RAIN

PROTOCOL PAPER

Risk Markets

"Games. Predictions. Insurance.

One trade. One protocol.*"*

VERSION 1.0

RAIN FOUNDATION · TECHNICAL SPECIFICATION

RAIN.ONE

RAIN Risk Markets

Protocol Paper — v1.0

RAIN Protocol · Rain.one

> ***"Humans trade risk in three forms: at random through games, for profit through predictions, and as protection through insurance. It is one trade. RAIN is the protocol for that trade."***

Abstract

Every risk-transfer industry performs the same operation: a **risk seller** (an underwriter, a vault, a house, a market maker) and a **risk buyer** agree on an uncertain event, and value moves when the event resolves. Gaming, prediction markets, derivatives, insurance and reinsurance differ only in *where the truth comes from* and *who is allowed to participate* — not in the trade itself.

This paper specifies **RAIN Risk Markets**: an open protocol for buying, selling, and trading risk itself, with trustless resolution and instant, pre-collateralized settlement. The protocol serves three verticals on one primitive:

- 1. **Games** — risk bought at random, for entertainment. **Live today** on Arbitrum One: blackjack and roulette over state channels, ~300 ms per action, provable fairness with on-chain fraud proofs.
- 2. **Predictions** — risk traded for profit. **Live** via the Rain prediction-markets side of the project (permissionless markets, order book, AI + dispute + appeal resolution, `rain-sdk-v2`).
- 3. **Parametric insurance** — risk sold away as protection. **Supported now** as a composition of the first two: a parametric cover *is* a prediction-market position bought as protection against a public-data event, underwritten by the risk vault and settled instantly (§P). Personal/individual insurance on private loss events is explicitly Stage 2 (roadmap).

The protocol also produces **secure, decentralized, trustless randomness** as a first-class service (§R), and is **token-agnostic by construction** (§T): any deployment chooses its own settlement token and vault.

A prediction market prices uncertainty about the world; a Risk Market prices uncertainty from **any declared source of truth**: a verifiable random function, a signed data feed, or a bonded oracle. All are instances of one activity: **the transparent buying and selling of risk at a mathematically defined price.** The intended builders are insurance companies, trading firms, and gaming operators — all on the same rails.

The protocol itself takes no view on what the risk *is about*. A distribution is a distribution: the probability a random draw exceeds a threshold, the probability a flight lands late, the probability an index falls outside a

band. Anyone may define a market on any distribution that passes the mathematical listing rules (§2); anyone may build a venue on top of it. One builder ships a high-frequency trading venue for synthetic risk; another ships parametric protection products for travelers or event organizers; a third ships classic prediction markets. **Same primitive, same vault, same settlement — different frontends for different publics.**

RAIN Risk Markets is a **clearing house for risk positions**:

- Every market is a published probability distribution with an immutable, machine-verified spread.
- Every outcome derives from a declared, verifiable truth source (VRF, signed data feed, or bonded oracle — §2.4).
- Every payout is collateralized on-chain before the round begins.
- Every participant — trader, liquidity provider, brand, affiliate, market creator — is paid programmatically, per round, from a single transparent value flow.
- Protocol revenue autonomously buys and burns RAIN.

The design goal is stated plainly: **win on price.** The protocol's cost structure (no intermediaries, no payment processing, no physical operations, no discretionary staff) permits a base spread of **1.25%** (governance-adjustable within [0.5%, 2%]) — the lowest generalized price of risk offered across a full catalog of synthetic risk products — while sustaining competitive returns for liquidity providers and a real, volume-driven business for distribution partners.

A second design goal follows from the first: **the deepest position capacity in the market.** Because the risk book is carried by an open vault rather than a single operator's balance sheet, position limits scale linearly and transparently with vault TVL. At $300M TVL the protocol supports positions of $600,000 on even-money outcomes — exceeding the limits available at any centralized venue, physical or digital.

---

1. The Primitive

1.1 One definition covers every product

A **Position** is a tuple:

```
Position = { stake, distribution, payoutTable, spread }
```

- `distribution` — a finite set of outcomes with fixed probabilities, resolved by verifiable randomness
- `payoutTable` — the multiple paid on each outcome
- `spread` — the price of risk: $1 - \Sigma\ (p_i \times \text{payout}_i)$

A binary even-money draw, a uniform multiplier, a 36-to-1 single outcome, a heavy-tailed multi-outcome ladder, a threshold on a weather index, a delay on a scheduled flight — the protocol does not know product names. It knows distributions, payout tables, and spreads. This indifference is the core architectural property: **the engine is product-agnostic and distribution-aware.** Product identity — what a distribution is called, how it is rendered, whom it serves — lives entirely at the frontend layer (§7).

