BREAKING STRUCTURE PRESENTS

Capital *Protection.*

Position sizing, drawdown discipline, and the math that keeps small accounts alive.

WRITTEN BY
Kai Kaijuka
FOUNDER, ALSTRUM AI

// $1K · $10K · $25K · $100K · $500K //

ALSTRUM.AI
BREAKINGSTRUCTURE.SUBSTACK.COM

CONTENTS

Inside this issue

SECTION 01

The asymmetric ruin problem.

> *Small accounts blow up because the math is asymmetric. A 10% loss requires an 11% gain to recover. A 50% loss requires a 100% gain. Most retail traders size positions as if the math were symmetric. It is not.*

The single most useful chart in this entire series is the next one. It is the math, drawn out. Every dollar lost takes a larger return to recover than the dollar that lost it. The relationship is mechanical: a drawdown of *x* percent requires a gain of *x / (100 - x)* percent to return to even. The implication is not philosophical. It is operational. Capital preservation is the dominant constraint, not return optimization.

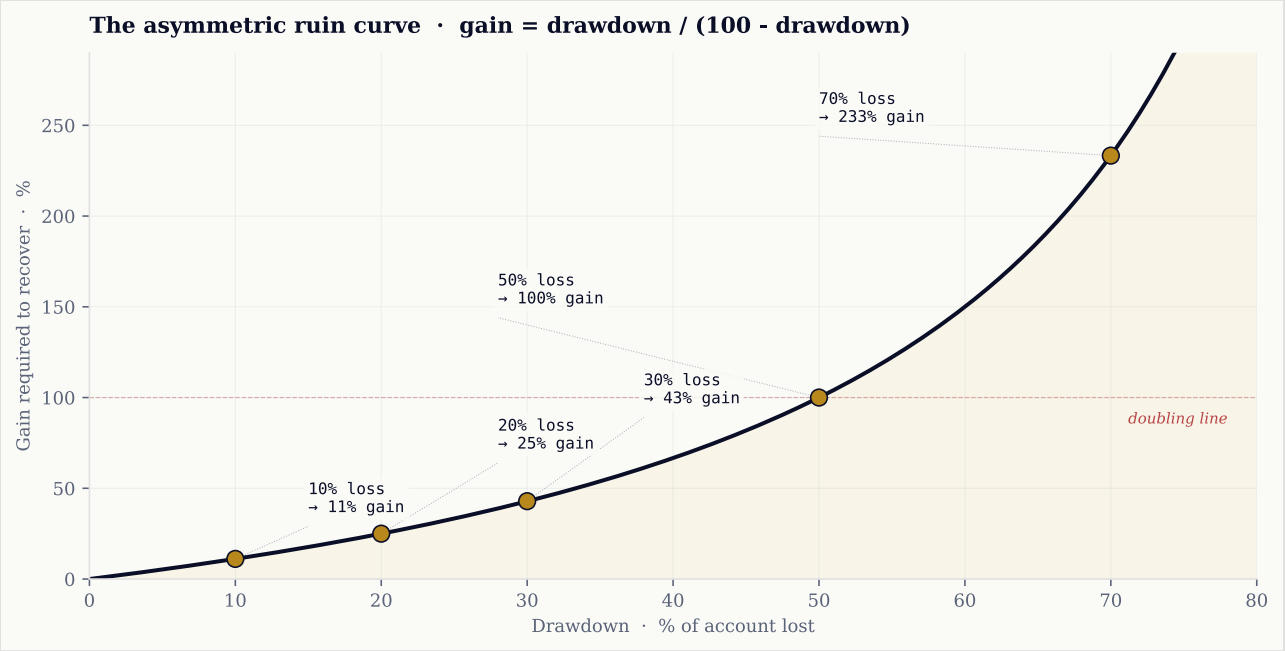

Figure 01 · The asymmetric ruin curve.

DRAWDOWN ALONG THE X-AXIS. REQUIRED RECOVERY GAIN ON THE Y-AXIS. THE FUNCTION IS GAIN = DRAWDOWN / (100 - DRAWDOWN). AT 50% LOSS, RECOVERY REQUIRES DOUBLING. AT 70% LOSS, RECOVERY REQUIRES MORE THAN TRIPLING. THE MATH IS MECHANICAL AND UNAVOIDABLE.

