ECONOMIC RISK RESEARCH

Vanna Protocol
Economic Risk Assessment

A quantitative assessment of the economic, market, liquidity, leverage, carry, incentive and systemic risks of a two-sided leveraged lending and on-chain margin protocol on Stellar / Soroban. Focus is on interest-rate and carry dynamics, liquidation economics, protocol solvency, and value-at-risk under stress.

Scope. Economic and quantitative risk only. This is *not* a smart-contract security audit, code-quality review, or vulnerability assessment; those are to be performed separately by a dedicated security auditor.

Independent Economic Risk Assessment

Prepared for governance, investor diligence, and protocol design review

Confidential draft • illustrative calibration

Reading note on the quantitative results

This report separates three classes of number. The distinction is load-bearing and the reader should keep it in mind throughout.

1. Exact / structural results [exact]. The interest-rate curve, the leverage bounds, the recursive-leverage series, the health-factor algebra, and the carry / equity-decay trajectories are *exact functions of the parameters embedded in the Vanna contracts* (the rate-model coefficients $c_1 = 0.1$, $c_2 = 0.3$, $c_3 = 3.5$ and the two health thresholds, 1.25 opening and 1.10 liquidation). They are not assumptions; they are what the protocol, as parameterised, computes, and each has been re-derived from the opt-5 code. These are stated as fact.

2. Illustrative / simulated results [illustrative]. The Monte-Carlo Value-at-Risk, Conditional VaR, bad-debt and liquidation-frequency figures are generated under an *assumed* price process (a correlated GARCH(1,1) diffusion with an assumed volatility surface, and in the stress regime an assumed flash-crash jump process) and *assumed* market depth. They are **not calibrated to live Stellar / Blend / Reflector market data.** They are directionally and structurally informative but the absolute magnitudes must be re-estimated against real data (Section 3.4) before they drive governance. Every such figure is labelled [illustrative].

3. Code-confirmed mechanism facts [code-confirmed]. Statements about how a mechanism actually behaves in opt-5 — for example that the debt term is interest-accrued, that liquidation reverts for Blend/LP accounts, or that the origination fee is borrower-borne — are read directly from the contracts and tagged [code-confirmed], with the corresponding security finding named where relevant.

No number in this report should be quoted to a third party without its accompanying class label.

Verification status (this edition, opt-5). This edition revises the prior assessment after a full opt-5 security audit; structural claims were reconciled against the opt-5 contract set. Reconciliation *confirmed* the rate-curve and leverage arithmetic; the two headline theses (structural negative carry; hollow / *failing* liquidation economics); and the design gaps carrying them (no collateral factors, no reserve, no caps, external-venue dependency, and a collateral double-count). It *refuted* four secondary claims an earlier draft had flagged for audit — the health factor is interest-accrued (no wrong-way signal on the debt side); opening and liquidation use distinct 1.25/1.10 thresholds (a runway exists and opening is capped $\approx 5\times$); the vToken NAV is struck against total assets (no utilisation-timing dilution); and spot pricing is TWAP-gated (not manipulable at the instant of action). It *corrected* one item the earlier draft had backwards — the origination fee is borrower-borne, not a pool drain (opt-5 NEW-M3). And it *escalated* two: liquidation for the flagship externally-deployed strategy *deterministically reverts* (REG-01), and the collateral double-count is *repeatable and effectively unbounded* via a swap-laundering path (NEW-11). The corrected items are recorded in Section 1.5; every consequential economic conclusion is tagged to a security finding.

Contents

CHAPTER 1

Executive Summary

1.1 Vanna in economic terms

Vanna couples a **pooled lending market** to an **isolated-margin leverage engine**. On the lending side, liquidity providers (LPs) deposit XLM, USDC or EURC into per-asset pools, receive `vTokens` representing a pro-rata claim, and earn interest paid by borrowers — mechanically similar to Blend. On the margin side, users open Smart Margin Accounts, post collateral, borrow the LPs' liquidity, and deploy it into external Soroban venues — **Blend, Aquarius and Soroswap** are all live in opt-5's collateral-valuation surface — to construct leveraged yield strategies. Each external position is represented inside Vanna by a **TrackToken** whose value is read live from the external venue's exchange rate multiplied by a Reflector oracle price. Solvency is enforced by a **health factor** HF = total balance/borrow balance with *two* thresholds — a borrow (opening) threshold at 1.25 and a liquidation threshold at 1.10 — and the debt term is interest-accrued [code-confirmed], so the debt side of the signal is unbiased.

From an economic standpoint the protocol is a leveraged carry machine: it manufactures leverage on an external yield source and intermediates the resulting credit risk between borrowers and a pool of LPs. Its risk therefore concentrates in four places, in decreasing order of economic importance: (i) the *carry* between the internal borrow rate and the external yield; (ii) the *economic and mechanical effectiveness of liquidation* — whether unhealthy positions can be closed at all, promptly, profitably for a keeper, and with enough buffer to cover slippage; (iii) *collateral-side signal quality* (full oracle value, no risk factor, *plus* a confirmed accounting double-count); and (iv) *liquidity and first-loss allocation* on the LP side. A fifth axis, elevated in this edition, is (v) *governance parameter risk* — the blast radius of unbounded, un-timelocked admin controls. This report quantifies each.

1.2 Economic risk posture

> **Overall economic-risk posture: the design is coherent, but as implemented in opt-5 it is not ready to custody value at meaningful leverage — and the reason is now sharper than "economic fragility." The liquidation path for the flagship, externally-deployed strategy is not merely under-incentivised; it *deterministically reverts*. The dominant tail is a mechanism / code artifact, fixable through parameterisation and code fixes rather than a market inevitability.**

Two economic theses dominate, both confirmed. First, the protocol's flagship strategy — leveraged external supply — is **structurally negative-carry** under realistic external yields: the borrow curve charges 17.5% at half utilisation against low-single-digit external supply APRs, so a leveraged position bleeds equity even when nothing goes wrong in the market (Chapter 5). Second, the **economics of liquidation now fail outright**: under an illustrative stress regime, the expected-shortfall ($CVaR_{99}$) bad debt from a liquidation process that does not execute ranges from roughly 0.49 to 0.73 per unit of debt across the health-factor spectrum, whereas a liquidation that *does* execute cuts that same tail to ≈ 0.0002 (Chapter 10). The gap is almost entirely attributable to liquidation economics — incentive, buffer, and latency — not to market volatility.

New in this edition: for accounts holding a Blend or LP position — the intended primary activity — liquidation *cannot execute* in opt-5 (finding REG-01). So the "no-working-liquidation" column is not a bracket endpoint; it is the *operating point* for the flagship strategy, and the ≈ 0.0002 "functioning" figure is currently unreachable for those accounts until REG-01 is fixed. A third confirmed defect — a *repeatable* collateral double-count (Chapter 6) — removes the leverage ceiling, so per-account bad debt is bounded only by pool liquidity.

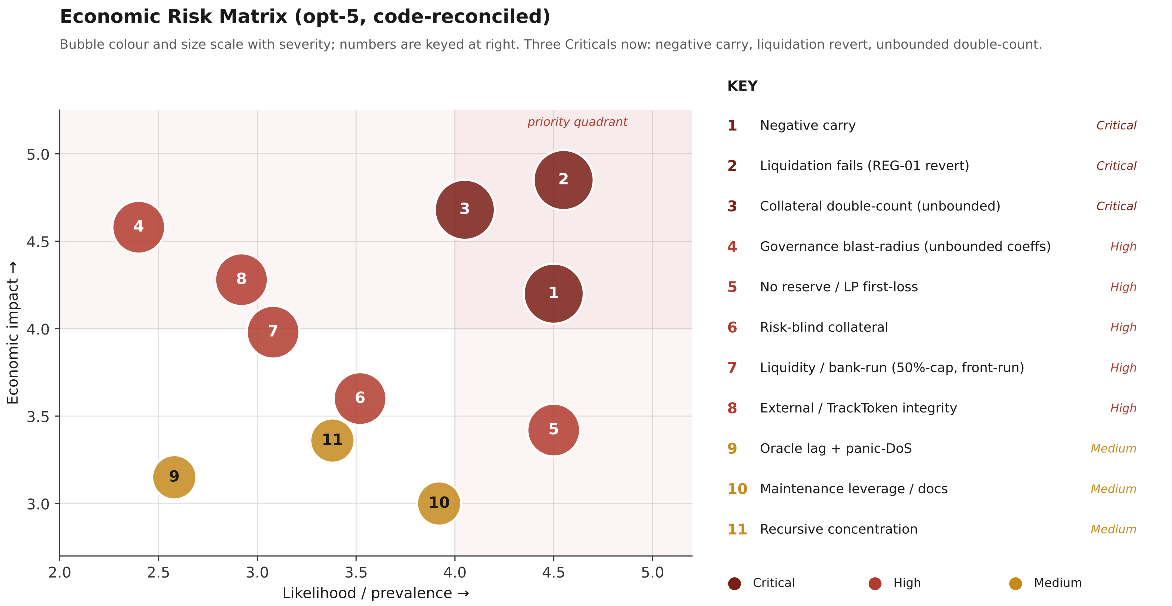

Figure 1.1: Economic risk matrix (opt-5, code-reconciled). Position and bubble size scale with composite severity. There are now three Criticals — negative carry, liquidation revert (REG-01), and the unbounded collateral double-count — and governance parameter risk enters as a low-likelihood, catastrophic-impact tail. [illustrative] for positioning; grounded in the quantitative chapters and named security findings.

1.3 Principal economic risks (ranked)

We frame findings as *economic risk themes* rather than code defects. Each is developed, quantified and given parameter guidance in the body.

#	Economic risk	Severity	Primary driver
E-1	**Structural negative carry.** Borrow rate (exact) exceeds realistic external yield at all utilisations; leveraged strategies lose equity over time. Adoption + bad-debt risk.	Critical	Rate curve vs external yield
E-2	**Liquidation economics — now a hard failure.** Thin 9.09% buffer, no bonus/penalty/close-factor; *and* the opt-5 unwind path deterministically reverts for Blend/LP accounts (REG-01), plus ≥ 5 other code paths that block liquidation.	Critical	Buffer, incentive + REG-01 revert

#	Economic risk	Severity	Primary driver
E-3	**Collateral double-count — repeatable/unbounded.** Deploying to Blend never decrements the underlying balance, and a swap-laundering path (NEW-11) re-injects it per cycle, so realised leverage is bounded only by pool liquidity, not $5\times$ or $11\times$.	Critical	WAD desync (NEWSOL-C1 + NEW-11)
E-4	**Governance parameter blast-radius.** `set_coefficients` is unbounded and un-timelocked; one admin write can make the next accrual insolvent for every borrower. Single-step admin.	High	Unbounded coefficients (NEW-05)
E-5	**Risk-blind collateral.** Collateral counts at full oracle value, no per-asset factor. (Debt side correctly accrued — exposure is collateral-side only.)	High	No collateral factors
E-6	**No reserve / LP first-loss.** 100% of interest accrues to LP value; LPs are the uninsured first-loss tranche.	High	No reserve factor
E-7	**Liquidity / bank-run.** On-demand LP claims vs perpetual non-callable loans, no utilisation cap; *plus* a hard 50%-redeem-cap panic and a redeem front-run path (NEW-12).	Med-High	Maturity mismatch; no caps; 50%-cap
E-8	**External / Blend dependency & TrackToken integrity.** External yield / exchange-rate risk feeds directly into collateral value and recovery; *plus* the TrackToken is initialised off-chain and front-runnable at deploy (NEW-03).	Med-High	Three venues; off-chain init
E-9	**Oracle — lag + panic-DoS.** HF pricing uses a ≈25-min TWAP (good) but the fallback lacks staleness validation, and several valuation paths `panic!` on bad inputs, converting oracle degradation into a hard freeze of borrow/withdraw/liquidate.	Medium	TWAP lag; fallback staleness; panic-DoS
E-10	**Maintenance leverage & docs.** Opening capped $\approx 5\times$ with a runway; drift ceiling $11\times$; docs cite $8\times$ (matches neither). Ceiling high given the thin buffer.	Medium	1.10 maintenance / drift
E-11	**Recursive concentration.** Re-pledging concentrates the levered stack on one venue/asset — a contagion multiplier — amplified by E-3.	Medium	Recursion; concentration

Table 1.1: Principal economic risks (opt-5-reconciled), ranked by composite severity. Severity reflects potential economic loss and likelihood, not exploit difficulty. Four earlier claims were resolved / corrected in code (Section 1.5).

1.4 Headline economic metrics

Metric	Value	Class
Borrow APR at 50% / 90% / 99% / 100% utilisation	17.5% / 32.8% / 115.2% / 175.0%	exact
Opening leverage cap (HF = 1.25, 80% LTV)	≈**5.0**×	exact
Maintenance / drift ceiling (HF = 1.1, 90.9% LTV)	**11.0**×‡	exact
Recursive-leverage limit (opening-gated, clean code)	**5.0**× equity‡	exact
Net carry, leveraged external supply @50% util (3% external)	−14.5%**/yr**	exact†
Time to economic insolvency, 5× (openable max) negative-carry position	≈**540 days**	illustrative
Collateral buffer at the 1.1 liquidation trigger	**9.09%** of collateral	exact
Liquidation slippage above which any liquidation loses	≈**9.1%**	exact
Stress $CVaR_{99}$ bad debt, no working liquidation, HF 1.5→1.1	**0.620** → **0.721** / unit debt§	illustrative
Same tail with a functioning liquidation	≈**0.0002** / unit debt§	illustrative
Origination fee attribution	**borrower-borne** (added to debt)	code-conf.

Table 1.2: Headline metrics (opt-5-verified). †Exact given the assumed 3% external yield; the decay is exact for any yield. ‡The confirmed double-count + swap-laundering (NEW-11) *removes* this ceiling — realised max leverage is bounded only by pool liquidity (Chapter 6). §For Blend/LP accounts the no-liquidation row is the *operating point* (REG-01), and the functioning row is currently unreachable until REG-01 is fixed (Chapter 7).

1.5 Claims verified and resolved against the code

Reconciling the report's structural claims against the contracts confirmed the core theses (negative carry; liquidation economics) but *refuted* four secondary concerns, *corrected* one it had backwards, and *escalated* two, relative to an earlier draft. They are recorded here for transparency.

Earlier claim	opt-5 verdict
HF omits accrued interest; "the signal moves the wrong way"	**Refuted.** Every debt read uses the interest-accrued borrow balance, so reported HF = true HF on the debt side. (But see the collateral-side double-count, Chapter 4, which re-introduces an *upward* reported-HF bias for self-deployed positions.)
No gap between opening and liquidation (IMR = MMR)	**Refuted.** Opening 1.25 (≈5×, 80% LTV) vs liquidation 1.10 (11×, 90.9% LTV); a runway exists and opening cannot be at the line. (The 11× maintenance ceiling remains real.)
vToken NAV struck against the liquid balance ⇒ dilution	**Refuted.** Priced against total assets (liquidity + borrows + offset); claims track interest; no timing arbitrage.
Spot pricing manipulable at the instant of large actions	**Largely refuted.** ≈25-min TWAP gated on spot-freshness; residuals are the opposite failure mode (lag; fallback staleness; panic-DoS) — retained as E-9.
Origination fee drawn from pool liquidity (LP drain)	**Corrected — borrower-borne.** `lend_to` records gross debt; the fee is repaid by the borrower and does not drain the pool (NEW-M3).
Liquidation *may* be unable to seize the external claim	**Escalated — it deterministically reverts** for Blend/LP accounts (REG-01), plus ≥ 5 sibling paths (Chapter 7).
Collateral double-count breaches the 5× cap (bounded)	**Escalated — repeatable / unbounded** via swap-laundering (NEW-11); realised leverage is bounded only by pool liquidity.

Table 1.3: Concerns resolved, corrected, or escalated by the opt-5 audit. The net effect narrows the report to the genuine, code-confirmed risks and sharpens two of them from "possible" to "certain."

1.6 How to read this report

Chapter 2 sets out the protocol's economics and architecture. Chapter 3 states methodology and the data required to calibrate the illustrative results. Chapter 4 derives the mathematical model. Chapters 5–8 are the core economic analyses (carry, leverage, liquidation, liquidity). Chapter 9 specifies the simulation framework; Chapter 10 reports stress results. Chapters 11 and 12 treat systemic and game-theoretic risk (including the phantom-collateral exploit); Chapter 13 covers oracle lag and panic-DoS, and Chapter 14 governance / parameter risk. Chapter 15 gives parameter optimisation and recommendations, and Chapter 16 concludes with a capacity statement. Appendix F records the deep numerical verification of every exact figure.

CHAPTER 2

Protocol Economics & Architecture

2.1 A two-sided leveraged-carry system

Vanna is best understood not as a money market and not as a synthetic dollar, but as a *leveraged-carry intermediary*. It takes passive LP capital, lends it to margin users, and those users lever it into an external yield venue. The protocol's revenue is the spread between what borrowers pay and what LPs receive; its risk is the credit exposure that arises when a leveraged position's collateral value falls below its debt faster than the protocol can close it.

Three economic properties follow immediately and shape everything downstream. First, because borrowed capital *leaves the Vanna perimeter* into an external protocol, the collateral backing a loan is not cash in a vault but a *marked claim* (the TrackToken) on a third-party position; its convertibility back into liquidity at the marked value is an assumption, not a guarantee. Second, because the same pool serves both on-demand LP redemptions and perpetual, non-callable margin loans, the system runs a structural *maturity mismatch*. Third, because leverage is manufactured against an external yield, the sign and size of the *carry* is the first-order determinant of whether the system creates or destroys value for its participants.

2.2 System architecture

Figure 2.1 shows the economic architecture: the flow of capital, the pricing inputs, and the points at which risk enters the health-factor computation.

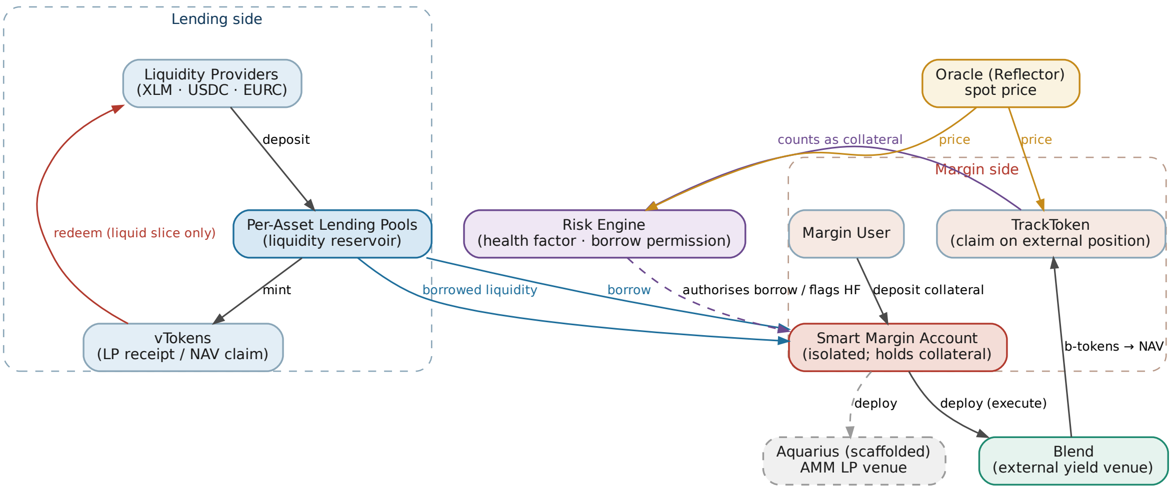

Figure 2.1: Protocol architecture and value flow. LP capital enters a single per-asset reservoir that simultaneously backs `vToken` redemptions and funds margin borrowing. Borrowed liquidity is deployed externally and returns to the risk computation only as a marked TrackToken claim, priced by the external exchange rate and a spot oracle.