1.1b Unified protocol architecture — the two halves of Rain

Rain is one project built by two teams, and the two halves compose:

- **The prediction-markets half** (the Rain markets team): permissionless market creation, trading, order book, liquidity, split/merge, and an AI + dispute + appeal **resolution layer** that serves as the protocol's oracle for world events. Shipped as `rain-sdk-v2` on npm, live on Arbitrum.
- **The risk-vault + state-channels half** (this paper's implementation): the underwriting vault, exposure engine, HouseBond, and RAIN Channels — instant, gasless, fraud-provable settlement of high-frequency risk positions, live on Arbitrum One with blackjack and roulette as the first products.

The unified SDK (`@rain/risk-markets-sdk`) exposes both halves behind one module map — `rain.games`, `rain.predictions`, `rain.vaults`, `rain.core` — wrapping `rain-sdk-v2` as an unmodified dependency. The three verticals map onto this architecture directly: games run on channels + vault; predictions run on the markets half; parametric insurance composes the markets half's resolution layer with the vault's underwriting capacity.

1.2 Relation to RAIN prediction markets

	Prediction Markets	Risk Markets
Source of uncertainty	Real-world events	Any declared truth source (§2.4)
Probabilities	Discovered by trading	Declared by the distribution
Settlement	Oracle / resolver	Truth-source proof
Counterparty	Order book / AMM	Risk Vault
Primitive	Risk transfer	Risk transfer

Same protocol family, same token, same brand/SDK distribution model, one thesis: **RAIN prices risk. Any risk.**

1.2b The risk continuum

Every risk-transfer industry — derivatives, insurance, reinsurance, prediction markets, gaming — performs the same operation: someone pays a spread over mathematical expectation to shed (or acquire) exposure to a distribution. The industries differ only in **where the truth comes from** and **who is allowed to participate**. The protocol organizes this as a continuum of truth sources, strongest first:

```
Cryptographic truth   (VRF proof)          → synthetic distributions — instant, exact
Parametric truth      (signed data feeds)  → indices: weather, flight status, prices
Economic truth        (bonded oracle game) → asserted real-world facts under stake
```

A position on a synthetic distribution converges in minutes; a parametric position settles the moment the index prints; an asserted fact settles through an optimistic assertion-and-challenge game (§2.4). The vault does not care: it prices spread against variance either way. **This is why the protocol layer speaks only in**

expectation, variance, and truth sources — any narrower vocabulary would silently exclude most of the addressable market.

1.3 Why the statistics guarantee convergence

For a product with per-unit edge e and per-unit standard deviation σ, after n positions the expected trader loss is $e \cdot n$ while the noise is $\sigma \cdot \sqrt{n}$. Expectation grows linearly; luck grows as a square root. The probability that aggregate trading ends profitable for the buy side falls below 5% when $n \approx (1.65 \cdot \sigma / e)^2$ and below 1% at $n \approx (2.33 \cdot \sigma / e)^2$.

At protocol scale — thousands of traders, millions of rounds — aggregate revenue is a deterministic function of volume:

```
Expected Vault PnL ≈ spread × volume
```

Every participant's share of that flow is therefore not a speculative bet but a **statistical annuity**, priced per round and paid per round.

---

2. Distribution Engine

2.1 MarketDefinition

Any market listed on the protocol is an on-chain object:

```
MarketDefinition {
  outcomes[]        // probabilities, summing to 1
  payoutTable[]     // multiple per outcome
  randomnessSource  // VRF config
  spreadCheck       // engine-verified: Σ(p·payout) ≤ 1 − minSpread
  exposureProfile   // σ and max multiple, computed by the engine
}
```

2.2 Listing rules — mathematics as the listing committee

A market is listable if and only if:

- 1. **Spread floor:** $\Sigma\ (p_i \times \text{payout}_i) \leq 1 -$ base spread (1.25% at launch; see §4)
- 2. **Closed distribution:** all outcomes and probabilities are fixed before the round; no in-round decisions affect expectation
- 3. **Exposure computability:** the engine can derive σ and max payout from the definition

No committee, no negotiation. If the math passes, the market lists.

2.3 Two tracks

- **Track A — Closed distributions (permissionless).** All outcomes and probabilities fixed before the round: binary draws, uniform multipliers, single-outcome long shots, multi-outcome ladders, threshold positions on external indices. Anyone may define and list a market. Creator incentives per §7.
- **Track B — Decision-path distributions (curated).** Products where in-round participant choices affect expectation. Expectation cannot be machine-verified from a payout table alone, so these are listed only by protocol governance, with pricing per §4.4.

2.4 Truth sources

Every `MarketDefinition` declares its **truth source** on-chain. A position buyer sees, before staking, exactly what will decide the outcome and how it is verified. Three classes are supported:

(a) Verifiable randomness — synthetic distributions.

