalid across Σ, and trajectories approaching Σ from the subcritical side in finite t-time exhibit $|\dot{P}_M| \to \infty$ on the cascade branch.

Proof. The system is already in singular-perturbation form with $\varepsilon = 1 - \varrho$. Fenichel's theorem requires ε to be a small positive constant in a neighbourhood of the slow manifold. Here ε is state-dependent and vanishes on Σ, so the hypotheses are violated on Σ. Blow-up $|\dot{P}_M| \to \infty$ follows from (42) when $\varepsilon \to 0^+$ with numerator $\theta(\bar{V} - P_M + f(L)) < 0$ (cascade condition). □

Remark 11.5 (Consequences of the canard classification)**.** Keeping $\varrho^* \ll 1$ preserves an $O(1)$ normal-hyperbolicity margin for $\mathcal{E}$. A fold would imply a structurally stable two-equilibrium picture near the threshold; the canard classification instead implies that trajectories crossing Σ exit the scalar-reduction regime and enter a fast two-dimensional layer dominated by LP withdrawal dynamics, consistent with the rapid collapse observed in Terra/Luna (May 2022) and Iron Finance (June 2021) [14, 15, 17].

12 Exogenous-Settlement Severance

The reflexive cascade of Section 10 operates only when clearing-layer quantities (margin, collateral, insurance fund) depend on P_M. When these are denominated in an exogenous asset S whose price is independent of L, the reflexivity coefficient of the clearing layer is identically zero.

12.1 The severance theorem

Definition 12.1 (Exogenous settlement asset)**.** A *settlement asset* is an asset S with price process $P_S(t)$ satisfying $P_S(t) \perp L(t)$ (the price is independent of the locked-liquidity process). The clearing layer is *denominated in S* if margin requirements, collateral valuations, and the insurance fund are all expressed in P_S-units.

Theorem 12.2 (Exogenous-settlement severance, first-order)**.** *Under Assumption 2.3 and Definition 12.1, suppose the clearing layer is denominated in an exogenous settlement asset S. Then, to first order in ∂_{P_M}, the clearing-layer reflexivity coefficient satisfies $\varrho_{\text{clearing}} = 0$ for all (P_M, L).*

The first-order feedback loop $P_M \downarrow \Rightarrow$ *margin calls* $\Rightarrow$ *liquidations* $\Rightarrow L \downarrow \Rightarrow V_{\text{end}} \downarrow \Rightarrow P_M \downarrow$ *is broken at the first step. A residual second-order channel via* $\mathrm{Cov}(P_S, P_M \mid \text{stress})$ *is characterised in Corollary 12.4.*

Proof. $\varrho = \tau\lambda L f'(L)/P_M$ measures the sensitivity of clearing-layer capital adequacy to the hub price P_M. Under exogenous settlement: collateral value in S-units is $C_{\text{eff}} = C \cdot P_S/P_S = C$, independent of P_M. Margin requirements in S-units are invariant to P_M. Insurance fund balance in S-units is similarly invariant. Hence the derivative of any clearing-layer quantity with respect to P_M vanishes, and $\varrho_{\text{clearing}} = 0$. □

Corollary 12.3 (Bounded contagion under exogenous settlement)**.** *Under Theorem 12.2, the worst-case clearing-layer loss on any hub-asset trajectory is bounded above by*

$$\mathcal{L}^{\max} \leq E_{\text{SIG}} + \text{IF} + \kappa \sum_k C_k + \sum_{k\in\mathcal{P}} \pi_k + \Delta^{\text{depeg}} + \Delta^{(2)}_{SM}, \tag{47}$$

where Δ^{depeg} *is the settlement-asset depeg tail term bounded in Theorem 13.4 and* $\Delta^{(2)}_{SM}$ *is the second-order* (P_S, P_M) *tail-covariance residual of Corollary 12.4. Under independent* (P_S, P_M), $\Delta^{(2)}_{SM} = 0$ *and the worst case reduces to a finite bounded-loss event.*

Proof. Theorem 12.2 removes the first-order reflexivity channel from the clearing layer; Corollary 12.4 bounds the residual second-order channel. The waterfall of Theorem 7.5 bounds the absorbable shortfall by the sum of defaulter collateral seizure, the CCP skin-in-the-game tranche E_{SIG}, the insurance fund, socialisation capacity $\kappa \sum_k C_k$, and ADL capacity $\sum_{k\in\mathcal{P}} \pi_k$. The settlement-asset depeg term follows from Theorem 13.4. □

Corollary 12.4 (Second-order severance residual)**.** *Under Theorem 12.2, the residual second-order contribution to clearing-layer loss from non-zero right-tail dependence between* P_S *and* P_M *under stress is*

$$\Delta^{(2)}_{SM} \leq \lambda^U_{SM} \cdot \sup_q \left|\mathbb{E}_q[P_S \mid P_M \downarrow]\right| \cdot C_{\text{total}}, \tag{48}$$

where $\lambda^U_{SM} \in [0,1]$ *is the upper tail-dependence coefficient of the* (P_S, P_M) *copula [10, 11, 40] and* $C_{\text{total}} = \sum_k C_k$ *is the aggregate collateral.* $\Delta^{(2)}_{SM} = 0$ *under* $\lambda^U_{SM} = 0$ *(independence of the extremes). Under* $\lambda^U_{SM} > 0$ *(e.g., correlated banking-channel failures affecting both* P_S *and* P_M *simultaneously under stress), severance is only first-order and Corollary 12.3's bound includes (48).*

Proof. Expanding the clearing-layer loss to second order in ∂_{P_M} introduces the bilinear term $\partial_{P_M}\partial_{P_S} \cdot \mathrm{Cov}(P_S, P_M \mid \text{stress})$. Under the Sklar representation of the joint distribution, the tail covariance is upper-bounded by λ^U_{SM} times the marginal-conditional expectations [11, Thm. 2.15]. The aggregate collateral C_{total} converts the covariance into a loss in settlement-asset units. □

Remark 12.5 (Allen-Gale correlated failure)**.** The second-order channel of Corollary 12.4 is the analogue, for the AMM-CCP system, of the correlated-failure account of financial contagion due to Allen and Gale [40]: severance operates at the level of counterparty-network topology, but systemic stress that affects P_S and P_M jointly (banking-channel contagion, regulatory action, custody failure) restores a partial coupling through the tail copula.

12.2 Residual risks and the counter-cyclical buyback

Exogenous settlement eliminates the dominant reflexive channel. Three residual contagion channels remain: (i) spoke-asset correlation (bounded by Theorem 9.1); (ii) settlement-asset depeg (addressed in Section 13); (iii) governance misconfiguration (structural, outside the model's scope).

A counter-cyclical buyback mechanism restores part of the exogenous funds to the hub asset during calm regimes. To preserve the reflexivity-severance property of Theorem 12.2 under a non-trivial buyback policy (so that reserves are not spent to acquire a declining endogenous asset), the buyback is gated on structural conditions.

Definition 12.6 (Gated counter-cyclical buyback)**.** Let $X \in (0,1)$ be a settlement-asset reserve threshold and $Y \in (0,1)$ a recovery-probability threshold. Let $\Delta_{\text{lag}} \geq 0$ be a decorrelation lag. The buyback amount at time t is

$$B(t) \;=\; \mathbb{1}_{G(t-\Delta_{\text{lag}})} \cdot r \cdot \left(\mathrm{IF}(t) - \mu_{\text{target}}\,\mathrm{IF}_{\min}\right)_+, \tag{49}$$

where $G(s)$ is the event $\{\mathrm{IF}_S(s)/\mathrm{IF}(s) \geq X\} \cap \{p_{\text{recovery}}(s) \geq Y\}$, $r \in (0,1)$ is the buyback fraction, $\mu_{\text{target}} > 1$ is the fund coverage target, and p_{recovery} is an estimator of $\mathbb{P}(\varrho_{t+1} < 1 \mid \mathcal{F}_t)$.