2.3 Capital and value flow

The protocol has one liquidity reservoir (the per-asset pool) feeding two claimants: LPs (senior in intent, but unprotected by any reserve) and margin borrowers (who remove liquidity and redeploy it externally). The critical economic feature, visible in Figure 2.2, is that the path from an external position *back* to pool liquidity — the unwind required to recover lender funds in a liquidation — is the fragile link: recovery presupposes that the external claim can be converted to underlying at or near its marked value, on demand, net of slippage. **In opt-5 that path is worse than fragile: for Blend/LP accounts it deterministically reverts (REG-01, Chapter 7). Correction vs the prior edition:** the origination fee is *added to the borrower's debt* (`lend_to` records gross; NEW-M3), so it is repaid by the borrower and does *not* drain pool liquidity — the treasury's fee income is borrower-funded, not LP-funded.

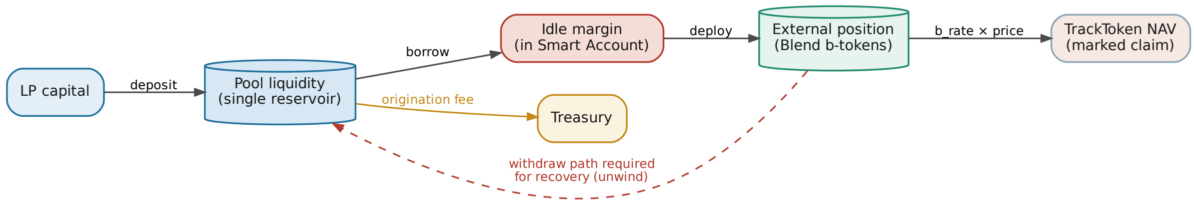

Figure 2.2: Capital and value flow. Solid arrows are the funding path; the dashed arrow is the recovery path (external position → pool) on which lender protection depends — and which reverts for Blend/LP accounts in opt-5 (REG-01). The origination fee is borrower-borne (added to debt), not a pool drain (NEW-M3).

2.4 Economic lifecycles

Capital (LP) lifecycle. Deposit → mint `vToken` at the prevailing claim price → accrue as borrowers pay interest into the pool → redeem. The economically important subtlety is that both the claim price and redemption availability depend on utilisation (Chapters 8 and 4).

Margin lifecycle. Create account → post collateral → borrow (opening health check at 1.25) → deploy externally → hold, during which carry accrues and the health factor drifts → repay/close, or breach → liquidation. The economic health of the system depends on the breach→liquidation transition being fast, incentivised and well-buffered.

TrackToken lifecycle. Minted on external deposit, valued live as external-exchange-rate × oracle price, burned on external withdraw. Its value inherits the integrity of both the external venue's accounting and the oracle.

Liquidation lifecycle. Figure 2.3 shows the economic state machine and the two exits that matter: a clean recovery when slippage is inside the buffer and the keeper is paid, versus bad debt when the buffer is exceeded, the incentive is absent, or the trigger is late.

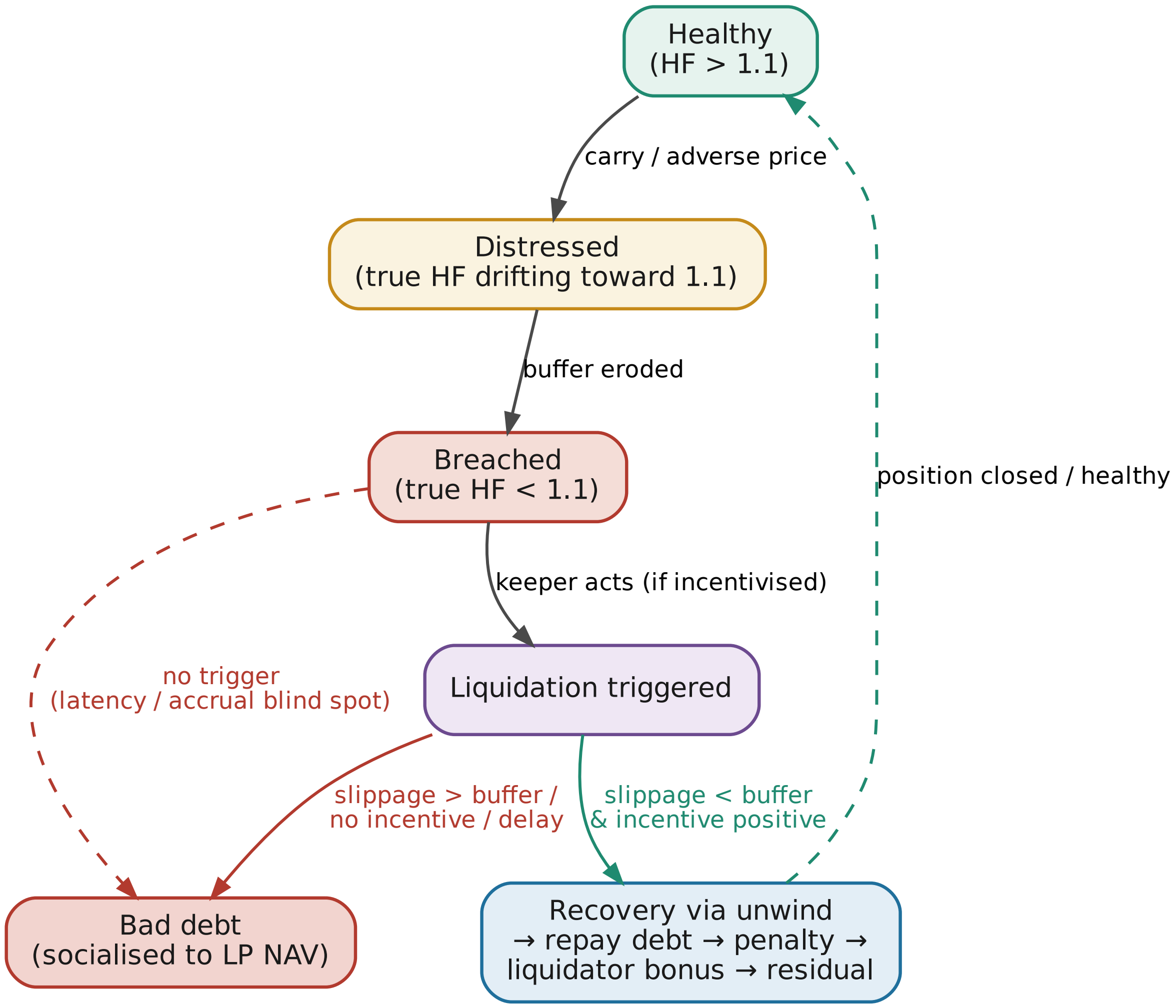

Figure 2.3: Liquidation economic lifecycle. The green exit (recovery) requires slippage below the collateral buffer *and* a positive keeper incentive; the red exits (bad debt) dominate when any of buffer, incentive, or timeliness fails. Bad debt is socialised to LP claim value.

2.5 Where value and risk accrue

Value accrues to LPs as interest (no reserve is retained) and to the protocol treasury as origination fees (borrower-borne, NEW-M3); risk accrues to LPs as the uninsured first-loss holders of a leveraged book. Margin users hold a levered exposure to the external yield — across *three* live venues (Blend, Aquarius, Soroswap), a wider dependency surface than a single external protocol — net of borrow cost. The asymmetry worth flagging up front: LPs bear tail credit risk but, absent a reserve, receive no first-loss protection. Aligning this — via a reserve financed from interest and from liquidation penalties — is a recurring recommendation.

CHAPTER 3

Methodology

3.1 Provenance and adaptation

The methodology follows, in reasoning rather than in headings, the practice of institutional economic-risk research (Chaos Labs, Gauntlet, Chainrisk, LlamaRisk). Two archetypes are relevant. The first is *simulation-driven parameter optimisation*: correlated stochastic price paths drive agent responses (borrowers, liquidators), producing distributions of bad debt and insolvency, from which risk parameters are chosen to bound a tail metric (VaR / CVaR / Liquidations-at-Risk) at a stated confidence subject to a capital-efficiency objective. The second is *empirical market / liquidity risk*: depth, slippage and redemption capacity assessed against market structure, yielding capacity statements.

Vanna requires a hybrid. Its collateral is a claim on an external yield venue, and its risk is dominated less by ordinary price volatility than by carry, by the mechanical economics of liquidation, and by external dependency. We therefore retain the simulation / VaR machinery but *foreground* mechanism-integrity economics and carry dynamics, which is where Vanna's risk actually concentrates. Where a reference method does not fit (e.g. Vanna has no funding-rate market and no insurance fund today), we derive the appropriate treatment rather than forcing the protocol into a template.

3.2 Exact versus illustrative, restated

As set out in the reading note, structural results (rate curve, leverage, recursion, carry, health-factor algebra) are exact functions of the protocol's parameters and are stated as fact. Stochastic results (VaR, CVaR, bad debt, liquidation frequency) are illustrative under the stated price and liquidity assumptions and are labelled accordingly. The purpose of the illustrative layer is to reveal *how the protocol's own mechanics transform market stress into loss*, and to rank parameters by their effect on the tail — not to publish a calibrated number.

3.3 Risk metrics

We use the standard institutional set, all expressed per unit of debt so they are scale-free:

- VaR_q — the q-quantile of the loss distribution over the horizon.
- CVaR_q (expected shortfall) — the mean loss conditional on exceeding VaR_q; the primary tail measure, because it is sensitive to the shape of the tail, not just its threshold.
- **Liquidations-at-Risk (LaR)** — here reported as the probability that a position of a given health factor is driven into the liquidation zone over the horizon.
- **Expected and tail bad debt** — loss that exceeds recoverable collateral, i.e. socialised to LPs.
- **Probability of insolvency** — probability that a position's debt exceeds its recoverable collateral value at the horizon.

3.4 Data required to calibrate

The simulation magnitudes are illustrative because the following inputs were not available in this engagement and **must** be supplied before the numbers drive governance:

1. **Price history** for XLM (and USDC / EURC) at minute resolution, at least twelve months including a stress episode, to fit the volatility surface and the jump process rather than assuming them.
2. **On-chain depth and realised slippage** for the XLM, USDC and EURC pairs across Stellar DEX / Soroswap / Aquarius, replacing the assumed slippage with an empirical price-impact function $s(\text{notional})$.
3. **External exchange-rate history** (Blend reserve rate) per market — realised supply APR, utilisation, and any historical socialised loss — to model TrackToken value drift and external-yield collapse empirically.
4. **Oracle telemetry** — update cadence, historical staleness or outage episodes, and deviation versus DEX mid — to parameterise oracle-latency and manipulation scenarios.
5. **Anticipated borrower / LP population** — leverage histogram, account counts, concentration (top-k share), and expected utilisation — to seed the agent-based layer with realistic populations rather than a swept grid.

Until these are provided, treat the stress magnitudes as *shape*, not *level*.

CHAPTER 4

Mathematical Model of Protocol Economics

Notation. p_i is the oracle price of asset i; bal_i a collateral balance; D the debt in asset units; U pool utilisation; L_{liq} the liquid pool balance; B outstanding pool borrows; Σv the `vToken` supply. Rates are annualised.

4.1 Health factor and the liquidation condition

A position is healthy iff

$$\frac{\sum_i \text{bal}_i \, p_i}{\sum_j D_j \, p_j} > 1.1 \quad \Longleftrightarrow \quad \text{HF} > 1.1 \quad (\text{liquidation}), \tag{4.1}$$

with a separate, stricter borrow (opening) threshold at 1.25 (verified in code, Section 4.3). There is no collateral factor $c_i < 1$ and no liability factor; every asset enters at full oracle value. For a single-asset position (collateral, external position and debt all in one asset) the price cancels:

$$\text{HF}_{\text{single}} = \frac{m + D_{\text{ext}} \, r_{\text{ext}}}{D}, \tag{4.2}$$

where m is idle margin (asset units), D_{ext} the deployed principal, $r_{\text{ext}} \geq 1$ the external growth factor, and D the debt. This is *price-independent* — a genuinely useful property, as single-asset positions carry no directional price risk — but it means single-asset risk is *entirely* carry- and mechanism-driven. Cross-asset positions (e.g. stablecoin collateral, XLM debt) retain full price exposure and are the ones exercised in the Monte-Carlo layer.

4.2 Interest accrual in the signal (verified)

The economic quality of HF as a solvency signal depends on debt and collateral being accrued consistently. **Code reconciliation confirms that the debt term is interest-accrued on every health path**: each debt read resolves to the interest-accrued borrow balance (principal plus accrued interest), not to raw borrow shares. Consequently the reported and true health factors coincide,

$$\text{HF}_{\text{reported}} = \frac{\text{collateral}}{D_{\text{accrued}}} = \text{HF}_{\text{true}}, \tag{4.3}$$

and the signal is *unbiased* with respect to accrued interest. An earlier draft hypothesised an upward bias *on the debt side* — debt measured without accrual — producing a health factor that “rises even as equity falls.” **That debt-side divergence cannot occur in the current code.** But the reported health factor is nonetheless biased upward for the protocol’s intended activity, through a different, confirmed channel on the *collateral* side.

The collateral-side double-count reintroduces an upward bias

A prior edition dismissed the collateral side as merely “stale, hence conservative.” That is true only for a position whose deposited collateral sits *idle*. For the protocol’s intended activity — depositing collateral and *deploying it* to Blend or an LP venue — the confirmed double-count

(NEWSOL-C1) makes reported collateral *double*: when the asset is supplied externally, its underlying balance ledger is never decremented (it stays, stale) *and* the tracking token is additionally counted. For one unit of deposited collateral deployed to Blend, reported collateral $= 1$ (stale underlying) $+1$ (tracking) $= 2$. The reported and true health factors then diverge by the deployed-collateral-to-debt ratio, which for a self-deployed position near the opening cap is ≈ 1 in HF units:

$$\mathrm{HF}_{\text{reported}} \;\approx\; \mathrm{HF}_{\text{true}} + \frac{\text{deployed collateral}}{D} \;\approx\; \mathrm{HF}_{\text{true}} + 1. \tag{4.4}$$

So a self-deployed position truly at $\mathrm{HF}_{\text{true}} = 1.25$ reports $\mathrm{HF}_{\text{reported}} \approx 2.25$. Because the liquidation trigger fires on the *reported* value, it does not fire until reported HF reaches 1.10, i.e. until *true* HF has fallen to ≈ 0.10 — collateral worth a tenth of the debt, catastrophically insolvent (Figure 4.1). **This is the surviving "wrong-way" distortion — collateral-side, via the double-count, not accrual-side — and it defeats the timely-trigger conclusion of Chapter 5 for exactly the strategy the protocol is built for.** The clean single-HF trajectory of Figure 5.4 holds only for the idle-margin configuration; it does *not* hold for self-deployed collateral. Fixing the WAD double-count is therefore not only a leverage-ceiling fix (Section 4.3) but a prerequisite for the solvency signal to be truthful, and it is P0.

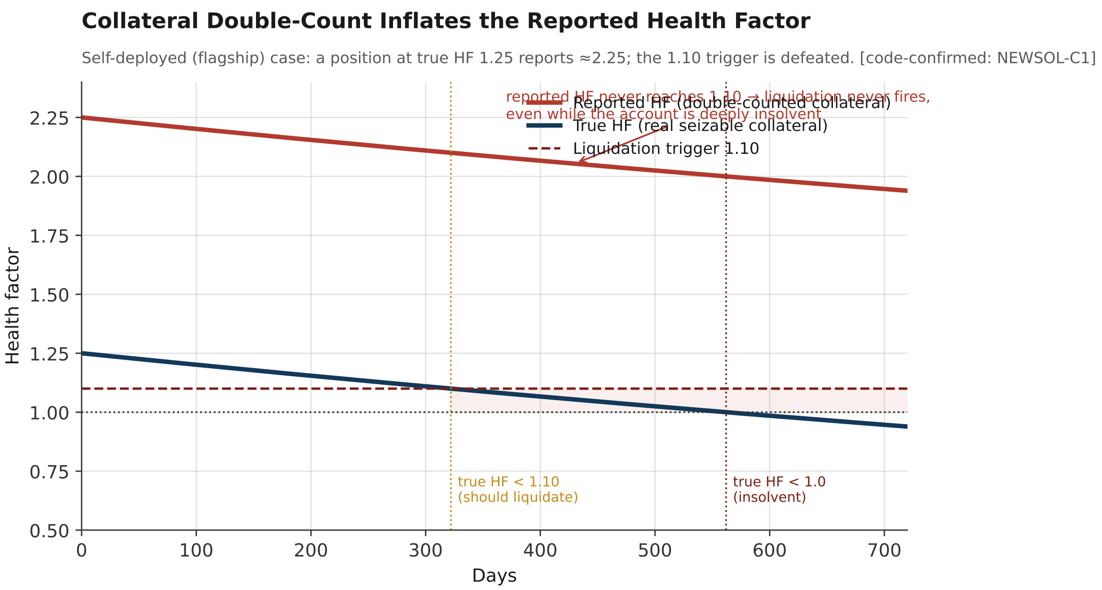

Figure 4.1: The collateral double-count inflates the reported health factor for a self-deployed (flagship) position: at true HF 1.25 it reports ≈ 2.25, so the 1.10 trigger is reached only when true HF is ≈ 0.10. The shaded band is "invisible insolvency" — true HF below 1.10 while reported HF is comfortably above it. [code-confirmed] (NEWSOL-C1); [illustrative] for the decay path.