- **Source:** Chainlink VRF v2.5 on Arbitrum. Each round's outcome is a deterministic function of a VRF output whose proof is verified on-chain.
- **Rapid mode:** VRF words are requested in batches and consumed sequentially, enabling 2–3 second rounds and multi-round positions (one transaction, N rounds).
- **MEV protection:** outcomes are bound to positions at request time (commit-then-reveal against the VRF fulfillment); no position can be placed or modified between randomness generation and settlement.

2.5 Randomness liveness & fallback

Verifiable randomness is a hard dependency, and the protocol treats its failure as a first-class scenario:

- **Invariant: no proof, no settlement.** A round may only settle against a randomness output whose cryptographic proof verifies on-chain. Under no failure mode does the protocol substitute unverifiable randomness.
- **Fulfillment timeout.** Every randomness request carries a deadline (measured in blocks). If the oracle fails to deliver within the deadline, all positions bound to that request are **voided and stakes refunded in full, automatically** — funds are never stranded on an unfulfilled request.
- **Buffer degradation, not security degradation.** Rapid mode consumes pre-verified batches. If the buffer empties (oracle latency or outage), new position placement **pauses** until fresh verified words arrive; open escrows follow the timeout rule above. Liveness degrades; integrity never does.
- **Redundant subscriptions.** The protocol maintains multiple independent VRF subscriptions and key hashes, so a single subscription or key failure does not halt the system.
- **Governed fallback provider.** A secondary verifiable randomness source (a threshold-signature beacon such as drand, wrapped behind the same on-chain proof-verification interface) is pre-approved by governance and can be activated **for liveness only**: activation applies to new rounds exclusively, never mid-round, and is announced by an on-chain event. Rounds are never migrated between randomness sources after placement.
- **Observability.** All timeout refunds, buffer pauses, and fallback activations emit events; the Anomaly Breaker (§3.2) operates independently of the randomness pipeline and remains active throughout.

(b) Signed data feeds — parametric truth.

Markets on external indices — weather readings, flight arrival times, published prices, official statistics — settle against **cryptographically signed data feeds** (Chainlink data feeds and equivalent attested sources). The market definition pins the exact feed, the reading window, and the threshold; settlement is mechanical. No claims process exists because nothing is claimed: the index either printed past the threshold or it did not. Feed liveness follows the same discipline as VRF: no attested reading by the deadline → positions void and refund automatically.

(c) Bonded assertion — economic truth (roadmap).

For facts that no feed publishes, the protocol specifies an **optimistic assertion game**: an asserter posts the outcome with a bond; a challenge window follows; an unchallenged assertion settles the market; a challenged one escalates to a stake-weighted panel of oracle stakers whose minority is slashed, with further escalation to wider panels at higher stakes. Truth becomes the cheapest equilibrium. This class is specified for completeness and gated behind governance activation — markets may only reference it once the assertion module has been separately audited and parameterized.

A market's truth-source class is part of its on-chain identity. Frontends may filter to whichever classes fit their product and public; the vault and the fee waterfall are identical across all three.

---

3. Risk Vault

3.1 Structure

An ERC-4626 vault denominated in USDC. Depositors mint vault shares and become the market's risk seller pro rata. The vault pays winners, collects from losers, and accrues the net spread. **No yield is promised. Depositors sell risk; realized results are whatever the distribution delivers.**

Capacity is capped — by design. The vault is capped at **$300M TVL** at launch. The cap exists because vault size serves exactly one purpose: underwriting position limits — and at $300M the protocol already offers the deepest limits in the market. Capital beyond that would dilute LP returns without adding trader utility. The cap is a **DAO parameter**: governance may raise it (e.g. to $500M) when volume justifies deeper limits.

Entry into a full vault — the share market. While the vault is at cap, new capital cannot mint shares; it can only buy them. Vault shares are freely transferable ERC-4626 tokens: an existing LP may list shares at any price (e.g. 1.05× NAV), and an entrant who values the seat pays the premium. This creates a secondary market for vault capacity — the scarcer and more profitable the book, the more a seat costs — without ever expanding risk capacity or diluting the spread. Withdrawals (48h delay to prevent result-timing) free capacity, which reopens minting at NAV until the cap is reached again.

3.2 Exposure engine — three rules

Every position must satisfy all three, evaluated at placement time:

Rule 1 — Volatility cap (per position):

```
stake ≤ 0.2% × TVL ÷ σ_position
```

where σ_position is the standard deviation of the position's payout multiple.

Rule 2 — Max single payout (tail cap):

```
max_payout ≤ 1.0% × TVL
```

A fixed, simple, publicly known ceiling: no single round can remove more than 1% of the vault.

Rule 3 — Global net exposure:

```
Σ NetExposure(open rounds) ≤ 5% × TVL
```

where `NetExposure(round) = max over outcomes [payouts to winners − stakes of losers]`. Offsetting positions on the same round (e.g. red vs black) net against each other: the engine measures true exposure, not gross stakes. Shared global rounds (one spin, many traders) are therefore also a capacity multiplier.