The closed-form sufficient bounds below make the gate's dependence on the underlying risk parameters explicit.

Proposition 12.7 (Closed-form bound on X)**.** *Let $\eta_{\text{safe}} > 0$ be a margin-of-safety parameter, $q \in (0, 0.05]$ a tail quantile for correlated settlement-asset depegs, $h_{\text{stress}} \in (0,1)$ a per-asset stressed haircut, and $\omega_{\text{non}-S} \in [0,1]$ the fund fraction held in non-settlement-asset (non-S) reserves. Write $\rho_{\mathrm{IF}} := \mathrm{IF}_{\text{balance}}/C^{\mathrm{IF}}_{\text{total}} \in (0,1]$. The smallest X such that post-buyback residual-S reserves cover the tail depeg shortfall with margin $(1+\eta_{\text{safe}})$ satisfies*

$$X \;\geq\; (1+\eta_{\text{safe}}) \cdot h_{\text{stress}} \cdot \omega_{\text{non}-S} \;+\; r \cdot \rho_{\mathrm{IF}} \;+\; f_{\min}, \tag{50}$$

where $f_{\min}$ is a structural primary-settlement-asset floor. The right-hand side is a closed-form monotone function of the haircut, depeg tail, and buyback parameters; no Monte-Carlo calibration is required.

Proof. Partition the fund into S-denominated and non-S-denominated subfunds: $\mathrm{IF} = \mathrm{IF}_S + \mathrm{IF}_{\text{non}-S}$. Pre-buyback, $\mathrm{IF}_S/\mathrm{IF} \geq X$. After a buyback consuming fraction r of $\rho_{\mathrm{IF}} \cdot \mathrm{IF}$ (of which a fraction is drawn from S), residual S fraction is $X - r\rho_{\mathrm{IF}}$. A subsequent tail depeg on non-S holdings (fraction $\omega_{\text{non}-S}$) imposes loss $h_{\text{stress}} \cdot \omega_{\text{non}-S}$; to cover this with margin $(1+\eta_{\text{safe}})$ and maintain the structural floor $f_{\min}$, we require $X - r\rho_{\mathrm{IF}} \geq (1+\eta_{\text{safe}}) \cdot h_{\text{stress}} \cdot \omega_{\text{non}-S} + f_{\min}$. Rearranging gives (50). □

Proposition 12.8 (Closed-form bound on Y)**.** *Let $W \in \mathbb{N}$ be a rebalance-window length and $q \in (0, 0.05]$ a tail quantile. Under an AR(1) observation model for ϱ, a sufficient threshold is*

$$Y \;\geq\; 1 - q/W. \tag{51}$$

Proof. By the Bonferroni union bound, $\mathbb{P}(\sup_{b<t\leq b+W} \varrho_t \geq 1 \mid p_{\text{recovery}} \geq Y) \leq W \cdot (1-Y)$. Requiring this to be at most q gives $Y \geq 1 - q/W$. □

Remark 12.9 (Bonferroni is loose under AR(1) persistence)**.** The Bonferroni bound ignores positive serial correlation in ϱ_t. Under an AR(1) with persistence $\phi \in [0,1)$, the W single-block tail events are positively correlated; an effective-sample-size correction gives the tighter bound $Y \geq 1 - q(1-\phi^2)/W$ (which recovers (51) at $\phi = 0$). Implementations that bound $\mathbb{P}[\sup \varrho > 1]$ via a direct variance estimator (Chebyshev-style, rather than a union over the W blocks) obtain this tightening without the AR(1) persistence dependence. Eq. (51) is valid as an upper bound on the required Y; it is not tight.

Proposition 12.10 (Closed-form bound on Δ_{lag})**.** *Let $\tau_\varrho > 0$ denote the autocorrelation timescale of ϱ_t and $\varepsilon_{\text{dec}} \in (0,1)$ a residual-correlation tolerance. A sufficient lag is*

$$\Delta_{\text{lag}} \;\geq\; -\tau_\varrho \log \varepsilon_{\text{dec}}. \tag{52}$$

For $\varepsilon_{\text{dec}} = 0.05$, $\Delta_{\text{lag}} \geq 3\tau_\varrho$.

Proof. An AR(1) process has autocorrelation e^{-k/τ_ϱ} at lag k. Requiring $e^{-\Delta_{\text{lag}}/\tau_\varrho} \le \varepsilon_{\text{dec}}$ gives (52). □

Remark 12.11 (Linear autocorrelation vs. concentration tail)**.** Eq. (52) bounds the *linear* autocorrelation of ϱ_t at lag Δ_{lag}, which is what the AR(1) moment estimator directly consumes. A stronger statement about the probability of adversarial-innovation lag-window gate bypass, i.e., $\mathbb{P}[\text{bypass}] \le \exp(-c\Delta_{\text{lag}})$, follows from a standard Chernoff argument under bounded-variance AR(1). A sub-Gaussian tail of the form $\exp(-c\Delta_{\text{lag}}^2)$ requires an explicit variance-proxy assumption on the innovation distribution and is not delivered by (52) alone.

Remark 12.12 (Closed-form versus simulation-based gating)**.** Monte-Carlo calibration of $(X, Y, \Delta_{\text{lag}})$ over specific parameter envelopes yields point estimates that depend on the simulation distribution; Propositions 12.7-12.10 give closed-form sufficient bounds that are transparent in the haircut, depeg-tail, window, and autocorrelation parameters and do not depend on unmodelled stress-scenario details.

13 Multi-Settlement-Asset Haircut Model

Theorem 12.2 assumes a single exogenous settlement asset. In practice, a CCP supports multiple settlement assets (a basket of stablecoins, government money-market instruments, or similar) to avoid single-issuer concentration risk. This section gives the multi-asset haircut model and proves the stablecoin-agnostic bounded-contagion theorem.

13.1 Haircut schedule and depeg isolation

Definition 13.1 (Haircut schedule)**.** Each accepted settlement asset $s \in \mathcal{S}$ carries a haircut $h_s \in [0, 1)$. Effective collateral in s is $C_s^{\text{eff}} = (1 - h_s) \cdot C_s \cdot P_s^{\text{ref}}$, with P_s^{ref} a median-of-oracles reference price.

Proposition 13.2 (Depeg isolation under independent failures)**.** *Suppose settlement assets in $\mathcal{S}$ depeg independently, with asset s^* depegging by magnitude $\delta_{s^*} > h_{s^*}$. The clearing-layer loss is bounded by*

$$\mathcal{L}_{\text{indep}} \le (\delta_{s^*} - h_{s^*})_+ \cdot C_{s^*}. \tag{53}$$

Pools settling in other assets are unaffected.

Proof. Haircut-adjusted collateral covers the first h_{s^*} of depeg; residual $(\delta_{s^*} - h_{s^*})_+$ is the unhedged loss on s^*. Independence of depegs isolates the loss to s^*. □

Remark 13.3 (Oracle-synchrony assumption)**.** Proposition 13.2 assumes the reference-median oracle P_s^{ref} is synchronous with the depeg event. In practice the oracle updates on a discrete schedule with lag $\tau_O > 0$; during $[t, t + \tau_O]$ the posted haircut reflects the pre-depeg price while the realised collateral value follows the depegged price. The bound (53) understates the instantaneous loss during the oracle-lag window by at most $\tau_O \cdot \sup_s |\dot{\delta}_s| \cdot C_{s^*}$. A depeg-detection circuit-breaker halt bounds this slippage outside the haircut model; the per-asset bound holds exactly at oracle-update time.

13.2 Correlated depegs and the copula bound

Independence is an idealization; banking-crisis, regulatory-action, and custody-failure channels correlate stablecoin depegs. We give a copula-based upper bound.