4.3 Leverage bounds

Code reconciliation confirms *two* distinct thresholds. Borrowing is permitted only while HF $>$ 1.25 (the opening constraint); liquidation occurs at HF $<$ 1.10 (the maintenance constraint). The corresponding equity multiples $\lambda = \mathrm{HF}/(\mathrm{HF}-1)$ are

$$\lambda_{\text{open}} = \frac{1.25}{0.25} = \mathbf{5}\times\ (80\%\ \text{LTV}), \qquad \lambda_{\text{maint}} = \frac{1.10}{0.10} = \mathbf{11}\times\ (90.9\%\ \text{LTV}). \tag{4.5}$$

A new position is therefore capped near $5\times$ and *cannot* be opened at the liquidation line: the interval $1.25 \to 1.10$ is a genuine initial-vs-maintenance margin gap and safety runway. An earlier draft, which assumed a single 1.1 threshold governed both, is corrected here. What remains true is that an *open* position can drift up to the 11× maintenance ceiling before liquidation triggers, and that the documented "8×" matches neither bound — so the 11× drift ceiling, not an opening-time excess, is the leverage concern (Chapter 6). **Crucially, opt-5 escalates this: the double-count is repeatable (swap-laundering, NEW-11), so neither the 5× opening cap nor the 11× maintenance ceiling actually binds — realised per-account leverage is bounded only by available pool liquidity** (Section 4.3 and Chapter 6).

4.4 Recursive leverage

Depositing collateral, borrowing, deploying, and re-pledging the resulting TrackToken as collateral permits a loop. Because each borrow is gated by the *opening* threshold, the implied per-loop loan-to-value is $\ell = 1/1.25 = 0.8$, and cumulative exposure after n loops is the geometric series

$$E_n = C \cdot \frac{1-\ell^{n+1}}{1-\ell} \xrightarrow{n\to\infty} \frac{C}{1-\ell} = 5C. \tag{4.6}$$

So the *intended* recursive ceiling is 5×, not 11× (Figure 6.2 in Chapter 6), consistent with the opening cap. Two economic hazards remain. First, recursion concentrates the entire levered stack on one external venue and (typically) one asset — a contagion multiplier if the external exchange rate or the asset feed misbehaves. Second, and more concretely, the security audit confirms a *collateral double-count*: when the asset is deployed to the external venue, the original collateral ledger entry is not decremented while the new external claim is also counted, so collateral is overstated and borrowing can exceed the 5× opening cap. The double-count, not the clean geometric series, is the real recursion risk.

4.5 Interest-rate model

With coefficients $c_1 = 0.1,\ c_2 = 0.3,\ c_3 = 3.5$, the annualised borrow factor is

$$R_b(U) = c_3\big(c_1 U + c_1 U^{32} + c_2 U^{64}\big) = 3.5\big(0.1\,U + 0.1\,U^{32} + 0.3\,U^{64}\big). \tag{4.7}$$

The high powers produce a smooth, kink-free curve with a very steep tail. With no reserve factor, the LP supply APR is simply $R_s(U) = R_b(U)\cdot U$. The exact schedule appears in Table 5.1 (Chapter 5). The economically decisive feature is not the steep tail — which usefully discourages running at 100% utilisation — but the *floor*: 17.5% at half utilisation already exceeds plausible external supply yields.

4.6 Carry and equity dynamics

Consider a single-asset position with initial margin 1 held idle in the account and borrowed principal $D = \lambda - 1$ deployed externally at yield y. In discrete daily compounding,

$$\text{collateral}(t) = 1 + D\Big(1 + \tfrac{y}{365}\Big)^t, \qquad \text{debt}(t) = D\Big(1 + \tfrac{R_b(U)}{365}\Big)^t, \tag{4.8}$$

$$\text{HF}_{\text{true}}(t) = \frac{\text{collateral}(t)}{\text{debt}(t)}, \qquad \text{equity}(t) = \text{collateral}(t) - \text{debt}(t). \tag{4.9}$$

The net carry is $y - R_b(U)$. When it is negative, equity declines monotonically and the position is a slowly-detonating short-carry trade; the time to any health threshold is obtained by solving $\text{HF}_{\text{true}}(t) = \text{threshold}$. Chapter 5 tabulates these times; at the openable maximum of 5× (opening HF = 1.25) with a 3% external yield the position reaches the 1.10 liquidation line in about 310 days and economic insolvency in about 540 days, and higher effective leverage (reachable only by drift or the collateral double-count) fails proportionally faster.

4.7 LP claim (vToken) economics (verified)

Code reconciliation confirms that the `vToken` claim price is struck against *total assets*, not the liquid balance:

$$\mathrm{mint}(a) = a \cdot \frac{\Sigma v}{L_{\mathrm{liq}} + B + v_0}, \qquad \mathrm{redeem}(v) = v \cdot \frac{L_{\mathrm{liq}} + B + v_0}{\Sigma v}, \tag{4.10}$$

where B is outstanding borrows (including accrued interest) and v_0 a small virtual offset. Because outstanding principal *is* included, the claim price grows as interest accrues and LP returns track interest earned, as intended. An earlier draft posited the denominator as the liquid balance L_{liq} alone, which would transfer value between LP cohorts by utilisation timing; **that mechanism does not exist in the current code** (an in-code comment records it as a fixed prior bug), and the utilisation-timing dilution and the associated deposit/withdraw game have been struck from this report. The one genuine LP exposure that remains is structural, not a pricing artifact: LPs are the uninsured first-loss holders of a leveraged book, because no reserve fraction is retained (Section 8).

4.8 Solvency and bad-debt conditions

At the portfolio level, solvency requires recoverable collateral to cover debt:

$$\sum \text{recoverable collateral} \ \geq \ \sum \text{debt}. \tag{4.11}$$

For a single liquidated position with liquidation threshold θ, execution overshoot $o \geq 0$ (adverse move between breach and execution) and slippage s, the recovered value per unit debt is

$$\rho = \theta\, e^{-o}\,(1 - s), \tag{4.12}$$

and bad debt is $\max(0, 1 - \rho)$. Setting $o = 0$ gives the break-even slippage $s^\star = 1 - 1/\theta$; at $\theta = 1.1$, $s^\star = 9.09\%$. This single expression links three levers — the liquidation threshold (buffer), execution latency (overshoot), and market depth (slippage) — and is the analytical backbone of Chapter 7.

CHAPTER 5

Interest-Rate & Carry Risk

Carry is the first-order economic risk in Vanna, so we treat it first. The question is simple: does a leveraged position earn more from the external strategy than it pays to borrow? For the flagship strategy the answer, at realistic external yields, is no.

5.1 The interest-rate curve

The exact borrow and supply schedule implied by the rate model (Section 4.5) is given in Table 5.1 and drawn in Figure 5.1. Two features matter economically. The curve is convex and kink-free, with a very steep tail (115% at 99% utilisation) that strongly discourages running the pool at extreme utilisation — a useful liquidity-retention property. But the *floor* is high: at half utilisation the borrow rate is already 17.5%, and because there is no reserve factor the full economic rent passes to LPs as supply yield ($R_s = R_b \cdot U$) with nothing retained to build a buffer.

Utilisation	Borrow APR	Supply APR
0%	0.00%	0.00%
50%	17.50%	8.75%
80%	28.03%	22.42%
90%	32.83%	29.54%
95%	43.97%	41.77%
97.5%	70.46%	68.70%
99%	115.21%	114.06%
100%	175.00%	175.00%

Table 5.1: Exact interest-rate schedule (no reserve cut). [exact]

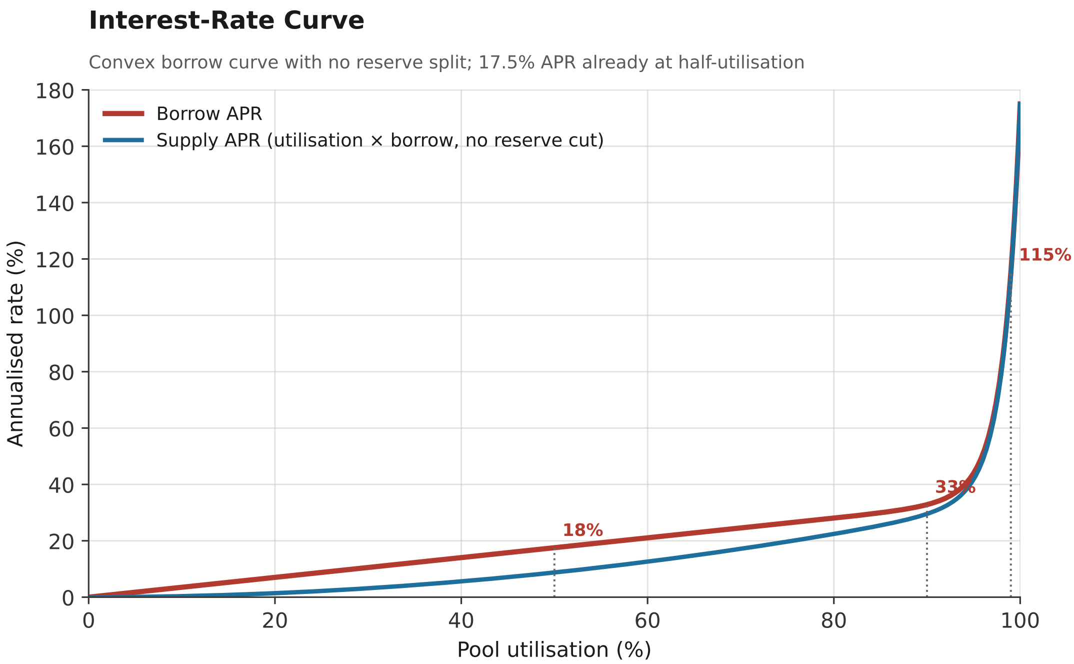

Figure 5.1: The interest-rate curve. The borrow floor of 17.5% at half utilisation is the root of the carry problem; the steep tail protects liquidity but the absence of a reserve split leaves no first-loss buffer. [exact]

5.2 The net-carry surface

Whether the flagship strategy is viable depends on the borrow rate relative to the external yield. Figure 5.2 maps net carry $(y - R_b(U))$ across the plausible region of external yields and utilisations. The economically viable (green) region requires external yields well above 15–20%; realistic external supply yields of 2–6% sit deep in the red across the entire utilisation range. In other words, under normal conditions the protocol's headline use case destroys borrower equity by construction.

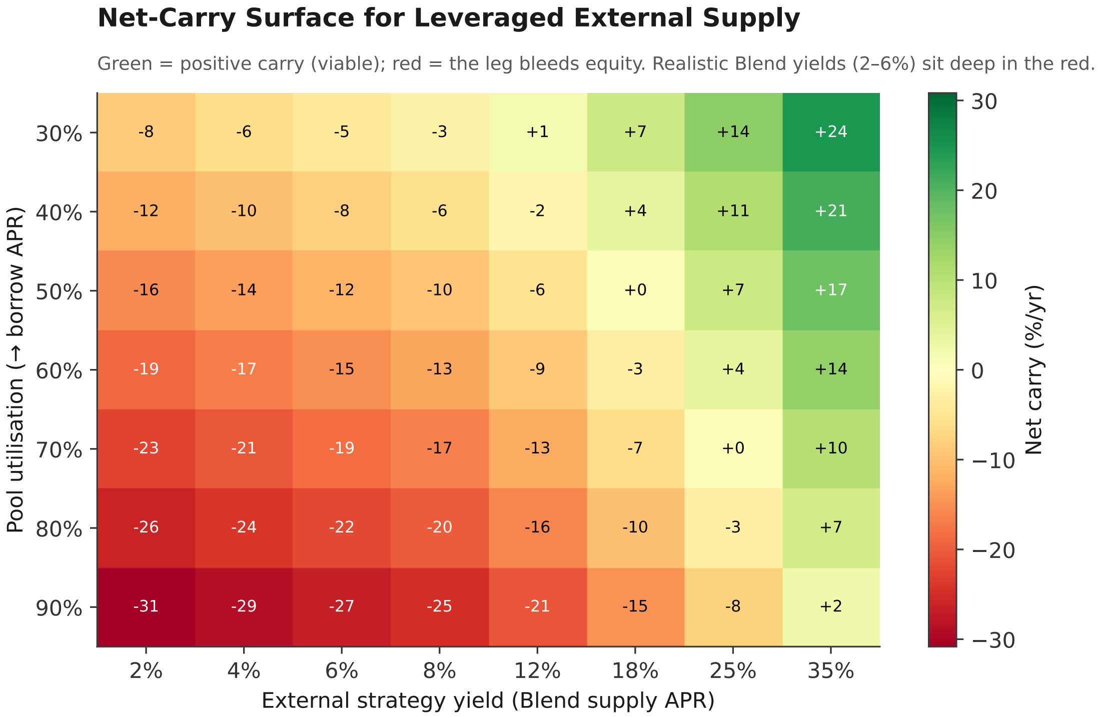

Figure 5.2: Net-carry surface for leveraged external supply. Positive carry (green) requires external yields far above what the external venue realistically pays; at 2–6% yields the strategy is negative-carry at all utilisations. [exact] given the rate curve; the yield axis is the free variable.

5.3 Equity decay and time-to-distress

Negative carry is not a static disadvantage; it compounds. Using the dynamics of Section 4.6 with $R_b(0.5) = 17.5\%$ and an assumed external yield of 3% (net carry -14.5%), Table 5.2 and Figure 5.3 show equity trajectories by leverage. At the openable maximum of 5×, the position reaches the 1.10 liquidation line in about 310 days and is economically insolvent in about 540 days on carry alone — before any adverse price move. Higher effective leverage (reachable only by drift toward the maintenance line or via the collateral double-count) fails faster: an 8× position would exhaust its equity within the year.

Leverage	Net carry	Days to HF 1.10	Days to insolvency	Equity @ 1yr
2× (open)	−14.5%	1368	1587	0.839
3× (open)	−14.5%	732	957	0.679
4× (open)	−14.5%	461	690	0.518
5× (open, max)	−14.5%	**310**	**540**	**0.357**
8× (drift only)	−14.5%	94	328	−0.125

Table 5.2: Carry-driven equity decay by leverage (assumed 3% external yield). Openable leverage is capped near 5×; the 8× row is a drift / double-count state. The decay is exact given the rate; the 3% yield is the assumption. [illustrative]

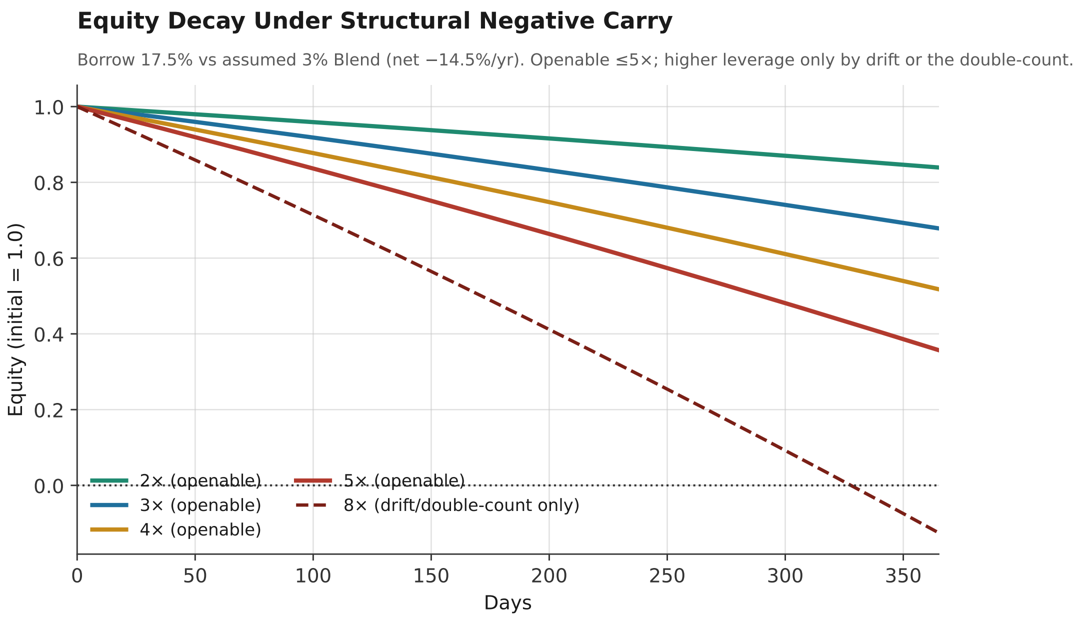

Figure 5.3: Equity trajectories under -14.5% net carry. Openable leverage ($\leq 5\times$, solid) decays steadily; the $8\times$ dashed curve is a drift / double-count state and crosses zero within a year. [illustrative]

Crucially, the solvency signal reflects this correctly *only for idle-margin collateral*. Because the debt term in the health factor is interest-accrued (Section 4.6), the reported and true health factors coincide when collateral is not deployed, and the health factor declines in step with equity, crossing the 1.10 liquidation line on time (Figure 5.4). An earlier draft hypothesised the opposite on the debt side — a reported HF drifting *up* as debt accrued — and the code refutes that. **However, for the protocol's intended self-deployed strategy, the collateral-side double-count (Section 4.7, Figure 4.1) does reintroduce an upward bias: reported HF $\approx$ true HF $+1$, so the trigger fires only near true HF ≈ 0.10.** In other words, carry produces a correctly-signalled path to distress *for idle margin*, but a *silently-signalled* one for deployed collateral until the double-count is fixed — and in both cases containment then depends on the liquidation economics of Chapter 7, which for Blend/LP accounts fail outright (REG-01).