Anomaly Breaker. Time-based stop-losses are statistically meaningless and are not used. Instead: if cumulative loss in a rolling window exceeds **5σ** of expectation given realized volume, trading halts pending review. A 5σ event (~1 in 3.5M) is not bad luck; it is evidence of a broken assumption — a randomness fault, a payout bug, or an exploit. The breaker is a security control, not a risk-management control.

3.3 The Anchor Reserve — full limits from block one

The headline limits (§3.4) are the protocol's strongest marketing asset for the highest tier of traders — and they require $300M of verifiable capacity from day one, while a young book's LPs need returns computed on capital that is actually working. The vault therefore runs in **two tranches under one pledge**:

- **Total pledge: $300M**, deposited on-chain at launch (founder/treasury capital). Position limits are a pure function of the **full pledge** — $600,000 even-money positions are live from the first block, verifiable by anyone.
- **Active tranche** (target: annualized volume ÷ 28, floor $30M): absorbs all round-by-round PnL. External LP deposits and the share market (§3.1) operate on this tranche only; LP returns are computed on it.
- **Anchor Reserve** (the remainder): a pledged backstop that is only drawn on tail events — a drawdown exceeding the active tranche's buffer or a payout beyond its capacity. For standing behind the tails it earns a **standby premium of 0.05% of volume**, paid from the vault residual before LP PnL (and paused, like all distributions, while the vault is under water).
- As volume grows, capital migrates automatically from reserve to active tranche until the reserve is exhausted at ~$8.4B/yr volume — at which point the vault is simply a $300M active book, and the DAO may vote to raise the cap.

This is the classic insurance/reinsurance structure, on-chain: the active layer prices and absorbs routine risk; the reserve underwrites catastrophe and is paid a standby rate for it. Traders see one number — the deepest limits in the market; LPs see another — returns on working capital only.

3.4 Limits table (illustrative)

Position limits scale linearly with TVL. At **$300M TVL**:

Distribution archetype	σ	Max stake	Max payout
Binary even-money (2×, p≈0.5)	~1.0	**$600,000**	$1.2M
Decision-path, curated (Track B)	~1.15	**$520,000**	~$2.1M (edge case)
Uniform multiplier @ 5×	~2.2	$270,000	$1.35M
Uniform multiplier @ 10×	~3.0	$200,000	$2M
Single outcome @ 35.44× (p=1/37)	~5.9	$82,000	$2.9M
Long shot @ 100×	~10	$30,000	$3M (Rule 2)
Heavy tail @ 1,000×	~32	$3,000	$3M (Rule 2)
Extreme tail @ 10,000×	~100	$300	$3M (Rule 2)

For reference, the deepest centralized venues — physical or digital — cap even-money positions at roughly $300K–$500K, and single-outcome high-multiple positions rarely exceed $40K anywhere. **At $300M TVL the protocol operates the deepest risk book in the market — publicly, programmatically, with no negotiation and no discretionary approval.** Limits are a pure function of TVL: the market itself decides capacity.

3.5 Vault risk profile (illustrative, $300M TVL, ~$1.5–2B monthly volume)

- Daily expectation: ~+$300K; daily σ: ~$2.5M (typical) to ~$4.5M (peak sessions of the largest position-takers)
- Bad day (2σ): ~−$5M (−1.7%); very bad day (3σ): ~−$9M to −$13M (−3% to −4.3%)
- Hard bounds per day: Rule 3 (5% open exposure) + Anomaly Breaker
- Bad month (2σ): ~−$2M to −$8M depending on volume mix; probability of a losing full year: low single-digit %
- Realistic multi-month max drawdown: ~5–10% peak-to-trough before expectation reasserts
- Expected annual return at this volume/TVL ratio: material double-digit; **not promised, engineered — and dominated in practice not by game variance but by smart-contract risk**, addressed via audits, bounties, and staged TVL growth

---

4. Fee Architecture

4.1 Base + Markup

```
Trader price = Base Spread (protocol-fixed) + Brand Markup (brand-chosen)
```

Base Spread: 1.25% (governance-adjustable within [0.5%, 2%]), decomposed:

Component	Share of volume	Recipient
Protocol fee	**0.15%**	RAIN buyback & burn (§6)
Base distribution	**0.25%**	Referring frontend/affiliate
Trader rebates (pool)	**~0.17%**	Volume-tiered rebates (§4.3)
Anchor Reserve standby	**0.05%**	Reserve tranche (§3.3)
Vault spread (net)	**~0.63%**	Risk Vault — residual claimant

Brand Markup: 0% to 2.5%, chosen per brand, disclosed on-chain and in the interface ("Protocol base 1.25% | Brand fee X%"). 100% of markup accrues to the brand. The ceiling protects the system-wide value proposition (worst-case total price ~3.75%, still below the effective pricing of legacy risk venues).