Theorem 13.4 (Stablecoin-agnostic bounded contagion)**.** *Let $\mathcal{K} \subseteq \mathcal{S}$ be the set of simultaneously depegging assets under a correlated-stress scenario with depeg magnitudes $\{\delta_k\}_{k\in\mathcal{K}}$ drawn from a joint distribution with copula C. The tail-expected clearing-layer loss satisfies*

$$\mathbb{E}_q[\mathcal{L}_{\mathrm{corr}}] \;\leq\; \sum_{k\in\mathcal{K}} \mathbb{E}_q[(\delta_k - h_k)_+] \cdot C_k \;+\; \delta_q^{\mathrm{tail}} \cdot \mathrm{IF}, \tag{54}$$

where $\delta_q^{\mathrm{tail}} := \mathbb{E}_q[\max_{k\in\mathcal{K}} \delta_k]$ is the joint-tail magnitude at confidence q, computable from C by standard copula-tail inequalities [10, 11].

Proof. Per-asset contributions sum by linearity; the joint-tail term δ_q^{tail} captures the simultaneous-depeg worst case. Copula tail-dependence coefficients [11] give upper bounds on δ_q^{tail} in terms of pairwise tail-dependence parameters $\{\lambda_{kk'}^U\}$. □

Corollary 13.5 (Stablecoin-agnostic bounded contagion)**.** *Theorem 12.2 and Theorem 13.4 together imply: for* any *exogenous settlement basket, the clearing-layer loss under a reflexive hub-asset trajectory is bounded by a closed-form expression involving the insurance fund, socialisation/ADL capacity, per-asset haircuts, and a joint-tail copula term. The bound is independent of which specific assets populate the basket, provided each is exogenous in the sense of Definition 12.1.*

Remark 13.6 (Diversification bounds)**.** Requiring the insurance fund to satisfy (i) a per-issuer concentration cap $\mathrm{IF}_k \leq \omega_{\max} \cdot \mathrm{IF}$, and (ii) a non-stablecoin reserve floor $\mathrm{IF}_{\mathrm{non-stable}} \geq \omega_{\min} \cdot \mathrm{IF}$, tightens the copula-tail bound: the joint-tail contribution is multiplied by $\omega_{\max} \cdot (1 - \omega_{\min})$, a strict improvement when $\omega_{\min} > 0$. The optimal $(\omega_{\max}, \omega_{\min})$ is left as Open Problem 1.

14 Adversarial Bounds

The preceding sections prove capital-efficiency and contagion bounds under stated structural assumptions. This section formalises the adversarial counterparts: for each defense, we (i) fix an adversary model, (ii) bound the adversary's payoff from above, and, where the model admits, (iii) construct a matching attack witnessing the lower bound. Adversarial statements are stated as propositions with explicitly labelled assumptions; proofs are given in sketch form with the reduction structure made explicit. Claims that we do *not* close in a matching sense are flagged explicitly as one-sided bounds.

14.1 Adversary catalogue

We distinguish four baseline adversary classes used below.

Definition 14.1 (Adversary classes)**.**

(a) *Rational single-account* adversary $\mathcal{A}_{\mathrm{RS}}$: a single account chooses a portfolio and trade sequence to maximise expected net profit, subject to the margin rules of §3 and the waterfall of §7. Non-adaptive within a block: chooses an action once per clearing step given public state.

(b) *Sybil-Rational multi-account* adversary $\mathcal{A}_{\mathrm{SR}}(K)$: controls K distinguishable accounts, each facing its own margin and KYC cost c_{KYC}. Collusion across accounts is free. Non-adaptive within a block.

(c) *Adaptive multi-block* adversary $\mathcal{A}_{\mathrm{AM}}(T)$: chooses actions block-by-block over horizon T, observing all public state at each block including realised prices, oracle feeds, and insurance-fund balance. Strategy may condition on history.

(d) *Correlated-event* adversary $\mathcal{A}_{\mathrm{CE}}$: controls both a portfolio and, at cost c_{ev}, the realisation of a single event outcome (or a correlated bundle) used as oracle input by the clearing layer.

All classes are computationally unbounded in the information-theoretic sense; all gains and costs are measured in the exogenous settlement asset S.

14.2 Reflexivity rate-limiter and adaptive-oscillation bound

Theorem 12.2 removes the primary reflexive channel when the clearing layer is denominated in S. An exchange-level control may nonetheless expose a residual reflexive channel through a policy parameter $\Phi(t)$ (for example, an LP-fee tilt or a dynamic haircut) whose movement induces cross-pool liquidity reallocation. We model this residual channel and prove that a rate-limiter on Φ caps the amplitude of any adaptive adversarial oscillation.

Definition 14.2 (Policy-driven reflexivity proxy)**.** Fix $\Phi : \mathbb{R}_{\geq 0} \to \mathbb{R}$ a scalar policy process with $|\Phi(t)| \leq \Phi_{\max}$. The *policy-reflexivity proxy* is

$$\varrho_\Phi(t) \;:=\; \tau\, \frac{\lambda_\Phi\, L(t)\, \partial_\Phi f(L(t), \Phi(t))}{P_M(t)},$$

with $\lambda_\Phi > 0$ a policy LP-sensitivity rate. ϱ_Φ measures the instantaneous sensitivity of locked liquidity to Φ.

Definition 14.3 (Rate-limiter)**.** A *rate-limiter* with cap $\delta_\Phi > 0$ is the constraint $|\Phi(t+1) - \Phi(t)| \leq \delta_\Phi$ imposed at every block boundary by the protocol.

Proposition 14.4 (Amplitude bound under rate-limiter; upper and lower)**.** *Let $\mathcal{A} \in \mathcal{A}_{\mathrm{AM}}(T)$ be an adaptive multi-block adversary controlling LP deposit and withdrawal flows against the policy Φ, seeking to maximise the supremum oscillation $\sup_{0\leq t\leq T} |\varrho_\Phi(t) - \bar{\varrho}|$ about an equilibrium $\bar{\varrho}$. Under the rate-limiter of Definition 14.3 and a controller natural frequency ω_n derived from the response of ϱ_Φ to Φ,*

$$\sup_{0\leq t\leq T} |\varrho_\Phi(t) - \bar{\varrho}| \;\leq\; \min\!\left(A_{\max}, \frac{\delta_\Phi}{\omega_n}\right), \tag{55}$$

where $A_{\max} := \sup_\Phi |\varrho_\Phi - \bar{\varrho}|$ is the absolute amplitude ceiling of the response. There exists an adversarial strategy $\mathcal{A}^\star$ achieving

$$\sup_{0\leq t\leq T} |\varrho_\Phi(t) - \bar{\varrho}| \;\geq\; \frac{\delta_\Phi}{\omega_n}\left(1 - e^{-\omega_n T}\right) - \varepsilon_{\mathrm{obs}}, \tag{56}$$

where $\varepsilon_{\mathrm{obs}} > 0$ is an observation-noise slack. The upper and lower bounds agree up to the noise slack and the transient factor $1 - e^{-\omega_n T}$, which tends to 1 for $T \gg 1/\omega_n$.

Proof sketch. Upper bound. Model the response of ϱ_Φ to Φ as a linear first-order filter with natural frequency ω_n; the rate-limiter restricts $|\dot{\Phi}|$ (in the continuous limit) to at most δ_Φ per block. The amplitude of the forced response is bounded by the ratio of input slew rate to natural frequency: $|\Delta\varrho_\Phi|_\infty \leq |\dot{\Phi}|_\infty/\omega_n \leq \delta_\Phi/\omega_n$. Combining with the ceiling $A_{\max}$ gives (55). A nonlinear extension via Lyapunov-style comparison (Khalil [27], §4) controls deviations from linearity on compact state sets; the linear sector bound is tight on the principal harmonic.