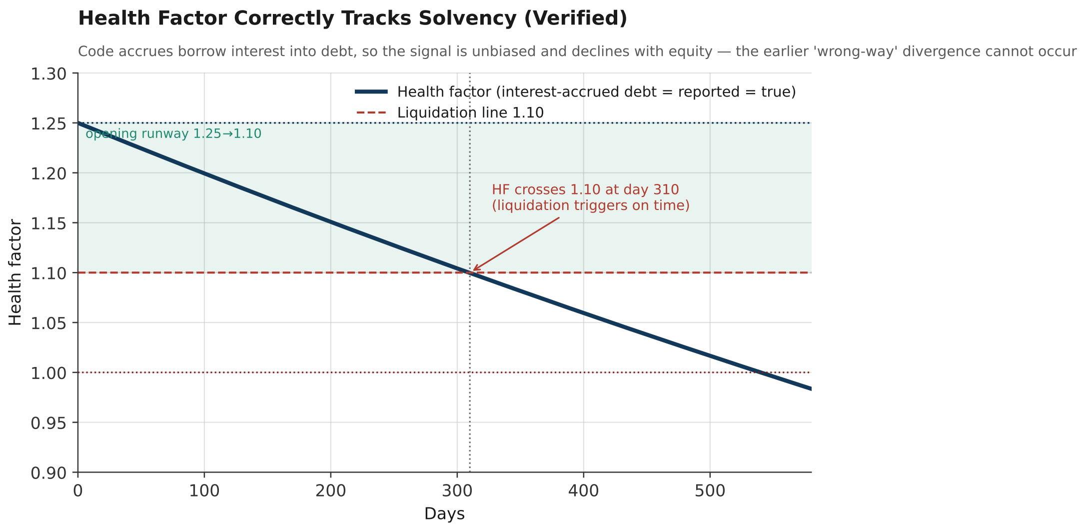

Figure 5.4: Health factor tracks solvency *for idle margin*. A 5× position opens at HF = 1.25; with interest-accrued debt the (unbiased) health factor declines through the runway and crosses the 1.10 line at day ≈310, so the trigger fires on time. **This holds only for idle-margin collateral;** for self-deployed collateral the double-count of Figure 4.1 defeats it. [illustrative]

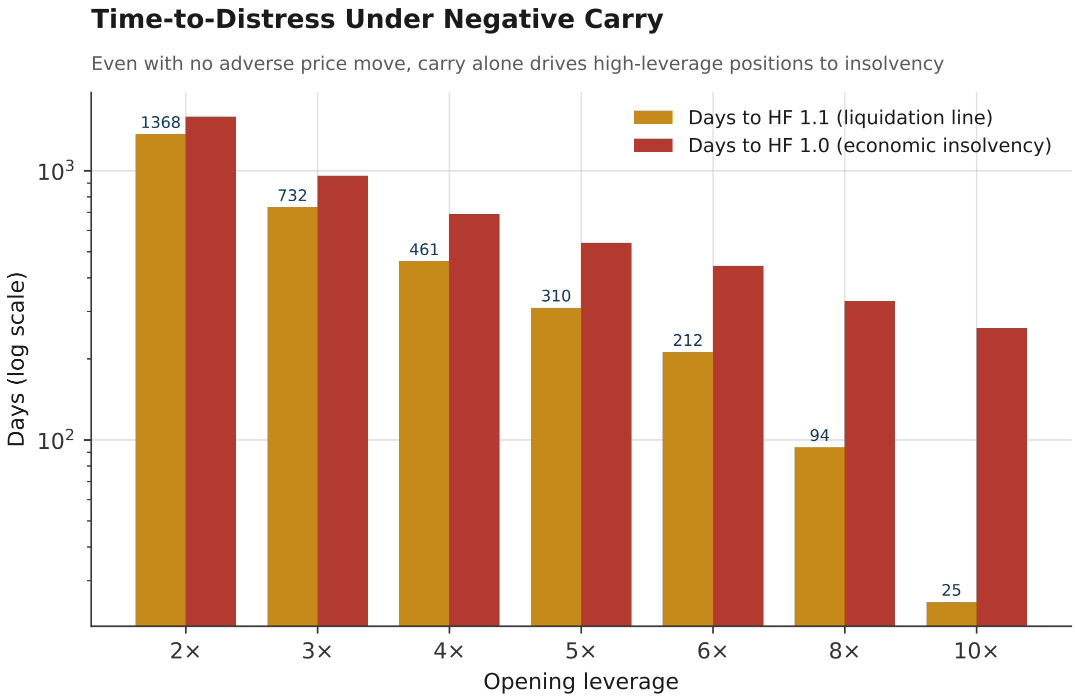

Figure 5.5: Time to the liquidation line and to insolvency, by leverage, driven by carry alone (no price move). Log scale; openable leverage is $\leq 5\times$. [illustrative]

5.4 Interest-rate risk assessment

The interest-rate configuration creates a twin economic risk. As an *adoption* risk, negative carry means rational borrowers either avoid the protocol or use it only to farm external incentives; sustained demand requires token subsidies that themselves must be funded. As a *bad-debt* risk, borrowers who do stay bleed toward insolvency; the health factor flags them correctly and on time (the signal is unbiased), so containment then rests on the liquidation economics of Chapter 7. The remedy is a recalibration of the curve to the realistic yield opportunity set — a lower base slope with a later, sharper kink — together with a reserve factor that retains part of the interest to build first-loss capital (Chapter 15). Strategies whose expected yield is below the borrow cost should be gated rather than offered.

CHAPTER 6

Leverage & Capital-Efficiency Risk

6.1 The leverage ceiling and the safety runway

Code reconciliation corrects an earlier assumption here. There are two thresholds (Section 4.3): opening at HF = 1.25 (≈ 5×, 80% LTV) and liquidation at HF = 1.10 (11×, 90.9% LTV), shown in Figure 6.1. So an initial-vs-maintenance margin gap *does* exist: a new position is capped near 5× and cannot be opened at the liquidation line, and the interval 1.25 → 1.10 is the borrower's safety runway and the protocol's reaction time. The earlier draft's "no runway, opened at the line" scenario is withdrawn.

Two leverage concerns nonetheless remain. First, an *open* position can still drift up to the 11× maintenance ceiling before liquidation triggers, which — combined with the thin buffer and hollow liquidation economics of Chapter 7 — is a high ceiling to permit. Second, the documentation cites 8×, which matches neither real bound (≈ 5× opening, 11× maintenance) and should be reconciled. The opening cap itself is well-placed and close to this report's recommendation; the action items are to lower the *maintenance* ceiling (raise the liquidation threshold) and to align the documentation.

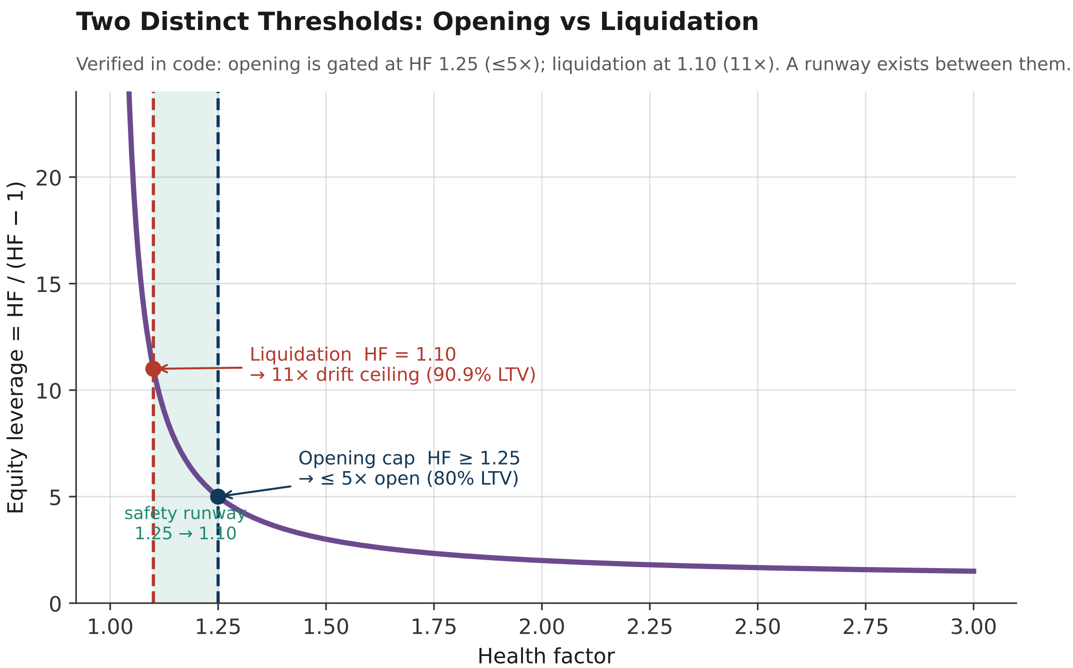

Figure 6.1: Effective leverage as a function of the health factor, with both verified thresholds. Opening is gated at HF = 1.25 ($\leq 5\times$); liquidation at HF = 1.10 (11×). The shaded band is the safety runway between them. [exact]

6.2 Recursive leverage and concentration

Because each borrow is gated by the 1.25 opening threshold, the intended recursive ceiling is 5× (per-loop LTV 1/1.25), not 11× (Figure 6.2). But opt-5 confirms this ceiling does not bind.

The economic hazard is *concentration*: the entire levered stack sits on one external venue and typically one asset, so a single external exchange-rate shock or asset-feed anomaly strikes the whole stack at once (the contagion topology of Chapter 11). The mechanical hazard, confirmed by the audit, is worse than the prior edition stated: the *collateral double-count is repeatable*. Deploying to Blend overstates collateral (Section 4.7), and a swap round-trip (NEW-11) re-credits the phantom WAD each cycle, so an attacker (or even an ordinary looping user) can inflate reported collateral without bound. Realised per-account leverage is therefore limited only by *pool free liquidity* — not 5×, not 11× (Figure 6.3). This converts the recursion analysis from a bounded geometric series into an unbounded exposure, and makes the WAD double-count fix (plus per-account exposure caps) a P0 item.

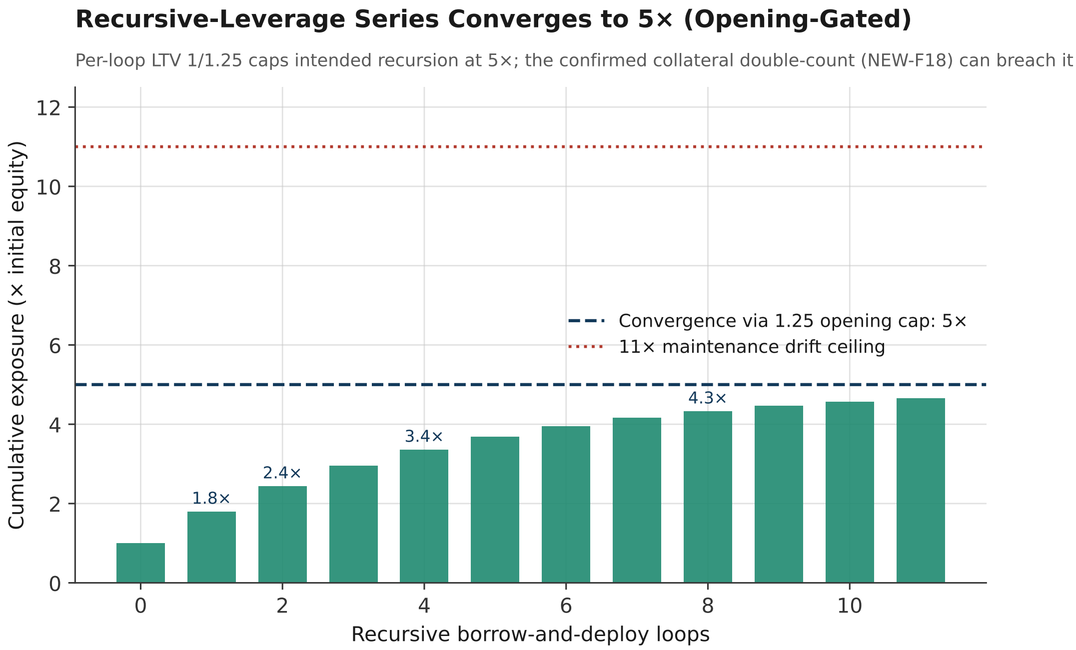

Figure 6.2: Recursive-leverage series under the 1.25 opening cap (per-loop LTV $= 1/1.25$), converging to 5× *in clean code*; the 11× maintenance ceiling is the drift limit. In opt-5 the repeatable double-count breaks both bounds (Figure 6.3). [exact] for the clean-code series.

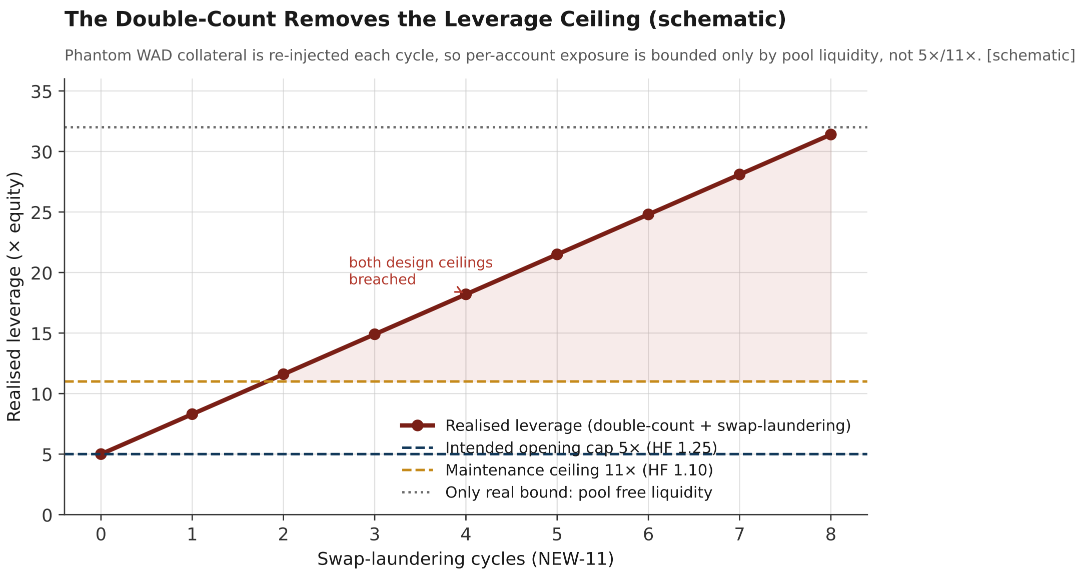

Figure 6.3: The repeatable double-count removes the leverage ceiling. Each swap-laundering cycle (NEW-11) re-injects phantom WAD collateral, so realised per-account leverage climbs past both the 5× opening cap and the 11× maintenance ceiling, bounded only by pool free liquidity. [code-confirmed] for the mechanism (NEWSOL-C1 + NEW-11); [schematic] for the per-cycle rate.

6.3 The capital-efficiency / safety frontier

Leverage is where capital efficiency and safety trade off most directly. Figure 6.4 and Table 6.1 map the stressed tail (as-implemented CVaR_{99} bad debt) against leverage. The relationship is steeply concave: most of the tail risk is already incurred by modest leverage, and the marginal efficiency gained above roughly 5× purchases disproportionate tail exposure. Encouragingly, the protocol's existing ≈ 5× opening cap (HF = 1.25) already sits at the efficient point this analysis would recommend — so the opening design is sound. The residual leverage risk is the 11× *maintenance* drift ceiling (and the double-count that can breach the opening cap), which argues for raising the liquidation threshold rather than tightening opening further, reserving any change only until the liquidation and reserve mechanisms are hardened and the results recalibrated on live data.

Max leverage	Stressed CVaR_{99} bad debt
2×	0.493
3×	0.620
4×	0.662
5×	0.683
6×	0.696
8×	0.710
10×	0.718
11×	0.721

Table 6.1: Stressed tail bad debt vs the leverage a position can reach (no working liquidation). The curve is concave: the safety cost of leverage is largely paid by $\sim 5\times$, which is where the opening cap already sits. [illustrative]

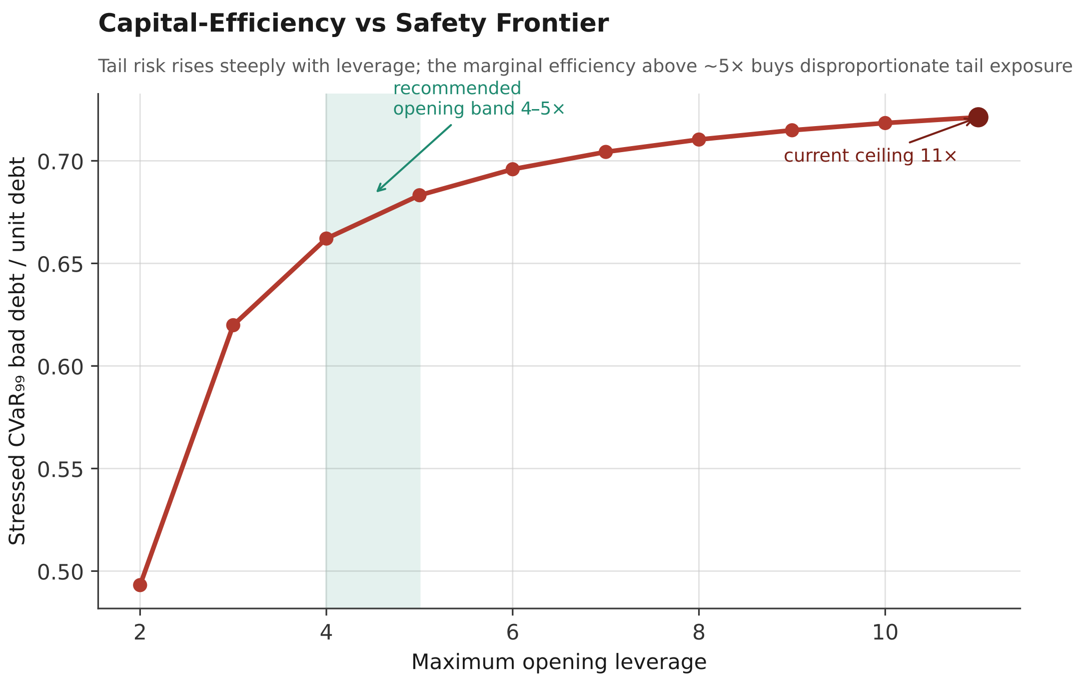

Figure 6.4: Capital-efficiency versus safety frontier. Tail risk rises steeply and concavely with leverage; the protocol's existing $\approx 5\times$ opening cap already captures most of the useful efficiency, while the $11\times$ maintenance ceiling is the outlier to address. [illustrative]

6.4 Capital-efficiency assessment

The opening design is largely sound: the $\approx 5\times$ cap and the $1.25 \to 1.10$ runway already embody the initial-vs-maintenance separation this analysis would recommend. Two frictions remain. The initial margin sits idle in the account and earns nothing, a drag that worsens the already-negative carry. And the $11\times$ *maintenance* ceiling, while nominally capital-efficient, is high given the thin buffer, the absent liquidation incentive, and the confirmed collateral double-count that can breach the opening cap. The remaining efficiency-for-safety trades are therefore targeted, not wholesale: raise the liquidation threshold (lower the maintenance ceiling), add collateral factors and exposure caps, and fix the double-count — rather than tightening the opening cap that is already well-placed.

CHAPTER 7

Liquidation Economics

Liquidation is the mechanism that converts a market move into either a contained loss or a catastrophic one. The prior edition assessed its *economics* (buffer, incentive, latency) and deferred implementation correctness to the security audit. That audit is now in, and it changes the conclusion in kind, not just degree: for the protocol's intended activity, liquidation does not merely lose money in the tail — **it cannot execute at all**. We therefore treat structural failure first, then the economics that apply once (if) it is repaired.

7.1 Structural failure: liquidation reverts for the flagship strategy

The opt-5 audit finds that `liquidate`, for any account holding a Blend or LP position, calls an auto-unwind routine that in turn invokes a *public* `execute` entry guarded by a trader-authorisation check and a reentrancy lock; invoked from inside liquidation it panics, and the entire liquidation transaction reverts (finding **REG-01**). Because the flagship strategy is precisely "deploy borrowed liquidity to Blend," this means the accounts that carry the most credit risk are the ones that *cannot be liquidated*. At least five further code paths independently block or corrupt liquidation:

- **REG-02** — an Aquarius skew assertion can abort the unwind for LP-collateral accounts.
- **NEW-08** — a Soroswap `RemoveLiquidity` mis-scaling returns the wrong underlying amount, mispaying the repayment leg.
- **V-H02 / NEW-07** — oracle valuation paths `panic!` on stale or malformed input, so oracle degradation freezes liquidation rather than degrading it gracefully.
- **NEW-C1 / NEWSOL-C1** — the collateral double-count inflates reported health, so an economically insolvent account reads as healthy and `liquidate` rejects it as not liquidatable (the same distortion as Figure 4.1).