The market self-segments: API/automation aggregators at 0%, lean community brands at ~0.5%, full consumer brands (concierge service, incentives, events — funded from their own margin) at 1.5–2.5%.

4.2 Why the base is 1.25%

- Low enough to be the headline: the lowest-priced full catalog in the market, roughly half the effective price of crypto-native incumbents once their loss-leader subsidies are averaged out, and a fraction of legacy venues (2.7–8%). One price governs the entire catalog.
- High enough that at volume/TVL ≈ 2.5–3 the vault earns a competitive double-digit expected annual return **after** funding rebates and the reserve standby premium, sustaining open liquidity without emissions.
- Calibration validated in the live simulator (Monte Carlo, 20K+ runs): at 1.0% the vault's risk-adjusted return under full launch limits was too thin relative to the tail risk it carries; 1.25% restores the correct hierarchy — the party bearing the risk earns the most.

4.3 Rebates

Volume rebates return 10–25% of a trader's generated spread, tiered by rolling volume, accrued per round, claimable on-chain. **Hard invariant: total rebates + incentives < spread, always.** (This closes the bonus-farming vector: no combination of incentives may turn expected value positive for volume-cycling strategies.)

4.4 Decision games (curated track)

Decision-path products (Track B — distributions where in-round choices affect expectation) carry ~0.5% structural spread (to the vault) against optimal play. The protocol adds a **notional fee of 0.10–0.15% per unit**

staked (initial stake and each in-round escalation) — certain revenue per action, rules untouched. Effective cost to an optimal participant ≈ 0.61–0.67%: the most transparent pricing available for this product class, and optimal (automated) strategies become high-frequency fee-paying volume rather than a threat.

4.5 Vault performance fee → burn

In addition to the volume-based protocol fee, the protocol charges a **performance fee on realized vault profits**, routed entirely to RAIN buyback & burn:

```
performance_fee = 10% × (net vault PnL above High-Water Mark)
```

- **High-Water Mark (HWM):** the fee applies only to net profits above the vault's all-time high (per share). After a drawdown the vault must first recover the full loss before any fee accrues again — LPs never pay a fee on recovered ground.
- **Accrual and settlement are on-chain, per epoch.** No discretion, no reporting: the fee is computed from vault share price and burned in a publicly visible transaction.
- **Alignment.** The protocol earns from the vault only when LPs earn. RAIN holders gain direct exposure to vault performance on top of volume exposure — the two token sinks (0.15% volume burn + 10% performance burn) compound.
- **Magnitude (illustrative, §3.4 profile):** at ~19% expected gross annual vault return on $300M TVL (~$57M/yr), the performance fee adds ~**$5–6M/yr of RAIN burn** while leaving LPs a net expected return of ~17%/yr. The rate is deliberately set at 10%, not higher: **the vault is the party carrying capital at risk and must remain the residual majority claimant on its own PnL.**
- **Not load-bearing.** The vault's solvency math (§3) is unchanged; the fee is taken from realized profits only and can never push the vault below its HWM.

4.6 Markup royalty → burn

Brands keep their markup as their core economics — but a brand charging 2% markup and a brand charging 0% should not contribute identically to the token sink. A deliberately light royalty applies:

```
markup_royalty = 10% × brand markup revenue
```

- Charged only on the markup itself, never on the 0.25% base distribution — a zero-markup brand pays nothing.
- At 1.1% average markup on 60% of flow, the royalty costs a brand ~0.07% of volume — small against their ~1.0% take — while adding a burn stream that scales with exactly the brands who monetize hardest.
- Routed to the same autonomous buyback & burn as all other protocol fees.

Token sinks, summarized: **0.15% of all volume + 10% of vault profit above HWM + 10% of brand markup.** Volume, vault performance, and brand monetization each feed the burn — without ever making RAIN load-bearing for solvency.

4.7 Locked credits (deposit bonuses)

Brands may issue locked credits with wagering requirements, funded from their markup. The contract displays the full arithmetic to the trader (locked balance, remaining wagering volume, expected cost at current spread). Transparency of the deal is mandatory at the protocol layer.

---

5. Value Flow

Per unit of volume (illustrative, base-only flow at 1.25% base spread):

```
Volume (100)
│
├── Protocol fee 0.15 ──────────→ RAIN buyback & burn
├── Base distribution 0.25 ─────→ referring brand/affiliate
├── Trader rebates ~0.17 ───────→ volume-tiered, on-chain
├── Anchor Reserve standby 0.05 → reserve tranche (§3.3)
▼
Risk Vault ~0.63 (mix-dependent) — the residual claimant
└── 10% of profit above HWM ──→ RAIN burn (§4.5)
```

Plus Brand Markup (0–2.5%) flowing entirely to brands, outside the protocol waterfall.