Lower bound (matching attack). $\mathcal{A}^\star$ drives $\Phi(t+1) := \Phi(t) + \delta_\Phi \cdot \mathrm{sgn}(\cos(\omega_n t))$, a rate-maximal resonant sign drive. The forced-response amplitude of a first-order filter under resonant sign input satisfies $|\varrho_\Phi - \bar{\varrho}|_\infty \geq (\delta_\Phi/\omega_n)(1 - e^{-\omega_n T})$ by direct integration of the filter impulse response. Subtracting the observation-noise slack $\varepsilon_{\mathrm{obs}}$ accounts for finite-sample discretisation of the sign function against the measured ϱ_Φ. □

Remark 14.5 (Design-time choice of δ_Φ)**.** Equation (55) converts the rate-limiter cap δ_Φ into a deterministic amplitude bound: given a maximum tolerated amplitude $\bar{A} < 1 - \bar{\varrho}$ (to preserve the $\varrho < 1$ safety margin of Theorem 10.3), δ_Φ must satisfy $\delta_\Phi \leq \omega_n \bar{A}$. The linearity in δ_Φ is tight: reducing δ_Φ by a factor of c reduces the adversary's achievable amplitude by the same factor, up to the observation-noise slack.

Remark 14.6 (Limitations of the upper bound)**.** Three limitations attach to (55). (i) The first-order-filter model is a linearisation; genuine resonance in a higher-order controller would require a phase-margin computation from the full loop transfer function. The linear bound remains conservative against such resonance whenever the natural frequency is conservatively underestimated. (ii) The noise slack ε_{obs} is a finite-sample artefact of measuring ϱ_Φ; it is small when the measurement process shares the block cadence of the control. (iii) The bound is an amplitude bound on ϱ_Φ; it does not speak to the adversary's monetary profit, which is treated separately in Proposition 14.7.

Proposition 14.7 (Oscillation-profit upper bound)**.** *Under the rate-limiter of Definition 14.3 and the adversarial strategy of Proposition 14.4, the adversary's per-block expected net profit from LP deposit/withdrawal flows is bounded by*

$$\mathbb{E}[\pi_{osc}] \;\le\; \frac{\delta_\Phi}{\omega_n}\cdot\sigma_V\cdot q_{max} \;-\; 2f_{AMM}q_{max} - c_{gas}, \tag{57}$$

where σ_V is the per-block return volatility induced by the oscillation, q_{max} is the adversary's per-block trade size, f_{AMM} is the round-trip AMM fee, and c_{gas} is the gas overhead.

Proof sketch. The adversary's gross per-block profit is at most $|\Delta\varrho_\Phi|_\infty \cdot q_{max}$ times the fraction of the price move the adversary captures through its LP position, which is at most σ_V. Subtracting round-trip AMM fees $2f_{AMM}q_{max}$ and gas c_{gas} gives (57). □

Corollary 14.8 (Break-even condition)**.** *The oscillation attack is unprofitable whenever $\delta_\Phi < \omega_n(2f_{AMM} + c_{gas}/q_{max})/\sigma_V$. This is a closed-form monotone condition on the rate-limiter cap δ_Φ, the natural frequency ω_n, the AMM fee schedule f_{AMM}, the per-block volatility σ_V, the adversary's trade size q_{max}, and the gas overhead c_{gas}; no calibration is required to state it.*

Proof. Set (57) non-positive and solve for δ_Φ. □

14.3 Gated-buyback manipulation path

Propositions 12.7-12.10 gave closed-form *sufficient* bounds on the buyback gate parameters. We now give the adversarial *necessary* bound: a lower bound on the cost any rational adversary must incur to bypass the gate, achieved by a matching multi-block construction.

Proposition 14.9 (Gate-bypass manipulation cost; matching bounds)**.** *Let $\mathcal{A} \in \mathcal{A}_{AM}(T)$ be an adaptive multi-block adversary. Let (X, Y, Δ_{lag}) denote the buyback gate parameters of Definition 12.6. Assume:*

(A1) The insurance-fund S-reserve observer is an unbiased estimator with sub-Gaussian observation noise $\sigma_{obs} > 0$.
(A2) The recovery estimator $p_{recovery}$ satisfies $\|p_{recovery} - \mathbb{P}(\varrho_{t+1} < 1 \mid \mathcal{F}_t)\|_\infty \le \varepsilon_p$.
(A3) The autocorrelation time τ_ϱ of ϱ_t satisfies $\Delta_{lag} \ge 3\tau_\varrho$.

Let $\bar{x}_0$ be the attacker's observed baseline $\mathrm{IF}_S/\mathrm{IF}$ ratio, and let C^{IF}_{total} and D^{clear}_M denote respectively the total fund size and the market-depth cost of moving the ϱ-estimate by unit magnitude. Then any $\mathcal{A}$ that causes a false buyback trigger with probability at least $p^\star \in (1/2, 1)$ satisfies

$$\mathcal{C}_{adv}(\mathcal{A}) \;\ge\; (X - \bar{x}_0)_+ \cdot C^{IF}_{total} \;+\; D^{clear}_M \cdot \psi^\star_{gate}(p^\star)\cdot(1 + \Delta_{lag}/\tau_\varrho), \tag{58}$$

where $\psi^\star_{gate}(p^\star) := \sigma_{obs}\cdot\Phi^{-1}(p^\star)/\sqrt{\Delta_{lag}}$ is the per-block forcing required to overwhelm observation noise and Φ^{-1} is the inverse Gaussian CDF. There exists a strategy $\mathcal{A}^\star$ achieving equality up to a universal constant $C > 0$:

$$\mathcal{C}_{adv}(\mathcal{A}^\star) \;\le\; C\cdot\Big[(X - \bar{x}_0)_+ \cdot C^{IF}_{total} + D^{clear}_M \cdot \psi^\star_{gate}(p^\star)\cdot(1 + \Delta_{lag}/\tau_\varrho)\Big]. \tag{59}$$

The upper and lower bounds match up to a constant.

Proof sketch. Upper bound (necessary cost). To flip the gate, the adversary must raise the observed $\mathrm{IF}_S/\mathrm{IF}$ from $\bar{x}_0$ to at least X; this requires acquiring and contributing $(X - \bar{x}_0)_+ \cdot C^{\mathrm{IF}}_{\mathrm{total}}$ in settlement-asset units, a sunk cost. Independently, the adversary must force p_{recovery} to cross Y during the Δ_{lag}-block window. Because observations of ϱ are sub-Gaussian and filter through a Δ_{lag}-block moving average with decorrelation factor $(1 + \Delta_{\mathrm{lag}}/\tau_\varrho)^{-1}$, the per-block forcing to achieve a desired miscalibration at confidence $p^\star$ is $\psi^\star_{\mathrm{gate}}(p^\star)$ by a Chernoff-tail inversion. Integrating over the Δ_{lag}-block window and weighting by the market-depth cost D^{clear}_M gives the second term of (58). The two costs are additive because they affect independent gates.

Lower bound (matching attack). $\mathcal{A}^\star$ pre-funds the S-reserve to X (sunk cost $(X - \bar{x}_0)_+ C^{\mathrm{IF}}_{\mathrm{total}}$), then at each block $t \in [t_0, t_0 + \Delta_{\mathrm{lag}}]$ places a directional trade of magnitude $\psi^\star_{\mathrm{gate}}$ against the price of M, paying market impact $D^{\mathrm{clear}}_M \cdot \psi^\star_{\mathrm{gate}}$ per block. By Chernoff's inequality applied to the moving average of observations, the estimated p_{recovery} exceeds Y with probability at least $p^\star$, and the gate triggers on the Δ_{lag}-lagged observation. The total cost is the sum of sunk and market-impact terms, matching (58) up to the universal constant C absorbing sub-Gaussian proxy-norm-vs-variance slack. □

Remark 14.10 (Chernoff versus Hoeffding)**.** The $\psi^\star_{\mathrm{gate}}$ bound uses sub-Gaussian tail inversion; for bounded observations, a Hoeffding tail can replace it with a constant-factor improvement. The asymptotic scaling $\psi^\star = \Theta(\sigma_{\mathrm{obs}}/\sqrt{\Delta_{\mathrm{lag}}})$ is unchanged; the constant depends on the noise model.