A small account starts at a disadvantage on this curve. A $5,000 account that loses $1,000 (a 20% drawdown) requires $1,250 in gains to recover. The same account that doubles in size to $10,000 only needs a 17% gain to recover from the same dollar loss. **The same drawdown is mechanically harder to recover from when capital is smaller.**

The framework that follows is built around this constraint. Every sizing rule, every drawdown band, every account-tier rule is downstream of the curve in Figure 01. Read it. Stay on the flat part.

SECTION 02

The 1% rule and why it is not enough.

> *"Risk 1% per trade" is the most-cited rule in retail trading. It is correct, necessary, and not sufficient. The reason: 1% on any single trade does not stop the account from losing 10% across a series of correlated trades.*

The rule, stated precisely, is: **risk no more than 1% of account NAV on any single trade, defined as the dollar loss from entry to stop, multiplied by position size.** The reason it works: at 1% per trade, a 10-trade losing streak (rare but realistic) costs 9.6% in compound terms. The account remains tradable. At 5% per trade, the same losing streak costs 40%, and the account is operationally dead even if technically alive.

The reason it is insufficient: **correlation**. Five Tier C single-name positions, each at 1% risk, are not five independent 1% bets. They share factor exposure (equity beta, sector beta, sometimes single-stock concentration). When the broader tape moves against you, they move together. The 5% aggregate exposure is more correlated than diversified.

The complete sizing rule has three components, applied in this order:

- **Per-trade cap:** 1% NAV at most, often 0.5-0.75% in higher-volatility tiers.
- **Per-correlated-cluster cap:** sum of risk across positions sharing the same factor (sector, theme, single direction) should not exceed 3-4% NAV.
- **Aggregate-book cap:** total risk-at-stop across all open positions should not exceed 8-10% NAV. Above that, the book is over-levered relative to its drawdown-recovery curve.

SECTION 03

Tier-based sizing.

> *Tiered sizing is the discipline that makes "1% per trade" actually work. Different instrument categories have different volatility, different liquidity, and different slippage profiles. One sizing rule across all of them is not risk control; it is averaging.*

The Three Layers framework (Issue 01) defined four tiers. Sizing maps to them:

- **Tier A · index and major instruments** (SPX, NDX, ES, NQ, major Treasuries). Highest liquidity, lowest single-event vol. Sizing typically 1-1.5% NAV per trade.
- **Tier B · sector ETFs and major futures** (XLE, XLF, CL, GC, ZN, 6E). Moderate vol. Sizing 0.75-1% NAV.
- **Tier C · single-name equities** (the convergence framework names from Issue 03). Higher single-stock vol, name-specific risk. Sizing 0.5-1.5% NAV depending on account size and conviction.
- **Tier D · crypto and exotic** (BTC, ETH, alts, exotic FX, low-liquidity futures). Highest vol, weekend gaps, exchange risk. Sizing 0.25-1% NAV maximum.

Position sizing matrix · account size × instrument tier

Cells show maximum risk-per-trade as % of NAV. Illustrative best-practice caps for retail traders.

		$1K-$10K	$10K-$25K	$25K-$100K	$100K-$500K	$500K+
Tier A	*index / major*	1.50%	1.00%	1.00%	0.75%	0.50%
Tier B	*sector ETF / fut*	1.00%	1.00%	0.75%	0.50%	0.50%
Tier C	*single-name eq*	0.50%	1.00%	1.00%	1.50%	2.00%
Tier D	*crypto / exotic*	*avoid*	0.25%	0.50%	0.75%	1.00%

Notes: smaller accounts cap Tier D (crypto/exotic) tighter due to liquidity and slippage risk.
Tier C single-name conviction-build allowed at higher %NAV for larger accounts due to multi-tranche scale-in capacity.
Cells are CAPS, not targets. Actual sizing conditions on conviction score and current drawdown band.

Figure 02 · Position sizing matrix.

RISK-PER-TRADE CAPS BY ACCOUNT SIZE AND TIER. SMALLER ACCOUNTS CAP TIER D TIGHTER DUE TO LIQUIDITY AND SLIPPAGE RISK. TIER C IS ALLOWED LARGER %NAV AT LARGER ACCOUNT SIZES BECAUSE MULTI-TRANCHE SCALE-IN IS OPERATIONALLY FEASIBLE. CELLS ARE CAPS, NOT TARGETS. EDUCATIONAL ILLUSTRATION. ALWAYS CONDITION ACTUAL SIZING ON CONVICTION SCORE AND CURRENT DRAWDOWN BAND.