Settlement is per round. There is no monthly reconciliation, no negative-carryover clause, no trust in reports: when the vault loses a round, there is nothing to distribute; when it earns, every claim pays instantly. All participants sit on the same boat at block granularity.

Recovery Gate. While the vault's cumulative PnL is below its High-Water Mark (i.e. LPs are under water), **distribution payouts to brands and trader rebates are suspended and accrue to the vault** until the drawdown is recovered. The 0.15% protocol fee and brand markup (the brand's own pricing) continue to flow. The party carrying the risk recovers first; the parties earning risk-free flow wait. This is enforced by contract, not by policy — and it is precisely what makes the per-round waterfall sustainable through drawdowns.

Affiliate/brand attribution is written on-chain at the wallet's first trade (default duration: 24 months, then evergreen at a reduced rate), making distribution revenue a composable, auditable — and potentially tradable — asset.

---

6. RAIN Token Integration

The protocol fee (0.15% of all volume) autonomously market-buys RAIN and burns it. The vault performance fee (10% of profits above HWM, §4.5) and the markup royalty (5% of brand markup, §4.6) are burned through the same autonomous mechanism.

- **Value accrual = f(volume) + f(vault performance) + f(brand monetization).** At $50M monthly volume: ~$75K/month burned; at $250M: ~$375K/month; at $1.5B (deep-liquidity steady state): ~$2.25M/month —

plus ~$450–500K/month from the performance fee at expected vault returns, plus ~$500K/month from the markup royalty at 1.1% average markup.

- The burn is an **autonomous protocol mechanism**, not a distribution of profits and not a promise of return. RAIN holders are exposed to protocol volume; vault LPs are exposed to spread versus variance. These are distinct instruments and are documented as such.
- The vault itself is USDC-native and requires no token emissions. Early-LP RAIN incentives from the existing treasury may be applied during bootstrap; they are additive, never load-bearing.

---

7. Ecosystem Roles

Role	Provides	Receives	Economics
Trader	Volume	The lowest generalized risk price on earth; on-chain rebates; trustless settlement; the world's highest limits	Pays 1.25% + optional brand markup
Risk Vault LP	Capital (sells risk)	Residual spread flow	No promises; expectation ≈ 0.5–0.6% of volume; drawdown profile per §3.4
Brand / Frontend	Users, UX, marketing, service	0.25% base + full markup	A $10K/mo-volume regular yields ~$175/mo at 1.5% markup — CAC pays back in weeks
Affiliate	Referrals	Share of base distribution, wallet-bound on-chain	An annuity on referred flow, verifiable per block
Market creator	New distributions (Track A)	Creator fee: 10% of protocol fee on their market	An open catalog that grows itself
RAIN holder	—	Burn exposure to total volume	§6

Go-to-market implication (from simulations): the top ~2% of traders generate roughly half of deposits and the bulk of volume (calibrated to public benchmarks of high-frequency retail risk venues: casual 78% at ~$12 avg position; regular 20% at ~$120; high-tier ~2% at ~$750; largest position-takers ~0.1% at ~$8K; top-tier globals on the order of 1 per 5,000 traders). The realistic scenario requires on the order of **hundreds of high-value traders**, not millions of casual users. The two magnets are engineered into the protocol: the lowest price and the deepest limits. Distribution economics favor organic, community-native channels and API-driven traders over paid acquisition; brands with markup can fund full-service, high-touch operations where justified.

Retention model (industry-calibrated). Acquisition follows the documented two-stage pattern of real-money platforms: **~40% of new sign-ups lapse within the first month**; the surviving base is sticky, churning at **~5%/month** (average established lifetime ~20 months). Organic referrals from the established base add ~4–5%/month of new sign-ups. Blended trader lifetime value converges to the **industry-standard ~$2,500** as cohorts complete their lifecycle (realized over ~18–24 months, not at sign-up).

Launch economics (illustrative, validated in the live simulator). At a blended CAC of **$700** per active trader (declining to ~$600 in year 2 and ~$500 from year 3 as the brand matures), a **$1M launch budget seeds ~1,400 traders** over a 90-day ramp. Growth thereafter is brand-funded: brands reinvest a declining share of their revenue into acquisition (≈50% in year 1 → 20% from year 4), bounded by an operational onboarding ceiling of ~25% base growth per month. An optional **max-growth posture** — brands reinvesting 90% of revenue and retaining 10% — roughly triples the two-year trader base at the cost of near-term brand profit; committed marketing accumulates in a war chest and deploys at the pace the onboarding ceiling allows. The cheaper channel remains indirect: recruiting brands and affiliates (who carry their own CAC for the 0.25% base + markup) scales the trader base at protocol cost near zero. One brand realistically serves ~600 traders (support, local marketing, payments), so the brand count grows naturally with the network.