The economic consequence is stark and simple: the "no-working-liquidation" regime of Chapter 10 is not a pessimistic bracket for these accounts — it is the *operating point*. The stressed CVaR_{99} bad debt of 0.49–0.73 per unit debt is what the protocol is exposed to today for its intended strategy, and the ≈ 0.0002 "functioning" figure is unreachable until REG-01 (and its siblings) are fixed. Repairing the liquidation path — so it can seize and unwind the external claim without reverting — is the single highest-priority action in this report, ahead even of the carry recalibration. Everything in the rest of this chapter (the economics of buffer, incentive, and latency) is necessary but *not sufficient* until this structural failure is resolved.

7.2 Buffer economics

Recall the recovery identity from Chapter 4: $\rho = \theta\, e^{-o}(1-s)$, with liquidation threshold θ, execution overshoot o, and slippage s. At $\theta = 1.1$ the collateral buffer over debt is only $1 - 1/1.1 = 9.09\%$. Setting overshoot to zero, the break-even slippage is $s^\star = 9.09\%$: below it a prompt liquidation covers the debt; above it, *every* liquidation loses money by construction, even with instant execution and even from a perfectly healthy starting point at the trigger.

Figure 7.1 shows how the tolerable slippage grows with the buffer — i.e. with a higher liquidation threshold or, equivalently, with collateral-factor haircuts. This is why buffer, penalty and collateral factors are not cosmetic: they are what buy the protocol slippage tolerance and the keeper a margin.

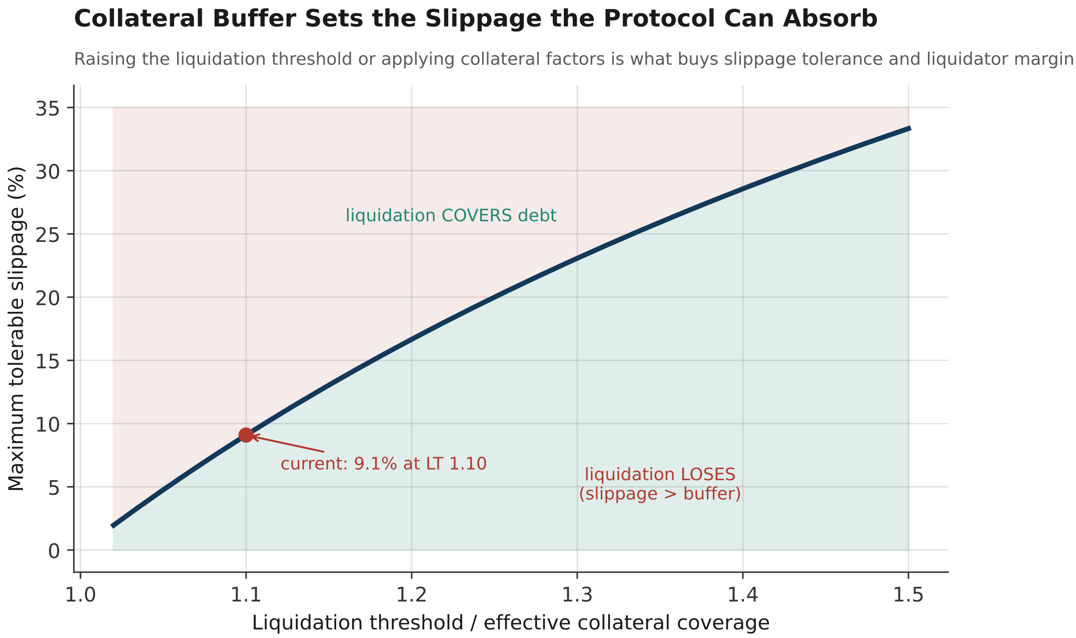

Figure 7.1: The collateral buffer sets the slippage the protocol can absorb. At the current 1.1 trigger the buffer is 9.09%; raising the effective coverage (higher threshold or collateral factors) is the lever that moves liquidation from loss-making to safe under realistic depth. [exact]

7.3 Slippage and market depth

On thin books the break-even band is reachable in stress. Figure 7.2 shows the stressed tail loss at $HF_0 = 1.15$ as a function of slippage: below $\sim 9\%$ a functioning liquidation rarely loses, but the loss climbs steeply beyond the buffer, reaching a $CVaR_{99}$ of 0.30 per unit debt at 35% slippage. Because Stellar-native depth for the relevant pairs is finite and deteriorates in stress, the protocol cannot assume the 9% buffer is comfortable; it must either widen the buffer or size a liquidation bonus above expected stressed slippage.

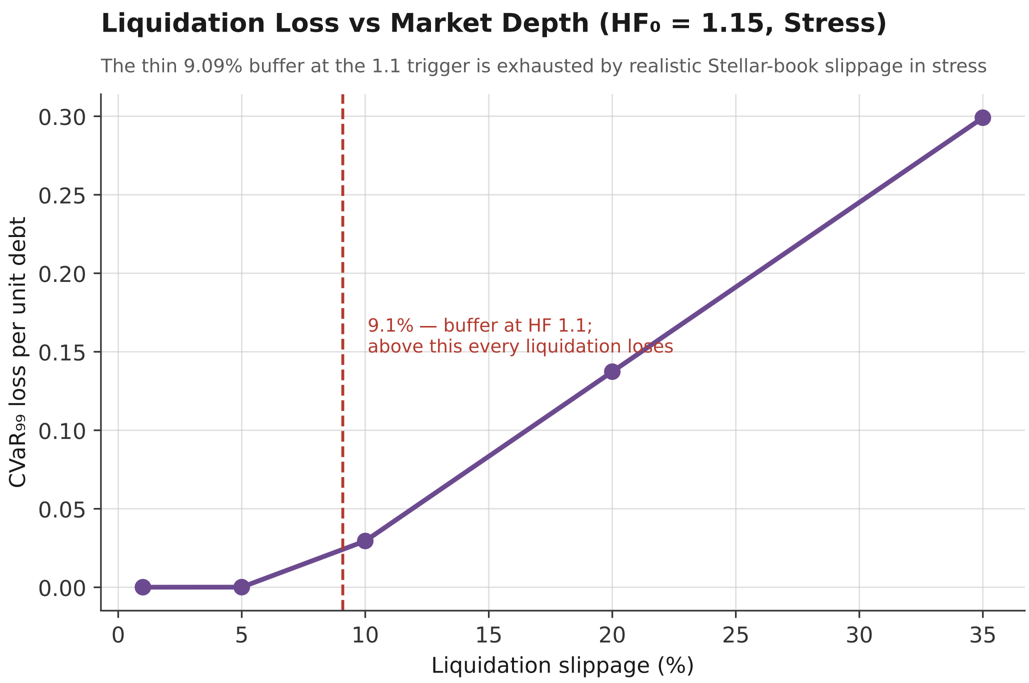

Figure 7.2: Liquidation loss versus slippage (stress, $HF_0 = 1.15$). The 9.09% buffer at the 1.1 trigger is exhausted by realistic stressed slippage, after which loss rises steeply. [illustrative]

7.4 Incentive economics

A liquidation mechanism has to pay for itself. Two prices are missing in the current design. A **liquidator bonus** (a discount on seized collateral) is what makes a keeper willing to spend gas and take inventory risk to close a position; without it, even a mechanically permissionless liquidation will not be *called* unless independently profitable, so enforcement becomes voluntary. A **borrower penalty** (retained by the protocol / reserve) is what deters maximal-leverage, strategic-default behaviour and what capitalises the first-loss buffer. With neither present, and with only a 9.09% buffer, the economics of liquidation are hollow even where it can execute: there is nothing to absorb slippage and nothing to reward the agent who must act. A **close factor** (partial liquidation) further reduces over-liquidation and its market impact. All three are standard and are recommended in Chapter 15.

7.5 Execution latency

Following the reference methodology, we treat the delay between a position becoming unhealthy and being closed as an economic parameter (the analogue of a bad-debt-recognition lag). Latency enters the recovery identity through the overshoot o: the longer the delay, the larger the adverse move that can accumulate — diffusive drift plus any jump landing inside the window — before execution, and the deeper the loss. Two things lengthen latency economically: a solvency signal biased high (so the trigger fires late or not at all, Chapter 5) and a missing keeper incentive (so no agent is watching). The liquidation lifecycle of Figure 2.3 makes the exits explicit: the benign exit requires buffer, incentive *and* timeliness simultaneously.

7.6 Liquidation risk assessment

Liquidation is the single highest-leverage place to improve Vanna's economic risk profile. The priority actions are economic, not just mechanical: introduce a liquidator bonus sized above expected stressed slippage; introduce a borrower penalty that funds a reserve; widen the effective buffer via an initial-vs-maintenance gap and collateral factors; and ensure the solvency signal accrues interest so the trigger is timely. The quantified payoff of these changes is the collapse of the stressed tail documented in Chapter 10.

CHAPTER 8

Liquidity, Redemption & Bank-Run Risk

8.1 The maturity mismatch

The pool runs a classic liquidity mismatch: LP claims are on-demand, whereas margin loans are perpetual, non-callable liens with no forced-repayment mechanism, and redemptions are served only from the liquid balance. At utilisation U, only the fraction $1 - U$ of deposits is redeemable before borrowers choose to repay (Figure 8.1). With no utilisation cap, U can reach 100% (at a punitive 175% borrow APR), at which point redemptions are fully blocked until borrowers voluntarily repay — which, under negative carry (Chapter 5) and *failing* liquidation (Chapter 7), they may never do. **opt-5 also hard-codes a single-transaction redemption cap of 50% of the liquid balance that `panic!`s rather than partially filling (NEWSOL-M / redeem-cap), so a large LP redemption reverts outright, and a front-runnable redeem path (NEW-12) lets an observer drain the liquid slice ahead of a queued redeemer** — both of which sharpen the run geometry beyond the smooth $1 - U$ curve below.

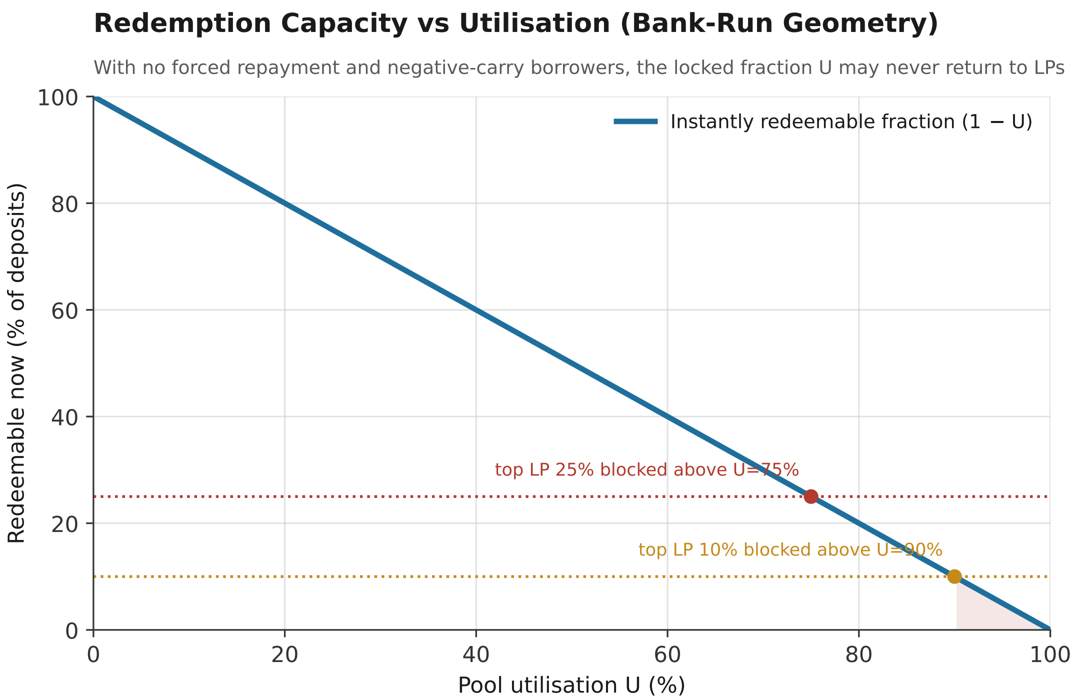

Figure 8.1: Redemption capacity versus utilisation. The instantly redeemable fraction is $1 - U$; a concentrated LP whose share exceeds $1 - U$ cannot exit at all. With no forced repayment, the locked fraction may not return. [exact] for the geometry.

8.2 Run dynamics and concentration

During a redemption rush, LPs race for the liquid slice first-come, first-served; the remainder is locked. Concentration sharpens this: a single LP holding more than $1 - U$ of the pool cannot exit even first in line, and the attempt itself pushes utilisation higher, worsening the position of everyone behind. (The claim price itself is *not* a source of distortion here: it is struck against total assets and tracks interest, Section 4.7.) The result is the canonical run-vulnerability, amplified by the absence of reserves and by the inability to liquidate for recovery.

8.3 LP first-loss and the missing reserve

LP economics deserve to be stated plainly. The LP is the uninsured first-loss holder of a leveraged book: they earn the full supply APR (because no reserve is retained) but, for the same reason, have no buffer between them and any bad debt, which reduces their claim value directly and instantly. The claim price is fair — it is struck against total assets and tracks interest earned (Section 4.7), so there is no utilisation-timing distortion — but the *first-loss* exposure is real and uninsured until a reserve exists and liquidation works. The remedies are a utilisation ceiling and/or withdrawal queue, per-asset supply and borrow caps, a reserve financed from interest and penalties, and — underpinning all of it — a functioning liquidation so that locked collateral can be converted back to liquidity for redeeming LPs.

8.4 Liquidity risk assessment

Liquidity risk is real but second-order relative to carry and liquidation, and it is addressed by the same reserve-and-caps package. The one liquidity-specific control that cannot be substituted is a utilisation ceiling (or explicit withdrawal queue): without a cap on U, the protocol has no way to guarantee any redemption capacity, and the bank-run geometry of Figure 8.1 is left entirely to borrower goodwill.

CHAPTER 9

Simulation Framework

9.1 Objectives

The simulation layer exists to answer one question: for a position (or book) of given health, leverage and collateral mix, what is the distribution of economic loss over a horizon, and how does that distribution respond to the protocol's own parameters? From the distribution we read VaR, CVaR, liquidation probability and bad debt (Chapter 3). The emphasis is on *sensitivity to parameters*, so that governance can see which levers move the tail — not on producing a single calibrated headline, which awaits the data of Section 3.4.

9.2 Price process

We follow the standard progression from the institutional literature, adapted to Vanna's assets.

Geometric Brownian motion is the canonical starting point, $S_t = S_0 \exp\big((\mu - \frac{1}{2}\sigma^2)t + \sigma W_t\big)$, but constant volatility understates both the clustering and the fat tails of crypto returns — precisely the features that dominate leverage risk.

GARCH(1,1) corrects this by making variance autoregressive. With zero drift,

$$S_t = S_{t-1}(1 + \varepsilon_t), \qquad \varepsilon_t = \sigma_t z_t, \qquad \sigma_t = \sqrt{\omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2}, \qquad z_t \sim \mathcal{N}(0,1). \tag{9.1}$$

We use $\alpha = 0.08$, $\beta = 0.90$ (high persistence, typical of crypto), with ω set from the target annualised volatility. The resulting daily-return distribution is visibly fat-tailed relative to a Gaussian of equal variance (Figure 9.1) — the regime in which leveraged positions actually fail.

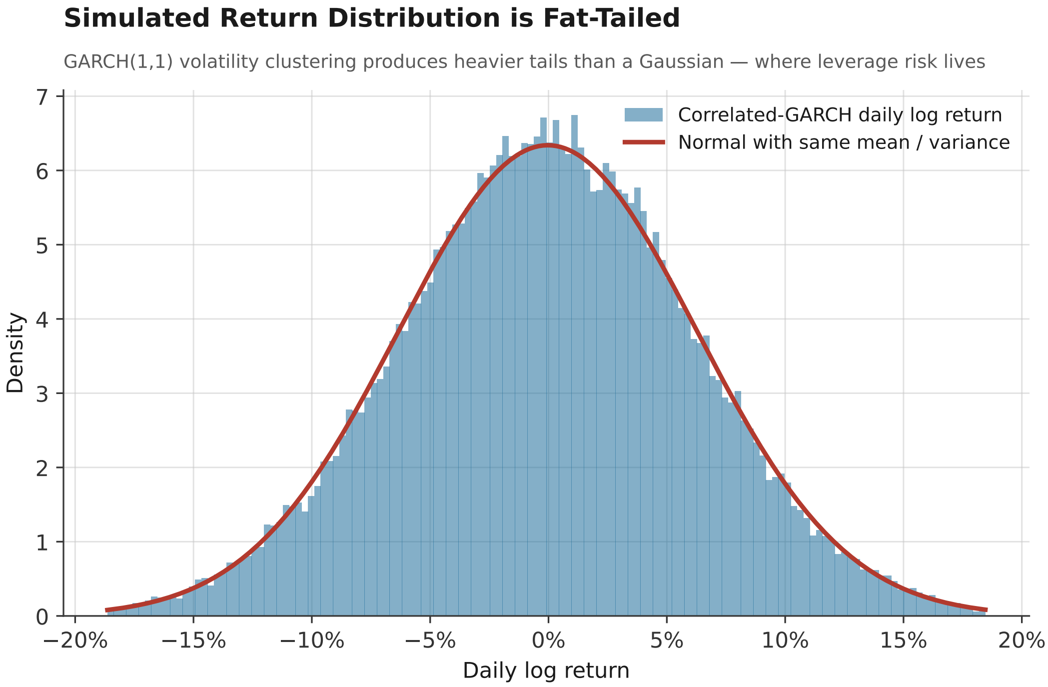

Figure 9.1: Simulated daily-return distribution (GARCH) versus a matched normal. Volatility clustering produces the heavy tails where leverage risk concentrates. [illustrative]

Correlated GARCH extends this to multiple assets by drawing the white-noise vector from a multivariate normal with the empirical correlation matrix, $\mathbf{z}_t \sim \mathcal{N}(\mathbf{0}, [\rho_{i,j}])$, so that stablecoin-collateral / XLM-debt positions inherit realistic co-movement. In the **stress regime** we superimpose a compound-Poisson jump: with intensity λ per day, a jump of assumed mean -15% and standard deviation 8% is added, modelling flash crashes. Representative paths appear in Figure 9.2.