The matrix is not a recommendation; it is a sketch of the design space. Real sizing for any particular trader depends on win rate, payoff ratio, and individual risk tolerance. The Kelly section that follows is the lens that turns these illustrative caps into something defensible for any specific edge.

SECTION 04

The Kelly fraction and quarter-Kelly.

> *The Kelly formula tells you the mathematically optimal fraction of capital to bet given a known edge. The output is almost always too large for any retail trader to actually risk. The practical adaptation is to size at 1/4 Kelly or smaller, and the reasons are concrete.*

The Kelly fraction for a binary win/loss bet is **k = (bp − q) / b**, where b is the payoff multiple on a win, p is the win probability, and q = 1 − p. The output is the fraction of NAV that maximizes long-run geometric growth assuming the edge is known exactly.

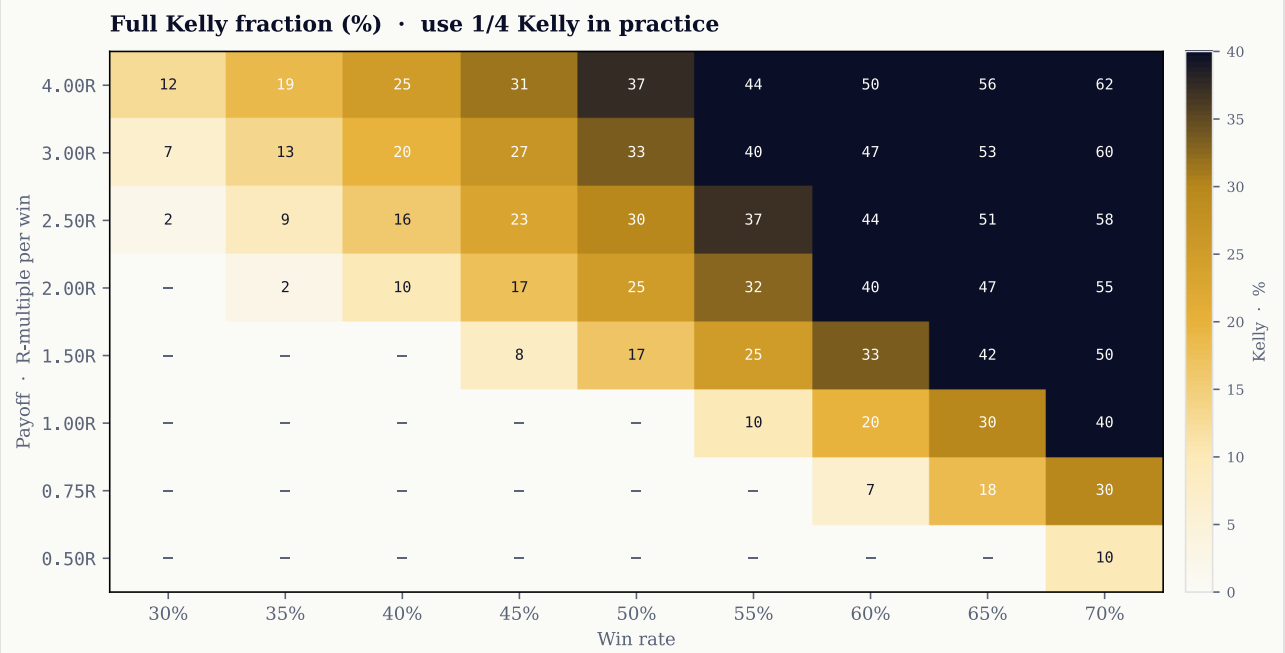

Figure 03 · Full Kelly fraction across win rate and payoff.

CELLS SHOW THE FULL-KELLY POSITION SIZE AS A PERCENT OF NAV. DASHES MARK ZONES WHERE THE EDGE IS NEGATIVE (DO NOT TRADE). THE DIAGONAL SHOWS WHERE MOST REALISTIC RETAIL EDGES LIVE, WITH KELLY OUTPUTS IN THE 5-25% RANGE THAT ARE TOO LARGE TO ACTUALLY RISK.