7P. The Parametric Insurance Model

A parametric cover is a prediction-market position bought as protection. No new machinery is required — the vertical is a composition of components that are live today:

- **The event** is public-data-verifiable: *flight AB123 lands ≥ 3 h late*, *rainfall at station X exceeds Y mm*, *price of Z closes below P*. No claims process exists because nothing is claimed: the data either printed past the threshold or it did not.
- **The risk buyer** is the protected party: they buy the YES side of the event as a hedge, sized to the loss they want covered.
- **The risk seller** is underwriting capacity: prediction-market counterparties and/or the risk vault, which prices spread against variance exactly as it does for any other distribution.
- **Resolution** runs through the prediction-markets resolution layer (AI resolution with a bonded dispute + appeal game) or directly against signed data feeds where available.
- **Settlement** is instant and pre-collateralized: the payout exists on-chain before the position opens. No adjuster, no cheque, no delay.

Worked example — flight-delay cover. A traveler with a €400 non-refundable connection buys €400 of the YES side of "flight AB123 arrives ≥3 h late" at the market's price (say 10% → €40 premium). Underwriters sell that risk for the ≈€36 expected margin. If the flight is late, resolution reads the public arrival data and the traveler is paid €400 the moment the market resolves; if not, the underwriters keep the premium. Same trade as any market on the protocol — the *intent* (protection vs profit) lives in the frontend, not the primitive.

The precise line: parametric now, personal later. What is supported **today** is parametric cover on public-data events. **Personal/individual insurance** — private loss events, per-claim assessment, quote-based underwriting, long-lived reserve management — is explicitly **Stage 2 / roadmap**: it requires parametric data sources per claim and underwriting workflows this paper does not yet specify. We state the line plainly rather than blurring it.

7R. Randomness as a Protocol Service

The games vertical forced the protocol to solve a problem every builder of risk products eventually hits: **secure, decentralized, trustless randomness at interactive speed.** The live solution — two-party commit-reveal hash chains (unbiased if *either* party is honest), dual-signed EIP-712 states, on-chain fraud proofs, with a Pyth Entropy session seed mixed in as defense-in-depth — produces verifiable random outcomes in ~300 ms with no per-outcome oracle dependency.

This capability is a protocol service in its own right: any builder who needs random outcomes with cryptographic fairness guarantees — gaming operators, lottery-style products, random allocation mechanisms, sampling — can consume the same channel randomness or the VRF-backed round pipeline (§2.4a, §2.5). The games are the proof of service quality: they are the most adversarial consumer of randomness that exists, and they run on it in production.

7T. Token-Agnostic Deployments

The protocol is **token-agnostic by construction**: `ChannelManager` takes the settlement token and the vault as constructor parameters. Any deployment may settle in USDC, USDT, or its own token; exposure math, the fee waterfall, bonds and limits are all denominated in whatever the instance chooses. **RUSD — the play-money token of the live demo — is one instance's choice, not a protocol requirement.** An insurer can deploy an instance settling in USDC with its own vault; a trading firm can run one against its own treasury token. The unified SDK reflects this: deployments are described by a validated config object `{chainId, token, channelManager, houseBond, verifiers, vault}` and the demo instance is just the default entry in the registry.

8. Simulations

Assumptions: blended spread ~1.0–1.05%, waterfall per §5, protocol opex $150–250K/mo, initial vault $10M (founder), scaling with adoption.

Scenario A — Pessimistic: $10M monthly volume

Protocol $15K/mo; vault expectation $52K/mo with monthly σ larger than the mean (≈35% losing months — normal and survivable at this scale); distribution $17K/mo. **Below protocol break-even; a proof-of-product phase with 6–12 months of acceptable burn, not a business.**

Scenario B — Realistic: $50M monthly volume (~5,000 active traders, ~250 high-volume)

Protocol $75K/mo (≈ break-even); vault expectation $260K/mo (~15–16% annualized at volume/TVL 2.5, ~29% losing months, converging annually); distribution layer $675K/mo including markup (weighted avg markup ~1.1%); a mid-size consumer brand at $3M monthly volume earns ~$52K/mo.

Scenario C — Optimistic: $250M monthly volume (~25,000 active, ~1,250 high-volume)

Protocol $375K/mo → ~$4.5M/yr burned; vault expectation $1.3M/mo on $30M TVL (~17% annualized, ~13% losing months); distribution $425K+/mo base plus markup. Protocol revenue at DeFi infrastructure multiples implies a $45–90M valuation for the extension alone — while total volume remains an order of magnitude below the largest centralized incumbents, indicating substantial headroom.