Corollary 14.11 (Linear-in-Δ_{lag} cost amplification)**.** *The adversarial cost in Proposition 14.9 grows linearly in $\Delta_{\mathrm{lag}}/\tau_\varrho$, whereas the defender's miscalibration risk decays exponentially in $\Delta_{\mathrm{lag}}/\tau_\varrho$ (Proposition 12.10). The upper bound on Δ_{lag} is therefore a recovery-window constraint $\Delta_{\mathrm{lag}} \le \tau_{\mathrm{recovery}}$ from the timescale on which the buyback must act, not a cost-amplification constraint (the latter is unbounded).*

14.4 Waterfall attacks

We formalise four attacks against the waterfall of §8: hedge-discount leg-removal, synthetic-directional cross-margin, insurance-drain via Sybils, and strategic default. Each is stated with an explicit adversary, payoff bound, and (where possible) matching construction.

14.4.1 Hedge-discount leg-removal

Portfolio-margin computations reward hedges via the correlation-structured formula (4). An adversary may post a hedged pair, collect the margin discount, then remove one leg immediately before a clearing step to capture the discount as profit.

Proposition 14.12 (Hedge-discount attack; tight bound)**.** *Let $\mathcal{A} \in \mathcal{A}_{\mathrm{RS}}$ post a hedged position with legs (A, B) of notionals Q and correlation $\rho_{AB} \ge \rho_0 > 0$. Let $\alpha_{\mathrm{lag}} \in [0, 1)$ be the intra-block recompute lag of the margin engine (fraction of a block during which the unhedged position is not yet re-margined). The adversary's expected leg-removal profit satisfies*

$$\pi_{\mathrm{hedge}}(\mathcal{A}) \;\le\; \alpha_{\mathrm{lag}} \cdot \big(1 - \sqrt{(1 - \rho_{AB}^2)}\big) \cdot Q \cdot \sigma_B \cdot \sqrt{\alpha_{\mathrm{lag}}} \;-\; 2 f_{\mathrm{AMM}}\, Q, \tag{60}$$

where σ_B is the per-block standard deviation of leg B's returns. A matching adversary $\mathcal{A}^\star$ (remove leg at block-start, hold through the adversarial price move, close within the same block) achieves the right-hand side up to a factor of $\sqrt{2/\pi}$ (a Gaussian-tail Jensen slack).

Proof sketch. Upper bound. The discount captured by the hedge is $(1 - \sqrt{1 - \rho_{AB}^2}) \cdot Q$ margin units (from (4) evaluated at unit standalone margins). Over the re-margin lag

α_{lag}, the unhedged leg accumulates P&L with volatility $\sigma_B\sqrt{\alpha_{\text{lag}}}$; the adversary's optimal closing time captures at most $\sigma_B\sqrt{\alpha_{\text{lag}}}$ of this volatility (Doob's martingale stopping-time inequality). Scaling by the discount factor and subtracting round-trip AMM fees gives (60).

Lower bound (matching attack). $\mathcal{A}^\star$ removes the hedge at block-start, holds for α_{lag} fraction of the block, and closes at the realised price. Expected directional P&L is $\sigma_B\sqrt{\alpha_{\text{lag}}/(2\pi)}\cdot Q$ (Gaussian-tail expected absolute value). Matching up to the Jensen factor $\sqrt{2/\pi}$ follows. □

Remark 14.13 (Defence: zero-lag re-margining)**.** The attack is blocked outright by enforcing $\alpha_{\text{lag}} = 0$: any margin-state transition arising from a leg removal recomputes portfolio margin before the next adversary action becomes valid. The AMM-deterministic price-impact kernel (Assumption 2.1) permits deterministic intra-block re-margining at block-commit time, realising $\alpha_{\text{lag}} = 0$ in the unified settlement model of §7.

14.4.2 *Synthetic-directional via cross-margin*

A second hedge-exploit arises when cross-margining across multiple spoke assets lets an adversary assemble a *synthetic directional* position at a cost below the standalone margin. We bound the advantage.

Proposition 14.14 (Synthetic-directional cross-margin upper bound)**.** *Let $\mathcal{A}\in\mathcal{A}_{\text{CE}}$ correlate the realisation of spoke asset B_j's event with a pre-positioned portfolio Π of spoke assets $\{B_i\}_{i\neq j}$. Let ρ_{pre} be the pre-event correlation estimate used for margining and ρ_{post} the realised post-event correlation. Under the parametric ES model* (3)*,*

$$\pi_{\text{synth}}(\mathcal{A}) \;\le\; (\rho_{\text{pre}}-\rho_{\text{post}})\cdot Q\cdot\sigma_B \;-\; c_{\text{ev}}, \tag{61}$$

where c_{ev} is the event-manipulation cost of Definition 14.1(d). A matching strategy (buy correlated synthetic, trigger correlation breakdown event, unwind) achieves the right-hand side up to AMM fees.

Proof sketch. The cross-margin discount scales with $\sqrt{1-\rho_{\text{pre}}^2}$; post-event, the realised portfolio variance scales with $\sqrt{1-\rho_{\text{post}}^2}$. The margin deficit at time of unwind is therefore $(\rho_{\text{pre}}-\rho_{\text{post}})\cdot Q\cdot\sigma_B$ to first order in the correlation differential. The event-manipulation cost c_{ev} subtracts by Definition 14.1(d). Matching follows from a corresponding pre-position + trigger construction. □

Remark 14.15 (Defence: correlation stress buffer)**.** Setting the margin-engine correlation input to $\max(\rho_{\text{pre}},\rho_{\text{stress}})$ for a stress floor $\rho_{\text{stress}}\ge\rho_{\text{post}}$ eliminates the upside: the attack captures zero whenever the margin is already computed at the stressed correlation. Setting $\rho_{\text{stress}} = 1$ (worst-case correlation) restores isolated-margin capital requirements and blocks all cross-margin discounts; intermediate values trade capital efficiency against attack surface monotonically.

14.4.3 *Insurance-fund drain via Sybils*

An adversary may open K small accounts, each with collateral C_{per}, and drive each to a small loss that triggers a per-account Stage-2 insurance-fund draw. If per-account identity cost is zero, the attack scales linearly and the insurance fund drains.

Proposition 14.16 (Sybil-drain lower bound on K)**.** *Let $\mathcal{A}\in\mathcal{A}_{\text{SR}}(K)$ open K accounts and drive each to a loss of magnitude $\varepsilon\cdot C_{\text{per}}$ with $\varepsilon\in(0,1]$. Let $\chi>0$ be the per-account concentration* *cap (fraction of insurance fund drawable per account per epoch) and c_{KYC} the identity cost per account. The attack is profitable* only if

$$K \geq 1/\chi \quad \text{and} \quad K \cdot (\varepsilon \cdot C_{\mathrm{per}} - c_{\mathrm{KYC}}) > 0. \tag{62}$$

Conversely, for any $K \geq 1/\chi$, the adversary $\mathcal{A}^\star$ (open K accounts, drive each to loss $\varepsilon \cdot C_{\mathrm{per}}$) achieves net profit $K(\varepsilon \cdot C_{\mathrm{per}} - c_{\mathrm{KYC}})$ provided $\varepsilon \cdot C_{\mathrm{per}} > c_{\mathrm{KYC}}$. The threshold $c_{\mathrm{KYC}} \geq \varepsilon \cdot C_{\mathrm{per}}$ eliminates the attack for all K.