Why quarter-Kelly

Three reasons retail traders should use 1/4 Kelly or smaller, not full Kelly:

→ **Edge uncertainty.** Full Kelly assumes you know p and b exactly. You do not. Your historical win rate over 50 trades has substantial standard error. Sizing at full Kelly on an over-estimated edge produces catastrophic drawdowns even when the edge is real.

- **Psychological tolerance.** Full Kelly produces 30-50% drawdowns even with positive edge. Most traders cannot execute through those drawdowns without overriding the system. Quarter-Kelly produces 10-15% drawdowns instead.
- **Correlated trades.** Kelly assumes independent bets. Most real trades share factor exposure. Sizing at full Kelly per trade aggregates to far more risk than the formula assumes.

The operational rule: **take the cell value from Figure 03 and divide by 4.** A 50% win rate at 2R payoff shows 25% full-Kelly. Quarter-Kelly is 6.25%. Compare that against the per-trade and per-tier caps in Section 03. Use the smaller number.

SECTION 05

Drawdown bands.

> *Sizing rules at trade entry are necessary but reactive. Drawdown bands are proactive: as the account loses capital, the rules tighten automatically, before judgment has a chance to deteriorate.*

The three trip-wires are unambiguous:

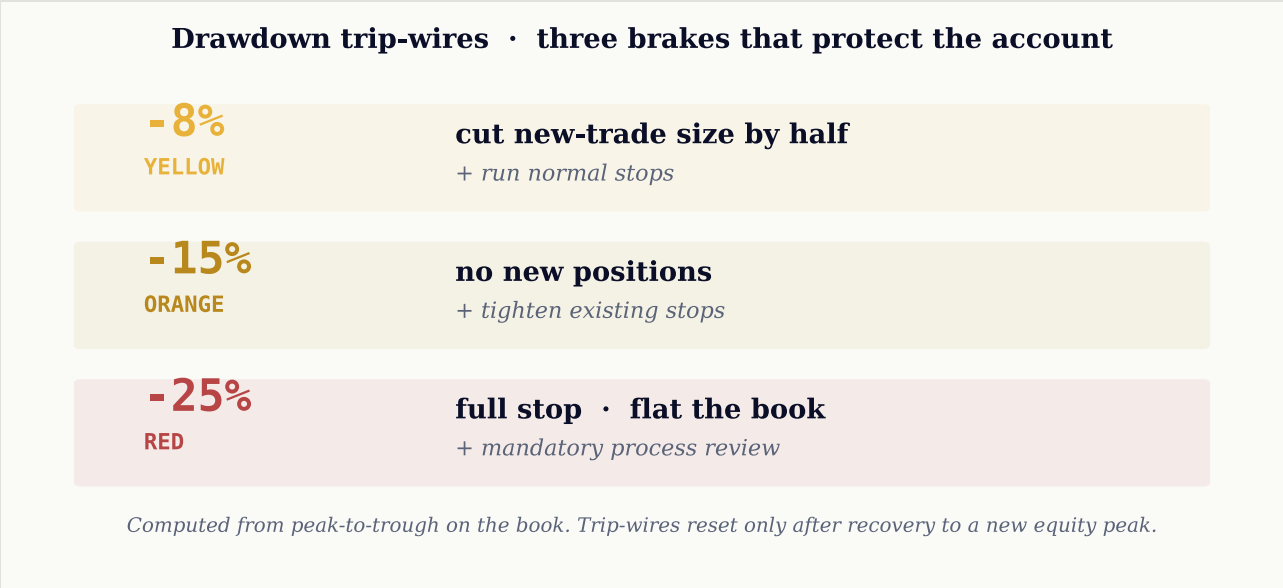

Figure 04 · Drawdown trip-wires.

THREE THRESHOLDS, THREE ACTIONS. TRIP-WIRES RESET ONLY AFTER RECOVERY TO A NEW EQUITY PEAK, NOT AFTER PARTIAL RECOVERY. THE MECHANISM IS AUTOMATIC TO REMOVE JUDGMENT UNDER STRESS.