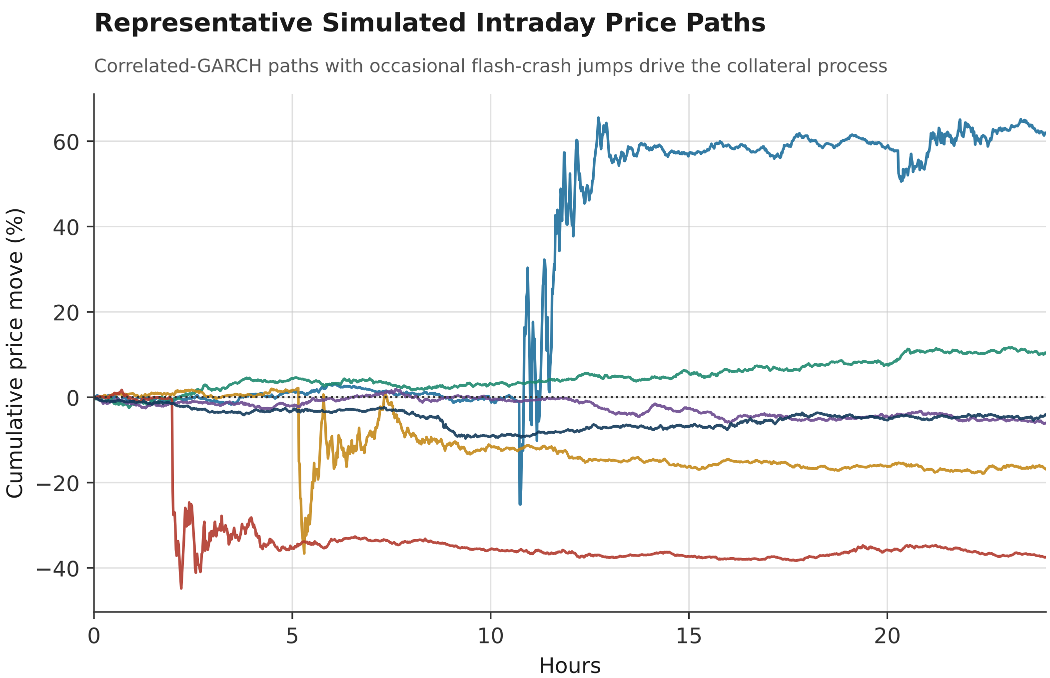

Figure 9.2: Representative simulated intraday price paths (correlated GARCH with occasional flash-crash jumps). These drive the collateral process in the Monte-Carlo layer. [illustrative]

9.3 Agent design

The economic actors are modelled as follows. **LPs** supply liquidity and redeem from the liquid slice; their redemption pressure interacts with utilisation. **Borrowers** hold leveraged positions initialised at a target health factor and are otherwise passive between rebalances (the conservative choice, since active de-risking would only reduce losses). **Liquidators** act on breached positions *only when doing so is profitable* given slippage and any bonus — which is exactly why the missing incentive (Chapter 7) matters: with no bonus, the rational liquidator population is small and slow. **Arbitrageurs** restore DEX prices after liquidation trades, bounding the depth available to the next liquidation.

9.4 Two liquidation regimes

Because the economics hinge on whether liquidation executes, every result is reported under two bracketing regimes. The **functioning-liquidation** regime closes a breached position a few minutes after breach, at market, net of slippage and overshoot (per the recovery identity of Chapter 4). The **no-working-liquidation** regime lets the position run to the horizon and marks the resulting shortfall as bad debt. These are not a prediction of the code's correctness — that is the security audit's remit — but an economic bracket: the true risk lies between them, and the distance between them is the economic value at stake in getting liquidation right.

9.5 Horizons and convergence

The base horizon is one day at minute resolution (1,440 steps), the window over which a leveraged position lives or dies in a volatility event; carry effects use the multi-month horizons of Chapter 5. Tail estimates are stable within the sample budget: Figure 9.3 shows the stress $CVaR_{99}$ converging well before the path count is exhausted, so the *magnitudes* reported are stable statistics of the assumed model (their dependence on the assumptions, not on sampling noise, is the caveat that matters).

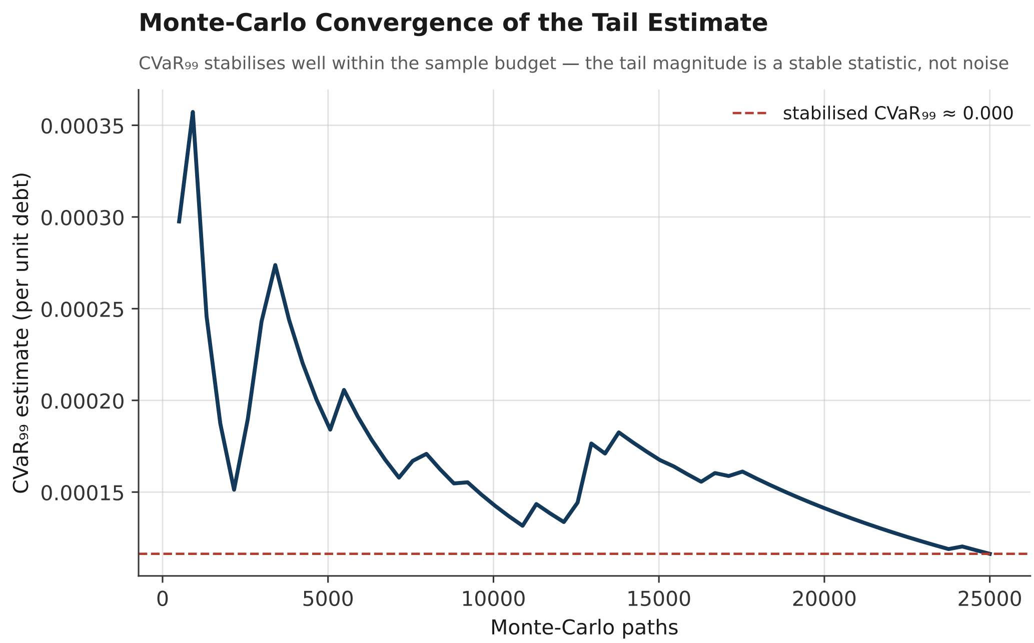

Figure 9.3: Monte-Carlo convergence of the stressed $CVaR_{99}$. The tail estimate stabilises within budget; residual uncertainty is in the model assumptions, not the sampling. [illustrative]

CHAPTER 10

Stress-Test Results

We report the base regime first (ordinary volatility, adequate depth), then the stress regime (elevated volatility, flash-crash jumps, thin depth), each under both liquidation regimes.

10.1 Base regime

At an assumed 80%/yr XLM volatility with a functioning liquidation, bad debt is negligible for any reasonably collateralised position; the risk that grows as health falls is the *frequency* of liquidation, not loss given liquidation. Table 10.1 shows liquidation probability rising from 3.4% at $HF_0 = 1.2$ to 98.6% at $HF_0 = 1.1$, while VaR/CVaR remain near zero — the signature of a system whose buffer is thin (so positions near the line are almost always liquidated) but whose losses are contained *provided liquidation works*.

Initial HF	P(liquidation)	VaR_{99} loss	$CVaR_{99}$ loss
2.0	0.0%	0.000	0.000
1.5	0.0%	0.000	0.000
1.3	0.02%	0.000	0.000
1.2	3.42%	0.000	0.000
1.15	27.18%	0.000	0.000
1.10	98.58%	0.000	0.000
1.05	100.0%	0.000	0.000

Table 10.1: Base regime, functioning liquidation (XLM vol 80%/yr, 1-day). Loss is contained; liquidation frequency is the variable that grows as the buffer thins. [illustrative]

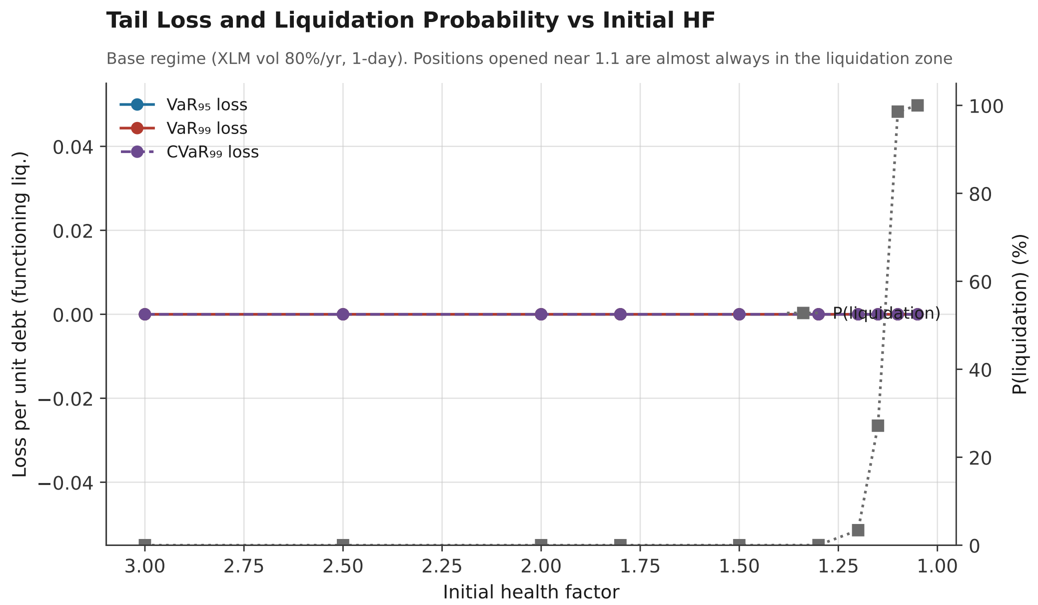

Figure 10.1: Tail loss and liquidation probability versus initial HF (base). A position opened near 1.1 is almost surely in the liquidation zone within a day. [illustrative]

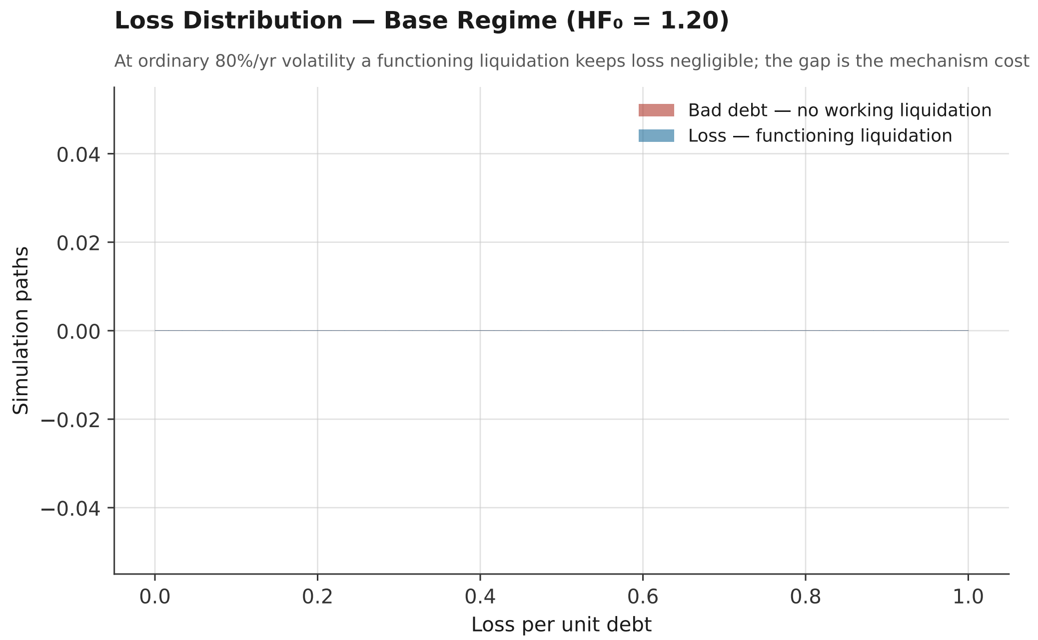

Figure 10.2: Base-regime loss distribution at $HF_0 = 1.20$. With a functioning liquidation the loss mass is negligible; the red (no-liquidation) mass is the cost of the mechanism failing even at ordinary volatility. [illustrative]

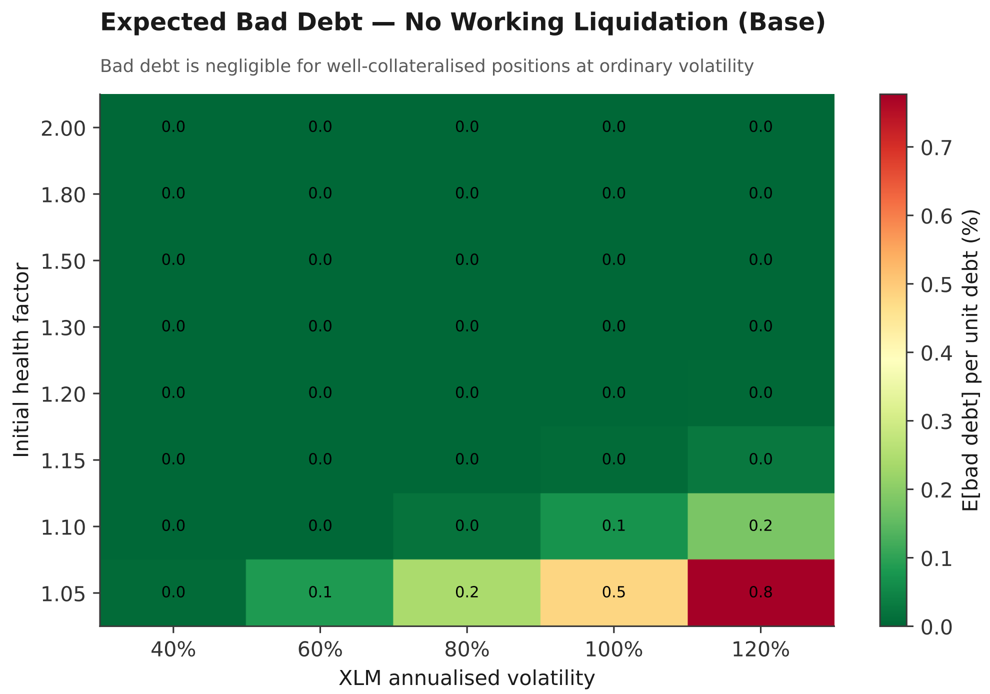

Figure 10.3: Expected bad debt over health factor × volatility (no working liquidation, base). Negligible for well-collateralised positions at ordinary volatility. [illustrative]

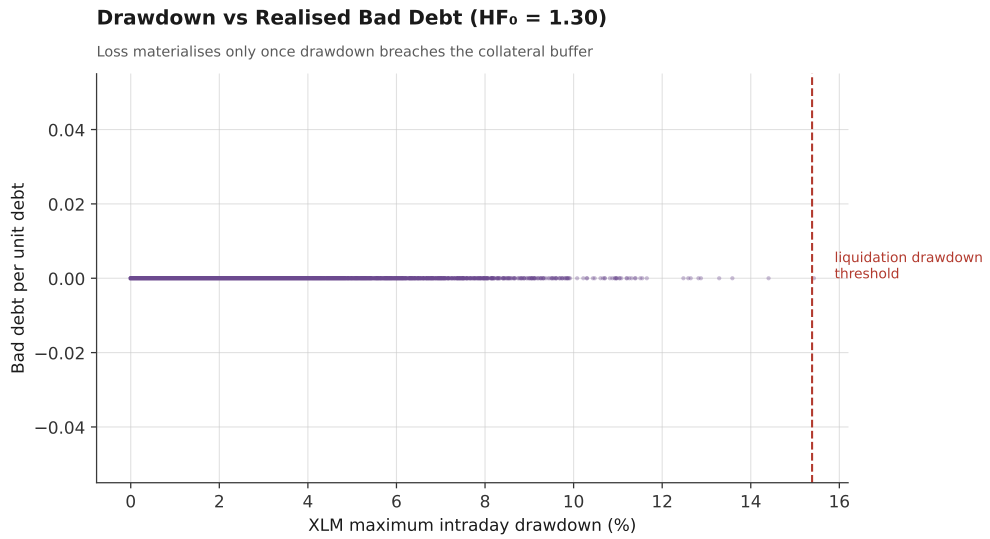

Figure 10.4: Maximum intraday drawdown versus realised bad debt ($HF_0 = 1.30$, base). Loss appears only once drawdown breaches the collateral buffer. [illustrative]

10.2 Stress regime

The picture changes qualitatively under stress. Table 10.2 and Figure 10.5 report both regimes across the health spectrum. With a functioning liquidation the tail stays near zero even under 120%/yr volatility with flash-crash jumps and 8% slippage ($CVaR_{99} \approx 0.0000$–0.0002). Without a working liquidation, the same shocks produce a $CVaR_{99}$ bad debt of 0.49 at $HF_0 = 2.0$, rising to 0.62 at 1.5 and 0.72 at 1.1 — that is, in the tail the protocol loses roughly half to three-quarters of the debt notional. The central result of this report follows directly:

The stressed tail is ≈ 0.49–0.73 per unit debt as-implemented versus ≈ 0.0002 with a functioning liquidation. The tail is therefore a *mechanism artifact*, not a market inevitability — which is precisely why it is fixable through liquidation economics and parameterisation.

The no-working-liquidation column is not a hypothetical bracket for the flagship strategy — it is the *operating point*. The opt-5 audit confirms (finding REG-01, Chapter 7) that `liquidate` *deterministically reverts* for any account holding a Blend or LP position, because the auto-unwind re-enters a trader-auth-guarded public entry behind a reentrancy lock. So for the accounts that carry the most credit risk, recovery does not merely lose in the tail; it does not happen. The realistic operating point for those accounts is the red (no-liquidation) column — CVaR_{99} bad debt of 0.49–0.73 per unit debt — and the ≈ 0.0002 "functioning" figure is unreachable until REG-01 and its siblings are fixed. This makes the liquidation-path repair the report's top-priority action.