Proof. The concentration cap χ restricts the per-account drawable fraction of the fund; draining the fund requires $K \geq 1/\chi$ accounts. The per-account economics is loss minus identity cost, so the aggregate is $K(\varepsilon \cdot C_{\mathrm{per}} - c_{\mathrm{KYC}})$. The converse is immediate. □

Remark 14.17 (Defence: identity gating)**.** The proposition isolates the two independent defences: (i) concentration cap χ forces $K \geq 1/\chi$ scale; (ii) identity cost c_{KYC} makes scale unprofitable. Neither alone suffices: a high χ permits a single account to drain the fund; a zero c_{KYC} admits unlimited accounts. Combined with $c_{\mathrm{KYC}} \geq \varepsilon \cdot C_{\mathrm{per}}$, the attack is ruled out regardless of χ.

14.4.4 Strategic default with friction cost

A rational defaulter may prefer to default and absorb the waterfall consequences if the realised off-exchange P&L exceeds the collateral seized. We account for the off-exchange hedge friction cost.

Proposition 14.18 (Strategic-default profitability condition)**.** *Let $\mathcal{A} \in \mathcal{A}_{\mathrm{SR}}(K)$ hold positions totaling notional Q across K accounts, each with per-account collateral C_{per}. Let $\|\Delta P\|$ be the realised adverse price move, σ_{basis} the per-block cross-exchange basis volatility of the hedge, and $p_{\mathrm{trigger}} := \mathbb{P}(|\Delta P| \geq \delta_{\mathrm{default}})$ the probability that the move crosses the default threshold. The expected net profit from a strategic-default attack is*

$$\mathbb{E}[\pi_{\mathrm{strat}}] = p_{\mathrm{trigger}} \cdot \left(Q \cdot |\Delta P| - K \cdot C_{\mathrm{per}} - C_{\mathrm{off}}(Q,T) - Q \cdot \sigma_{\mathrm{basis}} \cdot \sqrt{2/\pi} \right) - (1 - p_{\mathrm{trigger}}) \cdot c_{\mathrm{position}}, \tag{63}$$

where $C_{\mathrm{off}}(Q,T)$ is the off-exchange hedge cost over holding period T and c_{position} is the cost of holding positions through the non-triggering scenario. The attack is profitable only if the bracketed term exceeds the holding cost weighted by the failure probability $(1 - p_{\mathrm{trigger}})/p_{\mathrm{trigger}} \cdot c_{\mathrm{position}}$.

Proof sketch. The triggering scenario yields gross P&L $Q \cdot |\Delta P|$ less collateral seizure $K \cdot C_{\mathrm{per}}$ and hedge friction $C_{\mathrm{off}} + Q\sigma_{\mathrm{basis}}\sqrt{2/\pi}$ (Gaussian expected-absolute-value of basis noise). The non-triggering scenario yields a position-holding cost c_{position}. Taking expectations weights by p_{trigger} and $(1 - p_{\mathrm{trigger}})$. □

Remark 14.19 (Sensitivity of the profitability boundary)**.** The profitability boundary of (63) depends monotonically on σ_{basis}, $C_{\mathrm{off}}(Q,T)$, and the tail index of $|\Delta P|$. Under a Pareto tail with shape ξ and threshold $\delta_{\mathrm{default}}/x_{\min}$, p_{trigger} follows the standard power-law form; the sensitivity of the boundary to $(\xi, \sigma_{\mathrm{basis}}, KC_{\mathrm{per}}/Q)$ is characterised by direct substitution.

14.5 Waterfall cascade under worst-case concurrent defaults

Theorem 7.6 establishes that the waterfall is order-independent under simultaneous defaults when stages are applied level-by-level. We now bound the total consumable capacity under worst-case concurrent defaults and identify the capacity-exhaustion frontier.

Proposition 14.20 (Worst-case cascade capacity)**.** *Under Theorem 7.5, the total absorbable shortfall across any concurrent-default set $D \subseteq \mathcal{A}$ satisfies*

$$\sum_{k \in D} \Delta_k \;\le\; \sum_{k \in D} C_k \;+\; E_{\mathrm{SIG}} \;+\; \mathrm{IF} \;+\; \kappa \cdot \textstyle\sum_{k \notin D} C_k \;+\; \textstyle\sum_{k \in \mathcal{P}} \pi_k \quad\Longrightarrow\quad \Delta^{\mathrm{res}} = 0. \tag{64}$$

The five terms correspond to Stages 1 through 5 of the waterfall (defaulter margin, CCP skin-in-the-game, mutualised fund, capped socialisation, ADL). The inequality is tight: there exist default configurations saturating any single term while the others carry zero.

Proof. The implication follows from Theorem 7.5. Tightness is by construction: a concentrated default saturates Stage 1; a fully-collateral-depleting default against the SIG tranche saturates Stage 2; an aggregate deficit consuming IF exactly saturates Stage 3; a distributed residual absorbing $\kappa \sum_{k \notin D} C_k$ saturates Stage 4; a profit consumption equal to $\sum_{k \in \mathcal{P}} \pi_k$ saturates Stage 5. □

Corollary 14.21 (Capacity-failure frontier)**.** *The worst-case concurrent-default loss saturating $\Delta^{\mathrm{res}} > 0$ requires simultaneous exhaustion of all stages. Writing the aggregate capacity as $\mathcal{C}^{\mathrm{total}}(D) := \sum_{k \in D} C_k + E_{\mathrm{SIG}} + \mathrm{IF} + \kappa \sum_{k \notin D} C_k + \sum_{k \in \mathcal{P}} \pi_k$, the shortfall is $\Delta^{\mathrm{res}}(D) = \max(0, \sum_{k \in D} \Delta_k - \mathcal{C}^{\mathrm{total}}(D))$. The residual vanishes for every default set D satisfying $\sum_{k \in D} \Delta_k \le \mathcal{C}^{\mathrm{total}}(D)$, a linear constraint on the default-set configuration.*

Remark 14.22 (Reflexive-cascade elimination)**.** Under Theorem 12.2 (exogenous settlement), the terms C_k, IF, and π_k in (64) do not depend on the hub price P_M; $\mathcal{C}^{\mathrm{total}}$ is therefore invariant to any reflexive cascade on P_M. This is the quantitative statement of Corollary 12.3: the worst-case capacity of the waterfall is fixed at the clearing configuration and does not collapse with the hub price.

14.6 Completeness summary

Table 2 records the adversary model, attack budget, defence, upper bound, and lower bound for each adversarial proposition above. An entry marked *matching* indicates the upper and lower bounds agree up to a universal constant factor; an entry marked *one-sided* indicates that only one direction is proved (the other is either a named open problem or is ruled out structurally by the defence).

Proposition	Adversary	Defence	Upper bound	Lower bound
Prop 14.4	$\mathcal{A}_{\mathrm{AM}}(T)$	Rate-limiter δ_Φ	$\min(A_{\max}, \delta_\Phi/\omega_n)$	$\delta_\Phi/\omega_n \cdot (1 - e^{-\omega_n T})$
Prop 14.7	$\mathcal{A}_{\mathrm{AM}}(T)$	AMM fee f_{AMM}	$\delta_\Phi \sigma_V q_{\max}/\omega_n - 2 f_{\mathrm{AMM}} q_{\max}$	one-sided
Prop 14.9	$\mathcal{A}_{\mathrm{AM}}(T)$	$(X, Y, \Delta_{\mathrm{lag}})$ gate	$(X - \bar{x}_0)C^{\mathrm{IF}} + D\psi^\star(1 + \Delta_{\mathrm{lag}}/\tau_\varrho)$	matching up to const C
Prop 14.12	$\mathcal{A}_{\mathrm{RS}}$	Zero-lag re-margining	$\alpha_{\mathrm{lag}}(1 - \sqrt{1-\rho^2}) Q \sigma_B$	matching up to $\sqrt{2/\pi}$
Prop 14.14	$\mathcal{A}_{\mathrm{CE}}$	Correlation stress buffer	$(\rho_{\mathrm{pre}} - \rho_{\mathrm{post}}) Q \sigma_B - c_{\mathrm{ev}}$	matching up to AMM fe
Prop 14.16	$\mathcal{A}_{\mathrm{SR}}(K)$	χ-cap + KYC cost	$K(\varepsilon C_{\mathrm{per}} - c_{\mathrm{KYC}})$ if $K \ge 1/\chi$	matching by constructic
Prop 14.18	$\mathcal{A}_{\mathrm{SR}}(K)$	OI cap + basis risk	eq (63)	one-sided upper
Prop 14.20	default-set D	4-stage waterfall	$\mathcal{C}^{\mathrm{total}}(D)$	matching by constructic

Table 2: Adversarial-bound summary. All statements are proved in sketch form with adversary model stated explicitly.