The trip-wires are computed peak-to-trough on the equity curve, not from the original NAV at start of year. The reason: a 12% gain followed by a 9% drawdown is a 9% drawdown from peak, not a small gain. The yellow trip-wire fires. Sizing on new trades is halved until recovery.

The mechanism is automatic. Once the band is hit, the rule fires regardless of what the trader believes the next trade looks like. This is the entire point: in a drawdown, the trader's read of "the next trade looks great" is precisely the read that needs to be discounted, because by the structure of how drawdowns happen, that read is often where the drawdown deepens. Trip-wires defend against the trader's own optimism at the worst time.

SECTION 06

The $1K-$10K account.

> *The smallest account tier is also the most fragile. Liquidity, slippage, and per-trade fees consume a larger fraction of NAV than at any other size. The discipline that protects this account is different from the discipline at scale.*

Operational rules for the $1K-$10K tier:

- **Cash equity only.** No options, no futures, no crypto leverage. The bid-ask spreads and contract sizes do not work at this scale.
- **Three positions maximum.** Concentration is the constraint, not diversification.
- **Tier A and Tier B only.** Single-name (Tier C) and exotic (Tier D) instruments cap too tight to be meaningful at $1K-$10K.
- **1.5% risk per trade absolute cap.** Higher than larger accounts because absolute dollar risk is small. At $5,000 NAV, 1.5% is $75. Even that is barely worth the slippage on most retail brokers.
- **No leverage.** Margin at this scale magnifies operational risk faster than it grows returns.

The honest read: **this tier exists to learn process, not to compound returns.** The math says that even a 20% annual return on $5,000 is $1,000 before taxes, well under the value of a part-time job. The right framing is treating the small account as a paid education in process discipline, where the paid part is the small absolute dollars at risk. The framework, the discipline, and the journal mechanics are the actual deliverable from this tier. The capital is the medium of instruction.

SECTION 07

The $10K-$25K account.

> *The first tier where the framework becomes operationally meaningful. Tier C single-name positions start to be sizable enough to matter, and the account is large enough to absorb the realistic 5-10% drawdowns that come with single-name conviction trades.*

Operational rules for the $10K-$25K tier:

- **Cash equity + occasional debit verticals.** Option spreads at this size are tractable for defined-risk thesis trades. Naked options remain inappropriate.
- **Five positions maximum.** Enough breadth to absorb individual misses, few enough to monitor in retail-available time.
- **All four tiers permitted** at appropriate sizes per the matrix in Section 03.
- **1.0% risk per trade default.** Quarter-Kelly conditioned. 0.5% on Tier C convergence builds (multi-tranche scale-in).
- **No margin.** Cash account sufficient. Pattern Day Trader rule does not apply until $25K, which is its own consideration in the next tier.

The $10K-$25K tier is the most operationally important for new traders running the convergence framework. It is large enough that a 5/6 Tier C name at 1% NAV ($100-$250 risk-at-stop) produces meaningful dollar outcomes that test the discipline of holding through a 6-week thesis window. It is small enough that mistakes are recoverable. Most professionals who graduated from retail to institutional sizes did most of their learning in this tier.

SECTION 08

The $25K-$100K account · PDT threshold.

> *The $25,000 threshold is the regulatory inflection point in the US. Above it, a margin account holder can be flagged as a Pattern Day Trader, which removes the four-day-trade-per-week limitation. Below it, day-trading capacity is limited. For most convergence-framework traders the threshold matters less than is commonly assumed.*

Why the PDT threshold matters less than retail forums suggest: the convergence framework is a multi-day-to-multi-week holding pattern, not a day-trading pattern. The trades that survive Issue 03's six-source rubric are positioning trades held for weeks, not intraday names. The day-trading capacity unlocked at $25K NAV is operationally irrelevant to traders running the framework. **What matters at this tier is broader instrument access, not faster turnover.**

Operational rules for the $25K-$100K tier:

- **Options-eligible accounts** become viable for hedging applications. Defined-risk option structures (debit spreads, calendar spreads, ratio spreads) can be sized to add 0.25-0.5% NAV exposure with capped downside.
- **Eight positions typical, twelve maximum.** The book starts to be a book rather than a list of trades.
- **1.0% risk per trade, with conviction overrides allowed** on 5/6 or 6/6 convergence names up to 1.5% NAV per name.
- **Pattern Day Trader status optional.** Most framework traders should opt out of PDT to avoid the temptation to overtrade.
- **Modest margin usage** permitted for overnight positions, capped at 1.25x cash equity.