Initial HF	P(liq)	Functioning liq. CVaR_{99}	No-liq. CVaR_{99} (bad debt)
2.0	8.6%	0.0000	0.493
1.8	12.0%	0.0000	0.544
1.5	21.2%	0.0000	0.620
1.3	31.8%	0.0000	0.671
1.2	45.6%	0.0001	0.696
1.15	66.3%	0.0001	0.709
1.10	99.0%	0.0002	0.721
1.05	100.0%	0.0002	0.734

Table 10.2: Stress regime (vol 120%/yr + flash-crash jumps + 8% slippage, 1-day), both liquidation regimes. The two columns bracket the economic value of a working liquidation. [illustrative]

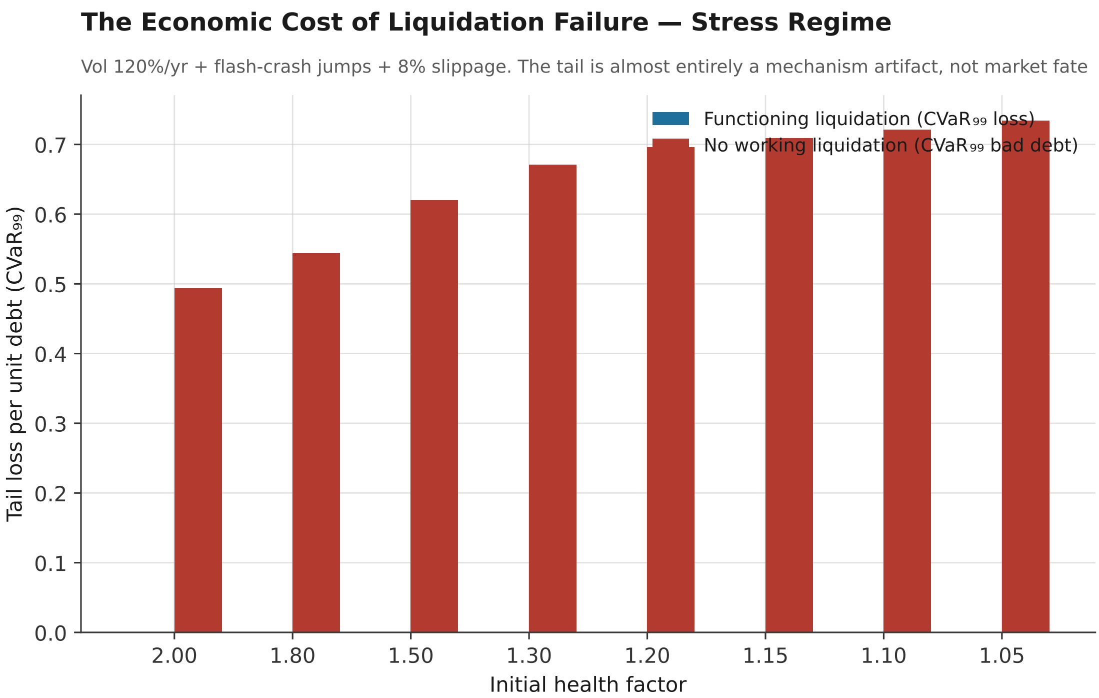

Figure 10.5: The economic cost of liquidation failure, stress regime. Blue (functioning liquidation) is near zero across the spectrum; red (no working liquidation) is 0.49–0.72 per unit debt. The gap is the value at stake in liquidation economics. [illustrative]

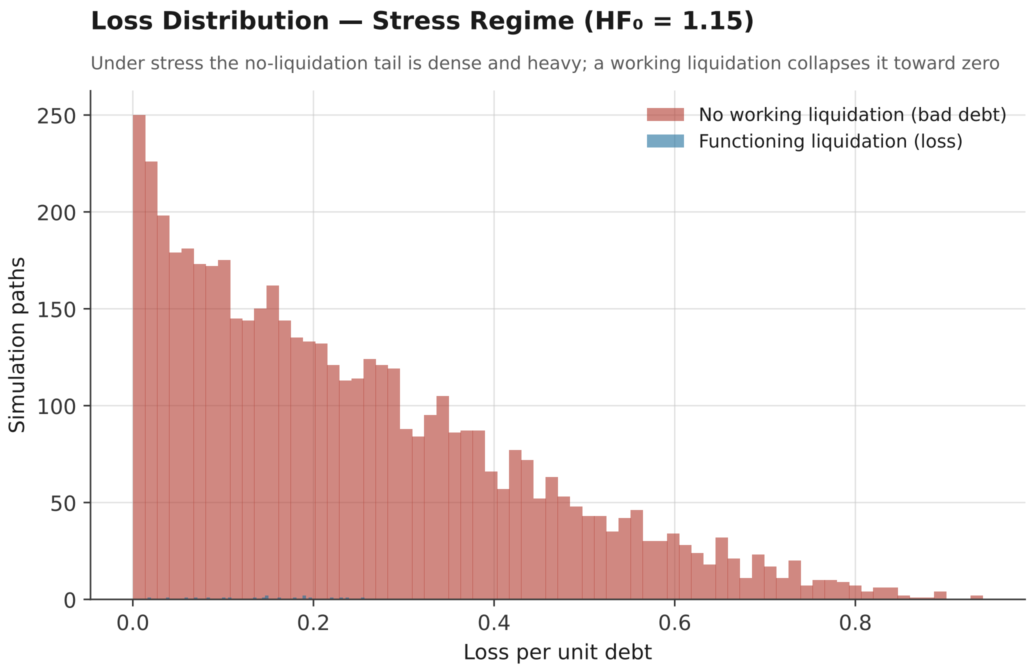

Figure 10.6: Stress-regime loss distribution at $HF_0 = 1.15$. The no-liquidation tail is dense and heavy; a functioning liquidation collapses it toward zero. [illustrative]

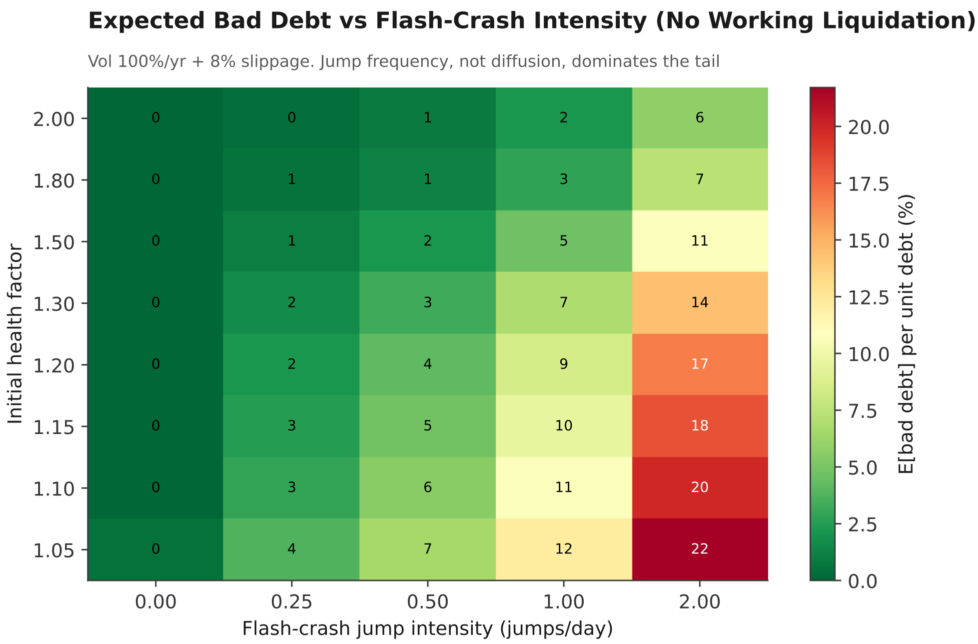

Figure 10.7: Expected bad debt over health factor × flash-crash intensity (no working liquidation, vol 100%/yr, 8% slippage). Jump frequency, more than diffusion, drives the tail. [illustrative]

10.3 Slippage sensitivity

Holding health at $HF_0 = 1.15$ and sweeping slippage (Table 10.3) isolates the buffer's fragility: inside the 9% buffer the functioning-liquidation tail is essentially zero, but it climbs to 0.14 at 20% and 0.30 at 35% slippage. This is the quantitative case for sizing a liquidation bonus above expected stressed slippage and for widening the buffer via collateral factors.

Slippage	$CVaR_{99}$ loss / unit debt
1%	0.0001
5%	0.0001
10%	0.0295
20%	0.1373
35%	0.2991

Table 10.3: Loss versus slippage at $HF_0 = 1.15$ (stress, functioning liquidation). The break in behaviour occurs at the 9.09% buffer. [illustrative]

10.4 Scenario matrix

Table 10.4 summarises the economic consequence of a set of named stress scenarios and identifies which mechanism governs the outcome.

Scenario	Economic consequence	Governing mechanism
XLM −20% intraday	Cross-asset positions near the line breach; contained if liquidation works, else material bad debt	Liquidation; buffer
Flash crash −40% (jump)	Overshoot exceeds buffer; loss even with prompt liquidation unless bonus/close-factor present	Buffer; incentive
External yield → 0	Every leveraged supply position turns/worsens negative-carry; equity decay accelerates	Carry; strategy gating
External exchange-rate anomaly	TrackToken mismarked; collateral over/under-valued; recovery impaired	Oracle/external dep.
Oracle stale / deviating	Liquidation fires late or on wrong price; overshoot grows	Signal; oracle
Utilisation → 100%	Redemptions blocked at 175% APR; run risk; borrowers may never repay	Liquidity; caps
Correlated de-risking	Simultaneous breaches exhaust thin depth; slippage spikes; tail fattens	Depth; close factor

Table 10.4: Economic scenario matrix. In almost every row the outcome is governed by liquidation economics and/or carry — the two priority areas.

CHAPTER 11

Solvency, Systemic & Contagion Risk

11.1 The solvency condition

Protocol solvency requires aggregate recoverable collateral to cover aggregate debt (Chapter 4). The economic fragility is that "recoverable" is not "marked": collateral is valued at oracle price and external exchange rate, but recovery requires actually converting the external claim to underlying, on demand, net of slippage. The wedge between marked and recoverable value — overshoot plus slippage minus buffer — is exactly the bad-debt term of Chapter 10, and it is why a protocol can be "solvent" on paper and insolvent in a liquidation.

11.2 Contagion topology and the death spiral

Vanna's design contains a reinforcing loop (Figure 11.1). A shock — an XLM crash, an external exchange-rate anomaly, or an oracle move — mismarks collateral; leveraged positions go underwater; if liquidation fails or loses on the thin buffer, bad debt is socialised to LP claim value; falling claim value triggers an LP redemption rush; redemptions push utilisation toward 100% (and the borrow rate toward 175%), locking exits; and the stressed, illiquid state makes the next liquidation even harder — closing the loop. Recursive leverage (Chapter 6) is the amplifier: because the levered stack is concentrated on one venue and asset, the initial shock hits a large correlated exposure at once.

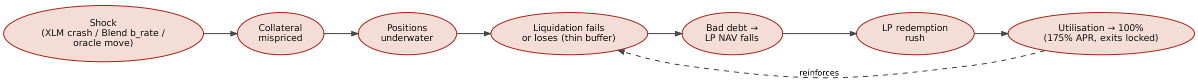

Figure 11.1: Contagion / death-spiral topology. Each arrow is an economic transmission; the dashed feedback (utilisation spike → harder liquidation) is what makes the loop self-reinforcing rather than self-correcting. [illustrative] for structure.

11.3 External-protocol dependency

Because borrowed capital is deployed into an external venue (Blend), Vanna inherits that venue's economic risk directly: a collapse in the external supply yield turns strategies negative-carry (Chapter 5); an anomaly in the external exchange rate mismarks every TrackToken; and any withdrawal friction at the external venue impairs recovery precisely when it is needed. This is a genuine, irreducible dependency, not a defect; the economic response is to bound it — per-venue exposure caps, a haircut on externally-deployed collateral, and monitoring of the external exchange-rate drift — rather than to assume it away.

11.4 Systemic risk assessment

Systemic risk in Vanna is not independent of the earlier chapters; it is their composition. The same three fixes — working liquidation economics, a reserve buffer, and leverage/exposure caps — break the death spiral at multiple nodes simultaneously: a working, incentivised liquidation stops the shock→bad-debt transmission; a reserve absorbs residual bad debt before it reaches LP claims; and caps limit the correlated exposure any single shock can hit. Contagion is therefore best treated not as a separate workstream but as the strongest argument for the priority actions of Chapter 15.

CHAPTER 12

Incentive & Game-Theory Analysis

Economic risk is not only about price paths; it is about how rational agents play the mechanism. Four games matter.

12.1 Strategic default under a thin buffer

With a 9.09% maintenance buffer, no borrower penalty, and an $11\times$ drift ceiling, the payoff to running high leverage is asymmetric: the upside accrues to the borrower while a fraction of the downside, once the buffer is breached in a gap move, is socialised to LPs as bad debt. Although opening is capped near $5\times$ (so a borrower cannot start at the line), a rational risk-seeking borrower is still incentivised to open at the cap and let the position drift toward the $11\times$ maintenance ceiling with minimal self-de-risking. A borrower penalty that makes liquidation strictly costly, together with a higher liquidation threshold (a lower maintenance ceiling), realigns this.

12.2 LP deposit / withdraw timing (resolved)

An earlier draft posited a timing game arising if the LP claim were priced against the liquid balance: deposit at low utilisation and redeem at high (or vice versa) to extract value from the other cohort. **Code reconciliation shows the claim is priced against total assets** ($L_{\text{liq}} + B + v_0$, Section 4.7), so claim value tracks interest and this arbitrage does not exist. The game is withdrawn; no action is required here beyond preserving the total-assets pricing.

12.3 The phantom-collateral exploit

The confirmed double-count (Chapters 4, 6) is not just a passive mis-accounting; it is an actively profitable strategy for an adversary, which is why it belongs in the game-theoretic analysis. The play is: deposit both legs of a pair as collateral; borrow against the (already correct) collateral; deploy the borrowed asset externally, which leaves the original underlying-balance ledger undecremented *and* credits a tracking token — reported collateral now exceeds real collateral; borrow again against the phantom excess; and repeat, using a swap round-trip (NEW-11) to re-credit the phantom balance each cycle. Two properties make it severe. First, it is *unbounded*: each cycle adds phantom collateral, so the only real limit is pool free liquidity (Figure 6.3). Second, the resulting position *cannot be liquidated* — the same inflated reported health that enabled the over-borrow also makes the account read as healthy, so `liquidate` rejects it, and even if it did not, the unwind reverts (REG-01). The expected bad debt per malicious account is therefore approximately

$$\mathbb{E}[\text{bad debt}] \approx \text{pool free liquidity} - \text{real seizable collateral}, \tag{12.1}$$

i.e. the attacker can extract close to the entire free liquidity of a pool. This single exploit chains E-3 (double-count), E-2 (liquidation failure), and E-6 (no reserve) into a drain, and is the reason the WAD double-count fix is ranked P0 alongside the liquidation-path repair.

12.4 Recursion incentives

Because recursion carries no explicit marginal cost (no per-loop penalty, no exposure cap), a yield-maximising agent is incentivised to loop toward the opening-gated ceiling ($5\times$) and, via the repeatable double-count above, without bound — in both cases maximising concentration on a single venue and asset. Per-asset and per-venue exposure caps, a cost to re-pledging, and the double-count fix blunt this incentive.

12.5 Oracle-timing and liquidation-timing games

An earlier draft treated spot valuation at the instant of a large action as manipulable. **Code reconciliation shows HF-critical valuation uses a $\approx$25-minute TWAP** gated on spot-freshness, with raw spot only a staleness-checked fallback — so the direct spot-manipulation game is largely closed. Two residual, opposite-signed points remain and are the real oracle agenda: the TWAP introduces a lag (an intentional manipulation-resistance tradeoff that can delay the trigger), and the fallback price lacks staleness validation. Adding staleness validation to the fallback, and monitoring TWAP-vs-DEX deviation, addresses these without reintroducing spot manipulability.

12.6 Incentive assessment

After reconciliation, the live incentive problems are three: the strategic-default asymmetry (thin buffer, no penalty, high maintenance ceiling); the recursion / concentration incentive; and — most serious — the *phantom-collateral exploit*, which weaponises the double-count and the liquidation failure into a pool drain. The LP-timing and spot-manipulation games are resolved in code. The corrective set for what remains is standard but urgent — fix the WAD double-count and the liquidation-revert (both P0), then add a borrower penalty, liquidator bonus, close factor, a higher liquidation threshold, exposure caps, and fallback-price staleness validation — and is consolidated in Chapter 15.

CHAPTER 13

Oracle Risk: Lag and Panic-DoS

The prior edition treated the oracle only through the manipulation lens (Chapter 12) and concluded, correctly, that the TWAP largely closes spot manipulation. The opt-5 audit surfaces a distinct and more economically consequential oracle risk that is not about manipulation at all: the valuation code *panics* on degraded input, so oracle problems freeze the protocol rather than degrade it gracefully.

13.1 The lag tradeoff

HF-critical valuation uses a ≈25-minute time-weighted price [code-confirmed]. This is the right design choice against manipulation, but it is a genuine economic tradeoff: in a fast move, the collateral is marked at a price up to ≈25 minutes stale, so a position can be materially underwater on spot while still reading healthy on the TWAP. Under negative carry and a thin buffer this lag is on the dangerous side — it delays the (already fragile) liquidation trigger — and it should be monitored via a TWAP-vs-DEX deviation alarm rather than tightened blindly (tightening re-opens manipulation).

13.2 Panic-DoS: degradation becomes a freeze

Several valuation paths `panic!` on stale, zero, or malformed oracle input rather than returning an error or a conservative bound (findings **V-H02 / NEW-07**, and the fallback path **V-M03**, which lacks staleness validation entirely). Because these paths sit under `borrow`, `withdraw`, and `liquidate`, an oracle outage or a single malformed feed does not merely pause pricing — it *reverts the core state-changing entry points*. The economic consequences compound the rest of the report:

- **Liquidation freeze.** Exactly when the oracle is stressed (a volatile market) is when liquidations are most needed; a panic on the valuation path means they cannot run, pushing the system further into the no-liquidation operating regime of Chapter 10.
- **Redemption freeze.** A panic under `withdraw` blocks LP exits, converting an oracle incident into a liquidity incident (Chapter 8).
- **Asymmetric availability.** If any state-changing path remains callable while valuation panics elsewhere, the protocol can be left in a partial state that favours whoever acts in the surviving path.

13.3 Oracle risk assessment

The fix is to make every valuation path *total*: return a typed error or a conservative (health-reducing) bound on stale/zero/malformed input, never `panic!`; add staleness validation to the fallback price; and expose the TWAP-vs-spot deviation as a monitored signal. Graceful degradation — freeze new borrowing but still permit liquidation and withdrawal under a conservative mark — is far safer than a symmetric panic that freezes the protective actions alongside the risky ones. This is a Medium economic risk in isolation but it interacts multiplicatively with the liquidation failure of Chapter 7, so its fix is P1 and gated on the same valuation-path review.

CHAPTER 14

Governance & Parameter Risk

Every quantitative result in this report is conditional on the protocol's parameters. The opt-5 audit shows those parameters can be changed in ways, and at a speed, that make the blast radius of a single administrative action larger than any market scenario modelled here. Governance is therefore an economic risk in its own right, not merely an operational one.

14.1 Unbounded, un-timelocked rate coefficients

The interest-rate coefficients are settable by admin through `set_coefficients` with *no bounds and no timelock* (finding **NEW-05**). Because the borrow rate enters every position's carry and health trajectory (Chapters 5, 4), a single write can, at the *next* accrual, make every borrower's debt jump discontinuously — rendering the entire book insolvent in one block, with no warning period in which borrowers could unwind. This is a strictly larger tail than any price shock in Chapter 10: it is a 100%-of-book event reachable by one transaction. Even absent malice, a fat-fingered coefficient (the values are raw fixed-point integers, easy to mis-scale) has the same effect.