Remark 14.23 (One-sided bounds are structural)**.** The two one-sided entries (Prop 14.7, Prop 14.18) are upper bounds without matching constructions. In both cases, the lower bound reduces to the unprofitability threshold, which is a calibration statement (parameters satisfy inequalities that zero the profit) rather than a construction (an adversary achieving a specific profit). The structural-impossibility form is recorded as a remark attached to each proposition.

15 Open Problems

Open Problem 1 (Optimal diversification bounds)**.** Derive the optimal issuer-concentration cap $\omega_{\max}$ and non-stablecoin floor $\omega_{\min}$ for the insurance fund, minimising the joint-tail bound of Theorem 13.4 subject to a capital-efficiency constraint. Characterise the Pareto frontier in $(\omega_{\max}, \omega_{\min})$.

Open Problem 2 (Optimal liquidation sequencing under coupling)**.** Given simultaneous liquidation volumes $\{q_i\}$ across pools with coupling exponents $\{g_i\}$, find the sequential liquidation order minimising total slippage. The deterministic AMM kernel makes this a finite-horizon dynamic program in the sense of Bertsimas and Lo [41]; characterise when the greedy policy (largest pool first) is optimal and when it is suboptimal.

Open Problem 3 (Optimal shrinkage intensity)**.** The blending weight w_s in Proposition 3.4 is exogenous. Derive the Ledoit-Wolf-type optimal w_s^* for structural priors with known sparsity pattern (a graph Laplacian), extending [9] to the graph-structured-prior case.

Open Problem 4 (Closed-form haircut optimisation under copula)**.** Given a copula C on joint stablecoin depegs, derive the optimal haircut schedule $\{h_k^*\}$ minimising the expected copula-tail loss of Theorem 13.4 subject to a total-haircut budget constraint. For Gaussian, Clayton, and Gumbel copulas, characterise when $\{h_k^*\}$ is unique and when ties occur.

Open Problem 5 (Canard invariant manifold extension)**.** Extend Proposition 11.4 to characterise the canard invariant manifold structure near Σ: the family of trajectories entering Σ in finite time from the subcritical side, their scaling law as a function of distance to $\mathcal{E}$, and the connection to the fast layer. The blow-up desingularisation of Krupa and Szmolyan [42] and the canard theory of Benoît et al. [20, 21] apply; the structural-endogeneity setting adds the constraint that $\varepsilon = 1 - \varrho$ is state-dependent.

Open Problem 6 (Endogenous settlement-asset contagion)**.** Theorem 12.2 assumes the settlement asset's price is exogenous to L. When the clearing system grows sufficiently large that its settlement flow affects the settlement asset's own stability, the exogeneity assumption fails. Derive the scaling condition under which this feedback becomes non-negligible, and characterise the induced second-order reflexive channel.

Open Problem 7 (Sharp bound on cascade depth)**.** Corollary 8.5 bounds the cumulative absorbed obligation across a correlated-default cascade. The number of rounds remains open. Given a density $\phi(C^{\mathrm{buf}})$ on the per-counterparty margin buffer above maintenance, derive a sharp bound on the expected cascade depth $\mathbb{E}[N_{\mathrm{rds}}]$ as a function of ϕ, the socialisation cap κ, and the concurrent-default size $|D|$.

Open Problem 8 (Tight adversarial lag-window tail)**.** Proposition 12.10 bounds the linear autocorrelation at lag Δ_{lag}; Remark 12.11 observes that a bounded-variance AR(1) yields a $\exp(-c\Delta_{\mathrm{lag}})$ tail on adversarial innovation sequences, while a sub-Gaussian $\exp(-c\Delta_{\mathrm{lag}}^2)$ tail requires an explicit variance proxy. Identify the minimal distributional assumption on the innovation under which the quadratic tail holds, or construct an adversarial innovation sequence saturating the linear bound.

Open Problem 9 (LP concentration versus fee-revenue scale)**.** Proposition 10.8 shows that a concentrated LP distribution can trigger a canard crossing through a single withdrawal. Fee-revenue scale encourages LP concentration [38, 39]. Jointly optimise the fee schedule, the per-LP withdrawal-rate limit, and the concentration cap $s_{\max}$ to maximise pool-level welfare subject to a guaranteed lower bound on the time to Σ under any single-LP withdrawal.

Open Problem 10 (Heavy-tailed ES aggregation)**.** Extend Theorem 3.2 to a Student-t elliptical family with explicit tail index ν, or to a copula-tail aggregation [10, 11]. Characterise the efficiency-ratio correction as a function of ν and the tail-dependence coefficients; identify the regime in which the elliptical bound is order-accurate and the regime in which it systematically understates margin.

Open Problem 11 (Porting versus liquidation frontier)**.** Proposition 8.3 gives the threshold under which porting dominates liquidation. Under surviving-member capacity process $C_{k'}^{\text{slack}}(t)$ and AMM slippage kernel $\mathcal{L}_{\text{slip}}(D, t)$, characterise the optimal porting-versus-liquidation split as a function of default-set size, hub-asset concentration in surviving members, and pool-depth dynamics. Relate to the EMIR Art. 48 porting framework [36].

Open Problem 12 (Dynamic correlation margin under event calendars)**.** Proposition 14.14 bounds the synthetic-directional attack under a static stress-correlation floor. Develop an event-aware dynamic correlation model in which the margin engine's ρ^{margin} updates against a scheduled event calendar, and prove a tighter attack bound under the event-conditional measure.

A Numerical calibrations for cross-margin efficiency

Equation (14) expresses $m^{\text{cross}}/m^{\text{iso}}$ as a closed form in $(N_{\text{br}}, \bar{\rho}, \delta_\rho, D^{sp}, D^{pE})$. Representative numerical evaluations, useful for cross-checking an implementation against the formula, are collected below.

Scenario	Inputs	$m^{\text{cross}}/m^{\text{iso}}$
Balanced portfolio, unshocked	$N_{\text{br}} = 10$, $\bar{\rho} = 0.20$, $\delta_\rho = 0$, $D^{sp} + D^{pE} = 0.08$	≈ 0.49
Balanced portfolio, shocked	$N_{\text{br}} = 10$, $\bar{\rho} = 0.20$, $\delta_\rho = 0.20$, $D^{sp} + D^{pE} = 0.08$	≈ 0.63
Tighter shock regime	$N_{\text{br}} = 10$, $\bar{\rho} = 0.20$, $\delta_\rho = 0.30$, $D^{sp} + D^{pE} = 0.08$	≈ 0.70
Concentrated hedge	$N_{\text{br}} = 10$, $\bar{\rho} = 0.20$, $\delta_\rho = 0.20$, $D^{sp} + D^{pE} = 0.04$	≈ 0.66

Table 3: Representative numerical values of the ratio (14) at equal per-spoke notionals and a non-binding floor. These are reference points on the sensitivity surface, not theorem outputs; they serve as a cross-check against implementations of Definition 3.5.

Empirical basis-volatility calibrations (used to set d_{sp} in (9)) should be drawn from the target exchange's observed basis series; representative industry ranges for liquid perpetuals are 99.5% ES basis of 2-4% of notional, and 8-12% in stressed thin markets. These calibrations are implementation parameters, not theorem outputs.