Scenario D — Deep-liquidity steady state: $300M TVL, $1.5–2B monthly volume

The maximum-capacity regime of §3.3–3.4. Vault expectation ~$9M/mo; protocol burn ~$2.25–3M/mo; daily risk profile as specified. This is the north-star scenario the limits flywheel targets:

```
More TVL → higher limits (linear, public) → largest traders arrive →
more volume → higher realized vault returns → more TVL
```

A unique property closes the loop: a large trader who believes the seller side always wins can *become* the seller — the same wallet can buy risk in the market and sell risk in the vault, each side deepening the other.

---

9. Competitive Position

	Legacy centralized venue	Crypto-native centralized venue	
Price (headline)	2.7–8%	1% on a narrow product subset	**1.25% on everything**
Price (effective, high-volume)	~2.5–7%	~0.85–0.9% (discretionary rebates)	**~0.75–0.85%, coded and public**
Fairness	Trust	Provably fair outcomes	Provably fair outcomes **and payments**
Solvency	Trust	Trust	**Collateralized pre-round, on-chain**
Limits	Negotiated, opaque	~$1M max payout	**Public function of TVL; $600K even-money at $300M**
Rebates	Discretionary	Discretionary	**Tiered, on-chain, guaranteed**
Distribution	Closed contracts	Closed affiliate program	**Permissionless Base+Markup; white-label without negotiation**

Three claims no centralized incumbent can copy without destroying its own economics:

- 1. **"One low price on every product."** Incumbents' cheap products are loss-leaders subsidized by expensive ones; here a single 1.25% base governs the catalog — below the incumbents' volume-weighted effective price.
- 2. **"We cannot hold your money — ever."** Not a policy. An architecture.

- 3. **"The deepest risk book in the market, by design."** A single operator carries tail risk on one balance sheet; an open vault distributes it across a capital market with a mathematically bounded profile.

---

10. Launch Plan & Parameters

V1 markets (Track A, synthetic truth): binary even-money, uniform multiplier (2×–100×), single-outcome 36×, multi-outcome ladders — the canonical closed distributions. **V1.5:** heavy-tail and extreme-tail archetypes; first parametric markets on signed data feeds (flight delay, weather indices). **V2 (Track B, curated):** decision-path distributions with per-action notional fee; bonded-assertion truth source subject to separate audit and governance activation.

Initial parameters:

Parameter	Value	Mutable by governance
Base spread	1.25%	Within [0.5%, 2%]
Protocol fee	0.15%	Within [0.1%, 0.25%]
Base distribution	0.25%	Within [0.15%, 0.35%]
Markup ceiling	2.5%	Within [1.5%, 4%]
Rule 1 constant k	0.2%	Within [0.1%, 0.3%]
Rule 2 payout cap	1.0% TVL	Fixed
Rule 3 exposure cap	5% TVL	Within [3%, 8%]
Anomaly breaker	5σ	Fixed
Withdrawal delay	48h	Within [24h, 96h]

Bootstrap: founder-seeded vault ($10M, growing to target); staged TVL caps during audit maturation; double independent audits + public bug bounty before uncapped deposits; early-LP RAIN incentives from treasury (time-decaying).

Legal: published by the RAIN Foundation (Panama). The protocol layer performs no marketing, no end-user relationship, and no custody beyond contract escrow. **What a market is called and to whom it is offered is a frontend decision, and each frontend carries the regulatory posture of its own product**: a venue offering high-frequency synthetic risk positions, a distributor of parametric protection products, and a prediction-market interface face different regimes in different jurisdictions — distribution-layer compliance (licensing, geo-restrictions, responsible-use tooling) is the responsibility of each brand, and the Foundation's official frontend applies its own geo-policy. This paper is a technical specification, not an offer of securities, insurance, investment advice, or a solicitation to trade.

Appendix A — Convergence thresholds (per §1.3)

Product	e	σ	n for 95% seller certainty	n for 99%
Even-money binary (1.5% spread)	0.015	~1.0	~12,100	~24,100
Uniform 2× (0.99%)	0.0099	~1.0	~27,800	~55,400
Decision-path, optimal play (0.65% eff.)	0.0065	~1.15	~85,000	~170,000
Heavy tail (0.99%)	0.0099	~8–10	~10^7	— (portfolio-level only)

High-σ products converge at the portfolio level, not the individual level — which is precisely why Rules 1–3 bound them tightly.

Appendix B — Glossary

Position — a stake on a published distribution. **Spread** — the price of risk embedded in a payout table. **LP PnL** — gross revenue accruing to risk sellers. **Net LP PnL** — after rebates and incentives. **Volume rebate** — programmatic, tiered fee return on traded volume. **Locked credit** — a programmatic incentive with disclosed release conditions. **Netting** — offsetting opposing positions on one round before measuring vault exposure.

RAIN.ONE