The economic controls are standard governance hygiene: hard on-chain bounds on each coefficient (so no admissible value can produce an economically-absurd curve), a timelock on parameter changes (so borrowers have a guaranteed window to react), and per-change rate-of-change limits. None is present today.

14.2 Single-step admin and privileged setters

Administrative control is single-step (no two-step ownership handover, no multisig requirement enforced in-contract; findings **V-L01 / V-L02**), and other privileged setters (oracle address, external-venue addresses, fee/treasury) share the same un-timelocked, single-key exposure. The economic risk is the familiar one: key compromise or a single erroneous transaction can redirect value flows (e.g. point the oracle or a venue address at an attacker-controlled contract) or brick the protocol. A two-step, timelocked, multisig-gated admin model bounds this.

14.3 Off-chain TrackToken initialisation

The TrackToken representing an external position is initialised off-chain and its init call is front-runnable at deploy (finding **NEW-03**): an observer can seize or mis-parameterise the token between deployment and initialisation. Economically this is a supply-integrity risk on the very instrument that backs collateral valuation — if the token's parameters or ownership are not what the protocol assumes, every position priced against it inherits the discrepancy. Initialisation should be atomic with deployment and permissioned.

14.4 Governance risk assessment

Governance risk is low-likelihood but catastrophic-impact — the upper-left region of the risk matrix (Figure 1.1). It cannot be parameterised away because it *is* the parameterisation layer; the remedy is structural: bounds, timelocks, multisig, and atomic/permissioned initialisation.

Because a single governance action can dominate every market scenario in this report, these controls are P0 for the coefficient bounds and timelock, and P1 for the remainder.

CHAPTER 15

Parameter Optimisation & Recommendations

The reference-report discipline is to *bound a tail metric at a stated confidence, then maximise capital efficiency subject to that bound*. Applied to Vanna, and given that the dominant tail is now a code defect (Chapter 10), the highest-value moves are to *fix the liquidation-revert and the double-count*, then to stop the carry bleed and recalibrate buffers on live data. The table below states each economic parameter, its current status, the economic issue, a recommendation, and the tradeoff.

Parameter	Current	Economic issue	Recommended	Tradeoff
Opening leverage (IMR)	$\approx 5\times$ (HF 1.25)	Already well-placed (verified)	keep $\approx 5\times$	—
Maintenance HF (MMR)	1.10 ($11\times$ ceiling)	Drift ceiling high; thin buffer	1.15–1.20	Earlier liquidation
Liquidation unwind path	**reverts (REG-01)**	Liquidation cannot execute for Blend/LP accounts	**fix (P0 code)**	—
Collateral double-count	**repeatable**	Unbounded phantom leverage; blocks liquidation	**fix (P0 code)**	—
Rate coefficient controls	**unbounded, no timelock**	One admin write can insolvent the book	**bounds + timelock (P0)**	Slower param updates
Collateral factors	none (1.0)	Risk-blind; volatile collateral under-margined	stables 0.85–0.90; XLM 0.55–0.70; TrackToken haircut	Lower borrow power
Liquidation penalty (borrower)	none	No deterrent; no reserve funding	2–5%	Borrower cost
Liquidation bonus (keeper)	none	No incentive to liquidate	4–10% ($\geq$ stressed slippage)	Paid from collateral
Close factor	full	Over-liquidation; market impact	25–50%	More transactions
Reserve factor	0%	LPs fully first-loss; no buffer	10–20%	Lower LP APR
Borrow base slope	high (17.5%@50%)	Negative carry	lower base (2–6%), later kink	Less low-util LP yield

Parameter	Current	Economic issue	Recommended	Tradeoff
Utilisation kink	implicit	High floor across range	target 80–90% kink	Curve shape
Oracle valuation paths	**panic on bad input**	Oracle degradation freezes borrow/withdraw/liquidate	make total (typed error / conservative bound)	—
Oracle fallback	no staleness check	Stale fallback price accepted	add staleness validation; monitor lag	—
TrackToken init	off-chain, front-runnable	Supply-integrity risk on collateral instrument	atomic + permissioned init	—
Redemption cap	hard 50% panic	Large redeem reverts; front-runnable (NEW-12)	partial-fill + queue	—
Supply / borrow caps	none	Concentration; liquidity risk	per-asset + per-account caps	Capacity limits
Admin model	single-step key	Compromise / fat-finger blast radius	two-step + multisig + timelock	Ops overhead
Utilisation ceiling / queue	none	Run risk; no guaranteed redemption	cap U or queue	LP flexibility

Table 15.1: Economic parameter recommendations (opt-5-reconciled). Rows already satisfied (opening leverage; interest-accrued HF; total-assets NAV; TWAP mid-price) are omitted; rows in bold are confirmed code defects with direct economic consequence, several now P0. Ranges are starting points to be refined on live-calibrated data (Section 3.4).

15.1 Prioritised economic roadmap

P0 — before custodying *any* meaningful value on mainnet. The gating items are now code defects, not tuning. *(i)* Fix the liquidation-revert (REG-01) and its sibling paths so liquidation can execute for Blend/LP accounts; *(ii)* fix the repeatable collateral double-count (NEWSOL-C1 / NEW-11) that enables unbounded phantom leverage and the pool-drain exploit; *(iii)* bound and timelock the rate coefficients (NEW-05) so no single admin write can insolvent the book. Only then do the economic P0s bite: recalibrate the borrow curve and gate uneconomic strategies (carry); add a liquidator bonus and borrower penalty; add a reserve factor; lower the maintenance ceiling toward 1.15–1.20; introduce collateral factors; and add per-asset and per-account caps. (Already in place and *not* P0: the opening cap, interest-accrued HF, total-assets NAV, and TWAP mid-price.)

P1 — before scaling. Make every oracle valuation path total (no `panic!`) and add fallback staleness validation; make TrackToken initialisation atomic and permissioned (NEW-03); convert the 50% redemption cap to a partial-fill/queue and close the redeem front-run (NEW-12); move admin to two-step + multisig + timelock; add a close factor / partial liquidation; re-estimate all illustrative magnitudes on the live-calibrated data of Section 3.4; and stand up the monitoring dashboard below.

P2 — ongoing. Introduce dynamic (utilisation- and volatility-responsive) parameters, per-venue exposure caps, scheduled stress re-runs, and a governance process for parameter updates with clear bounds.

15.2 Monitoring KPIs

The following should be tracked continuously and alarmed at thresholds: utilisation per pool; the distribution of open position health factors and aggregate leverage; realised net carry; bad-debt balance and reserve coverage ratio; liquidation success rate and latency (breach-to-close time); oracle staleness and oracle-vs-DEX deviation; top-k LP concentration; and external-venue exchange-rate drift versus expectation.

CHAPTER 16

Conclusion

16.1 Economic capacity statement

On the economics as built in opt-5, Vanna is **not ready to custody value**, and the opt-5 audit sharpens *why* from the prior edition's "economic fragility" to concrete, confirmed code failure. Several design elements are genuinely sound and verified, and deserve credit: opening leverage is sensibly capped near $5\times$ with a real safety runway; the debt side of the health factor is interest-accrued; the vToken NAV is struck fairly against total assets; the origination fee is borrower-borne, not an LP drain; and HF pricing uses a TWAP rather than raw spot. But three load-bearing problems now dominate. First, the flagship leveraged-supply strategy is **structurally negative-carry** under realistic external yields, destroying borrower equity absent any market move. Second, and decisively, **liquidation for that strategy does not merely lose in the tail — it deterministically reverts** (REG-01, plus ≥ 5 sibling paths), so the accounts carrying the most credit risk cannot be closed; the 0.49–0.73 per-unit-debt stress tail is the *operating point*, not a bracket. Third, a **repeatable collateral double-count** removes the leverage ceiling entirely and, chained with the liquidation failure and the absent reserve, yields a phantom-collateral exploit that can drain close to a pool's free liquidity. A fourth, structural risk sits above all of these: an **unbounded, un-timelocked admin** can insolvent the whole book in one transaction.

The constructive finding survives: none of this is a market inevitability. The stress tail collapses to ≈ 0.0002 once liquidation actually executes, and the exploit disappears once the double-count is fixed. The path is therefore concrete and *code-first*: fix the liquidation-revert and the repeatable double-count and bound/timelock the coefficients (all P0 code); then recalibrate the carry, add liquidation incentives and a reserve, lower the maintenance ceiling, add collateral factors and caps, make the oracle paths total, and harden governance and TrackToken initialisation; and only then re-estimate every illustrative magnitude against live Stellar / Blend / Reflector data before raising limits. The sound elements already in place show the protocol can get there — but the P0 code defects must be closed before any value is at stake.

16.2 Scope reminder

This assessment is economic and quantitative. It does not itself certify the correctness or security of the contract code; it consumes the findings of the opt-5 security audit and translates them into economic consequence. Its structural claims have been reconciled against the contracts (Section 1.5), and its most severe economic conclusions — the liquidation-revert (REG-01), the repeatable double-count, and the unbounded coefficients — are code defects that the security audit owns and that must be fixed, not merely parameterised around, before operating with user funds at scale.

APPENDIX A

Assumption Ledger

Every material assumption behind the illustrative results, with the direction in which it biases the reported loss. "Anti-conservative" means the true risk is likely *worse* than reported; "conservative" means the reported risk likely *overstates* the truth.

Assumption	Value used	Bias direction
External supply yield	3% (Blend)	Neutral; decay is exact for any yield
XLM volatility (base / stress)	80% / 120% per yr	Uncertain — calibrate to data
GARCH persistence	$\alpha = 0.08,\ \beta = 0.90$	Neutral (typical crypto)
Flash-crash jumps (stress)	$\lambda = 0.5$/day, -15% mean, 8% sd	Uncertain — calibrate
Liquidation slippage (base / stress)	1% / 8%	Anti-conservative if depth thinner
Execution delay (overshoot window)	2–3 minutes	Anti-conservative if latency longer
One liquidation per position; passive borrowers	—	Conservative (overstates no-liq loss)
USDC treated as hard peg	1.00	Anti-conservative (ignores depeg)
Liquidation regime for Blend/LP accounts	functioning *and* none, as a bracket	**Now anti-conservative: the "functioning" end is unreachable (REG-01), so the no-liq end is the operating point**
Horizon	1 day (vol) / months (carry)	Neutral

Table A.1: Assumption ledger (opt-5). The net bias is that the illustrative tail is if anything *understated* for the flagship strategy, because the functioning-liquidation regime — the benign end of every bracket — cannot execute for Blend/LP accounts (REG-01) until the code is fixed. Magnitudes must still be recalibrated on live data.

APPENDIX B

Reproducibility

The structural results are closed-form functions of the protocol parameters (rate coefficients $c_1 = 0.1$, $c_2 = 0.3$, $c_3 = 3.5$; health threshold 1.1) and can be reproduced directly from the equations in Chapter 4. The stochastic results use a correlated GARCH(1,1) price process (α=0.08, β=0.90, ω from the target annualised volatility) at minute resolution over a 1,440-step day, with 20,000 to 25,000 Monte-Carlo paths per configuration and fixed seeds for repeatability; the stress regime adds a compound-Poisson jump (λ/day, mean −15%, sd 8%) and a slippage haircut on liquidation proceeds. Loss is computed per unit debt under two liquidation regimes (functioning and none) via the recovery identity from Chapter 4, $\rho = \theta e^{-o}(1 - s)$. Convergence was checked (Figure 9.3).

APPENDIX C

Glossary

Health factor (HF) Ratio of total collateral value to debt value; liquidation at HF < 1.1.

IMR / MMR Initial / maintenance margin requirement — how much may be opened / when liquidation occurs. Vanna currently sets them equal.

Carry External yield minus borrow cost; negative carry destroys equity over time.

vToken LP receipt token representing a pro-rata claim on a pool.

TrackToken Vanna's internal claim on an externally-deployed position, valued live.

Collateral / liability factor Haircut (< 1) applied to collateral value (or add-on to debt) to reflect risk. Absent in Vanna.

Reserve factor Fraction of interest retained as protocol first-loss capital. Currently 0%.

Close factor Maximum fraction of a position liquidated in one action.

Overshoot Adverse price move between breach and liquidation execution.

Slippage Price impact of selling seized collateral into the market.

VaR$_q$ / CVaR$_q$ q-quantile loss / mean loss beyond that quantile (expected shortfall).

LaR Liquidations-at-Risk; here, probability of entering the liquidation zone over the horizon.

APPENDIX D

Economic Risk-Category Coverage Map

Risk category	Where treated
Interest-rate / carry	Chapter 5 (curve, net-carry surface, equity decay)
Leverage / capital efficiency	Chapter 6 (ceiling, recursion, unbounded double-count)
Liquidation / credit	Chapter 7 (structural failure REG-01; buffer; incentive); Chapter 10 (VaR)
Market / volatility / tail	Chapters 9–10 (GARCH, jumps, stress VaR/CVaR)
Liquidity / redemption	Chapter 8 (bank-run geometry; 50%-cap; front-run)
Oracle (lag + panic-DoS)	Chapter 13 (TWAP lag; panic freeze); Chapter 12 (timing)
Governance / parameter risk	Chapter 14 (unbounded coefficients; single-step admin; TrackToken init)
Contagion / systemic	Chapter 11 (death-spiral topology)
External dependency	Chapters 11, 14 (Blend / Aquarius / Soroswap)
Incentive / game-theory	Chapter 12 (default, phantom-collateral exploit, recursion)
Solvency-signal quality	Chapters 4, 5 (debt-side unbiased; collateral-side double-count)
Parameterisation	Chapter 15

APPENDIX E

References & Inputs

Reference methodology (style and quantitative practice). Chaos Labs, *Perpetual Futures Liquidity and Funding Rate Considerations for Ethena* (funding / carry, insurance-fund sizing, slippage, market-impact modelling); Chaos Labs, *Aave V3 Risk Parameter Methodology* (agent-based Monte-Carlo, GBM→GARCH→correlated-GARCH price modelling, VaR / bad-debt estimation, risk matrix).

Protocol under assessment. The Vanna Protocol Soroban contract set, optimisation 5 (the opt-5 `contracts` tree): rate model, risk engine, margin account, TrackToken / vToken accounting, and Blend / Aquarius / Soroswap integration, from which all structural parameters and mechanism behaviour in this report were read.

Security findings referenced. REG-01, REG-02 (liquidation revert / unwind); NEW-C1, NEWSOL-C1, NEW-11 (collateral double-count and swap-laundering); NEW-05 (unbounded rate coefficients); V-H02, NEW-07, V-M03 (oracle panic / fallback staleness); NEW-08 (Soroswap remove-liquidity mis-scale); NEW-12 (redeem front-run); NEW-03 (off-chain TrackToken init); NEW-M3 (borrower-borne origination fee); V-L01, V-L02 (single-step admin).

Required calibration inputs. As enumerated in Section 3.4: XLM/stablecoin price history, on-chain depth/slippage, external exchange-rate history, oracle telemetry, and anticipated borrower/LP population.

APPENDIX F

Deep Numerical Verification Log

Every [exact] number in this report is a closed-form function of opt-5 parameters and was re-derived independently. This appendix records the derivations so a reader can reproduce them without the code.

Interest-rate curve

$R_b(U) = 3.5\,(0.1U + 0.1U^{32} + 0.3U^{64})$. Evaluations: $R_b(0.5) = 3.5(0.05 + 0.05 \cdot 2^{-32} + 0.3 \cdot 2^{-64}) \approx 17.50\%$; $R_b(0.8) = 28.03\%$; $R_b(0.9) = 32.83\%$; $R_b(0.95) = 43.97\%$; $R_b(0.975) = 70.46\%$; $R_b(0.99) = 115.21\%$; $R_b(1) = 3.5(0.1 + 0.1 + 0.3) = 175.0\%$. Supply APR $R_s(U) = R_b(U) \cdot U$ (no reserve split): $R_s(0.5) = 8.75\%$, $R_s(0.9) = 29.54\%$.

Leverage bounds

$\lambda(\text{HF}) = \text{HF}/(\text{HF} - 1)$. Opening $\text{HF} = 1.25 \Rightarrow \lambda = 1.25/0.25 = 5.0\times$ ($\text{LTV} = 1/1.25 = 80\%$). Liquidation $\text{HF} = 1.10 \Rightarrow \lambda = 1.10/0.10 = 11.0\times$ ($\text{LTV} = 1/1.10 = 90.9\%$).

Recursive leverage (clean code)

Per-loop LTV $\ell = 1/1.25 = 0.8$; $\sum_{k \geq 0} \ell^k = 1/(1 - 0.8) = 5.0\times$. (Broken by the repeatable double-count; see Chapter 6.)

Collateral buffer and break-even slippage

Buffer at the 1.10 trigger $= 1 - 1/1.10 = 9.09\%$ of collateral. Break-even slippage with zero overshoot: $s^\star = 1 - 1/\theta = 1 - 1/1.10 = 9.09\%$.

Carry / equity decay (illustrative yield 3%)

Net carry $= y - R_b(0.5) = 3\% - 17.5\% = -14.5\%$. With idle margin 1 and borrowed $D = \lambda - 1$: solving $\text{HF}_{\text{true}}(t) = 1.10$ and $= 1.00$ gives, for $\lambda \in \{2, 3, 4, 5\}$: days to $1.10 = \{1368, 732, 461, 310\}$ and days to insolvency $= \{1587, 957, 690, 540\}$; equity at 1yr $= \{0.839, 0.679, 0.518, 0.357\}$. A drift-only $8\times$: $\{94, 328\}$ days, equity -0.125.

Double-count HF inflation

For one unit of deposited collateral deployed externally, reported collateral $= 1_{\text{stale}} + 1_{\text{tracking}} = 2$, so

$$\text{HF}_{\text{reported}} \approx \text{HF}_{\text{true}} + \frac{\text{deployed}}{\text{debt}} \approx \text{HF}_{\text{true}} + 1.$$

Hence true $1.25 \Rightarrow$ reported ≈ 2.25, and the 1.10 trigger is reached only at true $\text{HF} \approx 0.10$.

Illustrative Monte-Carlo tail (correlated GARCH; not calibrated)

Stress $CVaR_{99}$ bad debt with *no working liquidation*, by initial health factor:

HF_0	2.0	1.8	1.5	1.3	1.2	1.15	1.10	1.05
$CVaR_{99}$	.493	.544	.620	.671	.696	.709	.721	.734

With a functioning liquidation, $CVaR_{99} \leq 0.0002$ throughout. For Blend/LP accounts the no-liquidation row is the operating point (REG-01).

End of economic risk assessment. This document is an economic and quantitative analysis and is not a security audit, not financial advice, and not a recommendation to deploy; illustrative magnitudes require calibration on live data before use in governance.