References

[1] L. Eisenberg and T. H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236-249, 2001.

[2] L. C. G. Rogers and L. A. M. Veraart. Failure and rescue in an interbank network. *Management Science*, 59(4):882-898, 2013.

[3] R. Cifuentes, G. Ferrucci, and H. S. Shin. Liquidity risk and contagion. *Journal of the European Economic Association*, 3(2-3):556-566, 2005.

[4] R. Cont and L. Wagalath. Fire sales forensics: Measuring endogenous risk. *Mathematical Finance*, 26(4):835-866, 2016.

[5] V. V. Acharya and A. Bisin. Counterparty risk externality: Centralized versus over-the-counter markets. *Journal of Economic Theory*, 149:153-182, 2013.

[6] C. Pirrong. The economics of central clearing: Theory and practice. *ISDA Discussion Papers Series*, 2011.

[7] Chicago Mercantile Exchange. SPAN: Standard portfolio analysis of risk. Technical documentation, CME Group, 1988 (updated continuously).

[8] P. Jorion. *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill, 3rd edition, 2006.

[9] O. Ledoit and M. Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365-411, 2004.

[10] P. Embrechts, A. McNeil, and D. Straumann. Correlation and dependence in risk management: Properties and pitfalls. In *Risk Management: Value at Risk and Beyond*, pages 176-223. Cambridge University Press, 2002.

[11] H. Joe. *Dependence Modeling with Copulas*. Monographs on Statistics and Applied Probability 134, CRC Press, 2014.

[12] H. Adams, N. Zinsmeister, M. Salem, R. Keefer, and D. Robinson. Uniswap v3 core. Technical report, Uniswap Labs, 2021.

[13] G. Angeris, H.-T. Kao, R. Chiang, C. Noyes, and T. Chitra. An analysis of Uniswap markets. *Cryptoeconomic Systems*, 1(1), 2021.

[14] H. Uhlig. A Luna-tic stablecoin crash. *NBER Working Paper* 30256, 2022.

[15] A. Briola, D. Vidal-Tomás, Y. Wang, and T. Aste. Anatomy of a stablecoin's failure: The Terra-Luna case. *Finance Research Letters*, 51:103358, 2023.

[16] R. Clements. Built to fail: The inherent fragility of algorithmic stablecoins. *Wake Forest Law Review Online*, 11:131-154, 2022.

[17] Chainalysis. The collapse of TerraUSD: A forensic reconstruction. Chainalysis Crypto Crime Report, 2022.

[18] N. Fenichel. Geometric singular perturbation theory for ordinary differential equations. *Journal of Differential Equations*, 31(1):53-98, 1979.

[19] C. K. R. T. Jones. Geometric singular perturbation theory. In *Dynamical Systems (Montecatini Terme, 1994)*, Lecture Notes in Mathematics 1609, Springer, pages 44-118, 1995.

[20] E. Benoît, J.-L. Callot, F. Diener, and M. Diener. Chasse au canard. *Collectanea Mathematica*, 32(1-2):37-119, 1981.

[21] E. F. Mishchenko, Y. S. Kolesov, A. Y. Kolesov, and N. K. Rozov. *Asymptotic Methods in Singularly Perturbed Systems*. Monographs in Contemporary Mathematics, Consultants Bureau, 1994.

[22] Y. A. Kuznetsov. *Elements of Applied Bifurcation Theory*. Applied Mathematical Sciences 112, Springer, 3rd edition, 2004.

[23] L. Perko. *Differential Equations and Dynamical Systems*. Texts in Applied Mathematics 7, Springer, 3rd edition, 2001.

[24] M. Golubitsky and D. G. Schaeffer. *Singularities and Groups in Bifurcation Theory, Vol. I*. Applied Mathematical Sciences 51, Springer, 1985.

[25] M. K. Brunnermeier and L. H. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201-2238, 2009.

[26] T. Adrian and H. S. Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418-437, 2010.

[27] H. K. Khalil. *Nonlinear Systems*. Prentice Hall, 3rd edition, 2002.

[28] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

[29] R. Cont and T. Kokholm. Central clearing of OTC derivatives: Bilateral vs multilateral netting. *Statistics and Risk Modeling*, 31(1):3-22, 2014.

[30] D. Duffie and H. Zhu. Does a central clearing counterparty reduce counterparty risk? *Review of Asset Pricing Studies*, 1(1):74-95, 2011.

[31] D. Murphy. *OTC Derivatives: Bilateral Trading and Central Clearing*. Palgrave Macmillan, 2013.

[32] J. A. Cruz Lopez, J. H. Harris, C. Hurlin, and C. Pérignon. CoMargin. *Journal of Financial and Quantitative Analysis*, 52(5):2183-2215, 2017.

[33] S. Ghamami and P. Glasserman. Does OTC derivatives reform incentivize central clearing? *Journal of Banking and Finance*, 76:127-140, 2017.

[34] B. Biais, F. Heider, and M. Hoerova. Risk-sharing or risk-taking? Counterparty risk, incentives, and margins. *Journal of Finance*, 71(4):1669-1698, 2016.

[35] Committee on Payments and Market Infrastructures and Board of the International Organization of Securities Commissions. *Principles for Financial Market Infrastructures*. Bank for International Settlements, 2012.

[36] European Parliament and Council. Regulation (EU) No 648/2012 of 4 July 2012 on OTC derivatives, central counterparties and trade repositories (EMIR); Commission Delegated Regulation (EU) No 153/2013 on regulatory technical standards on requirements for CCPs. *Official Journal of the European Union*, L 201:1-59, 2012; L 52:41-74, 2013.

[37] U.S. Congress. Dodd-Frank Wall Street Reform and Consumer Protection Act, Title VIII: Payment, Clearing, and Settlement Supervision. *Public Law* 111-203, 124 Stat. 1376, 2010.

[38] A. Evans, G. Angeris, and T. Chitra. Optimal fees for geometric mean market makers. In *Financial Cryptography and Data Security 2021 Workshops*, Lecture Notes in Computer Science 12676, Springer, pages 65-79, 2021.

[39] J. Milionis, C. C. Moallemi, T. Roughgarden, and A. L. Zhang. Automated market making and loss-versus-rebalancing. *arXiv preprint arXiv:2208.06046*, 2022.

[40] F. Allen and D. Gale. Financial contagion. *Journal of Political Economy*, 108(1):1-33, 2000.

[41] D. Bertsimas and A. W. Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1-50, 1998.

[42] M. Krupa and P. Szmolyan. Relaxation oscillation and canard explosion. *Journal of Differential Equations*, 174(2):312-368, 2001.

[43] D. X. Li. On default correlation: A copula function approach. *Journal of Fixed Income*, 9(4):43-54, 2000.

[44] C. Acerbi and D. Tasche. On the coherence of Expected Shortfall. *Journal of Banking and Finance*, 26(7):1487-1503, 2002.

[45] M. Avellaneda and S. Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217-224, 2008.

[46] B. P. Welford. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419-420, 1962.

[47] M. Peters and J. Ye. A convex optimisation formulation of parimutuel mechanisms. Stanford University working paper, 2006.

[48] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620-630, 1957.

[49] M. J. Wainwright and M. I. Jordan. *Graphical Models, Exponential Families, and Variational Inference*. Foundations and Trends 1(1-2):1-305, 2008.

[50] I. Csiszár and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2nd edition, 2011.

[51] P. J. Davis and P. Rabinowitz. *Methods of Numerical Integration*. Academic Press, 2nd edition, 1984.

[52] D. T. Breeden and R. H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621-651, 1978.

[53] D. MacKenzie and T. Spears. "The formula that killed Wall Street": The Gaussian copula and modelling practices in investment banking. *Social Studies of Science*, 44(3):393-417, 2014.

[54] F. Salmon. Recipe for disaster: The formula that killed Wall Street. *Wired*, 17(3), February 2009.