SECTION 09

The $100K+ account · futures and hedging unlock.

> *At $100,000+ NAV, the entire instrument universe becomes operationally accessible. Futures contracts are sizable. Options structures can hedge entire books, not just single positions. Cross-asset hedging becomes meaningful. The framework operates at full capacity.*

The single biggest sizing shift at this tier is the move from per-trade risk to **book-level risk**. Issue 07 covers this in detail. The summary: at $100K+, the relevant question is not "what should I risk on this trade" but "what is the aggregate factor exposure my book carries, and where is the structural hedge that does not show up in any single trade." The trades-view of the book gives way to the factor-view of the book.

Operational rules for the $100K+ tier:

- **Full multi-asset access.** Futures, options, ETF complexes, FX, crypto at unconstrained sizes within tier caps.
- **Book risk capped at 8-10% NAV** aggregate risk-at-stop across all open positions.
- **Factor-level hedging budget of 2-4% NAV annually** dedicated to long volatility, tail protection, and cross-asset offsets per Issue 07.
- **0.5-1.0% per trade default**, with conviction overrides to 1.5-2% on high-conviction Tier C names.
- **Drawdown bands tightened to 6%/12%/20%** versus the 8%/15%/25% default, because the absolute dollar value of each percent is larger and the recovery curve steepens.

SECTION 10

Common mistakes.

> *Five patterns cause most account blow-ups. None of them are about being wrong on a trade. All of them are about sizing.*

1. Martingale (averaging down without thesis intact).

Adding to a losing position because the loss "feels temporary." The framework reads this as a thesis-failure event masked as a re-entry. The discipline: if the original thesis is broken (Section 11 of the failure-mode review), exit, do not add. Adding to losers is the single fastest path to ruin for retail accounts.

2. Revenge sizing.

Increasing per-trade size after a loss to "make it back faster." Mathematically the worst possible response: the drawdown band logic in Section 05 says exactly the opposite. After a 3% drawdown, the next trade should be smaller, not larger. Revenge sizing is what turns a 3% drawdown into a 15% drawdown.

3. Undersizing the winners.

The mirror failure. Convergence names that pass the 4+/6 threshold are sized identically to 2/6 watch names. The framework's edge is in differentiated conviction, but if sizing is flat across conviction scores, the edge does not compound. High-conviction names should be sized 1.5-2x baseline within tier caps.

4. Currency-of-record errors.

Trading a USD-denominated instrument while measuring P&L in a different currency, without explicit FX hedging. The FX move can be the entire P&L. Covered more fully in Issue 07.

5. Ignoring the journal.

The journal is what makes the sizing rules learn. Without it, the same five mistakes get made in different combinations every quarter. Section 11 covers the journal mechanics. The discipline is short: every closed trade gets logged with R-result, thesis-held flag, and one-line reason.

SECTION 11

The journal that enforces the rules.

> *Sizing rules without a journal are aspirations. The journal is the feedback loop that makes the rules executable over years rather than weeks.*

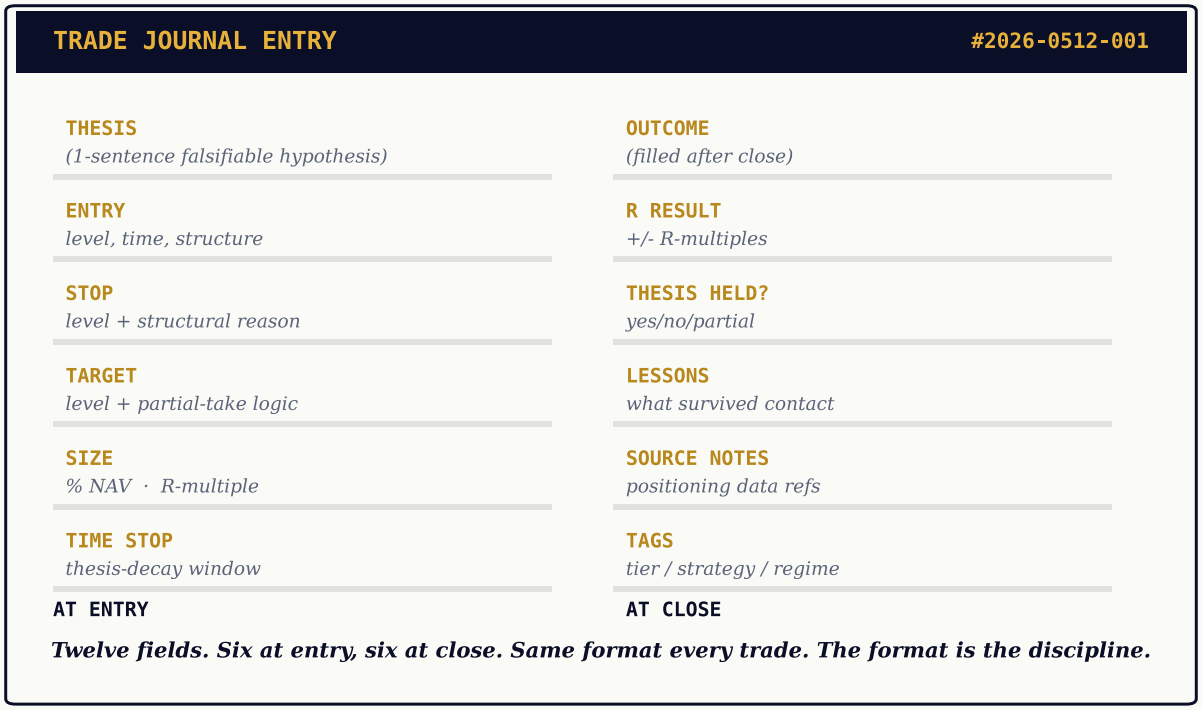

Figure 05 · The journal template.

TWELVE FIELDS PER TRADE. SIX AT ENTRY (THE PLAN). SIX AT CLOSE (THE OUTCOME). SAME FORMAT EVERY TRADE, NO EXCEPTIONS. PATTERN DETECTION BECOMES POSSIBLE ONLY WHEN THE FORMAT IS INVARIANT.

The journal is the mechanism that connects sizing rules at trade time to sizing rules at evaluation time. Without it, "I should size smaller on Tier D crypto" is a feeling. With it, "my last 30 Tier D crypto trades had a -0.4R average versus my +1.2R Tier C average" is a number that drives the next quarter's sizing matrix.

The discipline is one entry per closed trade. Two-line minimum. R-result, thesis-held flag, one-line lesson. Aggregated quarterly into a sizing review that updates the tier matrix from Section 03 based on observed performance. **The sizing matrix is not fixed. It evolves from journal data.**

SECTION 12

Resources and what is next.

FREE TOOLS

- **Position sizing calculator (any).** Most retail brokers and platforms include one. The calculator does the dollar arithmetic; the discipline is in the inputs.
- **Spreadsheet journal.** Google Sheets or Excel, formatted per the template in Figure 05. Pen and paper if you must.
- **R-multiple tracking.** Track the R-result on every trade. R = (exit - entry) / (entry - stop). This is the single most useful per-trade metric.

PAID TOOLS

- **Trade journal apps.** TraderSync, Edgewonk, TradesViz. The features that matter: tag-based filtering, R-multiple tracking, drawdown band visualization.
- **Risk management dashboard.** Most institutional traders use bespoke. Retail can approximate via Notion, Airtable, or a structured Google Sheet.

The Stryk angle on this issue.

Stryk's order ticket enforces per-trade and per-tier sizing caps before submit. Drawdown brakes engage automatically when the account hits the 8% / 15% / 25% thresholds, halving size or blocking new positions per Section 05. The journal auto-populates from executed trades with R-multiple and thesis-held fields. **The discipline is in the platform, not in the trader's willpower at the worst moments.**

Breaking Structure subscribers go to the front of the early access line by default. Until Stryk ships, the spreadsheet journal in Figure 05 is the manual version. *Build the habit now. The workflow ports.*

`breakingstructure.substack.com · alstrum.ai · @kjukas`

What is next in the series

Issue 06 covers **The Six-Field Plan**: the trade format that takes a sized position and turns it into a written, time-stamped, exit-defined plan. Sizing tells you how much risk. The six-field plan tells you which trade. The two together are the operating system.